MARKETSANDMARKETS™ | GRAPHWISE

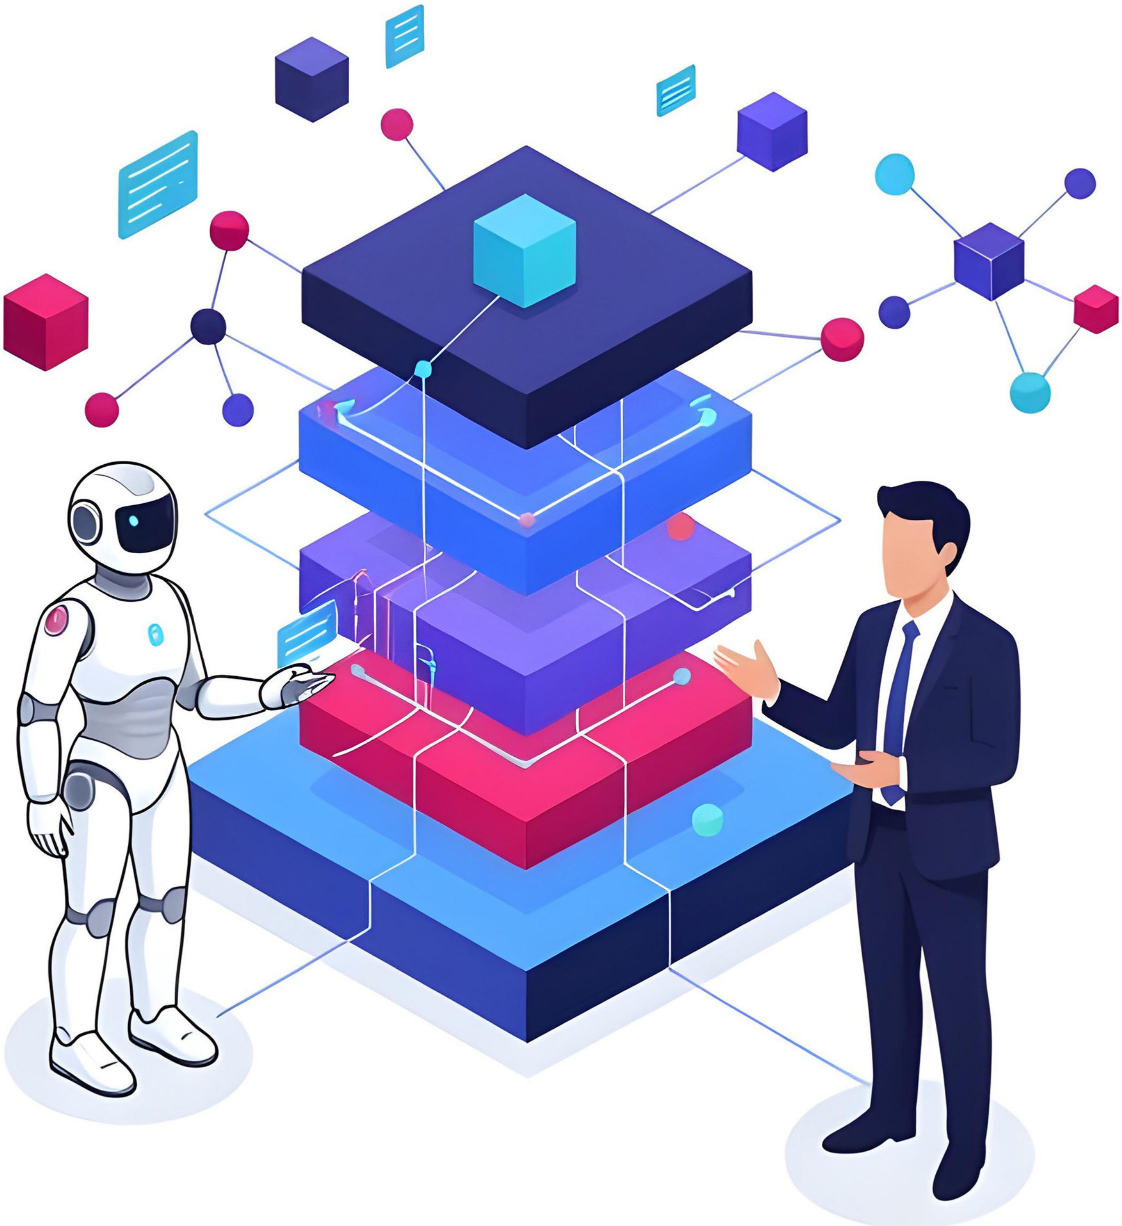

WHITE PAPER

The Enterprise Semantic Backbone

A Foundation for Reliable and Scalable Agentic AI

Table of Contents

1. Executive Summary

As AI becomes a central pillar of enterprise technology investment and business strategy, organizations are rapidly moving from experimentation to scaled deployment and broader adoption. Yet, despite this momentum, most enterprises continue to struggle to generate consistent, enterprise-wide value from AI, with many initiatives still stalled at the pilot stage or operating in fragmented environments.

To assess the state of enterprise AI adoption and architecture, MarketsandMarkets conducted a survey of 380+ enterprise leaders, including CIOs and technology executives, and combined this with analysis of market practices, vendor frameworks, and independent benchmarks. The objective was to identify the precise data and knowledge architectures required to successfully scale autonomous, agentic AI.

Key findings:

- ✓ The Scale Bottleneck: While 74% of organizations are investing heavily in AI and 90% of CIOs plan to increase spending by 2026, execution remains shallow. Only 28% have deployed an enterprise-wide AI strategy, while 95% report zero return on initial GenAI investments.
- ✓ System Fragmentation: Organizations manage an average of 957 applications, with only 27% integrated. As a result, nearly half of deployed AI agents operate in isolated silos because they lack unified, machine-executable context.
- ✓ The ROI Demand Signal: To address this challenge, 60% of executives rank a shared, reusable knowledge foundation among their top three priorities for improving AI ROI and scaling AI across the enterprise.

To move beyond the prototype plateau, the paper argues that organizations must shift toward a graph-based, vendor-independent knowledge infrastructure: the Semantic Backbone. Rather than treating knowledge as a temporary overlay on AI applications, enterprises must establish a governed, machine-executable foundation that provides persistent context, deterministic reasoning, and enterprise-wide interoperability. In this context, Graphwise is positioned as a leader in delivering this alternative.

- ✓ The Context Layer enhances AI with conversational memory and retrieval capabilities but remains probabilistic and limited to individual interactions.
- ✓ The BI Semantic Layer (anchored to the Open Semantic Interchange/OSI) provides a standardized metric layer but lacks ontologies, formal semantics, reasoning, and governance capabilities required for autonomous agent orchestration.
- ✓ Graphwise replaces these approaches with a standards-based Semantic Backbone—an enterprise semantic nervous system built on W3C graph standards such as RDF, OWL, and SKOS.
- ✓ The Context Graph functions as an operational “dashcam,” capturing reasoning paths, event histories, and workflow timelines, enabling agents to understand not only what happened, but also how and why decisions were made.
- ✓ By executing deterministic logic and SHACL constraints through platforms such as Graphwise GraphDB, this architecture delivers 100% explainable reasoning, achieves a 2x accuracy improvement compared with schema-less RAG approaches, and reduces AI hallucinations by 70% to 90%.

Enterprises must stop treating knowledge as an application-level enhancement and begin treating it as core infrastructure. Achieving production-scale AI success requires establishing a Semantic Foundation built on governed, standards-based knowledge assets and then activating enterprise intelligence through agent orchestration, reasoning, and automation. This roadmap provides the foundation necessary to deliver reliable, explainable, and scalable agentic AI at enterprise scale.

2. Methodology and Research Approach

This study sets out to answer a single strategic question: as enterprises move from AI experimentation to scaled, autonomous (agentic) AI, what architectural foundation makes those systems reliable, governed, and scalable? The question matters because investment in AI is accelerating while enterprise-scale value remains elusive-so decision-makers need a clear, evidence-based view of which data and knowledge architectures actually close that gap.

MarketsandMarkets led the research, writing, analysis, and production of this white paper. The study includes desk research, a survey of 384 enterprise leaders, interviews, and the resulting conclusions, all developed by MarketsandMarkets. Graphwise contributed supporting materials on request, including background on the Semantic Backbone concept, selected customer deployment references, and technical documentation.

Research design. The findings combine two complementary approaches:

- **Desk research** across market-research and industry sources, vendor technical material, and academic/independent benchmarks (for example, the Stanford HAI AI Index 2026) [1], used to establish trends, problems, and current solution patterns.
- **A company survey** of 384 enterprise leaders spanning three verticals-Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 respondents each)—used to validate the prevalence and ranking of adoption blockers, grounding technologies, and demand signals.

Logical flow of the desk research. The analysis follows a deliberate chain that mirrors the structure of this paper:

- **Identify market trends –** establish where enterprise AI investment and adoption are heading.
- **Surface the core problems –** distil the structural barriers that keep AI stuck at the "prototype plateau," then state the core problem explicitly.
- **Deep-dive into current solutions –** examine the first architectural responses the market has put forward: the Context Layer and the BI Semantic Layer.
- **Run a gap analysis, then introduce the Semantic Backbone –** test those responses against the problems defined earlier, then develop the standards-based foundation and contrast it with the BI semantic layer.
- **Establish the current domain state –** translate the architecture into concrete use cases and a phased implementation approach.

Research flow at a glance:

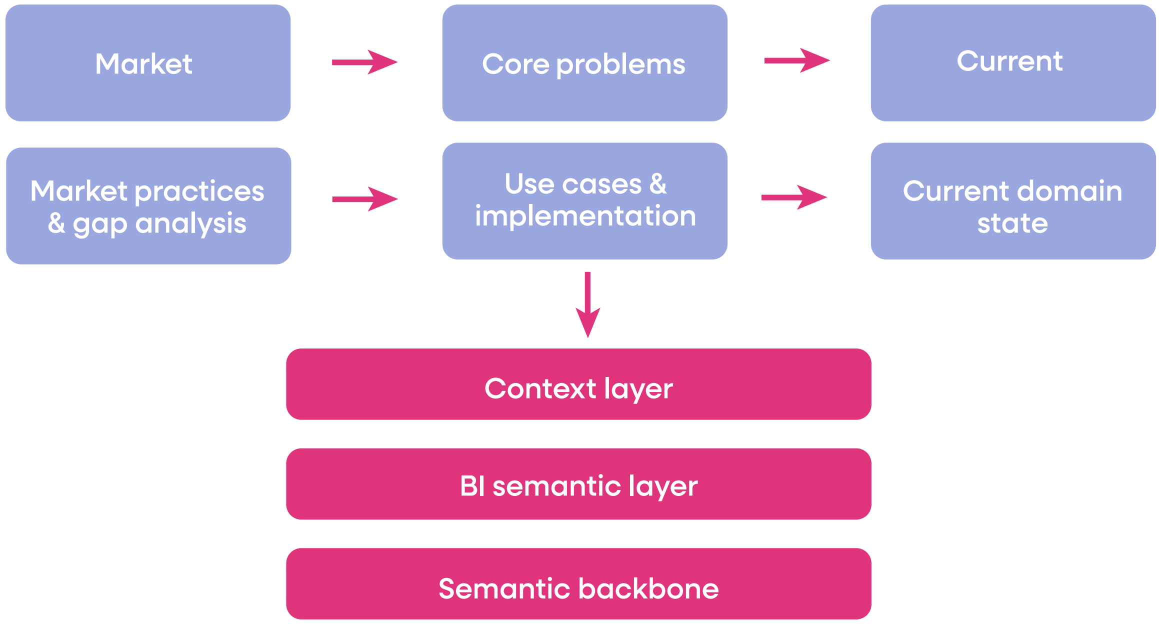

Survey sample composition (n=384 enterprise leaders)

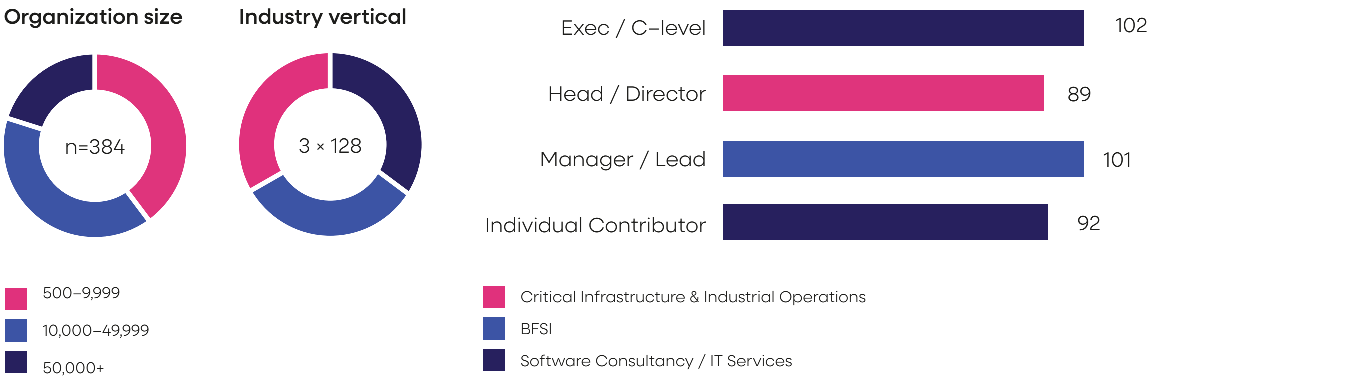

Figure 2.1 – Survey sample composition: the analyzed base is 384 enterprise leaders, balanced across three verticals (128 each) and spanning organization sizes and seniority levels.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3. Key AI Trends Overview

Enterprise AI adoption has entered a new phase–shifting from experimentation to scaled deployment, with AI increasingly becoming a core strategic capability underpinning business transformation. A synthesis of recent market research and industry reports highlights several defining trends shaping this trajectory. Industry framings position data, knowledge and orchestration–not the model alone–as the decisive layer for enterprise value.

Where Enterprises Expect The Strongest ROI

Where do you expect to see the strongest return on your AI investment over the next 2-3 years?

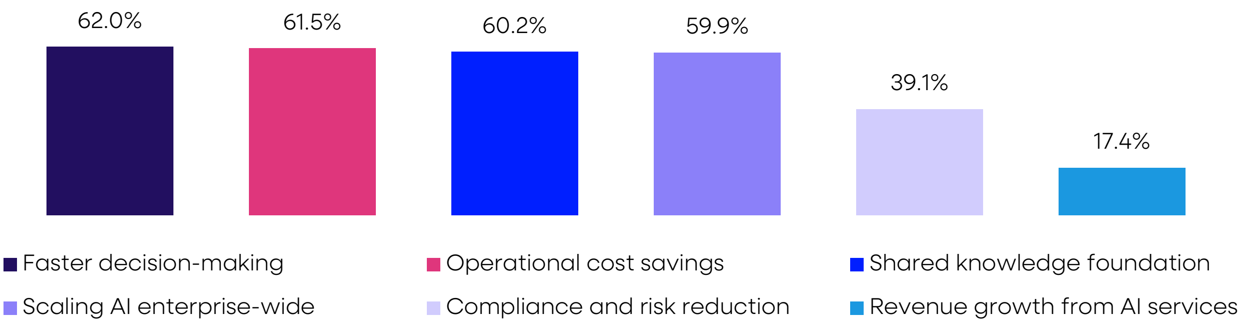

Figure 3.1 – The survey implied that the shared knowledge foundation is amongst the top 3 areas alongside faster decision-making, and operational cost savings, where the enterprises expect the strongest ROI.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3.1 AI is becoming the central pillar of enterprise technology investment

- AI is rapidly becoming the central pillar of enterprise digital strategy, with organizations reallocating budgets and resources toward AI–driven initiatives.
- In MarketsandMarkets' primary survey AI now sits firmly on the boardroom agenda: every respondent is involved in shaping AI, data, or analytics initiatives, and the objectives driving that investment are led by decision support and analytics (52.6%), cost optimization (51%), and customer–experience enhancement (51%)—ahead of process automation (49.7%), product and service innovation (48.4%), and revenue growth (47%) (see Figure 3.1).
- Around 90% of CIOs plan to increase AI spending by 2026, marking a decisive shift from experimentation to production; ~60% of firms are already running AI initiatives and ~90% have created dedicated budgets.
- This reflects AI's evolution into a strategic capability rather than a functional tool, reshaping enterprise architecture, operating models, and competitive positioning.

Primary Business Objectives driving AI Investment

What are the primary business objectives driving your AI initiatives?

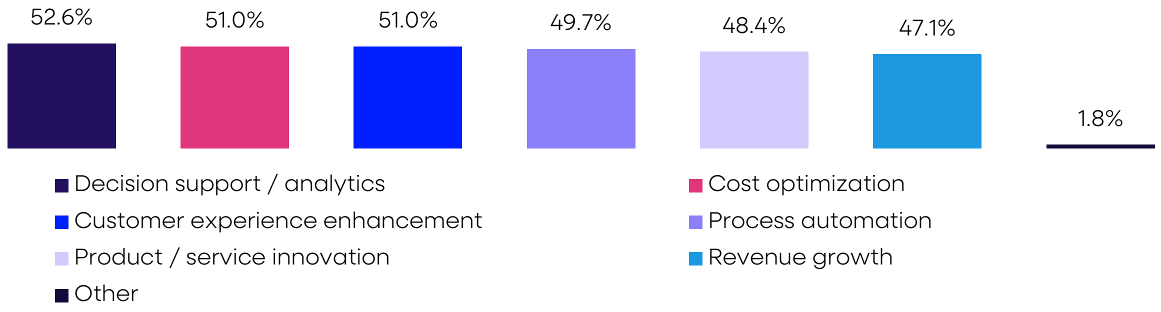

Figure 3.2 – Business objectives driving AI investment: Decision support / analytics, cost optimization, and customer experience enhancement lead a tightly-clustered set of objectives.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3.2 Rapid adoption of AI, but limited enterprise-scale impact

- AI has moved to widespread adoption, but adoption is still largely shallow, with most organizations struggling to move beyond pilots.
- In our survey, adoption has moved well beyond first experiments, but enterprise-scale value is still concentrated at the top: 64% of organizations have scaled AI—36% run multiple scaled use cases and 28% have executed an enterprise-wide AI strategy—while 36% remain in early pilots or proofs of concept (see Figure 3.3). The binding constraint has shifted from access to AI to the ability to scale it across enterprise systems, data, and workflows.
- The primary challenge has shifted from access to AI technologies to the ability to integrate, scale, and operationalize AI across enterprise systems, data environments, and workflows.

Stage of AI Adoption, By Industry

Which best describes the stage of AI adoption in your organization?

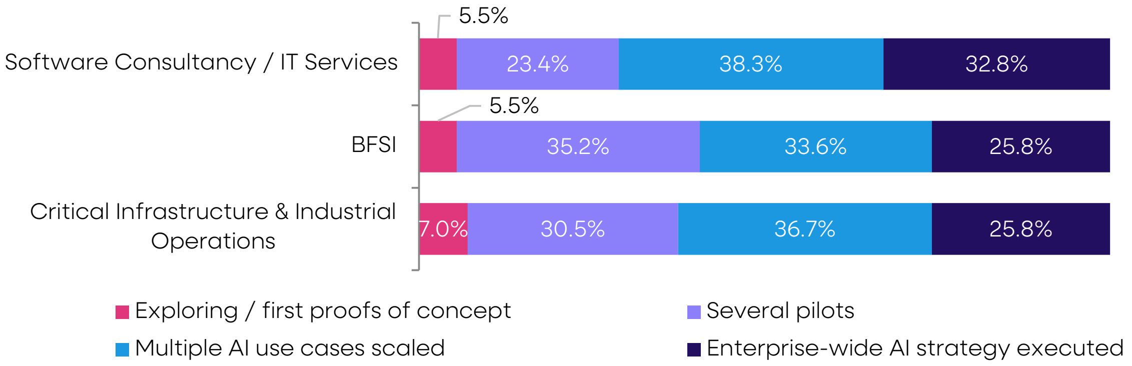

Figure 3.3 – Stage of AI Adoption, By Industry: Roughly two-thirds of organizations have scaled AI, but only ~28% have executed an enterprise-wide strategy, across all the verticals.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3.3 Rapid emergence and investment in Agentic AI

- ✓ A defining trend is the rise of agentic AI systems–autonomous, task–executing entities rather than static models.
- ✓ Agentic AI has already moved into mainstream deployment: half of organizations (50%) have deployed agentic or workflow–automation applications–the second most–common GenAI use case, just behind software–development and coding assistants (53%)–confirming the shift from static models to autonomous, task–executing systems (see Figure 3.4).
- ✓ Operational hurdles persist: about half of AI agents operate in silos and only ~27% of enterprise applications are integrated, creating fragmentation.

GenAI Applications Deployed, By Industry

Which of the following GenAI applications is your organization currently deploying or actively piloting?

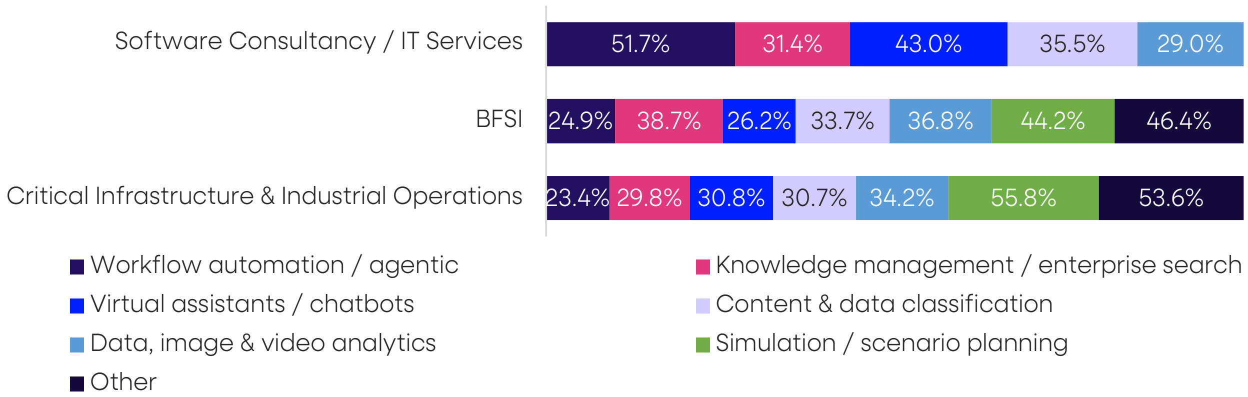

Figure 3.4 – GenAI Applications Deployed, By Industry: Workflow automation / agentic is the largest GenAI application deployed in Software Consultancy / IT Services, while Simulation / scenario planning is the highest amongst BFSI and Critical Infrastructure & Industrial Operations industry.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3.4 Data integration and architecture are the primary bottlenecks

- ✓ Enterprise infrastructure is not keeping pace with AI capability. Enterprises manage ~957 applications on average, with only ~27% connected; 64% of IT leaders cite integration challenges as a barrier.
- ✓ The Stanford HAI AI Index 2026 emphasizes a growing gap between AI capability and supporting systems such as governance, evaluation, and data infrastructure [1].
- ✓ Data interoperability, context management, and semantic consistency are now critical enablers for enterprise AI success.

Predominant AI Grounding Approach, By Industry

What technology or approach does your organization predominantly use to make AI / LLM answers more reliable and grounded in your own data?

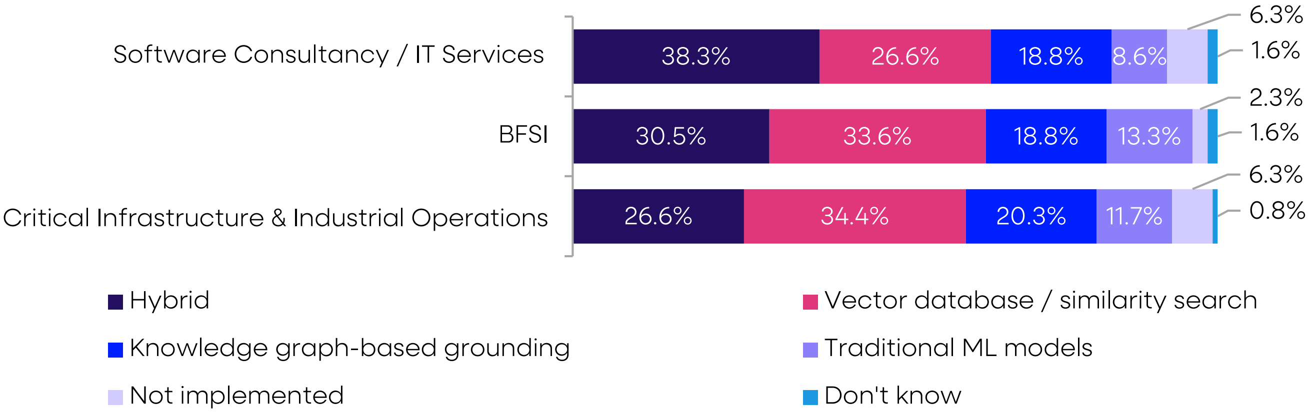

Figure 3.5 – Predominant AI Grounding Approach, By Industry : Although Hybrid is one of the highly adopted grounding approach, Vector database / similarity is largest across all 3 industries, followed by Knowledge graph-based grounding which is still in the early phase of adoption

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

3.5 Acceleration in AI capabilities and global investment

- Generative AI is reaching ~53% population adoption within three years and global corporate AI investment more than doubled in 2025; reported productivity gains of ~14%–26% in areas such as customer support and software development. [1].
- AI advancement is outpacing enterprise readiness, creating a gap between technological capability and the organizational systems required to support it.

3.6 AI is reshaping business models and competitive dynamics

- Organizations are shifting from services-based models to scalable AI-native platforms, a shift echoed in analyzes of professional-services transformation.
- In 2025, Amazon Web Services introduced Amazon Bedrock AgentCore at AWS Summit 2025–persistent memory, access controls, and system integration–with a $100M investment to accelerate adoption of agentic AI.

3.7 The ontology revival and the rise of the Semantic Backbone

A defining shift of late 2025 is that the industry has begun to converge on knowledge representation–ontologies, taxonomies, and knowledge graphs—as the missing foundation for production-grade agentic AI. Analysts and major platform vendors increasingly describe knowledge and context engineering, rather than raw model capability, as the decisive factor separating successful deployments from stalled prototypes.

- A wave of ontology-centric product announcements landed in Q4 2025: Microsoft introduced Fabric IQ; Palantir and NVIDIA partnered on ontologies; SAP added semantic capabilities to HANA Cloud and its Joule agents; and Snowflake, Salesforce, and others backed the Open Semantic Interchange initiative.

- ✓ The terminology is converging as well. Microsoft positions Fabric IQ–which combines an ontology, a BI semantic model, a native graph engine, and data and operations agents into a single semantic-intelligence system—as a semantic backbone serving the entire organization's teams, applications, and agents. That even the largest platform vendors now reach for the language of a semantic backbone underscores where enterprise architecture is heading, and frames the deeper question this analysis takes up: what such a backbone must actually contain to be reliable, governed, and scalable.

Expected Primary Impact of AI on the Business, By Industry

Looking at the bigger picture, what is the primary impact of AI on your organization's business over the next 2–3 years?

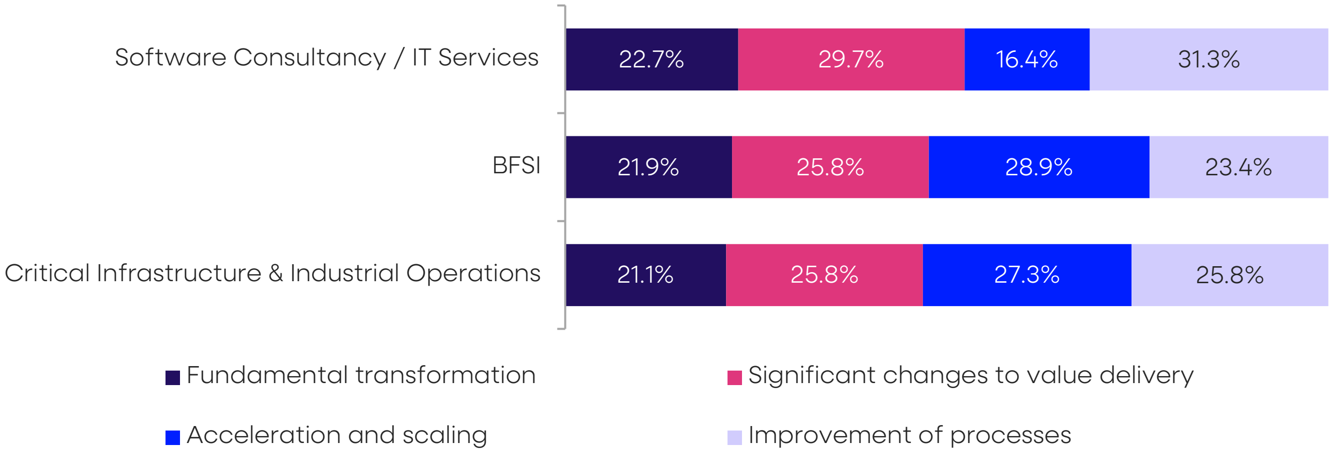

Figure 3.6 – Expected Primary Impact of AI on the Business, By Industry : A majority of enterprises expect AI to reshape — not merely improve — their business within 2–3 years; roughly half foresee fundamental transformation & significant change to how they create and deliver value. Software Consultancy / IT Services skews most towards improvement of processes, while BFSI and Critical Infrastructure & Industrial Operations lean towards acceleration and scaling.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

4. Key Challenges in AI and Agentic AI

While enterprise AI adoption is accelerating, organizations continue to face structural challenges that limit reliability, scalability, and enterprise-wide value realization. These challenges become more pronounced with agentic AI, where systems act autonomously, interact with enterprise environments, and make real-time decisions.

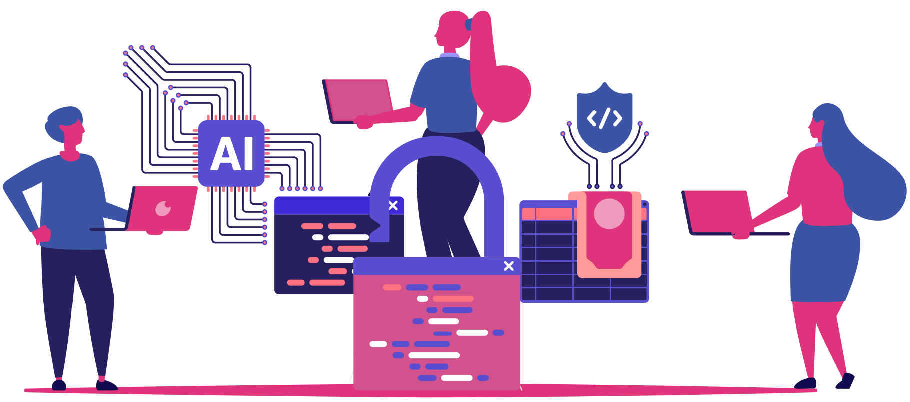

4.1 Low accuracy and reliability of AI outputs

Reliability remains uneven across real-world tasks; systems perform well on benchmarks yet fail in practical enterprise settings.

AI systems exhibit a "jagged frontier," solving complex problems yet failing simple tasks; improvements in one dimension (e.g., safety) can degrade others (e.g., accuracy).

Implication: AI cannot yet deliver consistent, enterprise-grade reliability, limiting use in mission-critical workflows

Hallucination Frequency, By Grounding Approach

1. How frequently do AI systems in your organization produce incorrect, fabricated, or misleading outputs (so-called "hallucinations")? &

2. What technology or approach does your organization predominantly use to make AI / LLM answers more reliable and grounded in your own data?

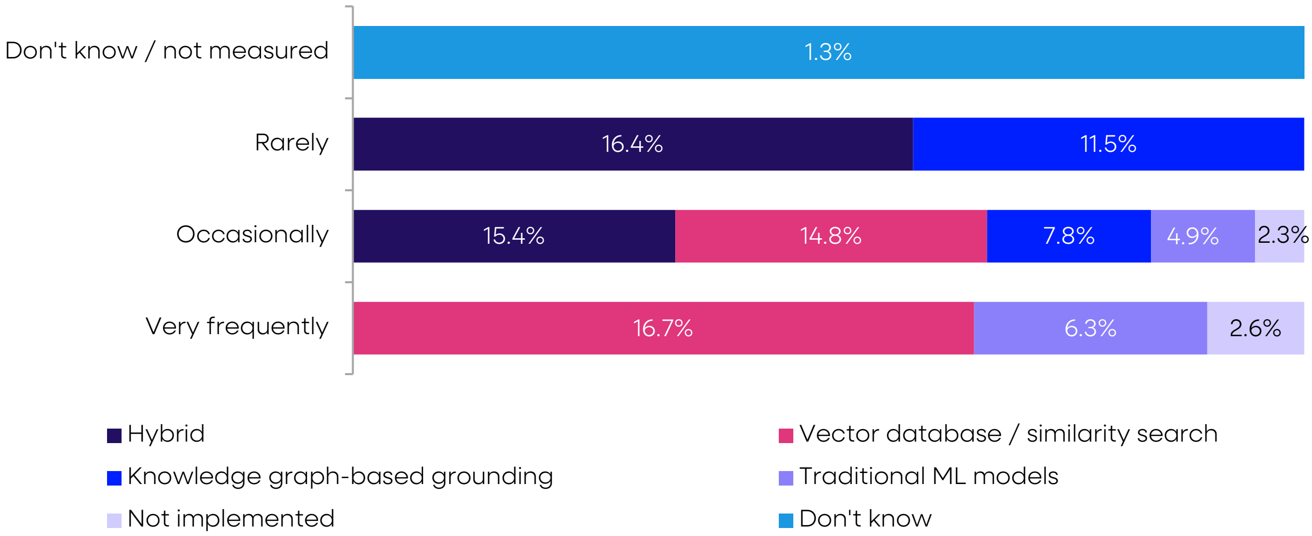

Figure 4.1 – Hallucination Frequency, By Grounding Approach: None of the organizations using knowledge graph-based grounding approach have reported frequent hallucinations

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

4.2 Lack of context awareness and learning capability

- Current systems struggle to retain context, adapt to enterprise-specific processes, and improve over time.
- Lack of contextual learning and brittle workflows are key reasons AI pilots fail to scale; a large share of organizations report negligible return on GenAI investment, highlighting the gap between capability and applicability.
- **Implication:** AI remains generic rather than enterprise-aware, leading to poor alignment with real business workflows.

4.3 Fragmented data ecosystems and integration challenges

- Disconnected systems, siloed data, and limited interoperability create a major barrier. Organizations manage ~957 applications on average with only ~27% connected; half of AI agents operate in isolation; 64% of IT leaders cite integration as a key barrier.
- **Implication:** Without unified data and integration layers, AI systems remain isolated and unable to scale.

Importance of Cross–System Knowledge Sharing, By Industry

How important is the ability to share knowledge and integrate data across different systems, teams and partners for scaling AI in your organization?

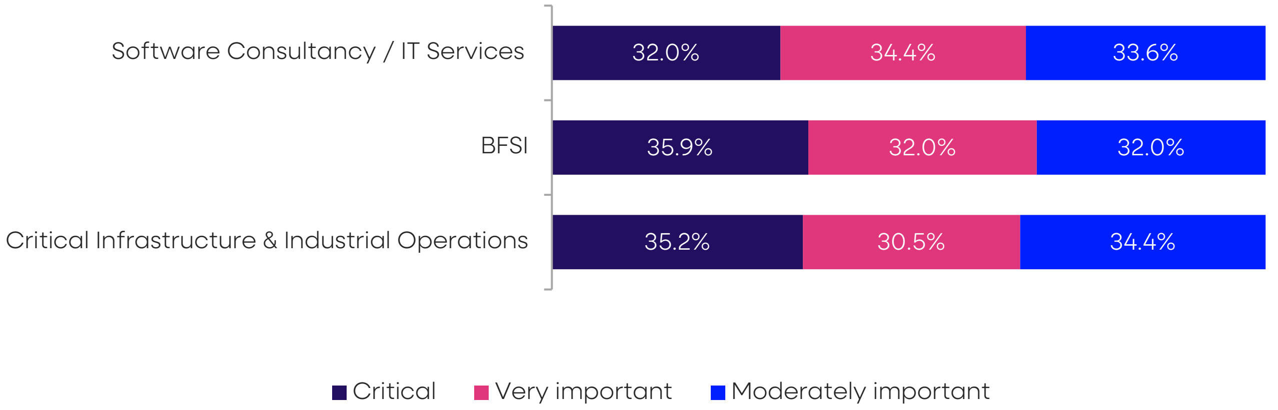

Figure 4.2 – Importance of Cross–System Knowledge Sharing, By Industry: More than 60% of respondents across all the 3 industries has rated cross-system knowledge sharing as critical & very important.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

Biggest Barrier to Sharing Data & Knowledge, By Industry

What are the biggest barriers to sharing data and knowledge across systems in your organization?

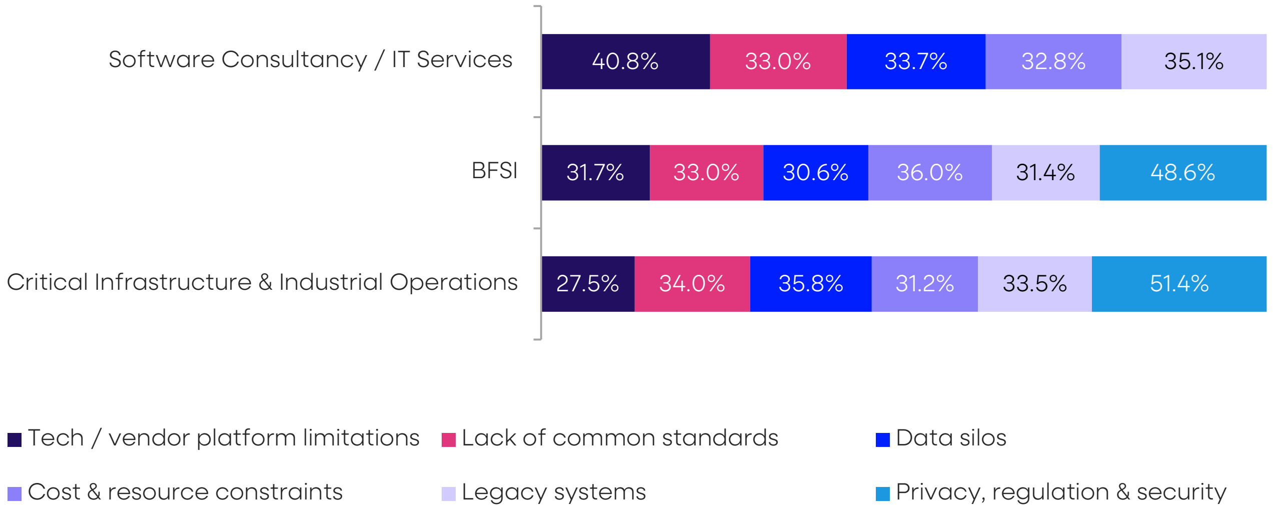

Figure 4.3 – Biggest Barrier to Sharing Data & Knowledge, By Industry : Tech / vendor platform limitations is the biggest barrier to sharing data & knowledge in the Software Consultancy / IT Services industry, while privacy, regulation & security is the biggest barrier within the BFSI and Critical Infrastructure & Industrial Operations industries.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

4.4 Low trust, governance gaps, and rising risk exposure

- As systems become autonomous, concerns around trust, governance, and risk intensify; inaccuracy and cybersecurity risks are top barriers to scaling agentic AI.
- Higher AI-governance maturity correlates with measurable gains (~23% in security metrics; ~20% in employee outcomes; ~18% improvement in adoption driven by greater robustness and confidence).
- **Implication:** Lack of trust limits adoption, particularly for high-stakes, decision-making use cases

Importance of Deterministic AI Behavior, By Industry

How important is it for your AI systems to be deterministic — that is, to give the same, traceable answer every time for the same input — compared to probabilistic AI that may vary?

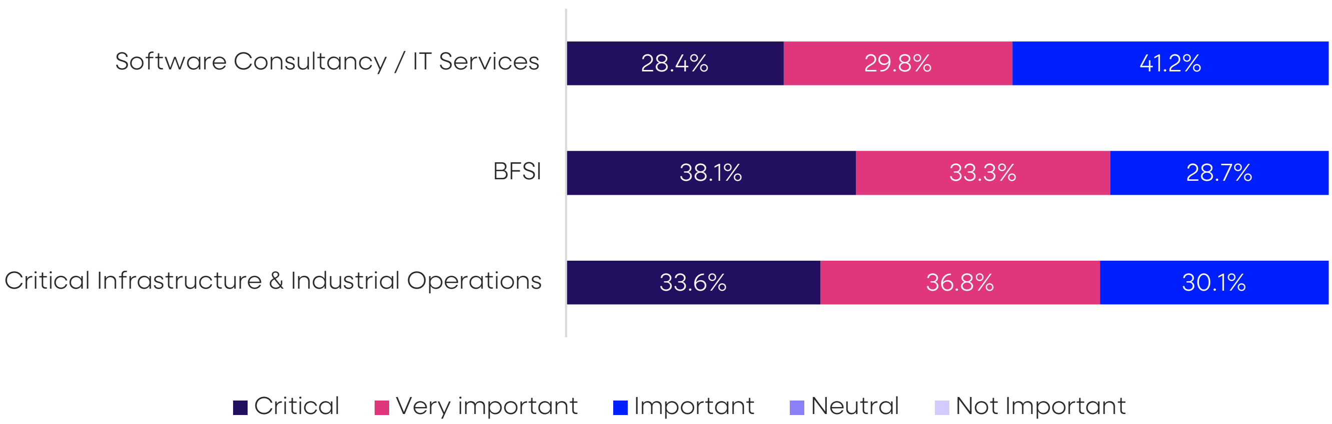

Figure 4.4 – Importance of Deterministic AI Behavior, By Industry: No respondent rated deterministic behavior as not important; it is "critical" or "very important" for most.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

Traceability of AI Output to Source Data, By Industry

To what extent can your AI outputs be traced back to the source data and the logic that produced them?

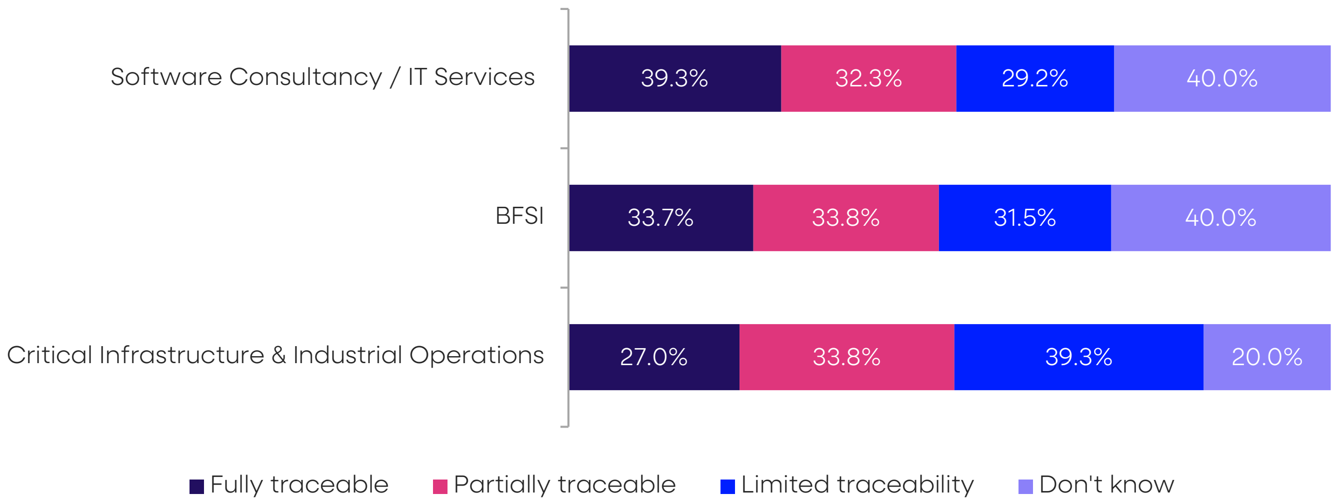

Figure 4.5 – Traceability of AI Output to Source Data, By Industry: Across all three industries, almost one third organizations cannot measure traceability at all.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

Measures That Would Improve Trust and Explainability The Most

Which measures would most improve trust and explainability in your AI systems?

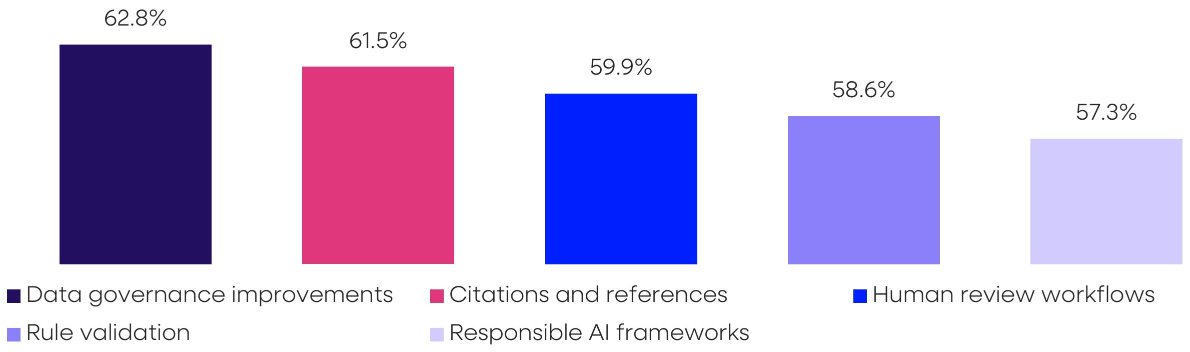

Figure 4.6 – Measures That Would Improve Trust and Explainability The Most: Data governance, citations and references, and human review workflows lead the measures to improve trust and explainability the most

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

4.5 Limited scalability from pilots to enterprise deployment

- While 80%+ of organizations explore AI, only ~40% reach deployment and ~5% achieve production-scale success-owing to misalignment with workflows, lack of integration, and inability to evolve with feedback.
- **Implication:** Enterprises face a "pilot-to-production gap."

4.6 Architectural limitations and lack of unified orchestration

- Current architectures are not designed for autonomous, multi-agent systems across distributed environments; only ~54% have centralized governance frameworks.
- **Implication:** AI transformation requires architectural evolution toward integrated orchestration and governance layers.

4.7 Evaluation, transparency, and capability-system misalignment

- As per the Stanford HAI AI Index 2026 [1], evaluation methods are not keeping pace—benchmarks are saturating and are often unrepresentative of real-world enterprise performance.
- Transparency is declining (limited disclosure on training data, architecture, and parameters), reducing visibility into how systems make decisions.
- A widening gap exists between AI capabilities and supporting systems (governance, evaluation, data infrastructure), limiting operationalization at scale.
- **Implication:** Addressing this imbalance requires strengthening foundational systems—governance, evaluation, and data infrastructure.

How Enterprises Evaluate/Benchmark AI Quality

How does your organization currently evaluate or benchmark the quality of AI outputs?

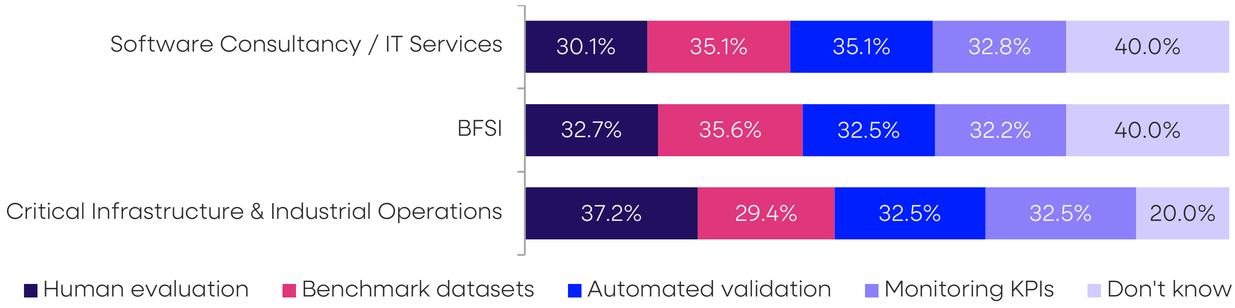

Figure 4.7 – How Enterprises Evaluate/Benchmark AI Quality: Benchmark datasets evaluation is the most prominent in Software Consultancy / IT Services and BFSI industry, whereas human evaluation is the preferred approach in Critical Infrastructure & Industrial Operations industry.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

5. The Core Problem and the Market's Responses

The trends and challenges set out so far converge on a single underlying question. Before examining individual architectures in detail, it is worth stating that core problem explicitly, setting out the main ways the market is responding to it.

The problem.

Enterprises can adopt AI but cannot reliably scale it. Models are capable; what they lack is governed, persistent, machine-executable context. Agents cannot consistently agree on what "revenue" or "customer" means, cannot find the authoritative data source, cannot reason deterministically across fragmented systems, and cannot show why they produced an answer–so outputs vary, hallucinate, and fail audit.

Concrete examples.

A data agent asked "what was revenue growth last quarter?" fails because the revenue definition lives in a stale YAML file, the right table is ambiguous (fct_revenue vs. mv_revenue_monthly), and fiscal-quarter conventions are not normalized. A compliance agent cannot trace which rule, evidence and control produced its answer. A maintenance agent cannot tell how a past decision unfolded because no system captured the inputs, policies applied and human overrides.

5.1 How the market is responding to this problem

Three architectural responses dominate the conversation. They are not equivalent, and the paper will treat them in turn.

- **The Context Layer – a memory theme.** This response is principally about how agents work with and retain context: long-term memory, awareness, and the situational state that surrounds a decision. Graphwise frames its contribution here as a "context graph."

- **The BI Semantic Layer – a metrics theme.** The thin BI/OSI camp throws around the term "semantic layer" in the Business–Intelligence sense–metric and analytics definitions, exemplified by the Open Semantic Interchange (OSI) initiative [6]. This is a thin, single–system construct and an unsettled concept that is not, on its own, capable of building robust, scalable, intelligent agentic systems.
- **The Semantic Backbone – a standards–oriented theme.** A more technical, standards–oriented camp approaches the same need from ontologies, formal semantics, and governance. This is effectively the transition toward the Semantic Backbone–the more advanced approach, and the concept the analysis develops from here. Graphwise (graphwise.ai) is a leading example of this camp.

Because "semantic layer" is not a standardized term, the same words describe very different things–which is itself a source of confusion the reader should hold in mind from the outset.

Where The Enterprise "Knowledge Layer" Lives Today

Where does the "knowledge layer" of your AI systems (definitions, business rules, ontologies, the model of your domain) primarily live today?

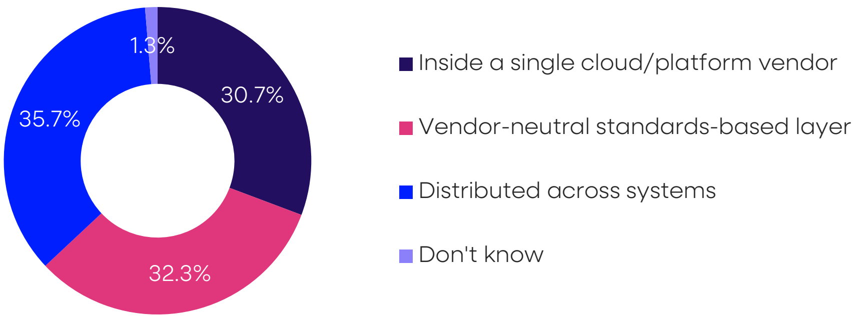

Figure 5.1 – Where the enterprise "knowledge layer" lives today: 35.7% of responding organizations have stated that the knowledge layer is distributed across systems followed by vendor-neutral standards-based layer.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

5.2 What this paper will analyze next

Having mapped the responses, the paper analyzes the pros, cons and concrete examples of each: the Context Layer and the context–graph idea; the BI/OSI semantic layer, including a gap analysis of why the thin interpretation does not solve the problem above; a consolidated gap analysis across both; and the Semantic Backbone.

6. Industry Perspective on the Context Layer

As enterprise AI evolves toward agentic systems, industry perspectives increasingly converge on the need for a dedicated context layer–an architectural component that lets AI operate with awareness of enterprise data, processes, and semantics. Crucially, companies see the "context layer" in different ways; this chapter first presents that range of views, then highlights the "context graph" introduced by Graphwise. The context layer is best understood as a memory theme–how agents retain and work with context–rather than as a substitute for the semantic foundation discussed later.

6.1 Defining the Context Layer

The context layer is a distinct architectural component that curates, integrates, and delivers the information AI agents require. It provides dynamic and persistent background knowledge so agents can interpret inputs, make informed decisions, and act in alignment with business objectives. It typically spans business entities and relationships, knowledge-graph representations, and decision traces-positioning it as the missing middle layer between raw data and AI models.

6.2 The range of market views on the context layer

There is no single, agreed definition. Across the market the same need is described as a “context OS,” “context engine,” “contextual data layer,” “ontology,” and more. The a16z article “Your Data Agents Need Context” (Cui & Li, 2026 [2]) argues that data and AI agents are essentially useless without the right context, and positions the context layer **as a superset of the semantic layer–**extending beyond metric definitions to canonical entities, identity resolution, business logic and workflows, governance and policies, and organizational “tribal knowledge.” Others (e.g., data-gravity platforms and dedicated context-layer companies) propose narrower, tool-bound versions. The common thread: tie an enterprise’s messy data together, add a contextual layer that helps agents understand business logic, and expose it to agents at run time.

This reflects a shift from data interpretation (BI) to decision execution (agentic AI): the context layer integrates structured and unstructured data, captures decision logic and workflows, and enables agent autonomy and reasoning across systems.

6.3 Graphwise’s “context graph”

Within this range of views, Graphwise has introduced the term “context graph” [4]. The context graph is an evolution of the knowledge graph, purpose-built for agentic AI. [4] Where a standard knowledge graph is a static “map” of facts, the context graph acts as a “dashcam” or “event clock” that captures dynamic reasoning: the situational variables, decision traces, procedural knowledge and continuous workflows that led to a given outcome.

It serves as long-term memory (LTM) for agents and directly tackles “context rot” by modelling temporal validity, causal chains of events and authority hierarchies. As a result, an agent can query exactly how a past decision unfolded-what inputs were used, which policies applied, where a human-in-the-loop override occurred, and the final outcome-understanding not just what happened, but how and why.

6.4 Core components and value of the Context Layer

- ✓ **Semantics –** organizes enterprise knowledge using ontologies and knowledge graphs; enables reasoning on meaning rather than raw data.
- ✓ **Operational state –** provides right-time data and system state from streaming and event-driven sources.
- ✓ **Provenance –** tracks data sources, decisions, and outcomes for auditability and compliance.

Together these let the context layer improve reliability (reducing hallucinations and noise), enable efficient knowledge retrieval, and support scalable agent orchestration-transforming AI from a stateless tool into a stateful, enterprise-aware system. Organizations that prioritize semantic, context-rich data report up to ~80% improvement in AI accuracy and up to ~60% reduction in operational costs.

6.5 Practical applicability and adoption considerations

The context layer is most applicable today in custom-built AI agents, enterprise-specific workflows, and high-value decision-making use cases. However, adoption is still evolving, with limited vendor standardization, integration challenges, and a need for organizational maturity. The context layer addresses memory and situational awareness-but it does not, by itself, supply the governed, deterministic semantic foundation that agentic systems also require.

7. Industry Perspective on the BI Semantic Layer

"Semantic layer" is not a standardized term, and vendors interpret it very differently. The discussion that follows uses it in a deliberately narrow sense: the "semantic layer" refers to the weak BI semantic layer–the Business-Intelligence construct popularized by the data and BI community and anchored to the Open Semantic Interchange (OSI) initiative [6]—and not the standards-based Semantic Backbone, which is a different construct altogether. The BI semantic layer is positioned by its proponents as a foundational enterprise capability–a strategic construct that defines how organizations structure, interpret, and operationalize meaning across data and applications. The urgency is real: in MarketsandMarkets' survey of 384 enterprise leaders, the most common way organizations keep the meaning of data consistent is still spreadsheets (42.4%) respondents–narrowly ahead of taxonomies and ontologies (41.9%), governance frameworks (39.6%), metadata catalogs (38.8%), and an enterprise knowledge graph (37.2%). Yet inconsistencies in the meaning of data remain frequent even where these mechanisms are in place, and they are most acute precisely where spreadsheets are the primary control (see Figures 7.2–7.3).

7.1 Role of the BI Semantic Layer

Proponents define the semantic layer as a unified, human- and machine-readable abstraction—a shared map of data, concepts, and processes. In this framing it acts as a bridge between data, processes, and organizational logic by enabling explicit modeling of relationships and hierarchies, standardization of business definitions, and transformation of raw data into contextualized knowledge. Some graph vendors push this further—using knowledge graphs to represent real-world entities and relationships for context-rich, machine-interpretable data—but, this is exactly the point at which the "thin" BI interpretation and the standards-based Backbone diverge.

The semantic layer is also positioned to support consistent interpretation (a single source of truth for meaning), governance and explainability (policy enforcement, traceability, standardized application of business rules), and AI orchestration (shared context, session memory, task state). Market activity reflects this: In November 2025, Salesforce completed its acquisition of Informatica (catalog, governance, MDM) [7], and in April 2025, KPMG acquired Metaphor's metadata-knowledge-graph platform—both supporting semantic-layer needs [8].

7.2 Benefits claimed for the BI Semantic Layer

- ✓ Reduces "AI debt" by replacing fragmented point-to-point integrations with a stable abstraction independent of tools and vendors.
- ✓ Enables interoperability by abstracting business meaning from underlying systems.
- ✓ Supports more coordinated AI by providing shared context across workflows.

7.3 The market split: two very different things called "semantic layer"

Because the term is not standardised, the market has split into two camps. One camp–the thin BI/OSI camp (metrics, analytics) and the OSI initiative–simply applies the label "semantic layer" to what is, in practice, a very thin layer: no schema mapping, no taxonomies and no vocabularies behind it; single–system and not scalable; and not sufficient for agentic orchestration. This thin interpretation is effectively "smoke and mirrors" and must be called out as such. A second set of (graph) vendors–Graphwise being a leading example–runs the same idea far more systematically on standards–based foundations, and has coined the Semantic Backbone, which genuinely fits agentic systems, orchestration and scaling. The figure below illustrates the thin, per–vendor BI interpretation.

The Semantic Layer Approach

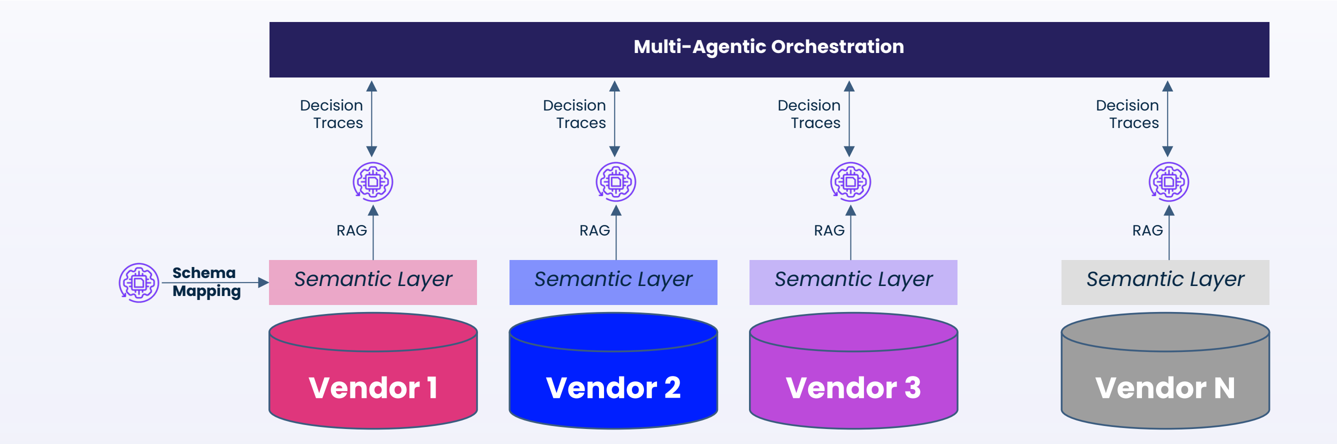

Figure 7.1 – The "thin" BI Semantic Layer approach: a separate semantic layer is bolted onto each vendor/system with only schema mapping and per–system RAG, leaving multi–agentic orchestration without a shared foundation.

7.4 Gap analysis: why the thin BI/OSI semantic layer does not solve the problem

A BI/OSI–style semantic layer does not resolve the fundamental problems that block agentic AI roll–out. Concretely, the public OSI specification (v0.1.1) shows why [6]:

- OSI is a YAML/SQL–centric model built around datasets, fields and metrics, with SQL dialects (ANSI SQL, Snowflake, MDX, Tableau, Databricks).
- Its "ai_context" is just free–text synonyms / instructions for LLMs.
- There are NO ontologies, NO controlled vocabularies/taxonomies, NO formal semantics (RDF/OWL/SKOS), NO reasoning and NO SHACL–style validation.

In other words, OSI is a BI metric–interchange format, not a true semantic standard–useful for consistent KPIs in a single BI context, but insufficient for agentic orchestration across systems. [6]

Without ontologies and enforceable logic the agent still lacks a governed definition of "revenue" (grounding gap); without formal reasoning it stays dependent on probabilistic LLM inference (reliability gap); being single–system and tool–bound, each new use case needs its own mappings (scalability gap); and with no validation or traceable reasoning paths it cannot satisfy audit (trust gap).

Mechanisms Used to Keep Meaning Consistent, By Industry

How does your organization currently make sure that the same business concept (e.g., "customer", "product", "contract") means the same thing across different systems and datasets?

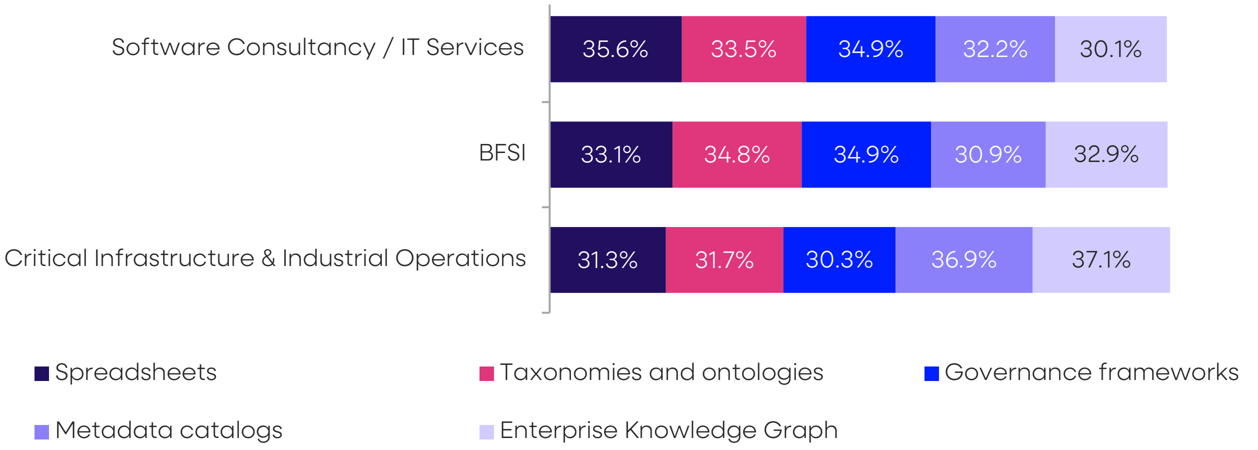

Figure 7.2 – Mechanisms Used to Keep Meaning Consistent, By Industry: Enterprise knowledge graphs remains the most preferred mechanism to keep meaning consistent in the Critical Infrastructure & Industrial Operations industry.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

Meaning Inconsistency, By Mechanism

1. How often do inconsistencies in the meaning of data (different definitions across systems, conflicting values, ambiguous terminology) negatively affect the quality of your AI or analytics outputs? &

2. How does your organization currently make sure that the same business concept (e.g., "customer", "product", "contract") means the same thing across different systems and datasets?

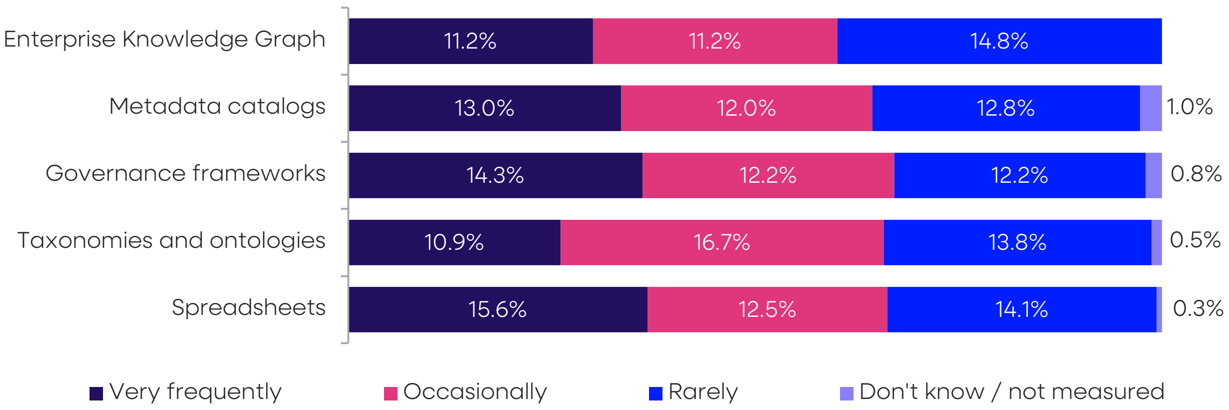

Figure 7.3 – Meaning Inconsistency, By Mechanism: Inconsistencies are most frequent where spreadsheets are the primary control and least frequent where an enterprise knowledge graph is used which is a clear evidence for ontology/KG–based semantic integrity over a thin BI/OSI layer.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

Maturity of Formal AI-Output Validation, By Industry

Does your organization have a way to validate that AI outputs follow business rules or definitions before they are shown to users?

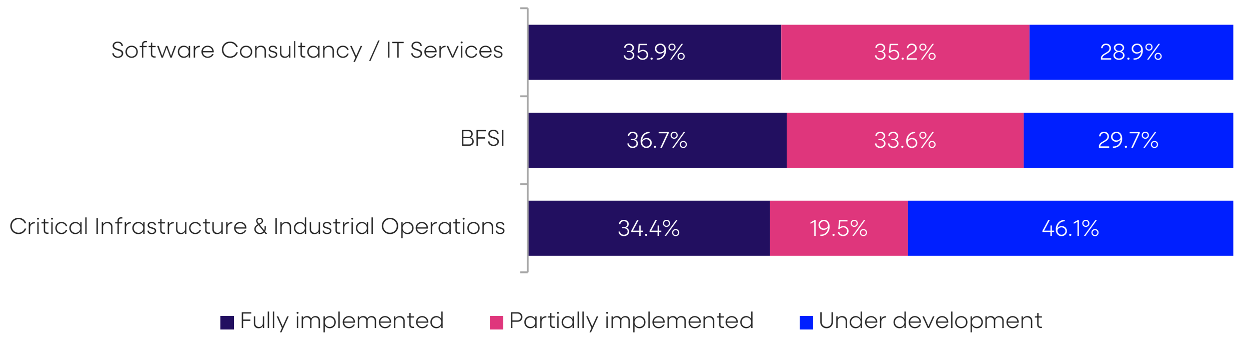

Figure 7.4 – Maturity of Formal AI-Output Validation, By Industry: Formal validation is still under development for a large share of organizations, reinforcing the need for built-in, standards-based validation (e.g. SHACL).

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

OSI in depth – why a BI/OSI semantic layer does not close the gaps. The Open Semantic Interchange (OSI) initiative is an industry effort, backed by BI and analytics vendors, to standardize how metric and analytics definitions are exchanged between tools. That is a useful goal for portable KPIs—but it is a metric-interchange contract, not a semantic foundation for autonomous systems. OSI leaves every gap open: it gives agents no governed, machine-executable definition of an entity such as “revenue” or “customer” (grounding gap); it carries no formal logic, so reasoning stays dependent on probabilistic LLM inference (reliability gap); it is single-system and tool-bound, so each new use case needs its own mappings (scalability gap); and it offers no validation or traceable reasoning paths, so outputs cannot satisfy audit (trust gap). OSI was never designed for agentic orchestration-there is no shared world model, no long-term memory, and no process or provenance layer for multi-agent systems. It is therefore a complement to, not a substitute for, the standards-based Semantic Backbone.

As “semantic layer” is not a standardized term, you find some agencies and companies the thin BI/OSI camp running it as a thin BI layer, while other standards-based camp-such as Graphwise-run it far more systematically on standards-based foundations. Once it is clear what the BI semantic layer is, how OSI contributes to the conceptual ambiguity around the term, and that this BI/OSI semantic-interchange approach is a weak, still-unsettled concept that cannot build real, robust, scalable, intelligent agentic systems, the natural next step is the standards-based answer. A true Semantic Backbone requires procedural knowledge, domain knowledge, expert-in-the-loop, schema/ontology mapping, standards-based representation, long-term memory and provenance-and the current industry shortcut (a simplified OSI-style layer plus a bespoke agent per system) does not work at enterprise scale.

8. Gap Analysis

8.1 From data access to autonomous AI systems

Enterprise AI architectures have evolved from traditional data stacks toward Context Layers and Semantic Layers to improve accessibility, interoperability, and analytics. Context layers enhance LLM performance by retrieving relevant enterprise data (e.g., RAG pipelines); semantic layers standardize definitions and provide abstraction, primarily for analytics. But as systems transition from assistive tools to autonomous, multi-agent systems, requirements shift fundamentally: from retrieval to reasoning, from static mappings to dynamic knowledge systems, and from human-driven workflows to autonomous decision-making. Many organizations remain stuck in a "prototype plateau."

8.2 Limitations of the Context Layer

- **Lack of deterministic reasoning –** outputs depend on probabilistic LLM inference; results vary across identical queries.
- **Query-dependent, fragmented context –** context is assembled per interaction; no persistent, reusable shared understanding.
- **Shallow grounding –** retrieved context is not governed by explicit business rules or ontologies.
- **Governance and traceability gaps –** limited auditability of outputs and reasoning paths.
- **Limited support for multi-agent systems –** no shared world model or long-term memory.

8.3 Limitations of the BI Semantic Layer

- **Thin abstraction ("bag of mappings") –** overlays that harmonize definitions rather than model deep domain knowledge. Enterprise-AI challenge: shallow grounding, weak contextual understanding, higher risk of incorrect interpretation.
- **Single-purpose design –** optimized for BI/reporting, not reuse across AI workflows. Enterprise-AI challenge: each use case needs separate mappings and tuning, raising cost and slowing scaling.
- **Absence of formal reasoning –** no native logical inference, multi-hop reasoning, or machine-executable rules. Enterprise-AI challenge: dependence on probabilistic LLM reasoning, weak explainability.
- **Limited semantic depth –** static definitions without process knowledge, decision flows, or temporal context. Enterprise-AI challenge: cannot understand how decisions evolve; weak long-term memory.
- **Static, non-evolving and tool-coupled –** does not adapt to changing enterprise state and is often tied to specific tools. Enterprise-AI challenge: outputs drift from current conditions; adoption fragments across systems.

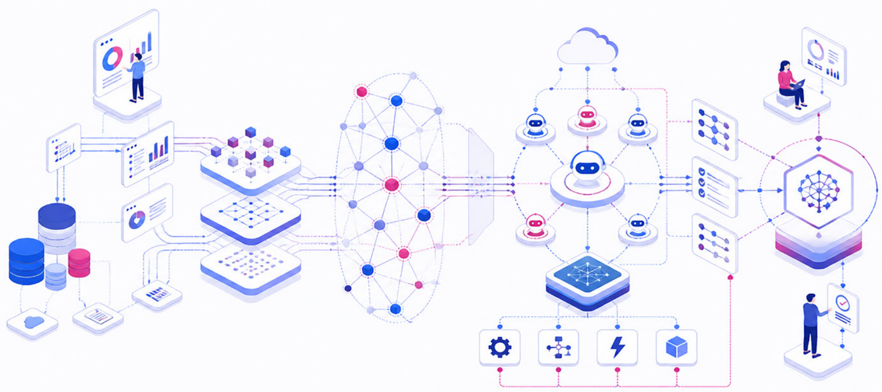

8.4 Core gaps in the context of Agentic AI

- **No shared world model –** inconsistent interpretations and actions.
- **Absence of long-term memory / context graphs –** no capture of decision history, workflows, or causal relationships.
- **Inability to enforce business logic deterministically –** logic embedded in prompts or code; fragile systems.
- **High hallucination risk –** reliance on probabilistic reasoning rather than governed knowledge.
- **Lack of explainability and auditability –** no traceable reasoning paths or lineage.
- **Scalability constraints –** each use case needs new pipelines, mappings, and prompt engineering.

8.5 Conclusion: The structural limitation of existing approaches

The Context Layer enhances data accessibility and the BI Semantic Layer improves interpretability, but both remain incremental solutions sitting on top of fragmented systems. Even in combination they do not provide a persistent, enterprise-wide knowledge foundation; fail to embed reasoning, governance, and explainability into the core architecture; and cannot support the coordination and scalability agentic AI requires. This creates a structural gap between AI capability and enterprise readiness—one that requires treating knowledge as infrastructure rather than as an overlay.

9. Introduction of the Semantic Backbone

The limitations above point to a fundamental architectural gap: current approaches improve access to data but do not provide a unified, governed, reasoning-capable foundation. The Semantic Backbone is that foundation.

A Semantic Backbone is not an extension of existing layers, nor merely a combination of context and semantic capabilities. As Graphwise defines it (graphwise.ai/fundamentals/what-is-a-semantic-backbone), it is a distinct architectural construct–a graph-based, enterprise-wide knowledge infrastructure (an Enterprise Knowledge Graph coupled with machine-readable governance) that acts as the "semantic nervous system" of the organization. It is therefore far more than a mapping layer: it is built as infrastructure rather than an overlay, serves as a tangible enterprise source of truth, and simultaneously powers multiple applications–domain knowledge hub, metadata hub, digital-twin model, and the contextual grounding layer for agentic AI–rather than a single function. [3]

A useful way to picture this is as a single enterprise knowledge graph presenting several faces at once. The same Backbone acts as a Knowledge Hub for precise retrieval and management of content; as a Semantic Layer that delivers unified views for analytics across siloed data; and as the semantic grounding for context graphs and Agentic AI. These are not separate products bolted together but distinct roles served by one governed foundation–which is precisely what distinguishes a Backbone from the thin, single-purpose layers examined earlier.

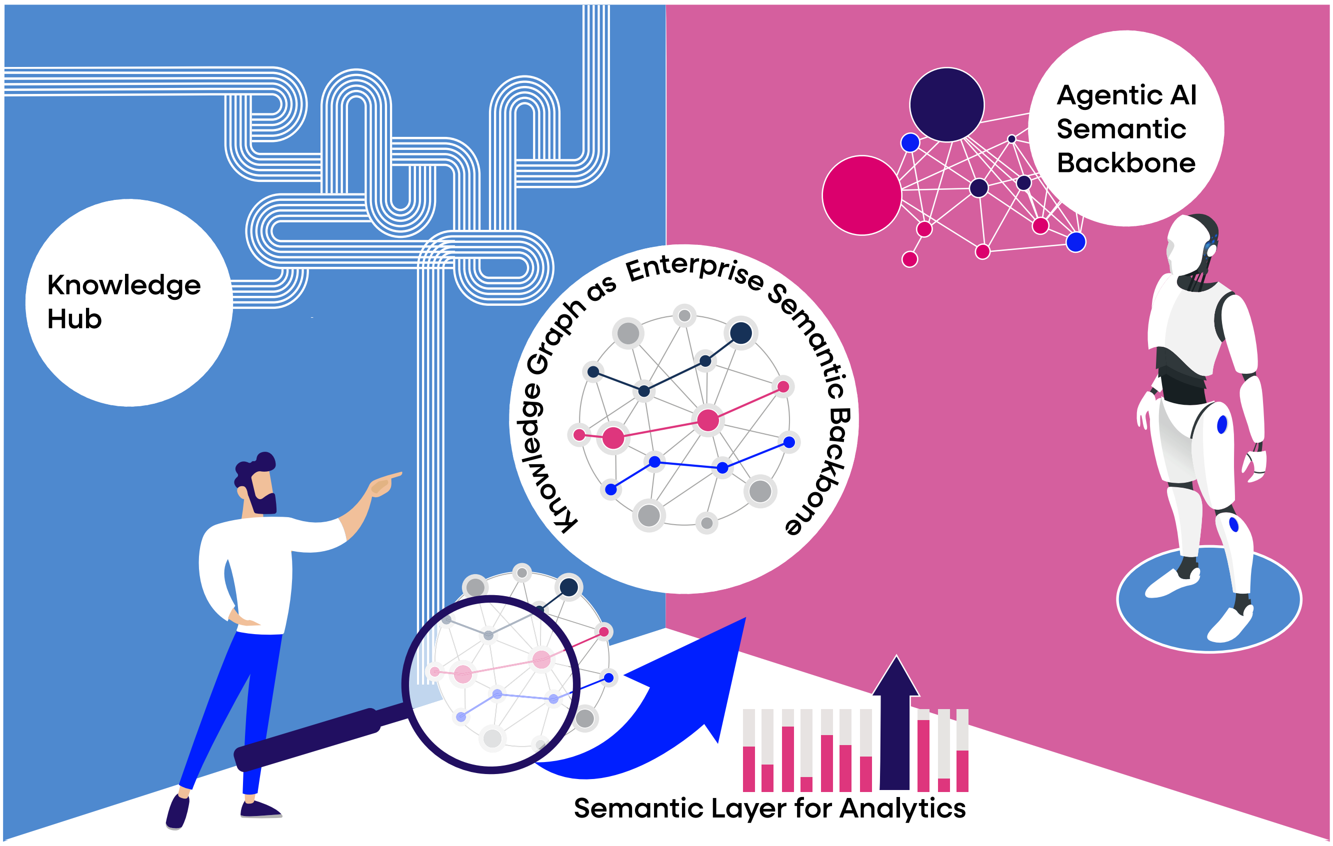

Figure 9.1 – The roles of the Enterprise Semantic Backbone: One enterprise knowledge graph simultaneously serving as a Knowledge Hub, a Semantic Layer for analytics, and the grounding layer for Agentic AI.

9.1 Fundamental characteristics of the Semantic Backbone

Graphwise frames the Backbone around four fundamental characteristics.

- ✓ **Deterministic, machine–executable logic –** business logic encoded in ontologies and executed via rule–based inference, enabling explainable multi–hop reasoning independent of LLM variability.
- ✓ **Knowledge as a trusted, interoperable asset –** entities, relationships and rules explicitly modelled and governed; a persistent, reusable foundation with validation, traceability, and policy enforcement embedded.
- ✓ **Consistent, unified access to data and process knowledge –** fragmented data integrated into one semantic model; shared world models, long–term memory, and orchestration for multi–agent systems.
- ✓ **Standards–based, vendor–lock–in–free representation –** enterprise logic decoupled from LLMs and applications for flexibility and portability.

9.2 RDF–based ontologies vs. LPG vs. proprietary graphs

The case is reinforced by comparing graph technologies. RDF–based ontology platforms—built on open W3C standards (RDF, OWL, SKOS–support data validation, provenance, multi–hop reasoning, precise semantic retrieval, domain/context awareness, and cross–platform interoperability. LPG graphs offer good analytics but limited governance metadata and standards support and require an additional semantic layer; proprietary graphs risk vendor lock–in and lower portability. Graphwise GraphDB, an RDF–based platform, reports strong benchmark performance and notes that ontologies reduce inaccurate answers by ~2× versus schema–less GraphRAG, with expert–in–the–loop tooling, explainable reasoning/provenance, and an all–in–one graph–AI platform (native graph engine, knowledge–modelling toolkit, low–code ingestion and GraphRAG).

9.3 Layered architecture of the Semantic Backbone

The Semantic Backbone is a multi–layered architecture whose six layers collectively enable semantic grounding, reasoning, and orchestration: a Vocabulary layer (taxonomies, master data, controlled vocabularies) that establishes a common language; a Logic layer (rule–based inference, data validation, multi–hop reasoning) that turns probabilistic interpretation into deterministic execution; a Data Linking layer (data catalogues, linked data, asset descriptions) that creates unified access across silos; an Intelligence layer (GraphRAG, semantic analytics, knowledge extraction) that bridges graphs with AI models; a Process Knowledge layer that captures decision flows, workflows, and event traces; and a Multi–Agent layer (shared world models, long–term memory, agent coordination) that supports orchestration of autonomous agents.

LAYER	INCLUDES	ROLE
Vocabulary Defines standardized identifiers, entities, and terminology	▪ Taxonomies ▪ master data ▪ controlled vocabularies	▪ Establishes a common language across the enterprise ▪ Eliminates semantic ambiguity across systems and agents
Logic Encodes business rules and domain knowledge using ontologies	▪ Rule-based inference ▪ Data validation ▪ multi-hop reasoning	▪ Transforms AI from probabilistic interpretation to deterministic execution of logic ▪ Ensures consistency and correctness in outputs
Data Linking Connects fragmented enterprise data sources through semantic metadata	▪ data catalogues ▪ linked data models ▪ asset descriptions	▪ Creates a single, unified access layer across silos ▪ Ensures consistent meaning across all data sources
Intelligence Enables precise retrieval, reasoning, and analytics	▪ GraphRAG ▪ semantic analytics ▪ knowledge extraction	▪ Bridges knowledge graphs with AI models ▪ Delivers accurate, context-aware, and explainable AI outputs
Process Knowledge Captures decision flows, workflows, and event traces	▪ a "dynamic memory" of enterprise operations	▪ Provides temporal and procedural context ▪ Enables systems to understand how and why decisions occur, not just static facts
Multi-Agent AI Supports orchestration of autonomous AI agents	▪ shared world models ▪ long-term memory ▪ agent coordination frameworks	▪ Enables collaboration across multiple AI agents ▪ Ensures consistency and alignment in decision-making

Figure 9.2 — Layered Architecture

Semantic Backbone Conceptual View

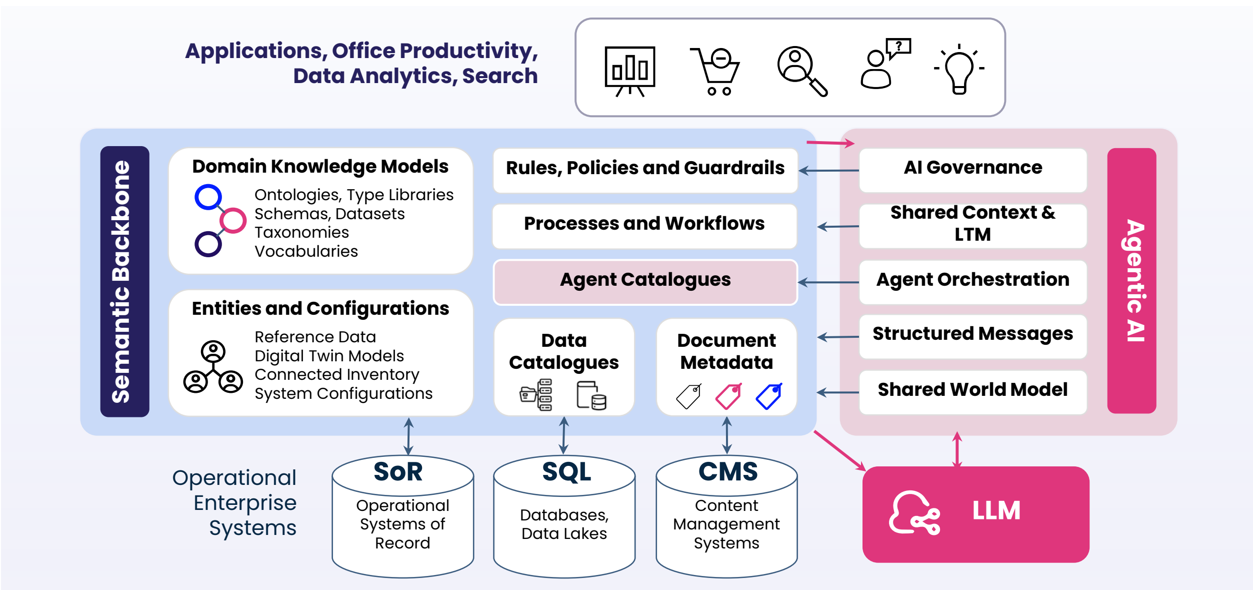

Figure 9.3 - Semantic Backbone Conceptual View

9.4 The Trust Layer: Enabling Deterministic and Governed AI

The Trust Layer should be understood as a **cross-cutting binder across the Semantic Backbone–not an additional seventh layer.** It connects and governs the Backbone's key enterprise roles–Knowledge Hub, Semantic Data Layer, Digital Twin, Data Portal, Metadata Hub, and Semantic Grounding–so that knowledge, data, process context, and AI outputs remain consistent, explainable, and reusable. In the Graphwise framing, the Trust Layer ties these roles together and delivers measurable business outcomes: a **70–90% reduction in hallucinations** through grounded reference knowledge; an **80% reduction in manual effort** by improving accurate retrieval for BI, CMS and RAG workflows; and **100% explainability** through governed validation and traceable reasoning paths–reframing trust as an embedded enterprise capability that helps organizations move from prototypes to governed, durable AI systems

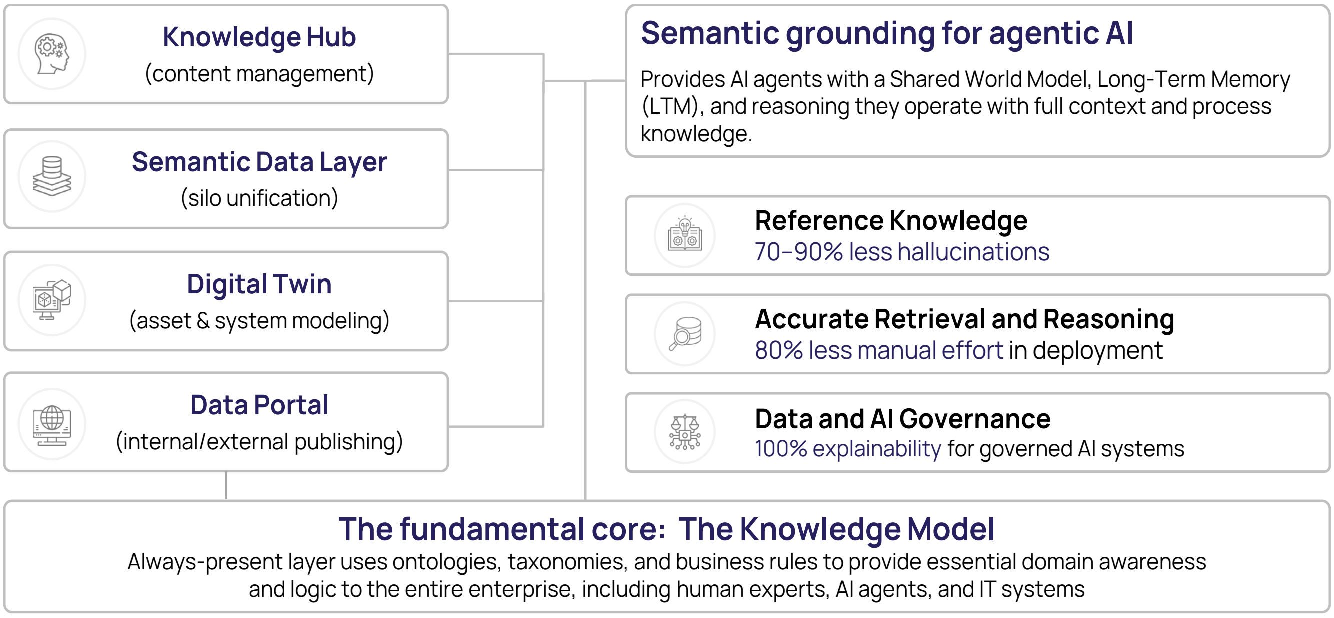

Understood this way, the Backbone is best read as knowledge graphs functioning as critical enterprise infrastructure—providing reference knowledge for humans, AI agents, and IT systems; accurate data retrieval for analytics (BI), search (CMS), and AI (RAG and context graphs); and governance of data and AI through semantic validation and business-logic enforcement. On that foundation, a single Backbone can serve a set of complementary roles across the enterprise:

- **Knowledge model –** domain awareness and business logic, expressed through ontologies, taxonomies, and business rules.
- **Knowledge hub –** precise retrieval and management of content enriched with semantic metadata.
- **Semantic layer –** unified views for analytics that consolidate fragmented, siloed data into a scalable data fabric.
- **Digital twin model –** configuration awareness for analytics through semantic models of devices, assets, and distributed systems.
- **Metadata hub –** publishing of master data and other data products for internal and external consumption.
- **Semantic grounding –** a reliable foundation for context graphs and Agentic AI.

Key components of the Trust Layer: deterministic reasoning (logic executed through ontologies and rules, not probabilistic inference); built-in governance and validation (semantic constraints (via SHACL) enforce data quality, structural integrity, and business-rule compliance, preventing incorrect or incomplete data from entering AI workflows); and explainability and traceability (every output traceable to source data, applied rules, and reasoning paths), enabling full auditability for regulated environments and 100% explainability.

The Semantic Backbone Approach

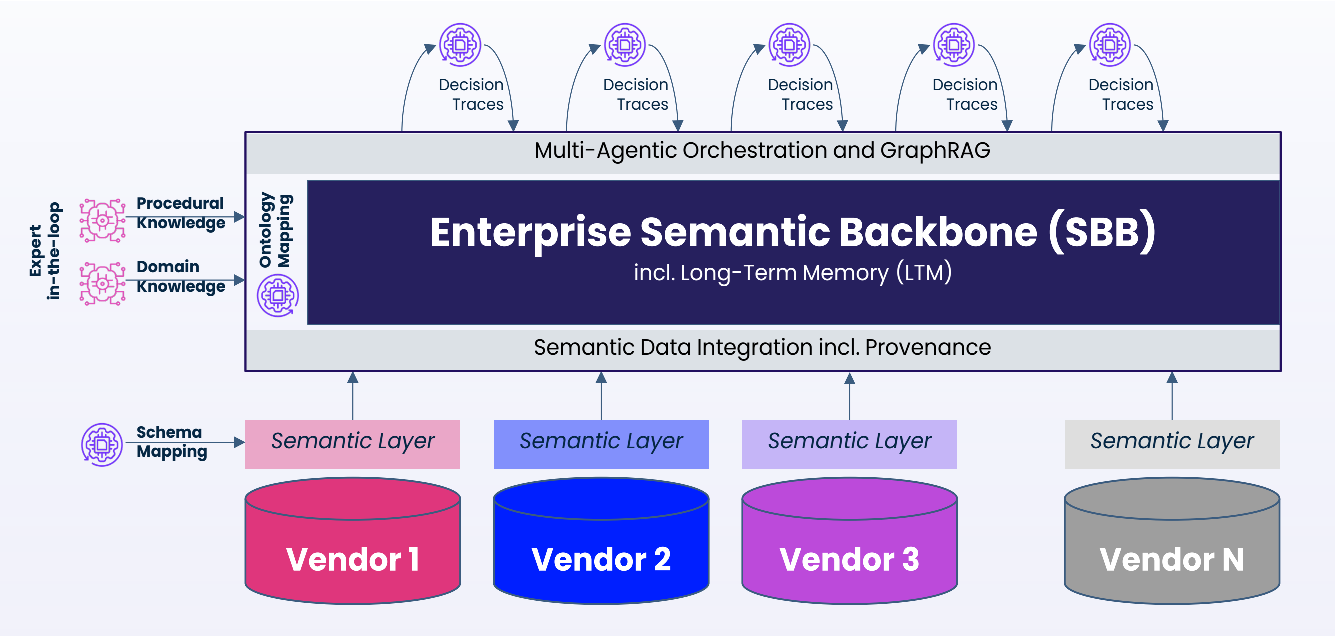

Figure 9.4 — The Semantic Backbone approach: Thin per-vendor semantic layers feed a single Enterprise Semantic Backbone (with long-term memory), unified by ontology mapping, expert-in-the-loop, procedural/domain knowledge and provenance, driving multi-agentic orchestration and GraphRAG.

9.5 Role of the Semantic Backbone in addressing the identified problems

- ✓ **Probabilistic → deterministic reasoning:** the Logic + Trust layers eliminate reliance on LLM inference for business logic.
- ✓ **Fragmented → unified knowledge:** Data-Linking + Vocabulary layers create a single source of truth.
- ✓ **Static → dynamic context:** the Process-Knowledge layer (context graphs) adds temporal and operational awareness.
- ✓ **Isolated agents → coordinated systems:** the Multi-Agent layer provides shared world models and orchestration.
- ✓ **Black-box → governed AI:** the Trust Layer ensures explainability, validation, and compliance.
- ✓ **Point solutions → scalable infrastructure:** the Backbone is a reusable foundation across use cases.

Importance of Vendor Independence Architecture, By Industry

How important is it for your AI and data architecture to remain independent of any single technology vendor (so that you can switch components without losing your knowledge model and data assets)?

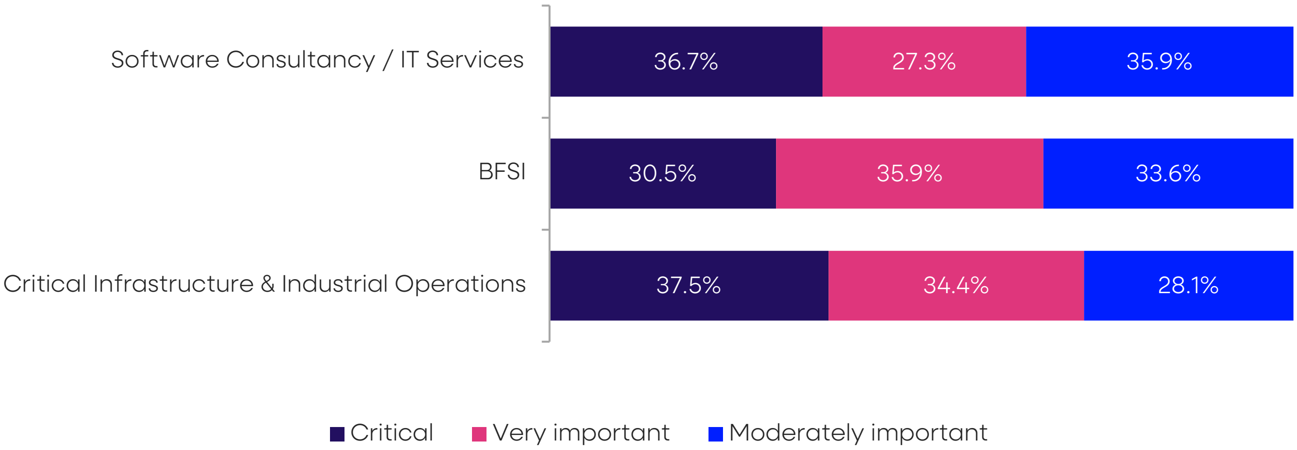

Figure 9.5 – Importance of Vendor Independence Architecture, By Industry: Vendor independence is rated critical or very important by the large majority–directly supporting a standards–based, lock–in–free Semantic Backbone.

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

When asked whether single-vendor dependence had ever limited their AI or data strategy, respondents described concrete, lived constraints. Thematic coding of all 384 responses surfaced nine recurring themes; the most prevalent were migration risk / switching costs (14.1%), cloud / platform lock-in (13.5%), portability / interoperability limits (12.0%), and vendor dependency as a gatekeeper (12.0%).

"Once our pipelines and models were built around one cloud ecosystem, moving to another environment became a major redesign exercise. The convenience of the initial setup turned into dependence because so many services and workflows were tied to that stack."

— Director, Software Consultancy / IT Services

"Our schemas and business definitions were effectively trapped inside the platform's structure. We could not reuse them across other tools without significant translation work, which limited the value of the semantic layer beyond the original implementation."

— C-Level Executive, Software Consultancy / IT Services

"Over time, the vendor became the gatekeeper for how fast we could change the data and AI stack. We could define the business need internally, but execution still depended on what the supplier supported."

— C-Level Executive, Critical Infrastructure & Industrial Operations

"The spend slowly shifted from innovation into platform maintenance. When we needed more storage, more environments, or more users, the commercial model became less attractive, and that made it difficult to build a broader AI strategy around one supplier."

— C-Level Executive, BFSI

"The migration risk was high because much of our process logic, mappings, and controls sat inside one ecosystem. Even considering a move meant a long business case, and the team often preferred to stay with the known platform rather than absorb that disruption."

— Vice President, Critical Infrastructure & Industrial Operations

9.6 Conclusion

The Semantic Backbone represents a fundamental shift—from fragmented layers and probabilistic systems to a unified, knowledge-driven, governed infrastructure that combines formal semantics and deterministic reasoning, unified data and context, multi-agent coordination, and embedded governance. Presented against the problem set rather than asserted, it is the response that most directly enables organizations to move beyond the prototype plateau toward production-grade, explainable, scalable AI.

10. Semantic Backbone vs. BI Semantic Layer

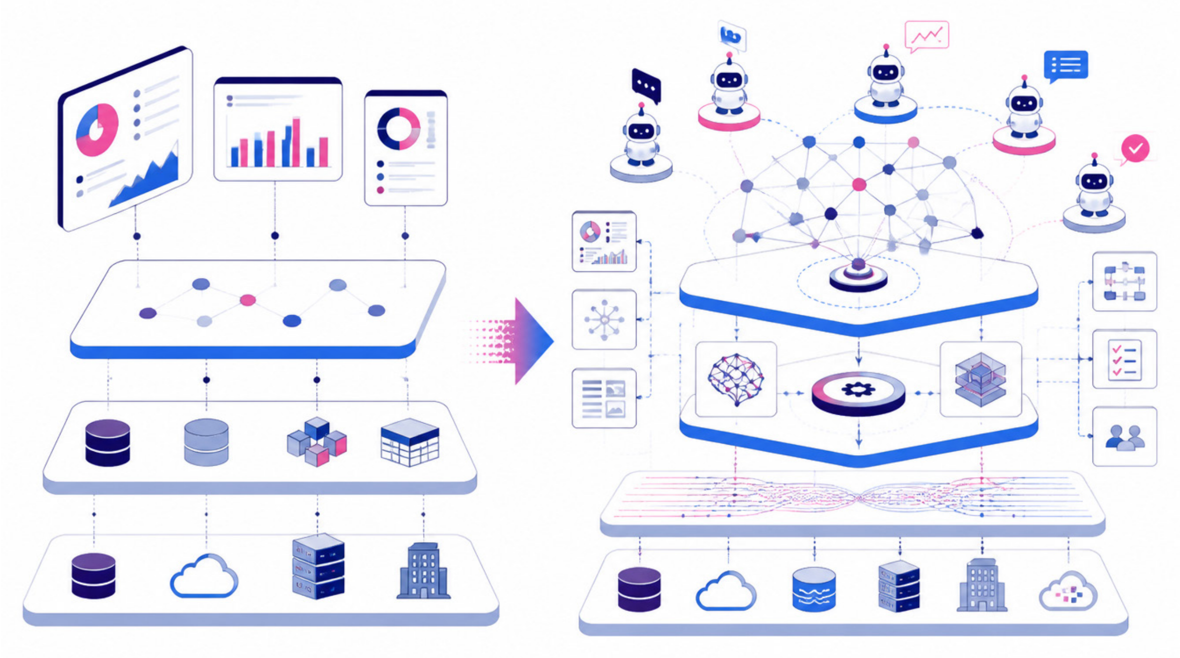

The distinction between a Semantic Backbone and a BI semantic layer is often referenced but rarely structured, causing confusion in architecture design. A typical semantic layer is "a thin layer on top… a pile of mappings," whereas the Semantic Backbone is developed as infrastructure and serves as a source of truth for multiple downstream applications. The tables below compare them across conceptual, architectural, and functional dimensions.

10.1 Conceptual differences

Aspect	BI Semantic Layer	Semantic Backbone
Core concept	Data abstraction layer	Enterprise knowledge infrastructure
Primary objective	Standardize data definitions for analytics	Establish a unified, governed understanding of enterprise knowledge
Design intent	Support specific use cases (e.g., BI)	Serve as a foundational, reusable enterprise asset
View of data	Data-centric (fields, metrics)	Knowledge-centric (entities, relationships, logic)
Role in enterprise	Supporting layer	Core system ("semantic nervous system")

10.2 Architectural differences

Aspect	BI Semantic Layer	Semantic Backbone
Position in architecture	Thin overlay on data systems	Foundational layer embedded across enterprise systems
Structure	Flat mappings and metadata	Multi–layered (vocabulary, logic, data linking, intelligence, process, agent coordination)
Knowledge representation	Limited schemas and mappings	Rich ontologies, taxonomies, knowledge graphs
Scalability model	Use–case specific, requires rework	Reusable infrastructure across applications
Interoperability	Often tool–dependent	Open standards (RDF, OWL, SKOS) enabling cross–system reuse
Context handling	Static and query–specific	Dynamic (context graphs, workflows, event traces)

10.3 Functional differences

Aspect	BI Semantic Layer	Semantic Backbone
Reasoning capability	No native reasoning; relies on queries/LLMs	Deterministic, rule–based, multi–hop inference
Grounding of AI	Partial, context–dependent	Deep grounding using governed knowledge
Support for agentic AI	Not designed for multi–agent systems	Shared world models, long–term memory, orchestration
Governance & validation	External or limited	Built–in semantic validation and rule enforce-ment
Explainability	Limited traceability	Fully explainable with traceable reasoning paths
Handling of hallucinations	High dependency on LLM be-haviour	Reduced via deterministic logic and grounded knowledge
Trust & reliability	Probabilistic and variable	Deterministic, governed, audit–ready

10.4 Conclusion

The BI Semantic Layer and the Semantic Backbone are not incremental variations of the same concept—they are two fundamentally different paradigms. The semantic layer suits data standardization and analytics; the Semantic Backbone is designed for enterprise–scale, trustworthy, autonomous AI through deep grounding, deterministic reasoning, built–in governance, and multi–agent orchestration

11. Use Cases

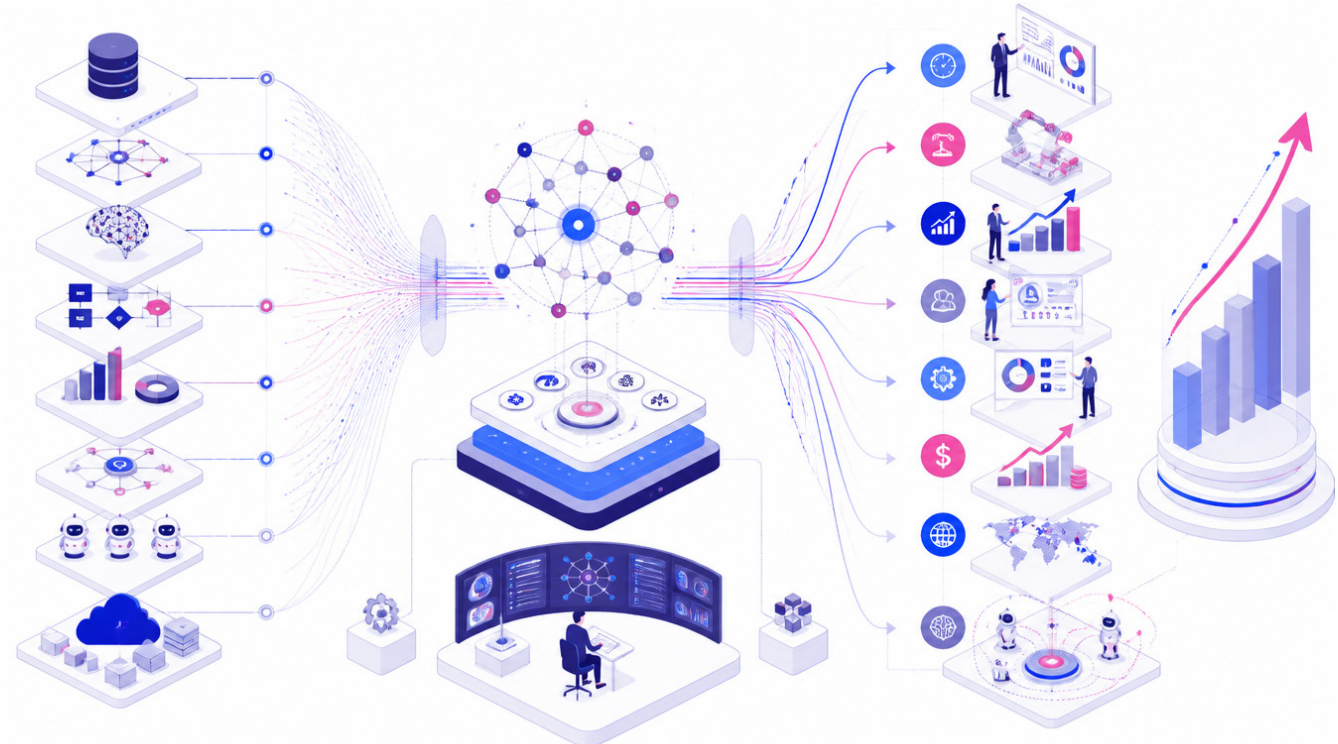

11.1 From Capability to Business Impact

The Semantic Backbone is designed as a production-grade foundation that directly enables high-value enterprise use cases. By providing semantic grounding (knowledge graphs, ontologies), deterministic reasoning (logic layer), unified data integration (data-linking layer), dynamic context (process-knowledge layer), and multi-agent orchestration (agent layer), it transforms isolated AI experiments into scalable, enterprise-wide solutions. The portfolio comprises Enterprise Knowledge Management, Technical Knowledge Management, the Semantic Digital Twin, Compliance Intelligence, Scientific Knowledge Management, and the Semantic Data Hub. [5]

11.2 Enterprise Knowledge Management

Enterprise knowledge is fragmented across documents, systems, and teams—difficult to access, inconsistent across functions, and hard to reuse. The Backbone addresses this: Vocabulary + Logic layers standardize terminology and encode domain knowledge; the Data-Linking layer integrates structured and unstructured data into a unified knowledge graph; the Intelligence layer (GraphRAG) enables accurate retrieval and reasoning; and the Trust Layer ensures validated, governed knowledge. At the business level this turns institutional knowledge into a reusable, governed asset-cutting search time, improving answer consistency, and accelerating onboarding—so knowledge management shifts from document storage to intelligent knowledge systems.

11.3 Technical Knowledge Management

Technical knowledge management turns fragmented technical content-documents, ticketing systems, engineering repositories, content platforms-into a reusable operational asset, addressing the whole content lifecycle rather than just storage and retrieval. (Graphwise: Technical Knowledge Management) Concrete capabilities include automatic preprocessing and chunking into a structured content graph, automated taxonomy-based tagging, ontology-driven content extension, and unified search across sources.

- ✓ Sandvik (engineering/mining) used Graphwise GraphDB with RWS, Kaleidoscope, and Ninefeb to build Smartmate, giving technicians and trainers one place to access equipment information and improving search, navigation, and contextual troubleshooting.

- ✓ EY integrated Graphwise into its Knowledge Discovery platform, improving discoverability and retrieval, reducing manual tagging, and delivering a 50–60% improvement in search performance.

The Data–Linking layer connects technical systems; the Process–Knowledge layer captures workflows and decision traces; context graphs retrieve similar past incidents; the Multi–Agent layer automates support workflows. At the business level this shortens troubleshooting cycles, reduces repeat incidents, and protects intellectual capital that would otherwise stay trapped in individual teams.

11.4 Semantic Digital Twin

A Semantic Digital Twin moves beyond static, table–based asset models by adding context, relationships, and reasoning–so organizations see dependencies clearly, run faster impact analysis, reduce false positives, and decide with business context. (Graphwise: Semantic Digital Twins) A natural–language semantic interface lets engineers query, for example, a power–grid digital twin without code; the Brick metadata schema is a concrete open–standard example that also strengthens the standards–based, vendor–neutral argument.

- ✓ Three leading BAS manufacturers integrated Graphwise GraphDB with the Brick metadata schema to create a unified Semantic Digital Twin, centralizing building assets and significantly reducing integrator implementation time and maintenance costs–replacing isolated silos with a "semantic backbone."
- ✓ Tietoevry used GraphRAG and a knowledge graph to ground AI outputs in industry–specific facts, enabling more confident automation for maintenance and nutrition decisions.

Vocabulary + Logic layers model assets, relationships, and constraints; Data–Linking integrates real–time operational data; the Process–Knowledge layer captures event flows; the Intelligence layer reasons across dependencies. At the business level this delivers 360° real–time visibility, predictive maintenance, lower maintenance cost and better uptime–turning static twins into reasoning systems.

11.5 Compliance Intelligence

Compliance is increasingly a data and knowledge problem, not just a policy problem: organizations must map regulations (e.g., DORA) to controls, track evidence, maintain traceability, and answer regulators quickly amid changing rules. (Graphwise: Compliance Intelligence) The Backbone reframes compliance from a cost–center into an intelligent asset by modelling relationships between rules, risks, and assets to create a 360° view that traditional databases cannot capture, with retrieval grounded in verified data for hallucination–free answers. The implementation pattern combines GraphRAG (LLM + semantic layer over knowledge graphs) and Talk–to–Your–Graph (TTYG).

- ✓ Reported impact: ~80% reduction in manual effort and a response–time improvement from 2 days to 1 hour for regulator inquiries.
- ✓ Cognizone used Graphwise TTYG and Graphwise GraphDB to let users query legislative archives in natural language, improving legal precision, reducing hallucinations, and cutting manual data–retrieval time. A leading pharmaceutical customer accelerated regulatory–authority response using AI–powered semantic similarity search.

The Logic layer encodes regulatory rules; the knowledge graph links regulations to processes and controls; the Trust Layer ensures validation, traceability, and auditability; the Intelligence layer enables explainable, rule–based checks. At the business level this lowers risk and operational burden while making controls, evidence, and obligations easy to trace and verify.

11.6 Scientific Knowledge Management

Scientific knowledge management connects disconnected lab results, documents, and expert reasoning into a machine-interpretable network-because discovery depends on meaning, relationships, and traceability, not keyword search. (Graphwise: Scientific Knowledge Management) FAIRification is becoming complex, error-prone, and prolonged across heterogeneous sources, and the Semantic Backbone is the way to operationalize FAIR (Findable, Accessible, Interoperable, Reusable) at scale; the PubMiner AI capability supports automatic knowledge extraction, with cross-industry positioning across biopharma, manufacturing, and finance.

- ✓ Oxford Drug Discovery Institute consolidated 30+ structured sources using Graphwise GraphDB and Linked Life Data Inventory, accelerating screening and target analysis and reducing R&D time from months to weeks or days.
- ✓ Queen's University unified 200+ datasets and 80 million publications into a searchable knowledge graph, cutting validation timelines from months to days and delivering 500% faster insight discovery. The FISH implementation uses Graphwise GraphDB and Semantic Objects to deliver a FAIR GraphQL API, FAIRifying animal-study data to enable reuse, ML/AI applications, and workflow automation.

Vocabulary + Logic layers model scientific entities; Data-Linking connects lab data, publications, and results; the Intelligence layer enables multi-hop reasoning; the knowledge graph ensures FAIR data principles. At the business level this accelerates R&D cycles and surfaces cross-study patterns, turning scientific knowledge into a connected, computable asset.

11.7 Semantic Data Hub

A Semantic Data Hub (SDH) is an advanced data architecture that unites the physical integration capabilities of a traditional Enterprise Data Hub (EDH) with the infrastructure of a **semantic backbone** built on ontologies and knowledge graphs. While a traditional EDH excels at consolidating and moving raw data for historical reporting and human-led business intelligence, it provides mainly technical metadata and lacks the rich business meaning required for complex automation.

By layering expressive, formal semantics over the data, an SDH translates granular data elements into a unified, machine-readable graph of data and content, eliminating **semantic inconsistencies** across different departments. Compared with a traditional EDH, this enables reasoning by preserving explicit relationships between entities, creates a highly traceable and explainable audit path for compliance, and provides the governed context needed to ground autonomous AI agents-drastically reducing hallucinations and increasing AI accuracy to 94-99%.

Physical vs. conceptual infrastructure. A traditional EDH provides the physical infrastructure to move and store data; the **semantic backbone** provides the conceptual infrastructure. An EDH tells a system where data lives and how it is formatted, whereas the semantic backbone tells the system exactly what the data means, abstracting complex data models into a unified view of business concepts and vocabulary.

Technical metadata vs. rich semantic context. A standard EDH manages technical metadata; a semantic backbone delivers rich semantic context-the business definitions of tables, how metrics like "revenue" or "retention" are calculated, which assets are trustworthy, and how entities relate logically. Where departments use different labels for the same concept ("counterparty," "external entity," "third party"), the ontology-backed knowledge graph maps them onto a single unified concept class so the enterprise "drinks from the same well" of truth.

Key benefits. Semantic interoperability (a "lingua franca" mapping diverse sources to one ontology); a flexible, scalable graph foundation that evolves without costly re-engineering; accelerated AI/GraphRAG adoption; democratized, natural-language data access; unified governance with data lineage; and neuro-symbolic reasoning with full explainability—so organizations can visually trace which facts, rules, and relationships led to an outcome.

Measurable KPIs. 40–60% reduction in time and cost to onboard new data sources/applications; faster time-to-market for AI (months to weeks); ~50% increase in data reuse across projects; and significant maintenance savings from retiring rigid ETL pipelines.

Proven at scale in pan-European data ecosystems (e.g., the European Railway Navigator), energy-sector interoperability (the Transparency Energy Knowledge Graph), and critical-infrastructure management (CIM-based integration at top-tier utilities). In short, the Semantic Data Hub provides "the semantic operating system for the enterprise, transforming a fragmented landscape of silos into a unified, intelligent data fabric that powers the AI strategy and ensures every system speaks the same language."

11.8 Cross-use-case Impact and Conclusion

Across all use cases the Semantic Backbone consistently delivers semantic grounding (accurate, context-aware outputs), deterministic reasoning (reliable, explainable decisions), a unified data foundation (elimination of silos), dynamic context (understanding of processes and workflows), and governance and trust (auditability and compliance). Organizations adopting a Semantic Backbone do not just improve individual use cases–they gain a systemic advantage: faster innovation, reduced operational risk, and scalable AI across the enterprise.

12. Implementation Approach

12.1 From Concept to Enterprise Deployment

Building a Semantic Backbone is not a one-time deployment but an evolutionary process integrating knowledge modelling, data integration, and AI enablement. The Backbone begins with a common foundation, then the deployment sequence branches into application-specific tracks depending on the first value target: knowledge management prioritizes retrieval and reuse, digital twins prioritize asset and process context, and the Semantic Data Hub prioritizes interoperability and governed exchange. This makes the implementation model both structured and flexible.

12.2 Step-by-step Implementation: Two Phases, Six Steps

Phase 1 – Establishing the semantic foundation. Build the core knowledge infrastructure for all downstream AI and analytics.

- ✓ **Step 1 – Define enterprise vocabulary (taxonomy layer):** standardize terminology with taxonomies and controlled vocabularies (W3C SKOS) for a common language. Outcome: consistent interpretation of data.
- ✓ **Step 2 — Develop ontologies (logic layer):** extend taxonomies into formal ontologies—relationships, constraints (via SHACL), and business rules—establishing a "contract on meaning" and machine-readable business logic for multi-hop reasoning.
- ✓ **Step 3 — Integrate and link enterprise data (data layer):** connect structured and unstructured sources via semantic enrichment, entity resolution, and metadata linking to create a single source of truth.

Phase 2 – Deploying intelligence and agentic capabilities. Build on the foundation to enable reasoning, context awareness, and automation.

- ✓ **Step 4 – Enable GraphRAG and semantic intelligence:** fuse the knowledge graph with LLMs for context-aware retrieval, explainable reasoning, and hallucination reduction (activates the Intelligence layer).
- ✓ **Step 5 – Build context graphs (process-knowledge layer):** capture workflows, decision traces, and event histories–a "dashcam" of how decisions are made over time.
- ✓ **Step 6 – Enable multi-agent orchestration:** deploy agents using shared world models, long-term memory, and semantic coordination for autonomous, cross-workflow operation.

Implementation Roadmap: Two Phases, Six Steps

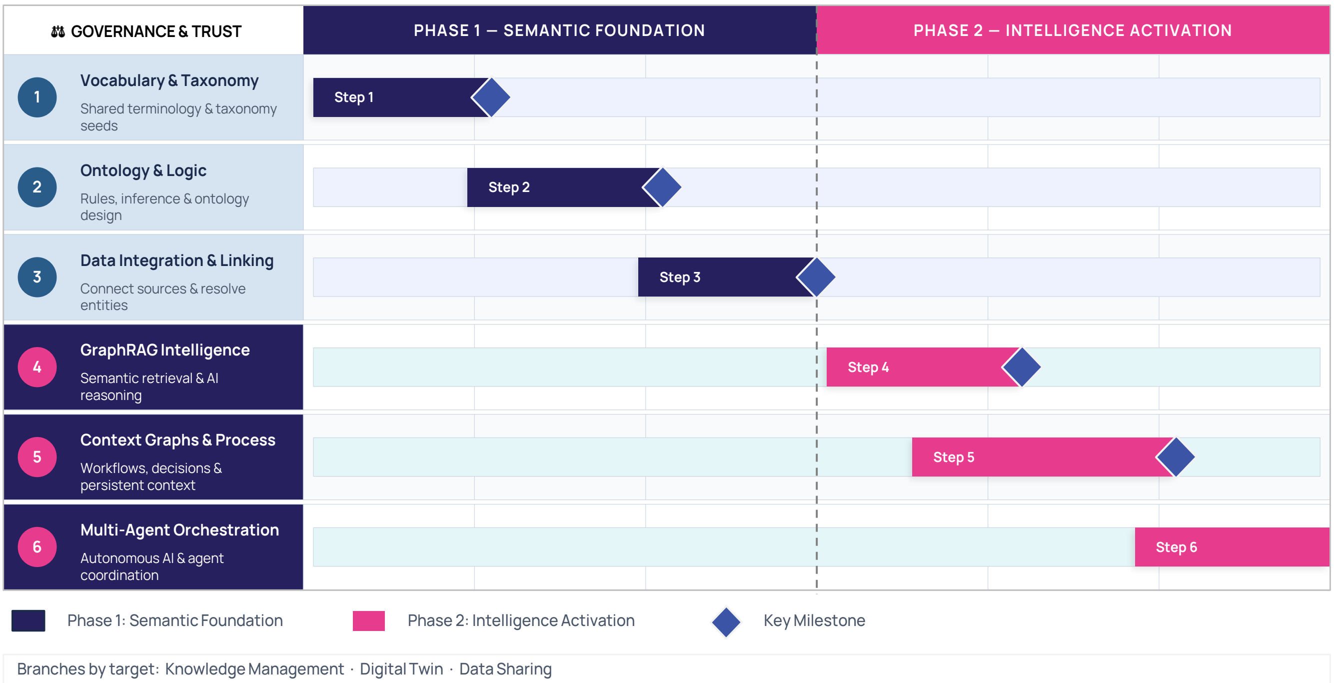

Phase 2 – Conclusion: Phase 2 transforms the Backbone into an active intelligence system enabling explainable AI, automation, and multi-agent collaboration.

12.3 Governance and Trust as a Continuous Layer

- ✓ Governance-as-code – encode business rules directly into ontologies and enforce validation using standards such as SHACL.
- ✓ Data quality and validation – prevent incorrect or incomplete data from entering the system.
- ✓ Explainability and auditability – ensure every decision is traceable.
- ✓ Scalability and evolution – allow continuous ontology updates without disrupting systems.

This keeps AI systems auditable, compliant, and production-ready. Governance is not an add-on–it is a core design principle enabling trust, compliance, and long-term scalability.

12.4 Implementation Considerations and Conclusion

- Start with high–impact use cases (compliance, knowledge management, digital twins).
- Adopt an incremental approach–build the backbone iteratively.
- Leverage open standards–RDF, OWL, SKOS and SHACL–for interoperability and future–proofing.
- Enable human–in–the–loop validation, combining domain expertise with automated extraction.
- Ensure integration with existing systems (CMS, databases, enterprise platforms) without disruption.

By following this phased, architecture–driven approach–semantic modelling, data integration, AI enablement, automation, and governance–organizations move from fragmented data and experimental AI to unified, governed, production–ready intelligence. The Semantic Backbone is the foundational capability that determines whether enterprise AI remains experimental or becomes scalable, trustworthy, and transformative.

Asked what single change would unlock the most value from AI, respondents converged on a unified, trusted data foundation (17.2%), better data quality and governance (14.6%), reduced vendor dependence / open architecture (13.5%), and reusable semantic models / a semantic layer (11.5%).

> "A single, trusted data foundation would make the biggest difference. Too much effort still goes into reconciling systems, aligning definitions, and fixing duplicates before we can use AI for anything meaningful."

> "Better data quality and governance would unlock more value than anything else. If the inputs are inconsistent, the outputs cannot be trusted, and that slows adoption."

> "I would reduce dependence on any single vendor. We need the freedom to move data and models across environments without rebuilding everything, otherwise the AI strategy will always be limited by someone else's roadmap and commercial terms."

> "A common semantic layer would save a lot of effort. If we could reuse business definitions across teams, we would spend less time aligning terminology and more time delivering AI-driven outcomes that people can use."
>
> **— Director, Software Consultancy / IT Services**

> "Faster access to approved data would have an immediate impact. The delay between identifying a use case and getting the right data is often longer than the model work itself."

Top Priorities for Improving AI Architecture

What are your organization's top priorities for improving AI capabilities over the next 2-3 years?

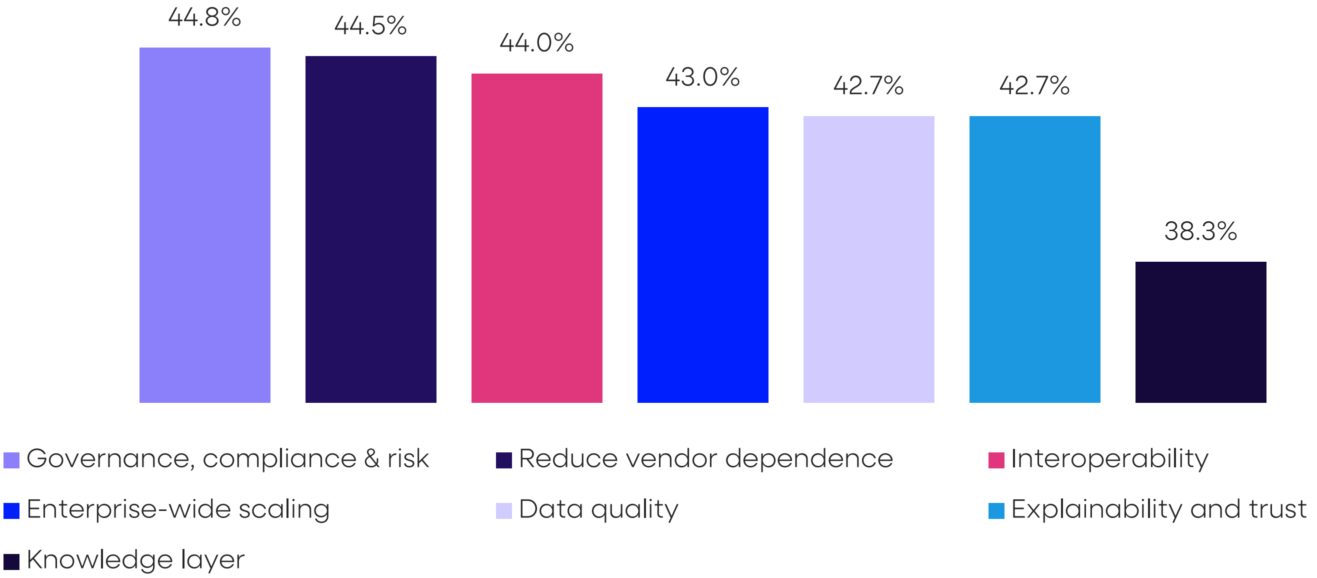

Figure 12.1 – Top Priorities For Improving AI Architecture: Governance, compliance & risk, vendor independence, and interoperability occupy the top 3 positions on the enterprise AI architecture priority list

Source: MarketsandMarkets Enterprise Semantic Backbone Survey, 384 enterprise AI/data leaders across Critical Infrastructure & Industrial Operations, BFSI, and Software Consultancy / IT Services (128 each), organizations of 500+ employees, May 2026.

13. Conclusion

The analysis in this whitepaper leads to one clear conclusion: as enterprises move from AI experimentation to agentic deployment, the limiting factor is no longer model capability alone, but the absence of governed, persistent, machine-executable context. AI has become a core investment priority, yet enterprise value remains uneven because organizations still struggle with fragmented data, limited interoperability, weak traceability, inconsistent grounding, and insufficient governance. In that environment, the traditional Context Layer and the thinner BI semantic layer are useful, but incomplete responses. They improve retrieval or standardize metrics, but they do not fully address the requirements of reliability, reasoning, validation, and enterprise-wide orchestration.

By contrast, the Semantic Backbone emerges as the more complete architectural foundation. Built on ontologies, controlled vocabularies, linked data, deterministic reasoning, provenance, and embedded governance, it is designed to make AI systems more reliable, more explainable, and more scalable across the enterprise. It shifts knowledge from being an overlay on top of applications to being infrastructure that can be reused across use cases, workflows, and agents. That is why it is better aligned to the core problems identified in the paper: hallucination risk, lack of context awareness, siloed data, weak trust, pilot-to-production friction, and the absence of unified orchestration.

The implication for enterprises is practical rather than theoretical. The path forward is to begin with a high-value use case, build a shared semantic foundation from vocabulary to ontology to linked data, and then activate intelligence through GraphRAG, context graphs, and multi-agent orchestration on standards-based foundations such as RDF, OWL, SKOS, and SHACL. Governance should not be added later; it should be embedded from the start. In that sense, the Semantic Backbone is not simply another layer in the stack. It is the foundation that determines whether agentic AI remains fragmented and experimental, or becomes trusted, scalable, and enterprise-grade.

14. References

This paper uses a single, consistent citation system: every source is listed once in the numbered entries below, and inline attributions in the body reference these entries by number (for example, [1] for the Stanford HAI AI Index). Findings labelled as MarketsandMarkets survey insights reflect MarketsandMarkets primary research: a survey of 384 enterprise leaders across three industry verticals (2026).

[1] Stanford HAI, The 2026 AI Index Report. Stanford Institute for Human-Centered AI.
https://hai.stanford.edu/ai-index/2026-ai-index-report

[2] J. Cui and J. Li, "Your Data Agents Need Context," a16z, Mar 10, 2026.
https://www.a16z.news/p/your-data-agents-need-context

[3] Graphwise, "What is a Semantic Backbone?" Graphwise Fundamentals.
https://graphwise.ai/fundamentals/what-is-a-semantic-backbone/

[4] Graphwise, "What is a Context Graph?" Graphwise Fundamentals.
https://graphwise.ai/fundamentals/what-is-a-context-graph/

[5] Graphwise use cases: Technical Knowledge Management, Semantic Digital Twins, Compliance Intelligence, Scientific Knowledge Management. https://graphwise.ai/

[6] Open Semantic Interchange (OSI) initiative and core specification (v0.1.1).
https://open-semantic-interchange.org/

[7] Salesforce Acquisition of Informatica,
https://www.salesforce.com/in/news/press-releases/2025/11/18/salesforce-completes-acquisition-of-informatica/;

[8] KPMG Acquires Metaphor's Technology Platform,
https://kpmg.com/us/en/media/news/kpmg-acquires-metaphors-technology-platform.html

About MarketsandMarkets

Founded in 2009, MarketsandMarkets (MnM) was built on a forward-looking vision: to identify and monetize the next wave of business opportunities embedded within global disruption. While others discussed megatrends like AI, 3D printing, IoT, and robotics, MnM saw a deeper insight—the unrealized revenue potential these shifts could unlock across industries.

MnM anticipates a USD 25 trillion revenue transformation in the B2B economy by 2030—driven not just by technology, but by macroeconomic shocks, policy shifts, climate pressures, and geopolitical flashpoints. Today, MnM tracks both emerging innovations and global dislocations—including landmark disruptions such as the 2025 US reciprocal tariffs—to help clients decode complexity, quantify exposure, and seize asymmetric growth opportunities ahead of competitors.

Over the past 13+ years, MnM has worked with 10,000+ clients across 30+ sectors, delivering over USD 140 billion in revenue impact. We've evolved from a market research publisher to a growth-enablement partner, powered by a 'GIVE Growth' philosophy and a 1,500+ strong team driven by outcome, not output.

Our commitment to strategic foresight and execution has earned us recognition from Forbes as one of America's Best Management Consulting Firms, and we remain the only India-origin firm on that prestigious list. What makes us a blue-ocean partner is our proprietary platform: Knowledge Store™—an AI-enabled, primary-research-driven market intelligence engine built to empower CXOs, strategy teams, and growth leaders to identify hidden inflection points and future profit pools in real-time.