



# THE DENODO AGENTIC AI MANIFESTO

The AI data layer for agents acting with live data, trusted context, and governed guardrails.

# EXECUTIVE INTRODUCTION

Agentic AI is ready to move from pilots to production but only when trust is engineered from the start.

CIOs, CDOs and CAIOs can trust autonomous agents when three criteria are met.

1. They act on real-time data
2. They carry the full context of the business process and
3. They operate within explicit guardrails

Together, these foundations create active context: live, governed, trusted enterprise context for the task at hand. With active context, agents don't just answer. They execute safely, explain decisions and improve outcomes.

The result is credible autonomy at enterprise scale, faster cycles, fewer escalations and consistent compliance.

Leaders must define where agents act, where humans remain accountable, and how trust is enforced before autonomy becomes systemic.

This manifesto defines the standard for trusted agentic AI and the architecture that enables it to operate autonomously, responsibly, and at enterprise scale.



# The promise and perils of agentic AI

The real promise of agentic AI is not task automation, but workflow transformation. McKinsey's first-year findings show that value emerges when organizations redesign end-to-end processes around autonomy, shifting humans to oversight, judgment, and exception handling. Gartner forecasts this shift will be profound, with autonomous decision-making becoming routine in core business operations.

The peril is scaling autonomy without readiness. Agents amplify whatever they inherit—broken workflows, fragmented data, or unclear accountability. Without explicit boundaries, governance, and observability, autonomy introduces drift, cascading failures, and loss of trust. Gartner warns that a significant share of AI initiatives will be cancelled due to weak governance and unclear value realization.

Agentic AI rewards clarity. Leaders must define where agents act, where humans remain accountable, and how trust is enforced—before autonomy becomes systemic.



# The moment: autonomy without trust breaks down today

When agents act without the right substrate, risk compounds. Stale copies, misread reality, fragmented silos starve decisions. Missing semantics conceals meaning.

Without active context, including live data, business meaning, and governed guardrails, agents drift, overreach or stall, and confidence erodes.

Dashboards tolerated latency and ambiguity. Autonomous workflows cannot. What fails is not ambition, but foundation.

Most data estates were built for reports, not decisions that trigger action to cross from demos to dependable production.



# The Trust Gap

Gartner forecasts that over **40% of agentic AI projects will be canceled by the end of 2027**, due to escalating costs, unclear business value, or inadequate risk controls.<sup>[1]</sup>

**66%** of organizations say AI data must be real-time, or no more than one minute old, to be trustworthy — yet traditional data architectures, designed for historical analytics, cannot support this sub-minute latency, creating a critical bottleneck.

There is a widening trust gap – leaders expect autonomous performance, yet agents often lack real-time data, complete process context and enforceable guardrails.

The result is missed SLAs, escalations and abandoned initiatives.

Closing the gap requires an architecture that can deliver live, governed enterprise context across distributed data without creating another layer of copies.

<sup>[1]</sup> Gartner Predicts over 40% of Agentic AI Projects Will Be Canceled by End of 2027.

# The standard: Active context for trusted autonomy

Live awareness: Agents act on current facts, not yesterday's copies. Direct access to distributed enterprise data keeps decisions aligned with what is happening now.

Business context: Semantics, relationships, lineage, and history clarify what data means, how it was derived, and why it matters, so decisions are more explainable and defensible.

Governed guardrails: Security, privacy, policy, and access controls remain consistently enforced as agents discover, reason over, and act on enterprise data.

## ARCHITECTURE:



## AN AI DATA LAYER FOR ACTIVE CONTEXT

An AI data layer provides a reusable architecture for delivering active context across distributed enterprise data.

It connects data across clouds, platforms, applications, and operational systems while minimizing unnecessary data movement and preserving governance.


Instead of rebuilding integration, semantics, and controls for every new agent or AI initiative, organizations can deliver trusted enterprise context through a common data foundation.

# How Denodo delivers active context



## Universal Connectivity

Connect distributed enterprise data across clouds, platforms, applications, SaaS, and operational systems through a common logical access layer.



## Zero-Copy Delivery

Deliver live data on demand while minimizing unnecessary movement and avoiding another copy for every agent, application, or use case.



## Governed Access

Apply security, privacy, policy, and access controls consistently as enterprise data is delivered to AI and other consumers.



## Semantic Trust

Give agents shared business meaning through semantics, relationships, lineage, and context so they can interpret enterprise data more reliably.

# Proof, outcomes, and ROI delivered



## A LEADING GLOBAL IT MANUFACTURER

Deployed enterprise-wide Agentic AI across **1,700+** governed business views, enabling **100,000+** employees to interact with trusted data.



## A NATIONAL FINANCIAL REGULATORY AUTHORITY

Modernized regulatory and financial oversight workflows using agentic AI grounded in live, governed data with full auditability.

## PERFORMANCE

Customers have delivered AI-ready data in one-third the time and achieved up to **10x performance** compared to traditional architecture.





# This is more than a manifesto; it is a movement.

We are committing to trusted autonomy: agents acting with active context, grounded in live enterprise data, shared business meaning, and governed guardrails. An AI data layer makes that possible across distributed enterprise data without requiring another integration and copy for every new agent. Denodo is the AI data layer, helping organizations prove value quickly and scale AI responsibly.

REQUEST A STRATEGY SESSION



Denodo is the AI data layer company. The award-winning Denodo Platform powers trustworthy agents and applications, transforming enterprise data into reliable insights for analytics and self-service. Organizations worldwide deploy Denodo alongside their data lakehouses to deliver AI-ready, business-ready data in a fraction of the time—achieving up to 4x faster time to insight, 345% ROI, and 10x better performance.

Visit [denodo.com](https://denodo.com) | Email [info@denodo.com](mailto:info@denodo.com) | Discover [community.denodo.com](https://community.denodo.com)

