


The Active Context Layer for Enterprise AI

How to deliver trusted, governed, and operationally aware context for AI agents



The Role of
Context in
Enterprise AI

01



Introducing
the Active
Context
Layer

02



Without
an Active
Context Layer

03



The
Requirements
for an Active
Context Layer

04



The Denodo
Active Context
Layer

05




Next Steps

06

CONTENTS

CHAPTER ONE

The Role of Context in Enterprise AI



Enterprise AI Needs More Than a Model

AI is moving from experimentation to execution.

Organizations are no longer just building chatbots or copilots. They are beginning to deploy AI agents that retrieve information, reason across systems, invoke tools, and support business workflows.

But as AI becomes more operational, the data challenge becomes more complex.

AI agents need more than access to data. They need trusted enterprise context: business meaning, the current operational state, governed access, lineage, and policy enforcement.

That is the role of the “active context layer.”

Why Context Has Become Critical

AI models are powerful, but they do not automatically understand your business.

THEY DO NOT
INHERENTLY
KNOW


- ⊗ Which data source is trusted
- ⊗ What business terms mean
- ⊗ Which data is current
- ⊗ Which policies apply
- ⊗ What each user is allowed to see
- ⊗ How operational systems relate to analytical data

Without this context, AI may generate answers that sound right but are incomplete, inconsistent, outdated, or unsafe.

Enterprise AI needs a trusted layer that helps agents reason with the right context at the right time.

CHAPTER TWO

Introducing the Active Context Layer



What Is an Active Context Layer?

An active context layer operationalizes trusted enterprise context for AI systems, applications, and users.

IT'S NOT JUST

- A semantic layer
- A metadata catalog
- A retrieval framework
- A data pipeline

An active context layer brings together **semantics, governance, live data access, metadata, lineage**, and **runtime policy enforcement** in a unified logical layer.

THE RESULT

AI systems can securely, consistently, and dynamically access trusted enterprise context across distributed environments.

Static Context Is Not Enough

AI agents operate in real time. They retrieve information, make decisions, invoke tools, and trigger workflows.

That means context cannot simply be documented. It must be applied during execution.

WITH AN
ACTIVE
CONTEXT
LAYER


- Business meaning is applied when data is retrieved
- Policies are enforced when data is accessed
- Sensitive information is protected before reaching the model

- Operational data reflects current business conditions
- Lineage and provenance are available for trust and auditability

For agentic AI, context must be active, governed, and operationally aware.

CHAPTER THREE

Without an Active Context Layer



Three Reasons Enterprise AI Fails to Scale

Without an active context layer, enterprise AI often fails in three predictable ways.

01

AI Without Context Cannot Be Trusted

AI responses become inconsistent, incomplete, or disconnected from real business meaning.

02

AI Without Governance Cannot Scale Safely

Policies fragment across tools, pipelines, models, and agents, increasing security and compliance risk.

03

AI Without Operational Efficiency Cannot Scale Economically

Each AI project builds its own integrations, retrieval logic, and governance controls, increasing cost and complexity.

FAILURE MODE 01

AI Without Context Cannot Be Trusted

Enterprise AI depends on more than data volume.

It needs trusted business meaning, current operational awareness, and consistent definitions across systems

WITHOUT THAT CONTEXT, AI AGENTS MAY

- ⊗ Misinterpret business terms
- ⊗ Use stale or incomplete data
- ⊗ Retrieve information from the wrong source
- ⊗ Generate inconsistent responses
- ⊗ Lose user trust

An active context layer gives AI systems a shared understanding of enterprise data, helping improve accuracy, consistency, and adoption.

FAILURE MODE 02

AI Without Governance Cannot Scale Safely

AI agents can access data, invoke tools, and act on behalf of users.

That makes governance a runtime requirement. Organizations need AI systems to only access the data each user, role, region, or process is authorized to use.

ENFORCE POLICIES AT THE POINT OF ACCESS

- ⊗ Row-level security
- ⊗ Data masking
- ⊗ Contextual authorization
- ⊗ Column-level restrictions
- ⊗ Attribute-based access control
- ⊗ Auditability and lineage

Governance must happen before sensitive data reaches the model.

FAILURE MODE 03

AI Without Efficiency Cannot Scale Economically

As AI use cases grow, many organizations risk creating a new layer of complexity.

This increases cost, latency, risk, and maintenance effort.


**EACH TEAM
MAY BUILD
ITS OWN**

- ⊗ Retrieval pipelines
- ⊗ Data integrations
- ⊗ Semantic definitions
- ⊗ Security controls
- ⊗ Governance workflows
- ⊗ Agent-specific data access patterns

An active context layer provides reusable trusted context across AI systems, reducing duplication and helping teams scale more efficiently.


CHAPTER FOUR

The Requirements for an Active Context Layer



What an Active Context Layer Must Provide

To support enterprise AI at scale, an active context layer should deliver:



Unified access to distributed enterprise data




Shared business semantics




Runtime governance and security




Live operational and analytical context



Metadata, lineage, and provenance



Reusable, governed data products



Flexible delivery through APIs, SQL, and AI frameworks



Performance optimization and cost visibility

Together, these capabilities help AI systems operate with trusted enterprise awareness. In the next few pages, we'll drill down into each of these capabilities in turn.

UNIFIED DATA ACCESS:

Connecting AI to the Full Enterprise

Enterprise data is distributed across cloud platforms, lakehouses, warehouses, SaaS applications, operational systems, APIs, streaming sources, and on-premises environments.

An active context layer provides a unified way to access this data without forcing everything onto a single physical platform. This is especially important for agentic AI.

AI AGENTS OFTEN NEED BOTH



AI agents need the full enterprise picture, not just the data already copied into one environment.

SEMANTIC CONSISTENCY:

Providing AI with a Shared Business Vocabulary


AI systems need to understand enterprise meaning.

A semantic layer translates complex technical data into consistent business concepts, metrics, definitions, and relationships.

THIS HELPS
AI AGENTS
AND USERS
UNDERSTAND

- What data means
- How entities relate
- Which metrics are trusted
- Which definitions are approved
- Which sources should be used

With shared semantics, organizations reduce ambiguity and help AI systems deliver more consistent, explainable answers.



RUNTIME GOVERNANCE:

Governing AI at the Point of Access



For enterprise AI, governance cannot be limited to policies documented in a catalog.

Policies must be enforced when data is accessed.

An active context layer helps ensure that AI agents only receive the data they are authorized to use. Sensitive information can be **masked**, **filtered**, or **restricted** before it reaches the model.

This supports safer AI adoption by applying consistent controls across users, agents, applications, and distributed data sources.

OPERATIONAL AWARENESS:

Keeping AI Connected to Current Business Conditions



Many AI use cases depend on a grasp of the current operational state:

A **customer service agent** may need the latest case status.



A **supply chain agent** may need current inventory data.



A **financial risk agent** may need recent transactions.



A **field operations agent** may need live service information.



PROVENANCE AND LINEAGE:

Making AI More Transparent and Auditable

Trust requires visibility.

Organizations need to understand which data was used, where it came from, how it was transformed, and which policies were applied.

An active context layer provides lineage, metadata, and observability so teams can:

- Audit AI responses
- Explain outcomes
- Track data usage
- Improve trust in AI-driven decisions
- Investigate issues
- Strengthen compliance

When AI becomes operational, explainability cannot be optional.

DATA PRODUCTS:

Packaging Trusted Context for Reuse

Data products make trusted enterprise context easier to find, govern, and reuse.

This creates a clean separation of duties.

Data teams manage the trusted context. AI builders consume that context to build agents and applications


A DATA PRODUCT CAN PACKAGE

- Trusted data
- Business definitions
- Ownership
- Documentation
- Quality expectations
- Access policies
- Lineage and metadata

AI builders should not have to become data engineers.


CHAPTER FIVE

The Denodo Active Context Layer



How Denodo Supports the Active Context Layer

The Denodo Platform provides the logical data management capabilities needed to support an active context layer for enterprise AI. Denodo brings together:



With Denodo, organizations can deliver trusted context across distributed environments without unnecessary data movement or fragmented point-to-point integrations.




CAPABILITY
1 OF 6

Federated Access Across Distributed Systems

Denodo connects to live data across hybrid, multi-cloud, SaaS, operational, analytical, and on-premises environments.

This gives AI systems a unified access layer across the enterprise.



Instead of building separate integrations for every model, agent, or application, organizations can use Denodo to provide governed access to trusted data wherever it resides.

This helps reduce duplication, simplify architecture, and support AI use cases that require both operational and analytical context.



CAPABILITY
2 OF 6

Universal Semantics for Trusted AI

Denodo provides a universal semantic layer that abstracts technical complexity and presents data in business-friendly terms.

This helps organizations create reusable definitions, relationships, classifications, and business views that can be shared across analytics, applications, and AI systems.

For AI, this improves context quality with data that is based on trusted enterprise meaning rather than raw technical structures.

Denodo also enables organizations to enrich data with semantic tags and attributes, including classifications for sensitive data such as PII, financial data, or regulated information. These semantic classifications provide the context needed to consistently apply global policies consistently across distributed sources.

Consistent semantics help AI systems to produce more reliable outcomes and provide governance teams with a stronger foundation for policy enforcement.



CAPABILITY
3 OF 6

Runtime Governance and Security

Denodo enforces governance at the point of access.

Using semantic tags, data classifications, user attributes, and policy rules, Denodo can consistently apply global governance policies across distributed data sources.

The platform supports fine-grained security controls, including:

- | | | |
|-------------------------|----------------------------------|--|
| ■ Row-level security | ■ Dynamic data masking | ■ Integration with enterprise identity systems |
| ■ Column-level security | ■ Execution restrictions | |
| ■ Cell-level controls | ■ Attribute-based access control | |

For AI use cases, this identifies, governs, masks, or restricts sensitive data before it reaches the model.



CAPABILITY
4 OF 6

The Denodo Data Marketplace

The Denodo Data Marketplace helps users discover, understand, and access trusted data products through a business-friendly, self-service experience.

It brings together rich metadata, business descriptions, ownership, lineage information, usage details, policy context, and governed data products in one place.

This helps organizations make trusted enterprise context easier to find, evaluate, and reuse across analytics, applications, and AI initiatives.


Business users

can quickly identify the right data for decision-making.




AI teams

can build applications using reusable, well documented data products.



AI agents

can be grounded in approved enterprise context instead of disconnected raw data sources.




By making governed data products easier to discover and consume, the Denodo Data Marketplace helps scale trusted data use across the enterprise




CAPABILITY
5 OF 6

The Denodo AI SDK and Assistant



The Denodo AI SDK

helps developers build AI-powered applications and agents that securely access governed enterprise data through Denodo.



Denodo Assistant

brings AI-powered assistance into the Denodo experience, helping users interact with data more naturally through intelligent recommendations, natural language capabilities, and guided support.

Together, these capabilities help teams build AI applications faster while keeping data access governed and trusted.



CAPABILITY
6 OF 6

Performance, Optimization, and FinOps

Enterprise AI must scale efficiently.

Denodo helps improve performance across distributed data environments with query optimization, caching, workload management, smart query acceleration, and monitoring capabilities.

It also helps reduce the work AI models need to perform. By providing a well-curated, governed location for trusted enterprise context, Denodo enables AI systems to retrieve ready-to-use data products instead of forcing models to interpret raw data, resolve joins, reconcile definitions, or reason across fragmented integrations.

This can reduce unnecessary token consumption, lower latency, and improve response quality because the model receives more precise, relevant, and business-ready context from the start.

Denodo also provides FinOps visibility into usage patterns, query behavior, workload performance, and resource consumption.

This helps organizations manage cost, improve efficiency, avoid unnecessary data replication, and scale AI adoption more sustainably.

What the Denodo Active Context Layer Enables

With Denodo as the active context layer, organizations can move from AI experimentation to trusted AI execution.

TRUSTED AI

Ground AI responses in governed, consistent, and business-ready context.

GOVERNED AI

Enforce policies at runtime and protect sensitive data before it reaches the model.

OPERATIONALLY AWARE AI

Provide agents with access to right-time operational and analytical context.

SCALABLE AI

Reuse trusted data products, semantics, and governance across multiple AI initiatives.

SUSTAINABLE AI

Reduce duplicated integrations, fragmented retrieval pipelines, and unnecessary data movement.

NEXT STEPS

Build the Trusted Operational Foundation for Enterprise AI

CONTACT US

Enterprise AI success depends on more than selecting the right model.


It depends on whether AI systems can access trusted, governed, and current enterprise context.

The Denodo Platform helps organizations establish the active context layer needed to support enterprise AI at scale.

With Denodo, organizations can unify semantics, govern access, deliver live operational context, expose reusable data products, and provide AI systems with trusted enterprise awareness across distributed environments.

As AI becomes more autonomous and embedded in business workflows, an active context layer becomes essential.

Denodo helps make enterprise AI trusted, governed, and ready to scale.



Visit www.denodo.com | Email info@denodo.com | Discover community.denodo.com

