Strategy₿

The Enterprise Semantic Layer Buyer's Guide

How to Evaluate 7 Critical Capabilities for a Semantic Layer Built for Analytics and AI

Selecting the Right Universal Semantic Layer
for AI-Ready, Future-Proof Enterprise Analytics

2026 Edition

Table of Contents

Executive Summary

Most enterprises today operate across multiple clouds, data platforms, and consumption tools. As a result, the semantic layer has become a core architectural component—translating raw data into governed, business-ready concepts that stay consistent across BI, applications, and AI. Yet many semantic layers are embedded within a specific warehouse or BI tool, creating duplicated logic, inconsistent KPIs, and costly lock-in.

A universal semantic layer addresses this by applying governed definitions, security, and performance consistently across heterogeneous platforms and tools. This guide provides a practical framework to evaluate solutions across seven critical dimensions.

Key findings from What 100 senior data leaders said (CIO Dive Studio, 2026)

95%+ reported that defining metrics separately in each tool is challenging. Nearly half cited multiple data sources as the top barrier.

70% say unified business definitions are the most critical factor for analytics and AI success.

68% of large enterprises are using or actively piloting an independent semantic layer.

87% rate visibility into how AI uses enterprise data as very or extremely valuable.

Measurable outcomes (UserEvidence study across retail, telecom, and financial services customers)

- **$3.4M** average net annual impact
- **551% ROI** with a two-month payback
- **Metric confidence** increased from **5 out of 10** to **9 out of 10**
- **Redundant metrics** decreased **44%**
- **AI hallucinations** decreased **22%**
- **End-user time** regained: **46%** previously spent on manual reconciliation

> Data & Analytics Leader
>
> Global consumer goods company
>
> *"Everyone can speak the same data language. Our semantic layer makes complex information universally understood and actionable."*

Who This Guide Is For

Semantic layer decisions typically involve two audiences with different success criteria.

For technical evaluators

- Prove KPI parity across at least two consumption surfaces and two user groups (permissions).
- Validate governance (row/column security, masking, auditing) across every access path, including AI interfaces where applicable.
- Confirm performance under realistic workloads and concurrency.
- Verify lifecycle operations (versioning, promotion, monitoring, and change control).

For executive and economic buyers

- Reduce KPI disputes, reconciliation time, and duplicated logic.
- Enable analytics and AI adoption with governed access, transparency, and auditability.
- Avoid lock-in and preserve portability as the stack evolves.
- Build a credible business case tied to time savings, cost efficiency, and risk reduction.

Why the Semantic Layer Matters

Modern enterprises run across multiple clouds, data platforms, SaaS applications, and on-prem systems. Teams use different BI tools, spreadsheets, applications, and increasingly AI assistants and agents. When business logic is defined separately in each layer, organizations end up with multiple versions of the same metric, inconsistent security rules, and growing confusion about which numbers are correct.

A semantic layer addresses this by centralizing business definitions, governed metrics, and access policies into a shared business context that stays consistent across tools and workflows. That consistency is essential for trusted analytics today and for scaling AI responsibly.

The CIO Dive research underscores the urgency. Nearly 80% of data teams at large enterprises spend more than half their time on data preparation rather than generating insights. Yet satisfaction with existing approaches remains mixed: among users of independent semantic layers, 79% report only slight or moderate satisfaction, suggesting the market is still searching for solutions that fully deliver on the promise.

> AI pilots stall not because models lack capability, but because existing data architectures cannot support them. Fragmented platforms, inconsistent business definitions, and governance models designed for static reporting all become constraints once AI systems begin operating autonomously.

The core challenges: Five persistent problems drive the urgency.

- **Inconsistent metrics:** The same KPI is calculated differently across tools and teams, leading to conflicting reports and eroded trust. In the CIO Dive survey, 42% cited this as a top barrier.
- **Duplicated logic and high maintenance:** Business rules are recreated across BI, data transformations, and spreadsheets, increasing rework and drift. Nearly half of respondents identified multiple data sources as the single greatest challenge.
- **Slow time-to-insight:** Teams lose time reconciling definitions and re-implementing metrics instead of producing insights. In the CIO Dive survey, analysts spent 38% of their time on work a semantic layer could eliminate.
- **AI readiness gap:** AI systems need governed context, consistent permissions, and auditability. Yet 61% blame overly complex infrastructure as the greatest AI implementation barrier, followed by poor data quality (36%) and too many data silos (35%).
- **Vendor lock-in risk:** Embedding definitions in one platform increases switching costs as enterprise database technology changes every 3–5 years.

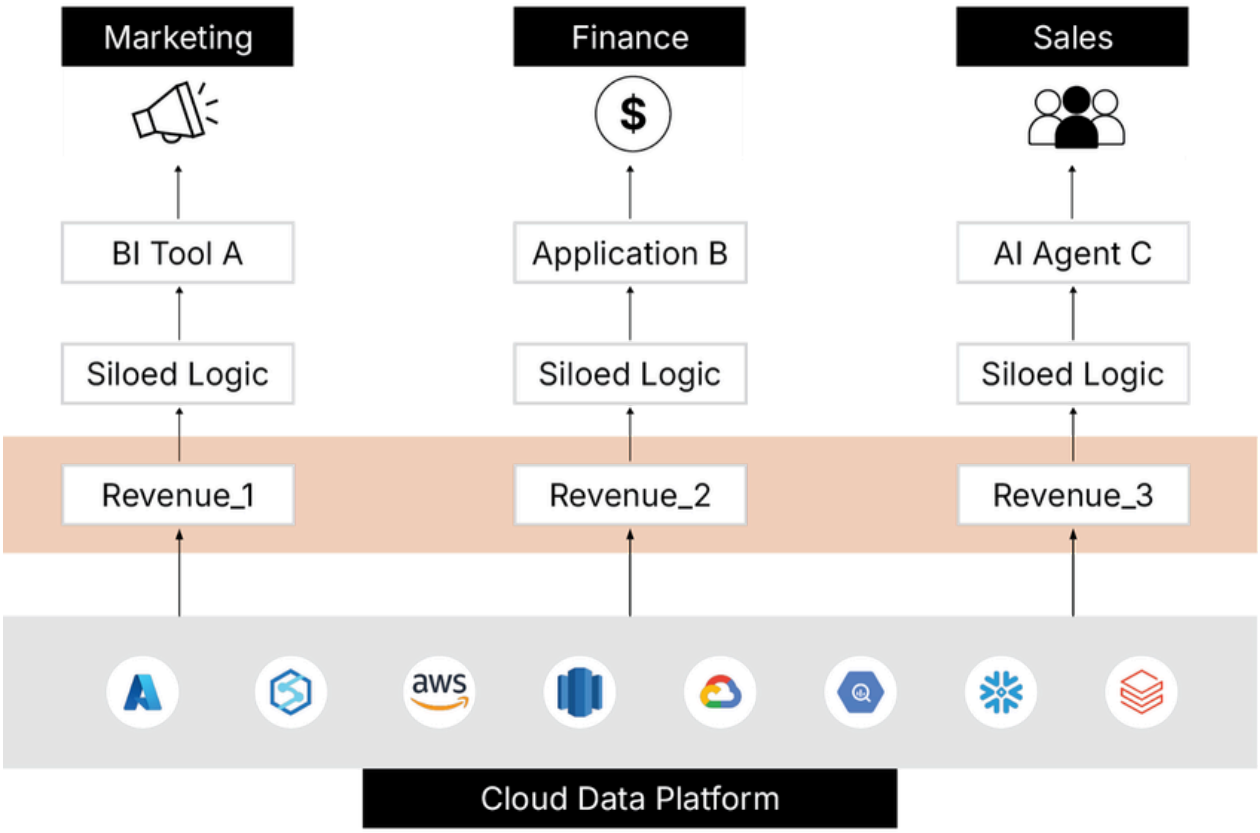

The solution: A universal semantic layer sits above data sources and tools to provide define-once/use-everywhere metric consistency, cross-platform federation, centralized governance that extends to AI agents, rich business context that reduces hallucinations, cost optimization through query abstraction, and self-service enablement. In the UserEvidence study, 67% of customers reported less IT dependency and two-thirds reported faster access to insights.

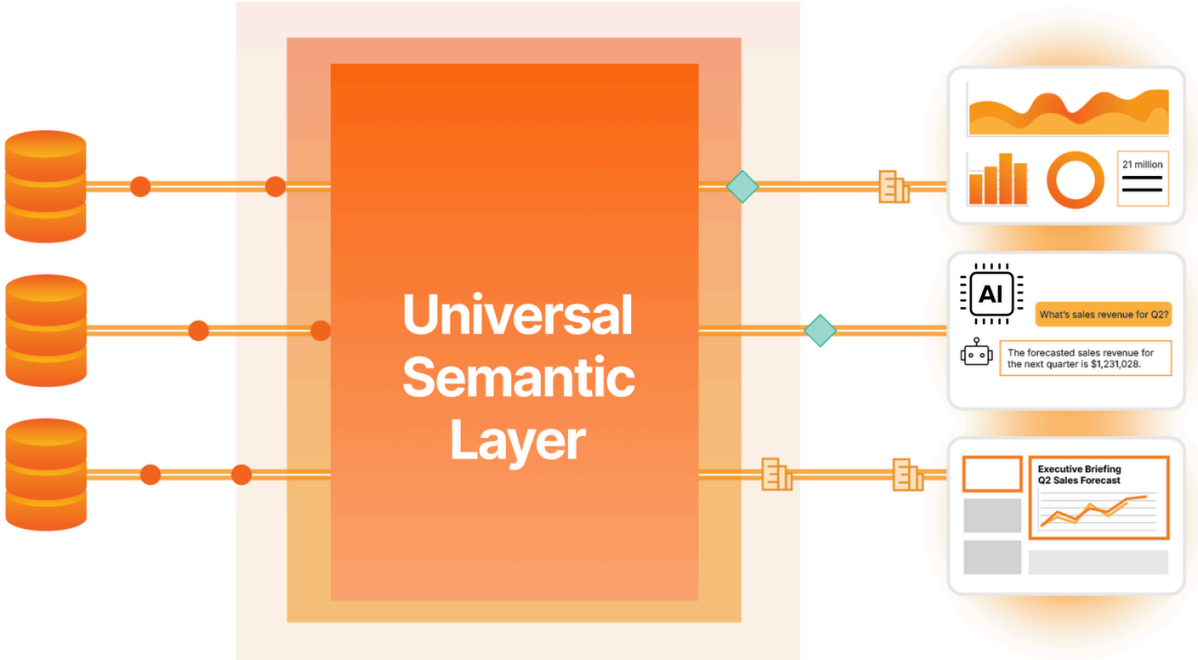

Market Context and Landscape

The semantic layer market sits at the crossroads of data virtualization (~$8B, ~20% CAGR), enterprise data management (~$99B, ~12% CAGR), and data warehouse platforms (~$34B, ~11% CAGR). This convergence means buyers encounter overlapping claims and inconsistent definitions. The term "semantic layer" can mean very different things depending on whether a vendor's roots are in BI tools, data warehouses, data virtualization, or purpose-built semantic modeling.

Two developments reshaping the landscape:

- **Open Semantic Interchange (OSI).** A cross-vendor industry initiative focused on improving semantic interoperability. Ask vendors how they support OSI and what practical interchange approach they offer to reduce rebuilds when you add tools or migrate platforms.
- **AI agents and Model Context Protocol (MCP).** MCP standardizes how AI agents interact with enterprise data through governed semantic context. When paired with a semantic layer, MCP enables AI requests to be authenticated, access-validated, logged, and contextually enriched, turning ad hoc AI integrations into governed, auditable systems.

The solution landscape broadly falls into three tiers, each with different tradeoffs for portability, AI readiness, and governance.

Types of Semantic Layer Solutions

The market can be broadly categorized into three tiers of capability:
Platform-embedded semantic layers are built into data warehouses or BI tools, such as warehouse-native semantic views or BI-tool metric layers. While convenient for single-platform environments, they create lock-in and fragmentation when organizations use multiple platforms. Definitions built in one tool do not automatically carry to another.

Specialized semantic layer platforms focus specifically on semantic modeling with varying capabilities, from code-first metric definition tools to more comprehensive modeling platforms. They generally offer better portability than embedded options but may lack integrated performance engines, AI-powered development, or comprehensive governance.

Comprehensive independent semantic layers combine semantic modeling with in-memory processing, multi-source federation, AI-assisted development, active governance, and open connectivity. They are designed to work across the full data ecosystem, independent of any single tool, warehouse, or cloud and therefore best match enterprise requirements for consistency and scale. Crucially, the strongest solutions in this tier are also built for fast time-to-value. With proper scoping, teams can deliver an initial governed domain within 6–8 weeks, then expand from there. That matters because the biggest risk to a semantic layer initiative isn't choosing the wrong vendor. It's an 18-month implementation that loses executive sponsorship before meaningful value appears.

Capability	Embedded	Specialized	Comprehensive independent
Works across multiple tools	❌	⚠️	✅
Semantic modeling	⚠️	✅	✅
Multi-source federation	❌	⚠️	✅
In-memory performance	❌	❌	✅
AI-assisted development	❌	⚠️	✅
Governance	⚠️	⚠️	✅
Open connectivity	❌	⚠️	✅
Fast time-to-value	⚠️	⚠️	✅

Evaluation Framework

Based on industry best practices and the evolving needs of modern enterprises, we recommend evaluating semantic layer solutions across seven critical dimensions.

Evaluation Criterion	Why It Matters	Weight
1. Platform Independence & Federation	Business logic must outlast any platform; query across clouds without data movement	**Critical (20%)**
2. AI-Powered Modeling & Data Readiness	Reduce modeling cycles from weeks to days; ensure data quality upstream	**High (15%)**
3. Semantic Depth & Richness	Handle hierarchies, time intelligence, non-additive measures, rich metadata	**High (10%)**
4. Open Standards & Universal Access	SQL, REST, MDX, DAX access; optimized BI connectors; OSI alignment	**High (10%)**
5. Active Security & Governance	Centralized RLS/CLS/RBAC enforced across all consumers including AI agents	**High (15%)**
6. AI Agent Readiness	MCP support, governed agent access, machine audit trails, hallucination reduction	**High (15%)**
7. Performance, Scalability & Cost	Sub-second on billions of rows; in-memory acceleration; predictable pricing	**High (15%)**

1. Platform Independence & Federation

Enterprise database technology changes every 3–5 years, and 90% of large enterprises use multi-cloud infrastructure. The semantic model must be fully decoupled from underlying storage and consumption tools—business logic, security, and metrics should survive a complete platform migration. Look for hybrid execution (passthrough + in-memory), intelligent query routing, real federation across clouds without ETL, and 100+ pre-built connectors. In the CIO Dive survey, 66% valued switching BI tools without rebuilding definitions, and 49% valued switching warehouses without costly rebuilds.

Watch out for: solutions that store definitions inside a specific warehouse or BI tool, or that require all data in one location.

2. AI-Powered Modeling & Data Readiness

Traditional semantic modeling requires weeks of manual effort. AI-powered modeling can reduce this dramatically while improving quality. The CIO Dive survey found 83% already use AI to accelerate data modeling. Look for automated discovery, relationship detection with confidence scoring, natural-language metric creation, and integrated data quality profiling. The best solutions can reduce modeling cycles from weeks to days while keeping humans in the loop.

Watch out for: solutions relying solely on hand-coded YAML/SQL, or that assume clean warehouse-ready inputs without addressing data quality.

3. Semantic Depth & Richness

A basic layer defines metrics and dimensions. Enterprise scenarios demand complex hierarchies (ragged, unbalanced), time intelligence (fiscal calendars, YTD), semi-additive/non-additive measures, parameterized metrics, and predictive analytics. The depth of the semantic model determines how effectively AI can reason about your data. Look for a true semantic graph with interconnected objects, rich metadata including AI-generated descriptions, and knowledge graph capabilities.

Watch out for: basic metric layers that only define simple aggregations.

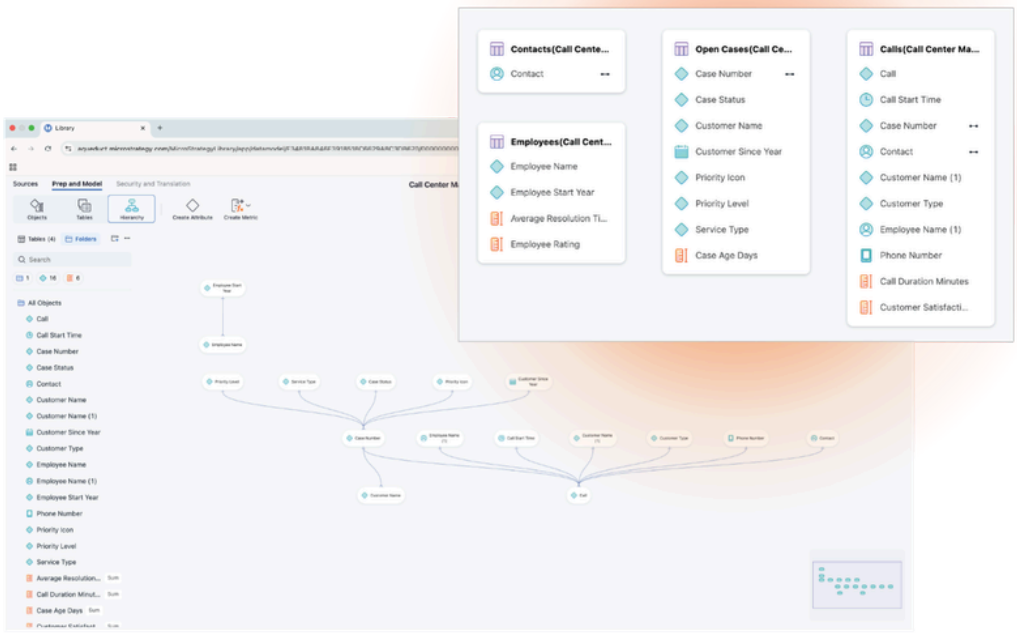

4. Open Standards & Universal Access

Your semantic layer must serve every consumer, including BI tools, AI agents, custom applications, and spreadsheets. Evaluate standard access methods such as SQL via JDBC/ODBC, DAX via XMLA, REST APIs, and SDKs, plus any optimizations that improve performance and usability.

Assess whether the vendor participates in open interoperability efforts like OSI and has a credible path to exchange or translate semantic definitions across platforms, rather than forcing you to rebuild logic per tool. The goal is a bridge, not a moat.

Watch out for: solutions that only work with specific BI tools, lack optimized connectors, or require you to manually recreate definitions that already exist in other semantic layers.

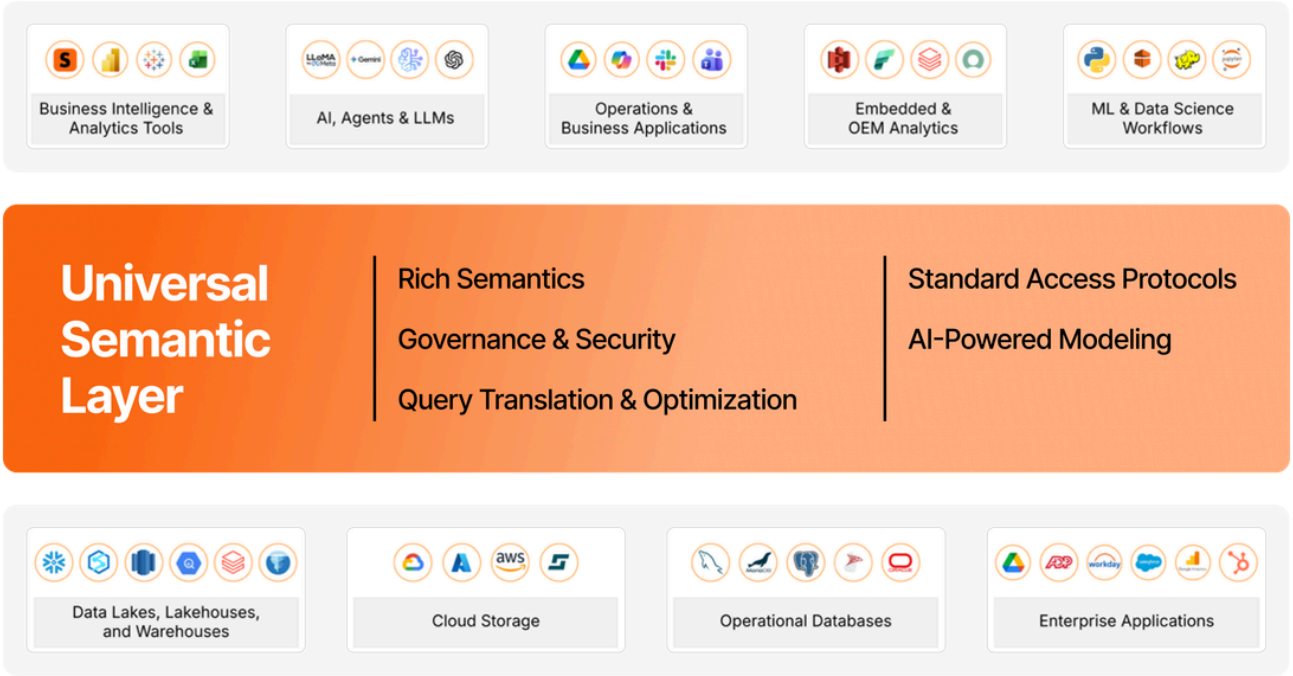

5. Active Security & Governance

Security must be enforced consistently across all access points, especially as AI agents become data consumers. Modern governance must be continuous and contextual. Policies should travel with the meaning of the data, not just its storage location. When AI systems interact through MCP and a semantic layer, governance is enforced at the point of execution with automatic audit trails, rather than reconstructed after the fact. Look for centralized RLS/CLS/RBAC, active monitoring with anomaly detection, lineage tracking, data protection for LLM interactions, and mature lifecycle management. On lifecycle specifically, the strongest platforms now offer native version control integration (for example, Git-based workflows) that stores semantic definitions as text-based objects, enabling branching, rollback, and CI/CD pipelines for model changes. This makes semantic model governance as rigorous as software development, not a manual, error-prone process. 82% of leaders rate governance and observability as very or extremely important.

Watch out for: static, configure-and-forget governance that fragments across databases.

6. AI Agent Readiness

The semantic layer is the governed bridge between AI and enterprise data. In the CIO Dive survey, 87% value visibility into AI's use of company data. The UserEvidence study showed a 22% decrease in AI hallucinations and 28% faster AI deployment with a governed semantic layer. Look for MCP integration with strong tenant isolation and enterprise security controls. Unlike standard API integrations, MCP enforces authentication, authorization, and audit logging as part of the protocol itself, not as an application-layer add-on. Also look for fine-grained access control for agents, comprehensive machine-to-machine audit trails, and architecture that prioritizes metadata-first patterns and minimizes data exposure to models.

Watch out for: basic SQL/API access without semantic richness, governance enforcement, or audit capabilities.

7. Performance, Scalability & Cost

A semantic layer that cannot deliver sub-second responses on large datasets will not be adopted. Consumption-based pricing can become difficult to forecast as adoption grows. Look for hybrid execution engines (pushdown + in-memory), semantic-aware caching with freshness policies, a cost-based optimizer, and pricing that remains forecastable as adoption grows. The UserEvidence study found a 20% decrease in downstream compute costs and organizations retiring an average of two BI tools. 99% of CIO Dive respondents consider predictable pricing at least somewhat important.

Watch out for: solutions relying entirely on underlying database performance, or consumption-based pricing that discourages broad access.

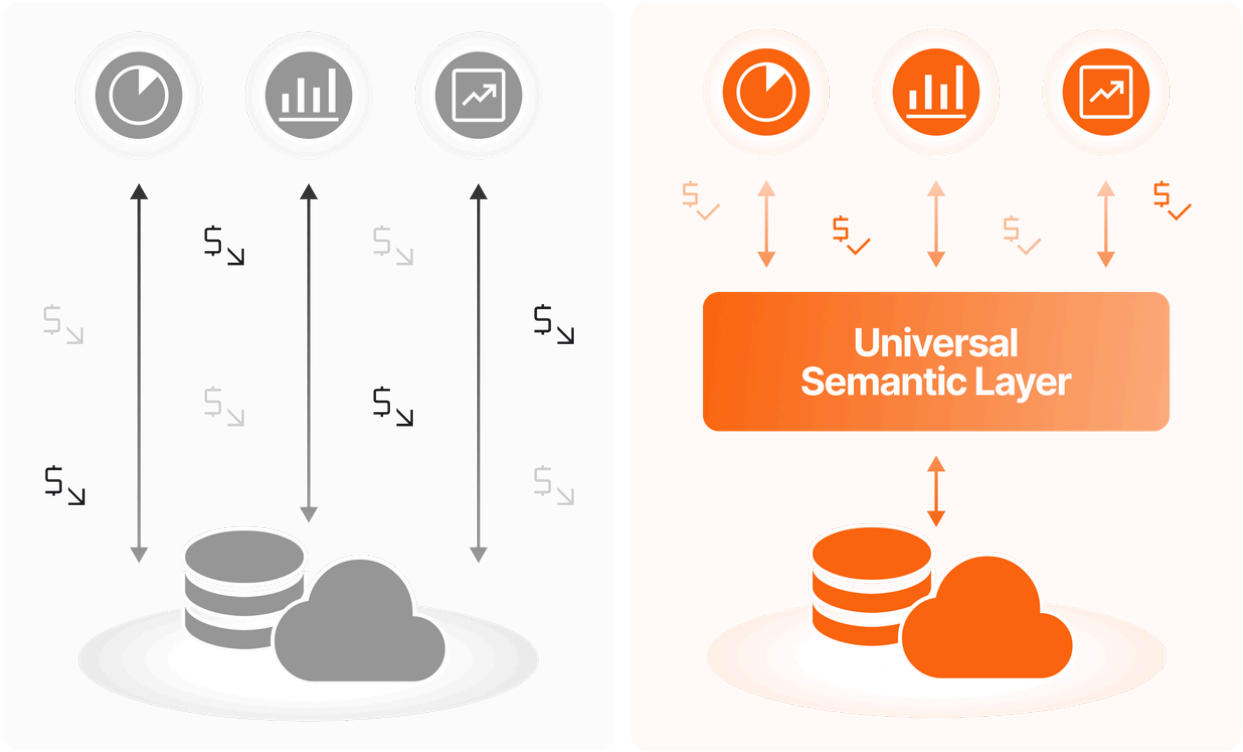

Semantic Layer Maturity Model

Assess where your organization stands and what to prioritize:

Level 1: Basic (Embedded and Fragmented)

Semantic definitions are embedded in individual BI tools or data warehouses, each with its own version of key metrics. Security is configured separately per platform with no centralized governance. Common symptoms include conflicting reports across departments, metrics that must be rebuilt when tools change, and AI initiatives that lack a governed data foundation. To move forward, identify a single high-impact KPI that is calculated differently across departments and use it as the basis for a proof of value.

Level 2: Intermediate (Standardized but Limited)

A dedicated semantic layer provides standardized definitions across some tools with basic governance (RLS/RBAC), but federation is limited and modeling is largely manual. Definitions are consistent within the layer's reach, but gaps remain for unconnected tools and sources. Governance tends to be configure-and-forget without active monitoring, and AI agents cannot easily consume governed semantics. The typical next step is extending federation to additional sources, introducing AI-assisted modeling, and enabling governed access for machine consumers alongside human ones.

Level 3: Advanced (Comprehensive and AI-Ready)

A universal semantic layer operates across all data sources, clouds, and consumption tools. AI-powered modeling accelerates development. Active governance continuously monitors and enforces policies across human and machine consumers. MCP and agentic AI support is integrated with full audit trails. In-memory acceleration provides independent performance with full portability across platforms. Target capabilities include broad connector coverage (100+ in mature platforms), AI-assisted semantic modeling, active governance with lineage and anomaly detection, MCP support, an optional acceleration layer, pricing that remains forecastable as adoption grows, and OSI alignment. With appropriate scoping, an initial governed domain is typically deliverable within 6–8 weeks. At this level, the focus shifts to continuous optimization: expanding agent coverage, refining governance policies based on usage patterns, and measuring business impact across the enterprise.

Recommendations

Select a semantic layer designed for analytics and AI at scale, not just BI metrics. Four principles should guide the decision:

- Favor a universal semantic layer that centralizes governed definitions and relationships, serves them consistently across tools, and supports federation so data can be unified without being consolidated.
- Look for an MCP-style integration pattern for AI and agents that supports authenticated, authorized, and auditable access.
- Prioritize continuous governance with near real-time visibility into access, lineage, compliance, and cost patterns.
- Ensure anomaly detection and policy enforcement operate where work happens, not after the fact.

Where to start based on your environment

Multi-cloud, multi-tool environments: Prioritize Level 3 capabilities. When data and consumption span multiple platforms and tools, metric drift and duplicated logic increase quickly. Choose a solution that supports federation, reuse, and governance so definitions remain consistent and changes propagate safely across teams.

Single-platform environments: Even with one warehouse, multiple BI tools quickly create duplicated definitions and metric drift. A Level 2 solution can standardize metrics across tools, and a Level 3 solution adds federation and broader portability as your stack evolves. Vendor independence protects your business logic if you add new tools or migrate platforms later.

Starting the journey: Begin with a high-impact cross-functional KPI defined inconsistently today (e.g., revenue, customer count). Use a proof of value to demonstrate impact, then expand. Prioritize solutions that grow from focused deployment to enterprise coverage without re-architecture.

AI-forward organizations: Prioritize active governance, AI agent readiness, and data quality. AI magnifies both the value and the risk of your data foundation. A semantic layer providing governed, auditable, semantically rich context for LLMs and agents is a prerequisite for trustworthy AI at scale.

Key Selection Criteria Summary

1. **Demand platform independence and real federation** — your business logic should outlast any data platform, and your semantic layer must span all your sources.
2. **Evaluate AI modeling and data readiness together** — manual-only development does not scale, and semantic quality depends on data quality.
3. **Verify semantic depth** — basic metrics are not enough for enterprise needs; look for rich metadata, hierarchies, and business context.
4. **Insist on open standards** — and ask about participation in OSI or similar initiatives, and the ability to interoperate with other semantic layers.
5. **Centralize and activate governance** — static governance is no longer sufficient, especially with AI agents as data consumers.
6. **Assess AI agent readiness** — including MCP support, machine audit trails, and hallucination reduction.
7. **Test performance at scale and model total cost** — proof of value with real data volumes, and full TCO including infrastructure savings and cost arbitrage.
8. **Validate time-to-value** — an initial governed domain should be deliverable within **~6–8 weeks** with appropriate scope. Long discovery phases before any user-visible outcome are a delivery risk.

Making the Case Internally

Selecting the right semantic layer is only half the challenge. The other half is building the internal case that moves a technical evaluation into a funded initiative.

Anchor to a known pain point. Start with a problem leadership already recognizes: a KPI that different teams calculate differently, a recurring data dispute that delays decisions, or an AI pilot stalled by inconsistent data. Frame the semantic layer as the solution to that specific problem, not as a platform purchase. Then use the proof of value to produce before-and-after evidence: time saved on reconciliation, reduction in conflicting reports, improvement in query response times.

Translate for the boardroom. Executive sponsors think in business outcomes, not technical capabilities. Reframe accordingly:

- Faster queries → faster decisions
- Eliminated metric duplication → reduced risk and audit exposure
- Federation across sources → infrastructure flexibility and lower switching costs
- Governed AI foundation → responsible scale rather than ungoverned proliferation

The UserEvidence ROI data in this guide—$3.4M average net annual impact, two-month payback, 46% of analyst time reclaimed—provides a credible external benchmark to anchor these conversations.

IT Leader

Premium automotive manufacturer

> *"We unified six departments into a single governed reporting ecosystem. Case resolution times improved 5–9% across support functions in the first year."*

Semantic Layer Evaluation Checklist

The following checklist maps to the seven evaluation criteria above. Use it during vendor demonstrations and proofs of value:

Platform & Data

□ Semantic model is platform-independent (not tied to one BI tool or warehouse)
□ KPIs and security remain reusable across cloud/platform changes
□ Can query/join across multiple sources without mandatory replication
□ Supports federation plus optional caching/acceleration
□ Supported connectors for your top data sources and tools

Modeling & Semantics

□ AI accelerates modeling (discovery, relationships, metric creation) with guardrails
□ Supports both visual and programmatic development workflows
□ Hierarchies, time intelligence, semi-additive/non-additive metrics
□ Rich metadata including AI-generated descriptions and lineage

Standards & Access

□ Governed semantics accessible across all BI tools and custom apps
□ APIs/SDKs available for automation, embedding, and AI agent access
□ No requirement to duplicate definitions per tool
□ Participates in open interoperability standards (OSI or similar)

Governance & AI

□ Central RLS/CLS defined once and enforced across BI and AI consumption
□ Active monitoring: lineage, anomaly detection, cost surveillance
□ MCP support with authenticated, governed agent access
□ Machine-to-machine audit trails for AI interactions

□ Agent request can be authenticated, authorized, executed, and logged end-to-end, including which semantic objects were used
□ Row/column security outcomes are validated to be consistent across BI and agent access paths for multiple user roles
□ Semantic model versioning (history, branching, rollback)
□ Dev/test/prod promotion for semantic changes
□ Usage and cost observability by semantic object (who/what drives spend)

Performance & Cost

□ Meets p95 latency targets on production-shaped data
□ Supports realistic concurrency (dashboards plus ad hoc)
□ Demonstrates acceleration beyond warehouse tuning
□ Costs are measurable and forecastable as adoption grows
□ Clear TCO model (license + infrastructure + operations)

Validation

□ KPI results match across at least two consumption paths
□ Drill paths reconcile from executive KPI to detail
□ Evidence captured (queries, results, screenshots)
□ First governed domain deliverable within ~6-8 weeks with appropriate scoping

Next Steps

Vendor evaluation

Score each vendor against the seven criteria weighted by your priorities. Request demos against your actual data sources—not canned demos.

Proof of value (POV)

Design a focused POV around a contested KPI, one that is currently calculated differently across tools or teams.

Key steps:

a. **Define the KPI** in the semantic layer and validate it renders identically across at least two BI tools and one AI agent interface, with different row-level security applied to each consumer.
b. **Test federation** by joining data from two or more sources without replication.
c. **Verify that governance policies** (RLS, CLS, masking) are enforced consistently across every access path, including programmatic and AI-driven requests.
d. **Measure query response** times under realistic concurrency.
e. **Document results** with implementation results on test data, query logs, and before-and-after comparisons—this evidence becomes the foundation of the business case.

Reference checks

Request references from customers who have migrated data platforms while preserving their semantic layer.

TCO analysis

Calculate full cost including licensing, infrastructure, operations, and impact on database compute. Factor in the cost of not acting: ongoing metric inconsistency, duplicated effort, and delayed AI initiatives.

The organizations that succeed with AI will not be defined by the number of models they deploy, but by the strength of the architecture that connects those models to trusted, governed data.

To see how Strategy Mosaic maps to this framework, request a scoped proof of value at https://www.strategy.com/software/strategymosaic

A. Implementation Risk & Change Management

A semantic layer evaluation rarely fails on technology. It fails on underestimated migration complexity, organizational change, and time-to-value expectations. Before selecting a vendor, pressure-test each of these dimensions.

Migration Complexity

Most enterprises carry years of business logic embedded across BI tools, spreadsheets, and data transformation pipelines. Moving to a universal semantic layer means inventorying, reconciling, and re-implementing that logic. It is not just a matter of installing new software.

- **Audit existing definitions first.** Before any vendor engagement, catalog how many distinct metric definitions exist and where they live. This becomes your migration scope and your before-state for measuring success.
- **Prioritize, don't try to do everything at once.** Start with one high-conflict KPI (typically revenue or customer count) and prove consistent results across two tools. Expand in 90-day increments.
- **Assume 25% longer than estimated.** Integration timelines routinely slip due to undocumented logic, data quality gaps, and security policy translation. Build buffer into your business case.
- **Validate connector coverage early.** Confirm that your specific BI tool versions, warehouse configurations, and data source combinations are supported, not just the product category.

Change Management

Technology adoption fails most often at the people layer. A new semantic layer changes how analysts find data, how developers define metrics, and how business users trust reports.

- **Identify a champion in each business unit.** Grassroots adoption requires a trusted voice in Finance, Marketing, and Operations who validates that the new definitions are correct—not just IT saying so.
- **Train for the new workflow, not just the tool.** Analysts need to understand why they should use the semantic layer instead of querying raw tables. The “why” matters as much as the “how.”
- **Expect a trust gap at launch.** When a unified metric first differs from a department’s historical number, stakeholders will question the semantic layer, not their old logic. Prepare narrative and evidence in advance.
- **Retire old definitions on a schedule.** Leaving legacy metric logic in place alongside the new layer creates a dual-source problem. Set a deprecation timeline with teeth.

Time-to-Value Expectations

Use this as a realistic planning benchmark, not a vendor promise:

Milestone	Realistic Timeline	Primary Risk
First governed KPI live	6–8 weeks	Undocumented legacy logic
Two BI tools consuming same layer	8–16 weeks	Connector and permission gaps
Governance policies enforced end-to-end	3–6 months	Cross-team policy alignment
AI agents consuming governed semantics	4–9 months	MCP configuration and audit setup
Full enterprise rollout	12–24 months	Change management at scale

Watch out for: vendors who quote go-live timelines without first reviewing your existing metric inventory and connector requirements. A vendor who can give you a fixed timeline without a discovery phase is guessing.

B. Tough Questions to Ask Vendors

The evaluation checklist tells you what to look for. These questions are designed to move beyond feature overviews and into specifics that reveal how a solution actually behaves under enterprise conditions. Use them in demos, technical deep-dives, and reference calls.

On Platform Independence

> *"If we change our data warehouse or replace a BI tool next year, what needs to be rebuilt in your solution, and how long does that take?"*
>
> - A strong answer walks through a specific migration scenario with minimal disruption to semantic definitions. An answer of "nothing breaks, it's fully portable" without supporting detail warrants a follow-up demonstration.

> *"Can you connect us with a customer who has migrated data platforms while keeping their semantic layer intact?"*
>
> - Customer references carry more weight than architecture diagrams. A vendor with mature portability will have this story readily available.

> *"Can you demonstrate a single governed metric that queries across two of our data sources in one scenario, without requiring a full data consolidation project? If caching or acceleration is used, explain what is cached, where it runs, and how it is governed."*
>
> - A strong answer should show one governed metric running across multiple live sources without full consolidation, with clear handling of execution, caching, and security.

On Governance & Security

> *"If an AI agent makes a request through MCP that the requesting user's permissions would deny, what happens? Please show the resulting audit log entry, including the policy decision and user context."*
>
> - This should be demonstrable, not described. If they cannot show you the log, governance for AI is not production-ready.

> *"How are governance policies stored and versioned? If a policy change causes a regression, how do you roll back?"*
>
> - Look for version control, promotion workflows, and rollback capability. "We test before deploying" is not a sufficient answer.

> *"Who in our organization can change access policies, and what controls prevent a privileged user from granting broad access without detection?"*
>
> - Enterprise governance requires separation of duties and monitoring. Ask for approval workflows and alerts for high-risk changes.

On AI & Agent Readiness

> *"Can you demonstrate that the same row-level security policy is enforced consistently for a BI user, a Python notebook, and an AI agent, using the same underlying semantic definitions with no per-tool policy duplication?"*
>
> - This exposes solutions that define security per tool rather than once at the semantic layer. A mature platform applies the same policy uniformly across all access paths without duplication.

> *"Are there any alternate execution paths that could bypass policy enforcement, including direct SQL? What controls exist to restrict them, and what alerts or logs would tell us an attempt occurred?"*
>
> - The answer should describe active enforcement and detection, not only policy configuration.

> *"What logging and monitoring do you provide for agent queries? Please walk through a real example showing user context, agent identity, policy decisions, and any semantic objects referenced."*
>
> - A dashboard view or export is fine, but it should be from a live environment.

On Implementation & Support

> *"What does onboarding look like for a company with our profile, including our approximate number of metrics, data sources, and BI tools?"*
>
> - A vendor with real implementation experience will ask clarifying questions before answering. A vague answer signals a generic process.

> *"What does your support model look like in year two, after the initial implementation team has rolled off?"*
>
> - Get specifics on named contacts, SLAs, escalation paths, and how governance changes are supported over time.

> *"What is the fastest a customer with our profile has gone from kickoff to first governed KPI in production? Can you connect us with that reference?"*

- This grounds the time-to-value conversation in reality rather than marketing timelines. A vendor who cannot provide a specific reference is estimating, not quoting from experience.

On Cost & TCO

> *"How does your pricing model work as adoption scales, for example if we double users or query volume in year two? Is cost tied to consumption, users, or a fixed model?"*

- Look for pricing that remains forecastable as adoption and usage grow.

> *"What infrastructure costs are not included in the license, and what do reference customers typically spend on infrastructure alongside the license?"*

- Total cost of ownership includes compute, storage, egress, and operational overhead. Get a full number, not just the license line.

> *"If our needs change and we need to migrate off your platform, what does that process involve and in what format can we export our semantic model?"*

- Portability commitments should be contractual, not verbal. Ask what format the semantic model exports to and whether that format is open.

Bring at least one question from each category to every vendor interaction.
How a vendor responds to specific, scenario-based questions reveals as much about their solution as any feature demonstration.

Glossary of Terms

ACL (Access Control List): Access permissions controlling which users and roles can access which objects.

Caching: Storing query results to improve performance and reduce warehouse compute costs.

CLS (Column-Level Security): Security controls restricting access to specific columns or fields.

Data Federation: Querying data across multiple sources without physically moving or copying the data.

In-Memory Processing: Data processing in RAM rather than on disk, enabling sub-second query performance.

Knowledge Graph: A programmable network representing entities and their interrelations, enabling inference and entity resolution.

MCP (Model Context Protocol): An open-source standard defining a secure protocol for AI applications to communicate with data sources and applications. Acts as a universal adapter for AI agent integration.

Metric Drift: When the same KPI is calculated differently across teams or tools over time, leading to inconsistent results.

OSI (Open Semantic Interchange): An industry initiative creating a vendor-neutral standard for semantic models. Uses a declarative YAML standard (MetricFlow) to define metrics, dimensions, and joins for interoperability across tools.

RLS (Row-Level Security): Security controls restricting data access at the row level based on user attributes.

Semantic Layer: A business abstraction layer translating technical data structures into business-meaningful terms and metrics.

Active Governance (Monitoring & Controls): A capability that continuously monitors and enforces policy across data access and usage, tracking lineage, security and compliance signals, and cost and consumption patterns in near real time.

Universal Semantic Layer: A vendor-independent semantic layer operating across all data sources, clouds, and consumption tools, serving both human and AI consumers.

Sources

CIO Dive Studio. “Data, AI & Analytics Trends Across Organizations in 2026.” Hosted by Strategy, October 2025. https://www.strategy.com/software/research-and-reports/data-ai-analytics-trends-across-organizations-in-2026
UserEvidence. “How Strategy’s Semantic Layer Delivers $3.4 Million & 551% ROI.” 2026 ROI Survey Report. https://www.strategy.com/software/research-and-reports/2026-roi-survey-report
DBTA Best Practices Series. “How to Build a Scalable Semantic Layer for AI and Analytics.” February/March 2026. https://www.dbta.com/DBTA-Downloads/WhitePapers/How-to-Build-a-Scalable-Semantic-Layer-for-AI-and-Analytics-14630.aspx

Market sizing references

[1] Verified Market Research (via PR Newswire). “Data Virtualization Market Size Worth $8.39 Billion Globally by 2028 at 20.9% CAGR.” February 2022. https://www.prnewswire.com/news-releases/data-virtualization-market-size-worth--8-39-billion-globally-by-2028-at-20-9-cagr-verified-market-research-301483610.html
[2] Research and Markets. “Enterprise Data Management Market Size, Share & Trends Analysis Report, 2023–2030.” January 2023. https://www.researchandmarkets.com/reports/5415471/enterprise-data-management-market-size-share-and
[3] Zion Market Research. “Data Warehouse Market by Type, Deployment, End-User: Global Industry Perspective, Comprehensive Analysis, and Forecast, 2025–2034.” December 2025. https://www.zionmarketresearch.com/report/data-warehouse-market
[4] LXA. “Cloud Computing: Stats and Trends.” June 2022. https://www.lxahub.com/stories/cloud-computing-stats-and-trends

Strategy Incorporated, 1850 Towers Crescent Plaza, Tysons Corner, VA 22182.

strategysoftware.com

Strategy₿