Solution Brief

The Denodo Active Context Layer for Agentic AI

Trusted enterprise context for scalable, governed, and effective AI agents

Organizations are rapidly shifting from generative AI (GenAI) experimentation toward agentic AI, where AI systems can reason, dynamically retrieve information, orchestrate workflows, and take action across the enterprise.

Traditional conversational AI focuses primarily on responding to user prompts. Agentic AI goes further by operating across business processes, interacting with enterprise systems in real time, coordinating workflows, triggering actions, and adapting dynamically as business conditions evolve.

This shift is changing how organizations think about enterprise AI architecture.

The success of agentic AI is no longer determined primarily by the model itself. Increasingly, success depends on whether AI agents can securely access trusted enterprise context across distributed systems while maintaining governance, operational awareness, and business consistency.

This is where many organizations begin to struggle.

66% of organizations say AI data must be real-time or less than one minute old to be trustworthy[1]

Enterprise data remains spread across cloud platforms, lakehouses, SaaS applications, operational systems, APIs, and on-premises environments. Governance policies may be defined centrally, but enforcing them consistently across systems with different platforms, owners, and control models is far more difficult in practice. Business meaning is often interpreted differently across teams and platforms, while operational data becomes stale as it moves through replicated pipelines. As AI agents become more autonomous, these challenges become harder to control.

Many vendors now position semantic layers, AI fabrics, metadata platforms, and context layers as the architectural answer for agentic AI. Most approaches, however, solve only part of the problem. Some organize metadata and business definitions well. Others provide strong AI capabilities, but only inside a single cloud ecosystem.

Denodo takes a different approach.

Denodo delivers a truly federated active context layer for agentic AI that unifies semantics, governance, provenance, and live enterprise access across the organization. Instead of forcing organizations to centralize everything into one platform, the Denodo Platform enables AI agents to securely access trusted context across hybrid, multi-cloud, SaaS, and on-premises environments.

The result is a trusted operational foundation for scalable agentic AI. In this solution brief, we'll review the challenges of implementing agentic AI and why existing approaches have so far been unable to address them. We then introduce the Denodo Active Context Layer solution and explain how it works to enable this critical foundation.

[1] Denodo, The AI Trust Gap Report

The Challenge: A New Agentic Reality

The first wave of enterprise AI focused largely on conversational interfaces. The next wave is centered on intelligent agents that can operate across business processes with increasing levels of autonomy.

These systems are no longer limited to answering questions. AI agents are beginning to retrieve information from multiple systems, evaluate changing business conditions, trigger workflows, interact with APIs, and coordinate multistep activities across the organization. As a result, enterprise AI is evolving from passive assistance toward the orchestration and automation of business processes with both humans and AI agents working together.

That shift introduces a very different set of architectural requirements.

AI agents cannot operate effectively on fragmented, stale, or inconsistently governed data. When agents operate on delayed or outdated information, they risk making decisions based on conditions that no longer reflect the current state of the business. In fast-moving operational environments, even small delays in data synchronization can lead to inaccurate recommendations, workflow errors, or poor business outcomes.

AI agents also become less efficient when forced to reason across disconnected systems, fragmented APIs, duplicated retrieval pipelines, and inconsistent business definitions. Instead of operating with clear contextual grounding, agents spend additional time and compute attempting to reconcile missing relationships and conflicting information.

> Flooding AI models with excessive or irrelevant information can reduce accuracy, increase token consumption, and drive up infrastructure costs.

As reasoning models perform more multistep planning and orchestration, retrieval efficiency becomes increasingly important. The challenge is not simply retrieving information from multiple systems. AI agents must also integrate and reason across that information in real time, often without an execution engine optimized for distributed data integration. As a result, agents can spend significant time and compute attempting to reconcile fragmented data within the model context itself, increasing latency and operational cost.

At the same time, more data does not necessarily produce better outcomes. When AI systems are flooded with excessive or irrelevant information, accuracy can decline as models struggle to identify the most relevant context for the task at hand. Without a trusted context layer guiding retrieval, integration, and reasoning, AI systems often generate unnecessary processing overhead, increased latency, higher token consumption, and rising infrastructure costs while still producing less reliable results.

Organizations quickly begin to encounter familiar challenges. AI responses become inconsistent because definitions vary across systems. Governance policies fragment as agents interact across multiple platforms. Retrieval architectures become more complex and expensive to scale. Over time, operational trust begins to erode because AI systems lack awareness of current business conditions.

The challenge is no longer simply generating AI responses. The challenge is enabling AI agents to operate safely, intelligently, and efficiently across the enterprise.

Why Existing Approaches to Supporting Agentic AI Have Been Falling Short

Many organizations initially attempt to support agentic AI using existing lakehouse, metadata, or semantic layer architectures. Each of these approaches delivers important value. However, none independently provides the full operational foundation required for enterprise-scale AI agents.

Lakehouse and cloud platform vendors such as Databricks, Snowflake, Microsoft Fabric, and BigQuery offer strong AI tooling and semantic consistency within their own ecosystems. These platforms work well when most enterprise data has already been centralized into a single environment.

[2] Denodo, The AI Trust Gap Report

The problem is that most enterprises are not structured that way. Critical enterprise context often remains distributed across operational systems, line-of-business applications, SaaS platforms, APIs, and multiple cloud environments. Enterprise AI initiatives now draw on an average of more than 400 original data sources, reinforcing how distributed and fragmented modern enterprise environments have become.[2]

As organizations expand AI initiatives, context outside the platform boundary becomes harder to govern consistently. Data movement increases, replication pipelines multiply, and maintaining real-time operational awareness becomes more difficult. Over time, this can increase both the cost and complexity of scaling AI initiatives across distributed environments while introducing latency that leaves AI agents operating on stale or incomplete business context.

Metadata and catalog vendors such as Atlan, Collibra, and Alation address a different part of the problem. These platforms help organizations organize business definitions, lineage, governance information, and semantic relationships. They improve discoverability and help establish a shared understanding of enterprise data.

However, metadata alone does not operationalize enterprise context for AI execution. AI agents still require live, governed access to operational and analytical enterprise data in order to reason and act effectively.

Most metadata-centric approaches do not sit in the runtime execution path, deliver live operational data, or enforce governance dynamically as AI agents interact with systems. Governance becomes informative rather than enforceable, especially once AI workflows begin operating across multiple environments. This can lead to fragmented retrieval architectures and inconsistent policy enforcement across AI systems.

Semantic layer vendors such as dbt, Cube, AtScale, and Looker solve another important challenge by standardizing metrics and reporting consistency for analytics workloads. These approaches improve BI governance and KPI alignment, but they were primarily designed for reporting and dashboard consistency rather than supporting AI agents operating across distributed enterprise systems. As AI agents become more operational and autonomous, organizations require broader contextual coordination across live enterprise environments.

Emerging AI integration frameworks and protocols help standardize connectivity between AI agents and enterprise systems, but connectivity alone does not create shared enterprise understanding. AI agents may still operate against fragmented local context where business definitions, identifiers, governance rules, and operational meaning vary between systems. Enterprise-scale agentic AI requires a broader contextual coordination layer that establishes consistent enterprise meaning across distributed environments.

Each of the technologies mentioned above addresses part of the challenge. However, enterprise-scale agentic AI requires a broader foundation that can unify trusted operational and analytical context, apply consistent governance, and provide live enterprise access across distributed environments.

The Denodo Difference: Introducing the Active Context Layer

Denodo approaches agentic AI from a fundamentally different perspective. Rather than focusing only on metadata, analytics semantics, or a single cloud platform, Denodo operationalizes trusted enterprise context across the broader enterprise ecosystem.

Denodo combines semantic abstraction, federated enterprise access, runtime governance enforcement, provenance, and live operational awareness within a unified logical data management architecture. This enables organizations to simplify AI retrieval architectures while improving the relevance, consistency, and efficiency of how AI agents access enterprise information. Rather than forcing models to reason across fragmented systems and duplicated pipelines, Denodo helps AI agents access the right operational and analytical context at the right time for the task at hand. The result is lower operational complexity, improved retrieval precision, and more consistent decision-making.

Because Denodo accommodates both hybrid and multi-cloud environments, organizations can leverage it to unify enterprise context without large-scale consolidation projects or disruptive migration efforts.

Existing investments in lakehouses, warehouses, metadata catalogs, governance platforms, APIs, cloud AI services, and operational systems remain valuable. Rather than forcing organizations to consolidate or replace those technologies, Denodo adds a federated contextual overlay that connects them together through shared semantics, runtime governance, and live access to operational and analytical enterprise data, for visibility into current business conditions, transactional systems, enterprise events, and other operational workflows, enhanced by historical, analytical context.

This becomes increasingly important as enterprises continue expanding across multiple cloud systems, SaaS ecosystems, operational applications, external data sources, and specialized AI platforms. AI agents increasingly need to reason across both internal enterprise data and external business context in real time, especially for such use cases as intelligent workflow orchestration, autonomous process automation, real-time customer engagement, supply-chain optimization, and operational decision making. The future of enterprise AI will not exist inside a single platform boundary, and organizations are recognizing that more quickly now than even a year ago.

Denodo acts as the operational coordination layer across distributed architectures, enabling AI agents to reason consistently and operate securely across the enterprise.

Delivering Trusted Context for AI Agents

Agentic AI requires semantic consistency, but it needs more than that. AI agents must operate against live enterprise conditions while maintaining alignment with centralized governance policies and identity-aware access controls. As agents increasingly act on behalf of users across enterprise workflows, organizations need to ensure that agents can only access the data and systems they are authorized to access. This becomes especially important in environments where sensitive operational and transactional data spans multiple enterprise platforms.

Denodo enables this by operating directly in the runtime execution path. Rather than relying solely on static metadata or replicated data pipelines, Denodo enables AI agents to securely access trusted enterprise data in real time while governance policies are enforced dynamically as information is retrieved. This helps organizations establish a governed data access surface for AI, so sensitive information such as PII, financial records, or regulated operational data can be masked, filtered, or restricted before it ever reaches the model.

67% of organizations struggle with AI data security and access controls across distributed environments.[3]

This includes capabilities such as:

- Attribute-based access control (ABAC)
- Row- and column-level security
- Dynamic policy enforcement
- Context-aware authorization
- Centralized governance across distributed environments

This approach becomes more important as AI agents evolve from simple assistants into multistep reasoning and orchestration systems.

In addition to governed access, AI agents also need to operate within a consistent business context across distributed systems. As agents interact across APIs, operational applications, cloud services, and business workflows, the quality and structure of enterprise context directly impact both AI effectiveness and operational efficiency.

Without trusted contextual grounding, AI agents often spend additional time and compute attempting to reconcile inconsistent definitions, fragmented retrieval paths, and incomplete operational awareness. Over time, this can increase latency, retrieval complexity, token consumption, and infrastructure cost.

AI agents should focus on reasoning and workflow execution, while a centralized context layer governs enterprise data access, integration, and policy enforcement.

Denodo helps address this challenge by delivering trusted enterprise context through a federated logical layer that unifies semantics, governance, and live operational access across the enterprise. Instead of forcing AI systems to reason across disconnected architectures, Denodo provides a more consistent and governed foundation for AI retrieval and orchestration.

By defining reusable, governed integration logic within the logical layer, Denodo enables organizations to establish deterministic integration paths across distributed enterprise data. Rather than requiring AI agents to repeatedly infer relationships, joins, and business context dynamically at runtime, Denodo delivers pre-defined and governed, contextual relationships that improve consistency, reduce reasoning overhead, and help prevent inaccurate or inconsistent AI outcomes.

[3] Denodo, The AI Trust Gap Report

Denodo also provides lineage, observability, and execution visibility across distributed environments, helping organizations strengthen auditability, compliance, and responsible AI governance initiatives while maintaining trust in AI-driven outcomes.

Business Outcomes and Value

Organizations that adopt the Denodo Active Context Layer for Agentic AI also gain improved control over their AI initiatives, so they become significantly more efficient to operate over time.

As enterprises scale agentic AI deployments, many discover that the largest challenge is not simply building intelligent agents. The larger challenge is controlling the growing operational complexity and cost associated with fragmented enterprise architectures. This becomes especially important as organizations move toward multi-agent environments where agents must coordinate tasks, share context, and operate consistently across business processes.

Denodo helps address this challenge by establishing a clean separation between agent process logic and enterprise data access logic. AI agents can focus on reasoning, orchestration, and workflow execution, while Denodo centrally manages governed access to trusted operational and analytical enterprise context.

When AI agents must retrieve information across disconnected systems, duplicated pipelines, inconsistent business definitions, and fragmented governance models, both infrastructure costs and operational inefficiencies grow rapidly.

Many emerging AI frameworks still lack a clear path into enterprise production environments. Organizations need AI foundations that are already proven at scale.

Reasoning models are computationally expensive by nature. When AI agents must retrieve information across disconnected systems, duplicated pipelines, inconsistent business definitions, and fragmented governance models, both infrastructure costs and operational inefficiencies grow rapidly. The challenge becomes even greater when multiple agents operate with different interpretations of enterprise data, governance policies, or business meaning. Without a shared contextual foundation, agents can produce inconsistent decisions, conflicting actions, and unreliable outcomes across workflows.

This problem becomes even more pronounced as AI agents evolve from simple question-answering systems into multistep reasoning and orchestration engines. Without trusted contextual grounding, models often generate excessive reasoning activity while attempting to resolve missing relationships, inconsistent semantics, or incomplete operational awareness.

Denodo helps solve this challenge by delivering trusted enterprise context through a federated logical layer that unifies semantics, governance, and live operational access across the enterprise. Because the platform has been deployed in complex production environments for many years, organizations can build agentic AI initiatives on top of an enterprise-ready foundation. This reduces reliance on emerging AI frameworks that may not yet provide a clear or supportable path into production operations.

This approach helps organizations:

- Improve the relevance and accuracy of AI responses
- Reduce unnecessary retrieval and reasoning overhead
- Lower token consumption and infrastructure costs
- Minimize duplicated integrations and retrieval pipelines
- Reduce operational complexity across AI environments
- Strengthen governance, compliance, and auditability
- Accelerate deployment of agentic AI initiatives
- Scale AI adoption more sustainably across the enterprise

By simplifying how AI agents access and reason across enterprise information, organizations can improve both AI effectiveness and the economics of scaling AI over time.

Perhaps most importantly, Denodo enables organizations to scale agentic AI without forcing large-scale data consolidation projects or introducing additional governance silos. Existing investments in lakehouses, cloud platforms, metadata systems, and operational applications remain valuable while Denodo provides the trusted context layer that connects them together. This also helps organizations enforce identity-aware governance policies consistently across AI environments, so agents can only access the enterprise data the initiating user is authorized to see. For CISOs and governance leaders, maintaining these guarantees becomes critical as AI agents increasingly interact with sensitive operational and transactional systems.

The result is faster AI adoption, stronger governance, improved operational trust, and a more cost-efficient path to enterprise-scale agentic AI.

The Future of Agentic AI Requires Federated Enterprise Context

The idea that organizations will centralize all enterprise data into a single platform is becoming increasingly unrealistic. Even in cases where large-scale consolidation efforts are successful, critical operational context often remains outside the platform boundary or is introduced through replicated pipelines that add latency and reduce real-time business awareness for AI systems.

Modern enterprises operate across multiple cloud systems, SaaS ecosystems, operational applications, regional environments, and specialized AI platforms. This fragmentation is not temporary. It reflects how enterprise technology environments are evolving.

The challenge is no longer simply consolidating data. The challenge is delivering trusted enterprise context efficiently enough for AI agents to operate accurately, securely, and economically at scale. Organizations that fail to address this challenge risk creating AI environments that become more expensive, fragmented, and difficult to govern as reasoning workloads grow.

> Even when organizations centralize data successfully, critical operational context often remains outside the platform boundary or arrives with latency that limits real-time AI decision-making.

This is where Denodo is uniquely differentiated.

The Denodo Active Context Layer for Agentic AI is a trusted operational foundation that unifies semantics, governance, provenance, runtime enforcement, and live operational context across the enterprise.

Unlike metadata-only approaches, cloud-specific AI platforms, or BI-centric semantic layers, Denodo enables organizations to operationalize trusted enterprise context for AI agents at enterprise scale.

Take the Next Step

Organizations that succeed with agentic AI do more than deploy powerful models. They establish a trusted operational foundation that delivers governed, live enterprise context across the business.

Denodo helps organizations build this foundation by enabling federated enterprise access, runtime governance enforcement, trusted semantic consistency, and operational AI readiness across hybrid and multi-cloud environments.

Request a tailored workshop or demo to explore how Denodo can help operationalize trusted enterprise context for agentic AI.