



Operationalizing Autonomous AI:

Architecting the Agentic
Enterprise on a Converged
Data Foundation



The market mandate centers on building autonomous agents capable of executing complex business processes. To survive this transition and achieve meaningful return on investment, organizations must abandon bolted-on AI applications and embrace a natively governed, multi-cloud data foundation that provides model optionality and secure, closed-loop execution.


The AI Reality Check: Mandates vs. Infrastructure Limitations

Enterprise leaders demand more from their artificial intelligence investments than passive insights and document summarization. Over the past two years, organizations poured capital into large language models (LLMs) to automate knowledge work, yet these systems consistently fail to execute meaningful business workflows reliably. Decision-makers recognize that the models themselves are rapidly commoditizing and converging in capability. The durable strategic value of these tools – whether frontier-scale or highly specialized smaller models – lives entirely in the underlying data foundation and the semantic context grounding them.

C-suite directives have decisively shifted away from experimental applications toward autonomous operational workflows. The enterprise requires an "Agentic Control Plane," a framework where the rigid delineations between application stack elements collapse into a single continuum of execution and control. This environment allows human professionals to assume supervisory roles, acting as human-in-the-loop orchestrators who validate, steer, and audit fleets of autonomous software.

According to the Futurum Research IH 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report, 44.5% of enterprises plan to increase their semantic layer spend over the next 24 months to support this exact control plane (see Figure 1). This investment highlights a vital market reality: a universal semantic framework is a non-negotiable prerequisite to grounding large language models and mitigating hallucinations. Autonomous agents cannot make mission-critical decisions if the underlying data is scattered across disconnected clouds, rigid legacy systems, and unverified schemas.

Figure 1: Expected Top Trends in Data Management & Analytics



Enterprise AI maturity is accelerating. Today's IT decision-makers are moving beyond foundational generative AI investments, looking instead toward autonomous, agentic frameworks capable of orchestrating complex analytics and workflows.

Q: "Which of the following do you expect to be a top trend in data management and analytics between 2026 and 2029?"
Base: IH 2026 (n=818)

Source: Futurum Research, IH 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report.

This ambition clashes directly with a harsh operational reality. While the vision for the agentic enterprise is clear, legacy infrastructure remains wholly unsuited for the task. Historically, data platforms functioned as passive repositories where data landed before being queried by human analysts. Attempting to build autonomous agents on top of these passive, fragmented systems guarantees high latency, brittle integrations, and unpredictable outcomes. Infrastructure must evolve to execute AI workloads securely in situ, collapsing the distance between where data resides and where models reason.





The Hidden Costs of Bolted-On AI Architectures


Treating AI as an application-tier bolt-on creates debilitating technical debt. The default architectural approach of the past two years relied on extracting data from secure warehouses, transforming it, and shuttling it into specialized, external vector databases to support retrieval-augmented generation (RAG) applications. This fragmented strategy forces data teams to maintain complex integration pipelines. Consequently, data duplication runs rampant across multiple environments, expanding the enterprise security surface area.

This friction exacts a heavy operational tax. As data moves across the network to external agent frameworks, it sheds the deep role-based access controls and governance policies established at the source. Furthermore, this approach severely limits the functional utility of AI. Most agentic systems operating in this application-tier model suffer from strict read-only limitations. They analyze historical data and provide recommendations, but they lack the secure integration necessary to execute a transaction or update a record.

Organizations pinpoint this inability to write back to systems of record as a severe barrier to operational automation. According to the Futurum Research IH 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report, 24.6% of organizations identify the inability to write back to systems of record as their primary architectural bottleneck. This limitation is felt most acutely in highly transactional, regulated verticals. Within the bottleneck cohort, Financial Services professionals identify it as their primary roadblock, recognizing that without the ability to write back securely, agents remain advisory novelties rather than true operational engines.

Allowing autonomous software to interact with sensitive corporate data presents an existential threat if improperly managed. According to the same Futurum Research primary data, accuracy and hallucinations remain the top reservation regarding generative AI replacing traditional business intelligence at 24.9% (see Figure 2). This substantial trust deficit underscores why validating the intent and permissions of an autonomous agent requires rigorous, multi-level guardrails that application-tier AI cannot provide.

Figure 2: Primary Reservation Regarding Generative AI in Analytics



If an executive cannot trust a model to calculate gross margin correctly on a read path, they will never authorize an agent to execute a transaction on a write path. Bridging this gap requires deterministic grounding layers.

Q: "What is your primary reservation, if any, regarding the use of Generative AI/LLMs to replace traditional BI Dashboards?"
Base: IH 2026 (n=818)

Futurum Research, IH 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report.

These challenges dictate the success or failure of digital transformation initiatives. A Senior Director of Cloud Operations at a mid-sized technology company – tasked with managing a highly complex hybrid multi-cloud setup spanning AWS, GCP, and Azure – realized that layering AI over decades of accumulated technical debt would invite disaster. The company required a unified governance foundation capable of supporting its entire infrastructure before launching revenue-generating AI products. To achieve this goal, the company turned to Snowflake.


"Zero – the baseline – was compatibility with our hybrid multi-cloud stack... Initially, we wanted a common governed layer for all our data. Snowflake supports all three clouds... The ability to run containers, code, and native AI features made us realize the system could be used much further than our initial project. Running model serving closer to the customer data provides a governance layer and removes latency."

By recognizing the limitations of the application-tier compromise early, enterprise leaders can avoid the integration trap and build architectures designed specifically for governed, in-situ AI processing.

Building a Native, Unified Agentic Data Core

To resolve these architectural constraints, Snowflake fundamentally rearchitected its data platform to support an agentic control plane natively. Rather than forcing customers to extract data into external AI environments, the company brings machine learning operations directly to the data. This converged foundation eliminates the severe security vulnerabilities inherent in data movement while enforcing unified governance across a multi-cloud reality. The architecture rests on key pillars that collapse silos, prioritize in-situ processing, and introduce crucial cost management capabilities (see Figure 3).

Figure 3: Market Drivers of Architectural Transformation Map



This architectural transformation map connects the foundational pillars of the Agentic Control Plane directly to how companies seek to modernize storage portability and enable closed-loop execution.

Q: "What is your primary reservation, if any, regarding the use of Generative AI/LLMs to replace traditional BI Dashboards?"
Base: 1H 2026 (n=818)

Futurum Research, 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report.

Unified Business Semantics and Governance

Semantic layers represent a major market driver of investment. When different departments define "revenue" or "active customer" using conflicting logic, AI agents inevitably hallucinate or return contradictory insights. Snowflake attacks this directly through Polaris Catalog and Horizon Context, a foundational semantic and governance layer. Horizon Context defines and governs the truth, allowing organizations to establish a core metadata specification. By defining business logic centrally and openly, human analysts and AI agents operate on a single, deterministic set of facts.

This approach anchors model outputs to a rigidly governed, centralized semantic catalog. According to the Futurum Research 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report, only 18.0% of organizations have managed to deploy advanced Graph RAG patterns, with 59.4% remaining limited to basic or hybrid search implementations. A unified semantic framework acts as the direct technical catalyst for advanced context engineering, advancing organizations from basic semantic search to robust, relational reasoning.

Preserving Model Optionality and Controlling Cost

The financial reality of generative AI dictates that enterprises cannot afford to run every workload on a single, massive proprietary model. Organizations demand dynamic model optionality – the ability to route tasks to the most cost-effective and performant model for the specific job, whether that is an expansive foundational LLM or a highly efficient, task-specific open-source model.


A converged data foundation, as with Snowflake, fundamentally enables this flexibility through Cortex AI. By managing models adjacent to the data, Snowflake allows engineering teams to swap underlying LLMs without re-architecting complex data pipelines or moving data to new environments. This in-situ approach intrinsically supports granular cost monitoring and management. Because model execution happens directly within the governed data boundary, IT leaders can accurately track compute consumption per agentic workflow, ensuring AI deployments yield a verifiable return on investment rather than generating unpredictable, spiraling cloud costs.

Open Interoperability and Closed-Loop Execution

A modern agentic foundation must honor data gravity. Moving petabytes of data to accommodate different compute engines is financially and operationally unsustainable. Snowflake embraces open interoperability through comprehensive support for Apache Iceberg v3 and the Polaris catalog. This capability allows external compute engines to securely read and write back to Snowflake-managed tables through open standards with zero data duplication, neutralizing the lock-in risks associated with legacy data warehousing models.

The true differentiator of a converged foundation lies in establishing closed-loop execution. By running AI workloads natively within the platform, Snowflake creates an environment where agents move beyond read-only analysis to take authorized action. Security policies, including dynamic data masking and granular role-based access controls, strictly govern this autonomous software. If an employee lacks the clearance to view a specific financial record, the AI agent acting on their behalf is similarly blocked. This in-situ processing ensures a secure, auditable trail for every agentic action.

A Managing Director of Cloud and Data Integration at a tier-one global financial institution requires rigorous, multi-level guardrails before allowing AI to operate on highly sensitive unstructured data. He chose a converged, native foundation to reduce the immense friction of implementing these controls across fragmented systems:



"Getting from a whiteboard to production takes 6 to 7 months. It has nothing to do with the vendors. It is driven by our own risk, compliance, testing, and lengthy pilots to ensure we face no reputational risk... We are not in the business of developing our own LLMs from scratch or building AI plumbing. We want to leverage vendor capabilities so we can focus on business value."


By shifting the operational burden of AI plumbing to the underlying platform, organizations can redirect their engineering resources toward building specialized agents that directly address business outcomes.



Democratizing the Governed Control Plane

Snowflake's architectural approach ensures advanced agentic power remains accessible across the organization within a strictly governed boundary. By layering Cortex Sense – an automated ontology and skills mapping engine – on its Horizon Context layer, Snowflake guarantees that all users interact with an AI system that deeply understands the specific nuances of the enterprise data estate.

This democratization is fully realized in Snowflake CoWork. CoWork serves as a highly capable personal work agent, bypassing the limitations of standard chatbots. It conducts deep research across the enterprise's structured and unstructured data, parsing historical contracts, generating interactive analytical artifacts, and providing verified answers with direct citations. Crucially, CoWork democratizes this power securely. It strictly enforces existing role-based access policies. If a marketing manager lacks authorization to see specific financial records, their personal CoWork agent faces identical restrictions. This capability allows enterprises to safely unleash autonomous capabilities to non-technical teams, eliminating their dependence on traditional IT reporting without risking data exposure.



"Sales and marketing no longer have lengthy lead times waiting on IT to generate reports. They use natural language to query SQL databases themselves."


– Managing Director of Cloud and Data Integration, Global Financial Institution

On the engineering side, Snowflake CoCo operates as an agentic coding partner designed to accelerate development. Developers use CoCo to rapidly build multi-agent workflows, optimize data pipelines, and deploy specific AI profiles tailored to discrete tasks. CoCo generates self-healing code that spots and corrects its own errors, drastically compressing migration and deployment timelines.

Extending Governed Access to Third-Party Systems

Snowflake recognizes that enterprise context does not live exclusively within its own databases. Operational workflows require agents to interact with external systems of record, from CRM platforms to ERP environments. To prevent the creation of a new walled garden, Snowflake embedded native support for the Model Context Protocol (MCP), allowing CoWork and CoCo to pull real-time context from external applications including Salesforce and Microsoft 365 (see Figure 4).

To further solidify this capability, Snowflake's acquisition of Natoma (May 2026), brings Natoma's enterprise MCP platform capabilities directly into the Snowflake platform for secure, enterprise-grade agent interoperability with Cortex AI Gateway. Cortex AI Gateway provides a secure conduit for agents to execute closed-loop business processes across external environments. Instead of relying on brittle API integrations that strip away access controls, Cortex AI Gateway ensures that an agent's actions in an external system adhere to the exact same governance policies that secure the core data warehouse.



Without a converged semantic and storage layer, the majority of enterprises remain stuck in basic or hybrid search patterns, lacking the advanced Graph RAG or integrated working memory required for autonomous closed-loop execution.

Figure 4: Current Enterprise RAG and Search Implementations

Q: "What is your primary reservation, if any, regarding the use of Generative AI/LLMs to replace traditional BI Dashboards?"
Base: 1H 2026 (n=818)

Futurum Research, 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report.



Conclusion and Recommendations

Achieving sustained success with agentic AI depends heavily on the quality, governance, and interoperability of the underlying data foundation, rather than the raw parameter count of the newest language model. The era of treating AI as an experimental science project is over. To move past pilot purgatory and operationalize autonomous software securely, enterprise decision-makers must take decisive architectural action. We recommend the following pragmatic steps:



Audit The Semantic Foundation: Establish a governed, centralized semantic layer utilizing open standards. Ensure human executives and AI agents work from the exact same mathematical truth before authorizing any automated workflows.



Demand Model Optionality: Architect your environment to decouple the data from the specific AI model. Deploy platforms that allow you to swap LLMs dynamically to optimize for cost, latency, and performance without undertaking massive data migrations.




Eliminate Operational Friction: By running models directly where data lives, companies can enforce existing role-based access controls and drastically reduce their exposure to security vulnerabilities.



Prioritize Open Standards and Governed External Action: Safeguard your data gravity by mandating infrastructure that natively supports open cataloging and table formats such as Apache Iceberg. Furthermore, ensure that any agentic action taken in third-party systems via tools such as Cortex AI Gateway adheres **strictly to your internal governance policies**.



Align AI Strictly to Business Pain Points: Define the measurable operational challenge first – whether that involves accelerating supply chain reconciliation or automating compliance audits – and then deploy agentic solutions targeted explicitly at solving that workflow.



"Keep the technology aside for a minute. Identify your pain points first. What is your North Star? Once you know that, look for a technology enabler... Don't be a hammer looking for a nail."

– Managing Director of Cloud and Data Integration, Global Financial Institution

The emergence of the agentic enterprise represents the rigorous evolution of software automation, demanding excellent engineering over speculative hype. Organizations winning in this era eschew the distraction of conversational novelty. Instead, they focus relentlessly on building a scalable, secure, and semantically grounded data infrastructure capable of letting autonomous systems act with absolute trust, measured cost efficiency, and deterministic precision.

Important Information About This Report

AUTHORS

Brad Shimmin

Vice President & Practice Lead, Data Intelligence,
Analytics, & Infrastructure | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and
The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts,
but must be cited in context, displaying author's name,
author's title, and "The Futurum Group." Non-press and
non-analysts must receive prior written permission by The
Futurum Group for any citations.

LICENSING

This document, including any supporting materials, is
owned by The Futurum Group. This publication may not be
reproduced, distributed, or shared in any form without the
prior written permission of The Futurum Group.

DISCLOSURES

The Futurum Group provides research, analysis, advising,
and consulting to many high-tech companies, including those
mentioned in this paper. No employees at the firm hold any
equity positions with any companies cited in this document.



ABOUT SNOWFLAKE

Snowflake is a cloud-based data platform that separates storage and compute, letting organizations scale each independently while querying data across AWS, Azure, and Google Cloud. Originally built for data warehousing, it has expanded into a broader data cloud supporting data engineering, data lakes, and data sharing. More recently, Snowflake has pushed into agentic AI, introducing tools like Cortex Agents and Snowflake CoWork that let AI agents reason over enterprise data and execute multi-step tasks with minimal human prompting. This positions Snowflake as a notable case study in how data infrastructure is evolving to support autonomous, AI-driven enterprise workflows.

Futurum

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

Futurum

CONTACT INFORMATION: The Futurum Group LLC | [futurumgroup.com](https://www.futurumgroup.com) | (833) 722-5337