Futurum

Informatica

salesforce

Elevate AI Agent Quality with Active, Intelligent Context

How Informatica Plugin Grounds Claude Reasoning in Trusted Data

Sponsored by Informatica, a Salesforce Company

SEPTEMBER 2026

The Trust Deficit in Enterprise Agentic AI

Enterprise generative AI is undergoing a critical architectural transition, moving from conversational interfaces to autonomous, task-oriented agents. This shift alters the primary deployment bottleneck. Model reasoning capability is no longer the primary constraint; operational trust, data governance, and output reliability now dictate production readiness.

Recent enterprise data highlights this operational tension (Figure 1). While 55.4% of enterprise technology leaders identify agent reliability and hallucination management as their primary hurdle to scaling generative AI,[1] application teams recognize that this is fundamentally an upstream data problem: 51.2% of organizations building internal AI applications report persistent data quality and availability challenges.[2]

Figure 1. The Agentic Trust Paradox

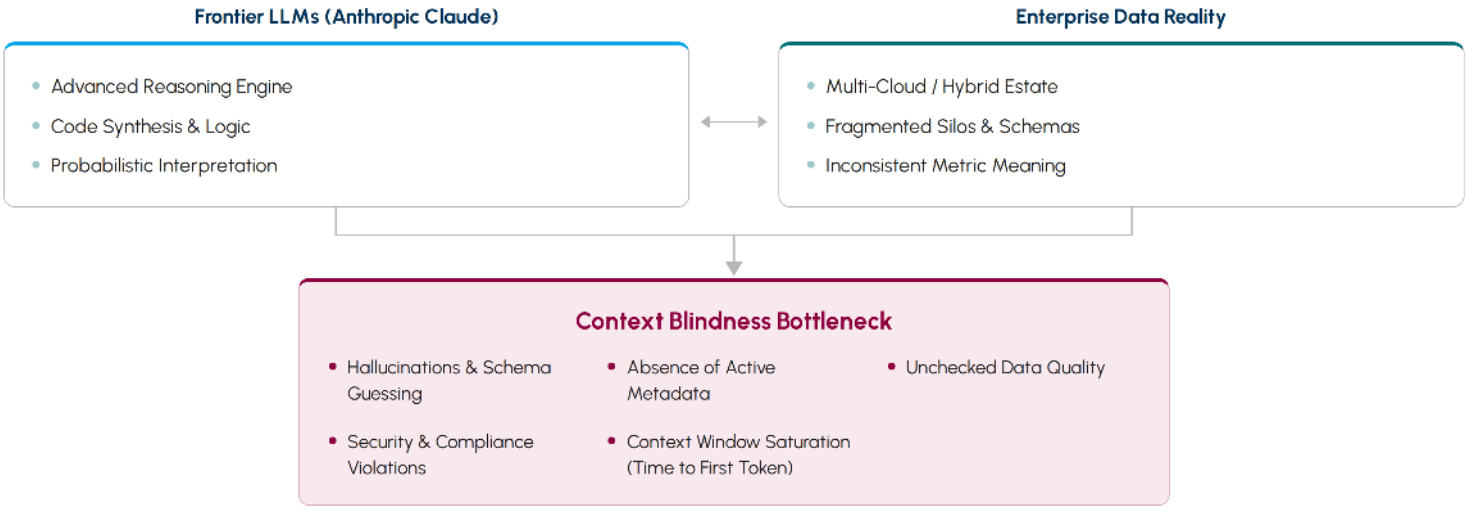

This reliability barrier stems directly from *context blindness*. Frontier large language models (LLMs) parse syntax and reason over unstructured text with precision, but they have no native awareness of enterprise relational hierarchies, system-of-record authority, data lineage, real-time business definitions, or active metadata. Of these, active metadata plays perhaps the most important role, enabling an agent to understand what enterprise data means, how it relates to other data, and what it can ultimately do with that data. If a data pipeline fails or an API is unavailable, active metadata can enable an agent to fix or troubleshoot the error rather than failing outright. Overcoming that kind of blind spot requires organizations to decouple deterministic enterprise context from probabilistic model weights.

As enterprise capital shifts toward infrastructure capable of grounding AI in verifiable corporate truth,[3] IT leaders are turning to an active metadata control plane to bridge the gap between static storage layers and autonomous reasoning agents.

1. Futurum Research, *1H 2026 Artificial Intelligence Platforms Decision Maker Survey*
2. Enterprise Technology Research (ETR), *AI Product Series Theme 3*, March 2026.
3. Futurum Research, *1H 2026 Data Intelligence, Analytics, & Infrastructure Market Sizing & Five-Year Forecast Report*

The Context and Meaning Dilemma: Why Unconstrained Retrieval and Application-Tier AI Fail

Early enterprise generative AI architectures relied heavily on naive Retrieval-Augmented Generation (RAG). While vector search effectively matches unstructured text fragments based on semantic proximity, it falls short in relational awareness, transactional state understanding, and governance controls. Ingesting uncurated source documents or raw database extracts inevitably causes metric divergence across operational silos. For instance, a financial query routed through basic vector retrieval can surface contradictory revenue calculations across enterprise resource planning (ERP) billing systems and customer relationship management (CRM) pipelines. Figure 2 displays RAG and other retrieval architecture pathways.

Figure 2. Retrieval Architecture Pathways

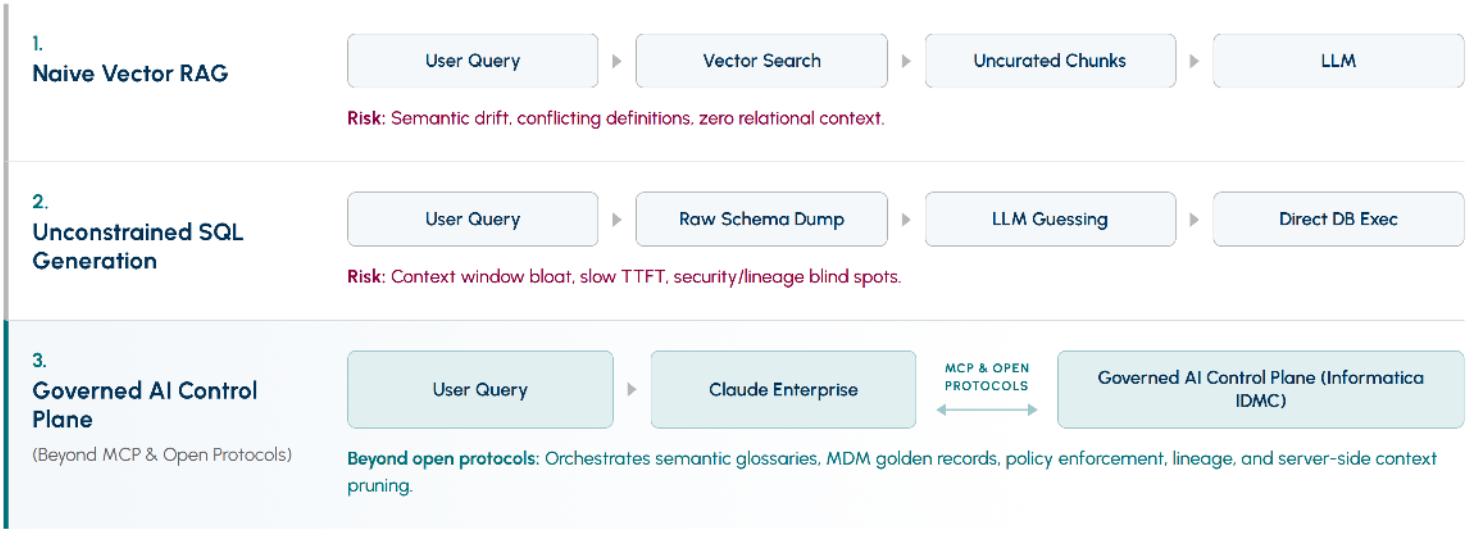

To work around these limitations, engineering teams frequently connect runtime models directly to raw operational schemas via unconstrained text-to-SQL generation or ad hoc APIs. This approach introduces significant performance, reliability, and security penalties:

- **Context Window Saturation:** Injecting uncurated database schemas into prompt payloads consumes limited context windows, inflating token costs and extending Time to First Token (TTFT) latency.
- **Semantic Guesswork:** Without certified business glossaries, the model must guess whether a column named REV_AMT corresponds to GAAP revenue, recognized bookings, or unbilled receivables.
- **Security and Lineage Gaps:** Point-to-point application connectors often bypass centralized policy enforcement, creating data leakage risks and leaving compliance teams without an auditable execution chain.
- **Absence of Active Metadata**: Direct, static connections lack runtime awareness of data freshness, schema evolution, and operational dependencies.
- **Unchecked Data Quality**: Ingesting dirty, unverified, or fragmented records leads directly to flawed operational decisions and automated downstream errors—turning minor source-data anomalies into high-impact agentic failures.

Enterprise survey data reflects these operational roadblocks (Figure 3). Technology leaders rank agent reliability, data privacy, and ROI uncertainty well ahead of compute infrastructure constraints.

Figure 3. Primary Bottlenecks to Enterprise Generative AI Production

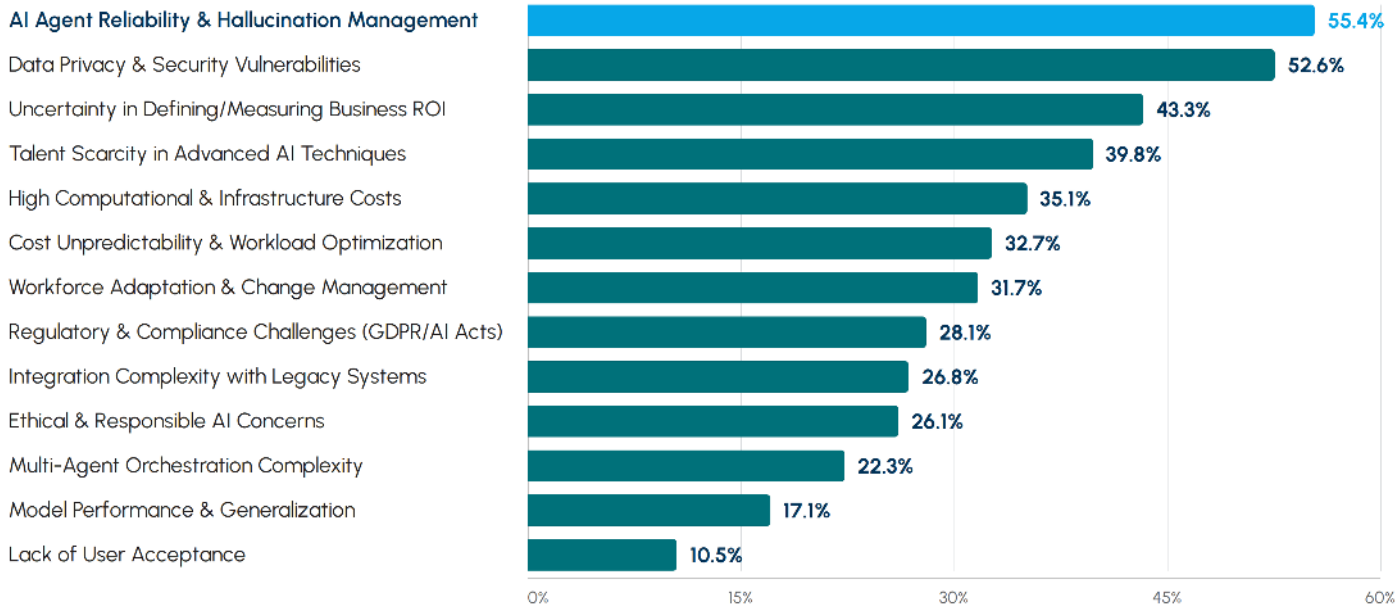

Source: Futurum Research 1H 2026 Artificial Intelligence Platforms Decision Maker Survey (n=820 enterprise decision-makers).[4]

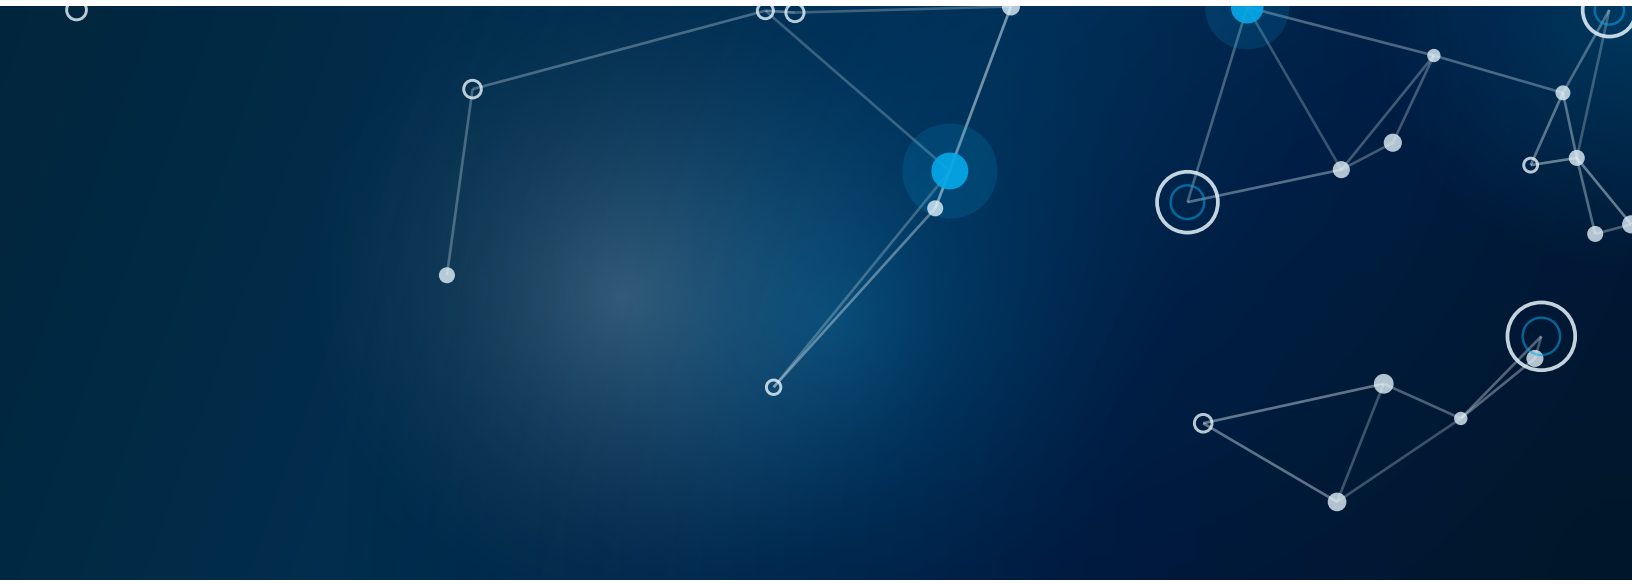

4. Futurum Research, *1H 2026 Artificial Intelligence Platforms Decision Maker Survey*

Architectural Mechanics: Grounding Claude Enterprise via Informatica's Headless Plugin

To resolve context blindness without brittle, point-to-point code, Informatica is extending its Intelligent Data Management Cloud (IDMC) directly into Anthropic's Claude Enterprise and Claude Code environments (Figure 4). Building on its headless integrations across cloud ecosystems, Informatica has built a headless plugin capable of delivering catalog discovery, master data resolution, and governance services via Anthropic's open Model Context Protocol (MCP) standard.

Figure 4. Informatica Headless Plugin Architecture

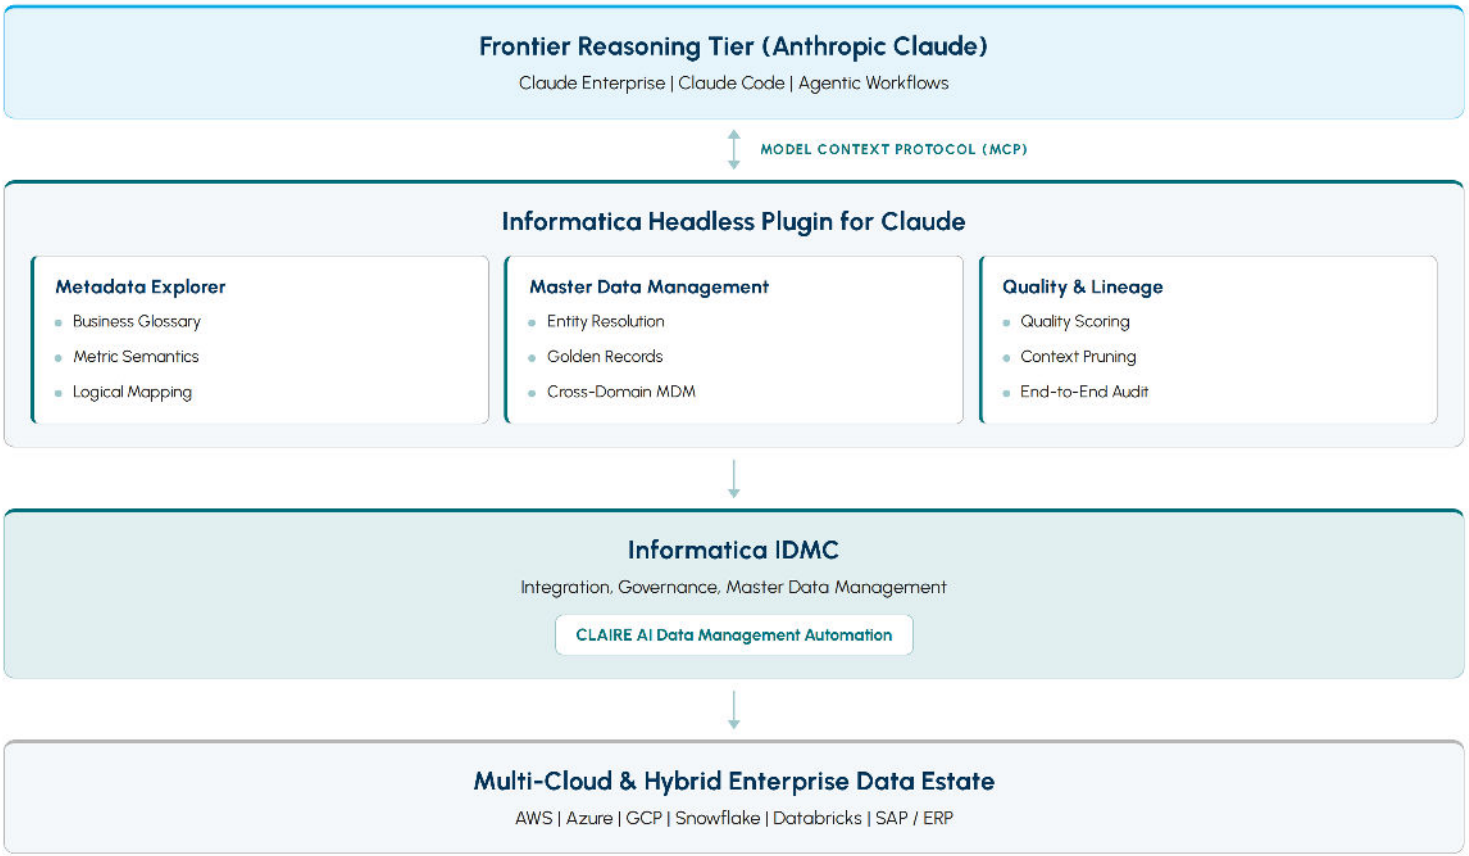

MCP is an open, standardized client-server interface that allows foundation models to query structured operational systems without custom application middleware. The Informatica Headless Plugin for Claude serves as a complete, working toolkit that addresses several use cases such as data discovery, lineage impact analysis, profiling, etc., deployed through several specialized MCP servers supported by Informatica's CLAIRE AI engine:

1. **Metadata Explorer MCP Server:** Dynamically supplies certified business glossaries, semantic metric definitions, and physical-to-logical data mappings. Claude interprets enterprise metrics deterministically instead of inferring them probabilistically.
2. **Master Data Management (MDM) MCP Server:** Executes server-side entity resolution across customer, product, and vendor domains. This prevents agents from taking actions based on duplicate accounts, mismatched IDs, or fragmented operational records.

3. **Point-of-Entry Quality and Lineage Engine:** Intercepts data retrieval at runtime to evaluate quality scores, strip unverified records, and attach an end-to-end lineage graph to the agent's output for regulatory auditing.

This headless approach delivers token efficiencies. By performing schema pruning, entity resolution, and catalog lookups inside IDMC before passing context to Claude, the system avoids dumping raw database schemas into prompt windows. Delivering concise, pre-governed context tokens reduces inference expenses and lowers TTFT latency.

Enterprise adoption metrics highlight this rapid movement toward standardized agentic interfaces (Figure 5): 76.2% of organizations with active agentic AI interest are running MCP in production, piloting it, or strongly considering it for future projects.[5]

Figure 5. Enterprise Adoption of the Model Context Protocol (MCP)

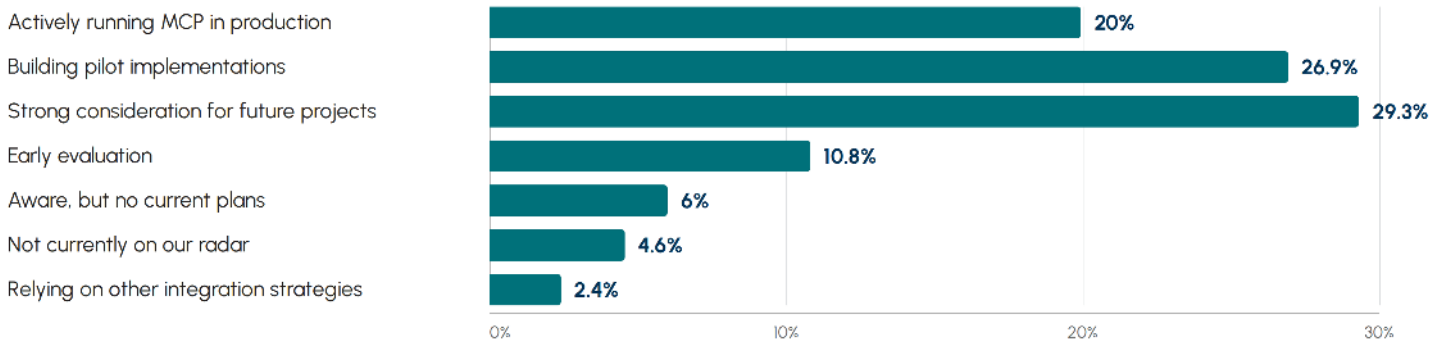

Source: Futurum Research 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey (n=416 of 818 enterprise data/AI leaders, asked of those prioritizing generative and agentic AI investment).[6]

5. Futurum Research, *1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey*
6. Futurum Research, *1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey*

Control Plane Neutrality Across Fragmented Multi-Hybrid Cloud and Multi-Platform Estates

Grounding AI context is significantly more complex in modern multi-cloud IT environments. Managing data across multiple public clouds is becoming very common, and according to survey data, 45.5% of organizations operate hybrid environments that span hyperscalers (AWS, Microsoft Azure, Google Cloud), modern analytical engines (Snowflake, Databricks), and transactional platforms (SAP, Oracle, on-premises databases).[7] Model choice compounds the problem: 56.2% of organizations deploy multiple foundation model providers simultaneously.[8]

Single-vendor data ecosystems cannot solve this problem across heterogeneous environments. Hyperscaler-native catalogs only govern data within their own infrastructure boundaries; they cannot unify lineage, enforce global access policies, or resolve master entities across competing cloud platforms.

Informatica IDMC operates as an independent context fabric above the storage and compute tiers (Table 1). It organizes operational context across distributed endpoints without requiring organizations to consolidate or replicate data into a single repository.

Table 1. Architectural Comparison — Naive Agentic Retrieval vs. Informatica Headless Data Management

Dimension	Naive Vector RAG and Direct APIs	Informatica IDMC	Strategic Operational Impact
Context Retrieval	Unstructured text search; blind to relational schemas and data integrity.	Governed metadata, unified MDM entities, and certified glossaries.	Reduces hallucinations caused by uncurated, conflicting data sources.
Data Governance	Point-to-point security rules coded into individual applications.	Centralized policy enforcement, object-level RBAC, and verifiable lineage.	Provides consistent compliance and access audits across multi-cloud environments.
Token Efficiency	Large schema dumps fill prompt windows with irrelevant metadata.	Concise, curated context payloads generated server-side.	Cuts inference costs, reduces TTFT latency, and preserves context window space.
Estate Coverage	Confined to single-vendor clouds or standalone SaaS platforms.	Neutral control plane spanning hybrid, multi-cloud, and legacy systems.	Enables cross-platform visibility without costly data replication.

7. Futurum Research, *1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey*
8. Enterprise Technology Research (ETR), *AI Product Series Theme 3*, March 2026.

Technical Strengths and Operational Realities

Pairing Anthropic's Claude with Informatica's headless plugin capabilities provides a clean separation of responsibilities: Claude supplies the reasoning and code synthesis, while Informatica provides deterministic governance guardrails. Rather than relying on guesswork when generating queries or interpreting business logic, Claude works from certified data definitions and verified master records. This allows the model to handle complex tasks—such as reconciling billing discrepancies across ERP and CRM systems—while maintaining an auditable data trail (Figure 6).

Figure 6. Operational Implementation Tradeoffs

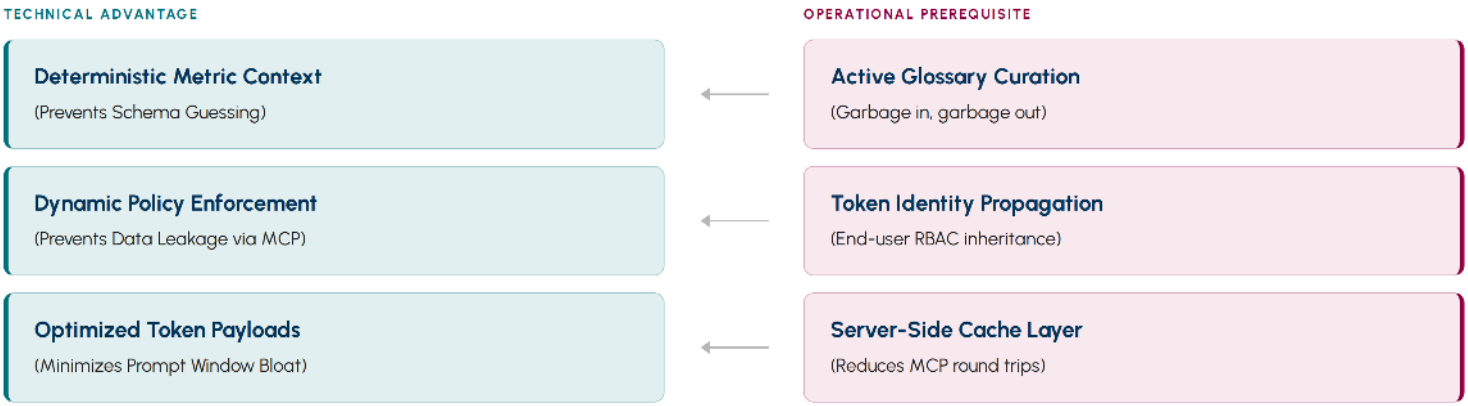

However, implementing this architecture introduces several practical operational challenges that enterprise engineering teams must plan for:

- **Underlying Data Quality Remains Essential:** IDMC provides the pipeline, catalog, and governance layers, but it does not automatically fix neglected source tables. Connecting an uncurated database to an MCP server will simply deliver incorrect data to downstream agents faster. Baseline data cleansing must occur upstream.
- **Tool-Level Security and Identity Propagation:** IT security teams must ensure that every MCP tool execution dynamically inherits the invoking user's credentials and object-level permissions. Passing generic service-account credentials through MCP creates significant data exposure risks.
- **Orchestration Latency in Multi-Step Workflows:** Complex workflows requiring multiple sequential MCP calls can introduce network and execution latency. Engineering teams should implement server-side metadata caching within IDMC to maintain interactive response times.

Blueprint for Enterprise Agentic Readiness

To maximize the return on governed context integration, enterprise technology leaders should follow a phased deployment plan (Figure 7):

1. **Standardize Semantic Definitions and Master Data Entities:** Audit corporate business glossaries and verify that core enterprise metrics (such as ARR, churn, and active customer definitions) are certified in IDMC before exposing MCP endpoints to agents.
2. **Benchmark Governed MCP Against Unconstrained Retrieval:** Run controlled proof-of-concept tests comparing standard vector RAG implementations against IDMC-governed MCP servers. Quantify output accuracy, hallucination rates, token consumption, and TTFT latency on standard analytical queries.
3. **Implement Identity-Aware Authorization at the Protocol Layer:** Configure dynamic identity propagation across all MCP connections, ensuring tool calls strictly inherit the active permissions and security entitlements of the individual user.
4. **Establish FinOps and Observability Controls:** Deploy continuous telemetry across MCP server endpoints to monitor execution times, cache hit ratios, and token utilization rates, keeping infrastructure costs predictable as agentic workloads scale.

Figure 7. Agentic Readiness Implementation Phases

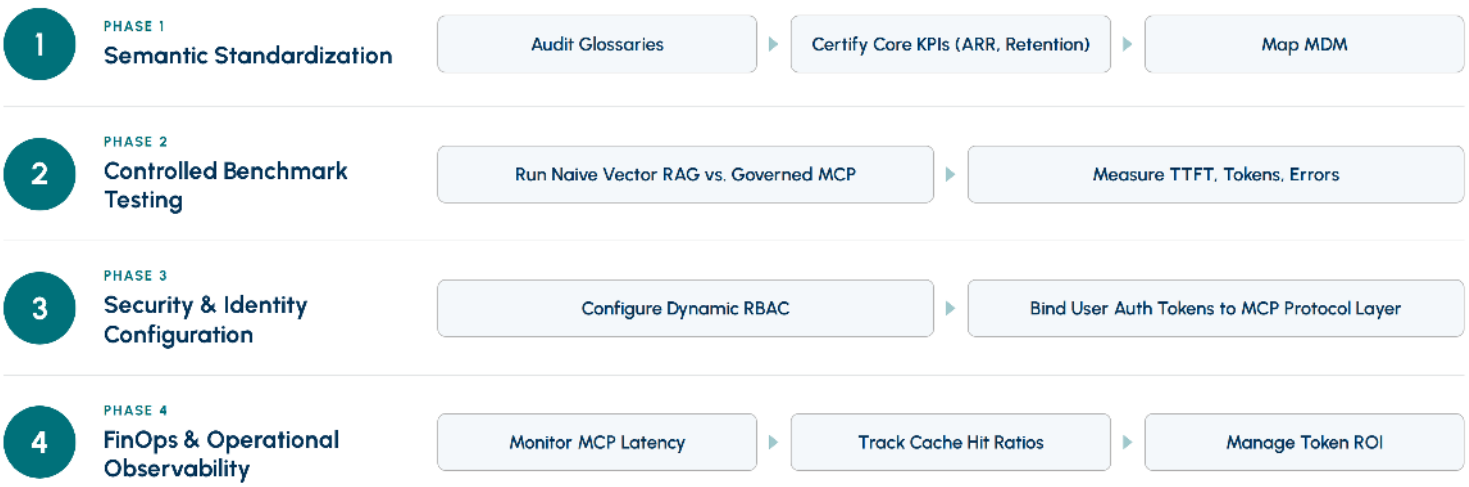

Analyst Perspective

Enterprise AI is moving past unconstrained experimentation into an era defined by operational governance and architectural discipline. Reasoning engines require structured, authoritative business context to function reliably in production workflows. Without deterministic metadata, active quality validation, and unified master data, even the most capable models will produce inconsistent and untrustworthy results.

Grounding Anthropic's Claude Enterprise through Informatica IDMC's neutral, headless MCP architecture directly addresses the enterprise context gap across hybrid and multi-cloud estates. By establishing a platform-agnostic control plane that delivers verified semantic definitions and audit-ready lineage directly to runtime agents, organizations can move generative AI from open-ended chat interfaces to reliable, production-ready enterprise workflows.

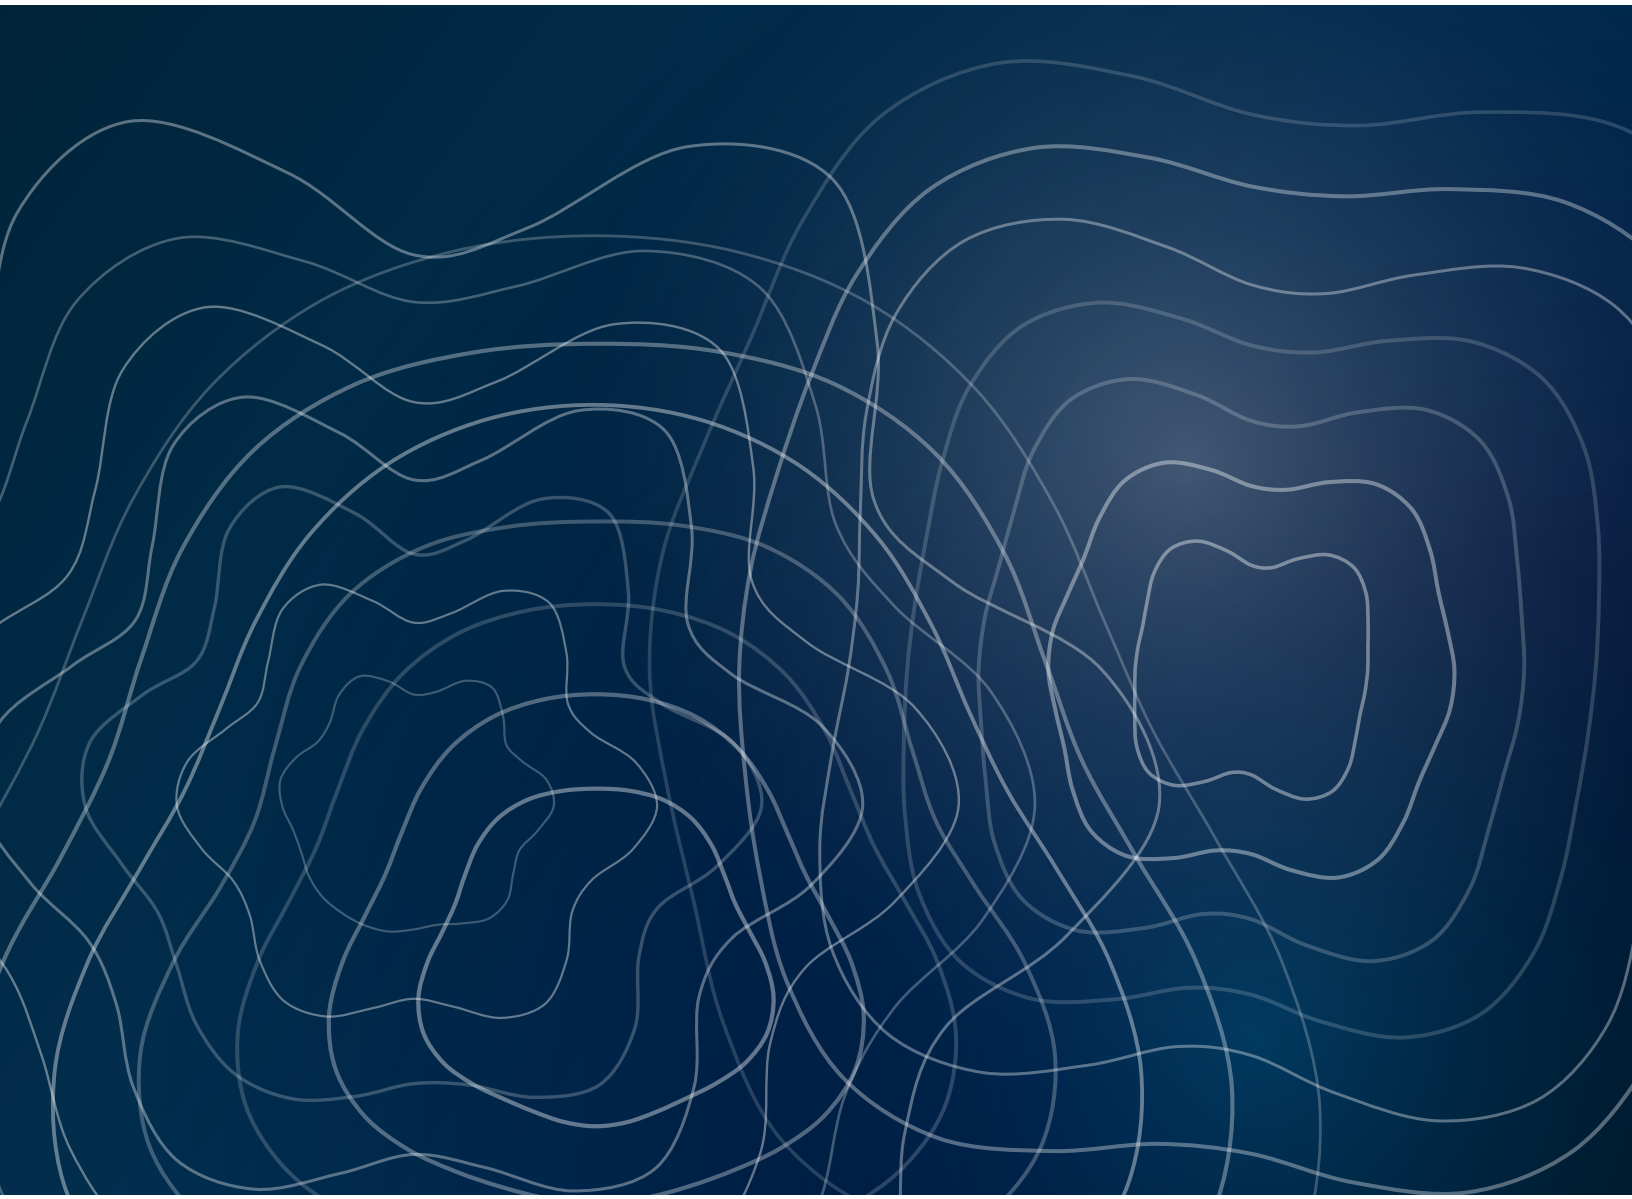

About This Brief

PROJECT TEAM

Brad Shimmin · Lead Author
Vice President & Practice Lead, Data Intelligence, Analytics, & Infrastructure | Futurum Research

Benjamin Brown · Contributor
Vice President, Custom Research | Futurum Research

PUBLISHER

Futurum Research

INQUIRIES

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

CITATIONS

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

LICENSING

DISCLOSURES

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

Informatica

ABOUT INFORMATICA

Informatica, a Salesforce company, is a leader in AI-powered enterprise cloud data management. Its Intelligent Data Management Cloud delivers an end-to-end platform of integrated solutions that connect, manage, and unify data across any cloud, hybrid, or multi-cloud environment, powered by CLAIRE AI. Informatica serves approximately 5,000 customers in nearly 100 countries, including more than 80 of the Fortune 100.

salesforce

ABOUT SALESFORCE

Salesforce helps organizations of any size become agentic enterprises — integrating humans, agents, apps, and data on a trusted, unified platform to unlock unprecedented growth and innovation. Visit www.salesforce.com for more information.

Futurum

ABOUT THE FUTURUM GROUP

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