




The Universal Semantic Layer Advantage

How Strategy Mosaic unifies your data, accelerates AI,
and future-proofs analytics with a Universal Semantic Layer

Start



The cost of fragmented data

The modern enterprise has more data than it can handle. Disjointed ecosystems, inconsistent metrics, fragmented security, and vendor lock-in stall analytics and cause AI initiatives to fail, raising costs while eroding trust in every KPI.

What is the real problem?

The problem lies in disconnected environments: legacy systems, multiple clouds, and tool sprawl, which slow delivery and inflate cost.

Disjointed data and inconsistent definitions: Teams using different tools define business terms and other core KPIs differently, so leaders see competing answers to the same questions.

Security and compliance risks: Security is duplicated and inconsistent across platforms, increasing compliance risk and making policy enforcement brittle.

Vendor lock-in: When data strategy tethers too tightly to a single vendor, change becomes expensive, blocking modernization and throttling innovation.

Net result: Analytics slows to a crawl and AI initiatives never reach scale because the foundation isn't unified or trusted.

Solution: Organizations need a common layer for business definitions, access, and governance, portable across tools and clouds, so everyone calculates, secures, and uses data the same way.



Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Why traditional paths fall short

Organizations typically try four routes to solve their data connectivity problem:

- 1 **BI tool-centric** approaches multiply definitions and duplicate logic inside each tool.
- 2 **Data virtualization** reduces movement but struggles with complexity and performance, staying SQL-centric.
- 3 **Data warehouses** centralize data yet increase cost, require heavy ETL, and deepen vendor lock-in.
- 4 **Standalone semantic layers** promise consistency but often lack integrated AI capabilities and a high-powered engine to accelerate queries across sources.

Each path addresses symptoms, not the core enterprise need: consistent semantics, performance, governance, and portability across tools and clouds. The alternative is a **Universal Semantic Layer** that standardizes business definitions, routes queries intelligently (live or cached), and enforces security centrally, without forcing a re-architecture of your data stack.

A universal layer unifies metrics, governance, and speed—without rebuilding your stack.

Self-service that meets users where they work

Give everyone intuitive, governed access to data with AI-powered modeling and natural-language queries, embedded in the tools they already use.

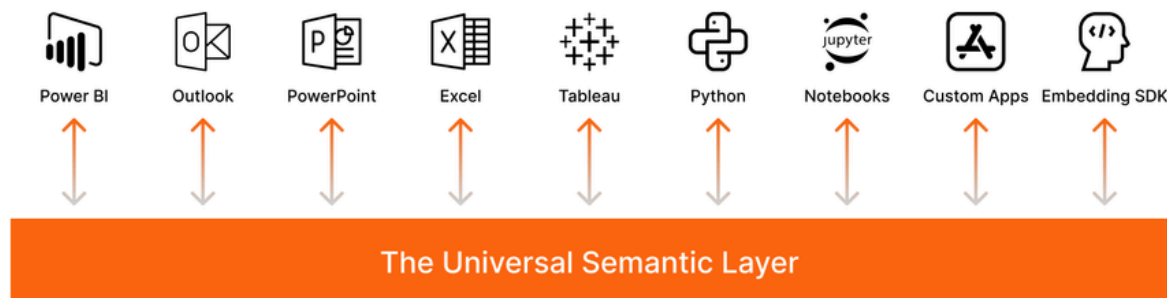


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Universal Semantic Layer

A Universal Semantic Layer is a vendor-neutral layer that standardizes business definitions and business logic, speeds queries, and governs access. It's portable across tools and clouds, and accelerated by AI-assisted modeling.

Mosaic is a Universal Semantic Layer that **decouples business logic and security from underlying databases and clouds**. With Mosaic, definitions and controls travel with you, not with a single platform.

The result is **consistent metrics everywhere** (Power BI, Tableau, Excel, apps, and AI) via open interfaces like SQL, DAX/XMLA, REST, and Python.

One semantic model, many tools: consistent answers, faster performance, centralized governance.

Where the layer lives

Decouple definitions and security from data stores; plug into BI, apps, and AI.

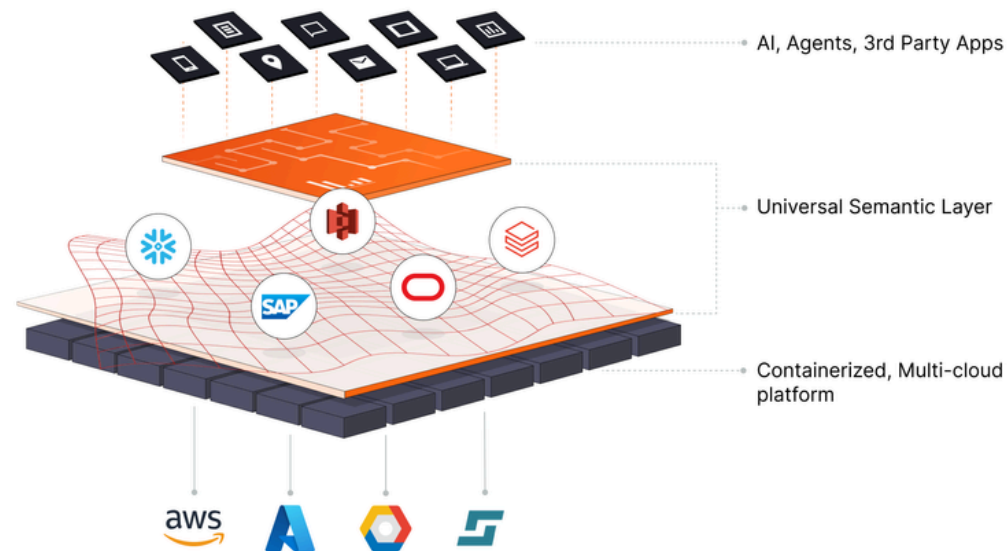


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Critical capabilities at a glance

A Universal Semantic Layer must ensure extensive data connectivity to unify diverse data sources, coupled with performance optimization and data capabilities.

The must-have checklist:

- **Unified definitions:** Consistent across platforms and tools for trusted insights
- **Cross-platform compatibility:** Connecting 200+ sources joining across clouds, and working with your preferred BI
- **Performance at scale:** Hybrid execution, in-memory acceleration, semantic-aware caching, and a cost-based optimizer
- **Centralized governance:** Row/object security, masking, auditing, and LLM-safe access
- **AI-readiness:** Exposing structured, governed semantics to LLMs and tracking lineage for auditability
- **Self-service:** Via business objects and natural-language metric creation
- **Portability:** Across clouds and databases without the need for re-architecture
- **Cost optimization:** By minimizing movement/egress and arbitraging platforms as needs change

Mosaic checks all the boxes providing universal data connectivity and performance optimization at scale.

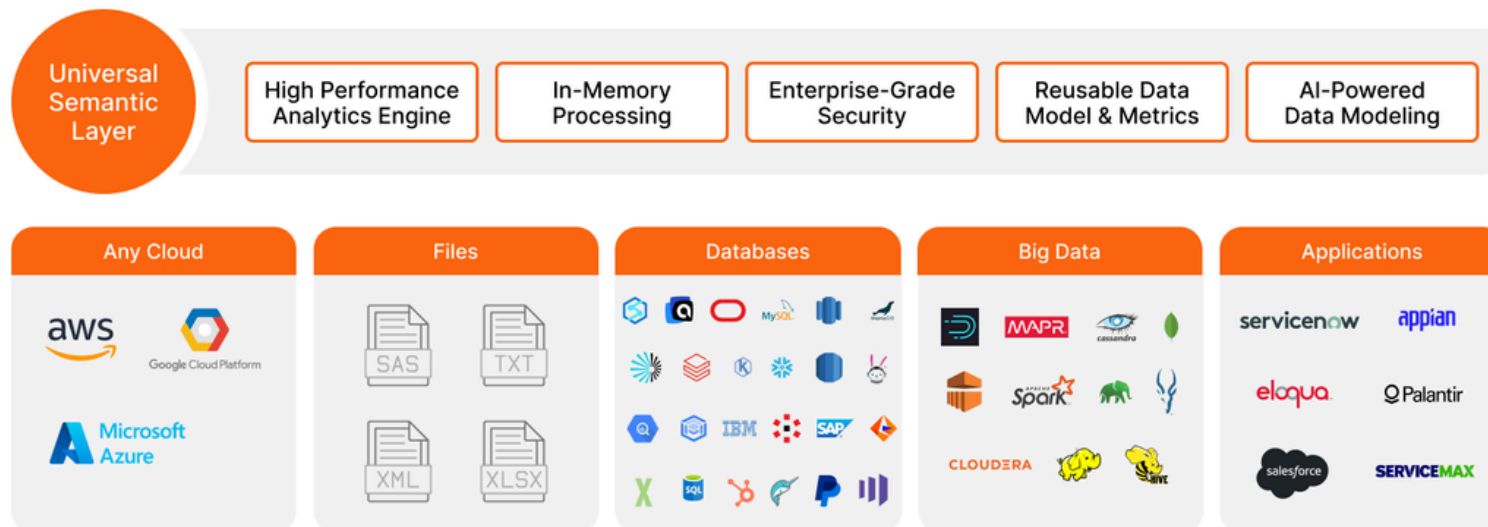


Table of Contents

- 1 The cost of fragmented data
- 2 Why traditional paths fall short
- 3 Universal Semantic Layer
- 4 Critical capabilities at a glance
- 5 Vuori: One source, faster insights
- 6 Inside the semantic layer: Unified definitions
- 7 Mosaic Studio: AI for data experts
- 8 Built for reliable AI applications
- 9 Multi-cloud agility, zero vendor lock-in
- 10 Connectivity, performance, and security

Vuori: One source, faster insights

Vuori implemented the semantic layer technology, which Mosaic is built upon, to fix a familiar problem: conflicting metrics and slow time-to-insight. The result was a single source of truth that aligned definitions across teams, faster executive reporting cycles, and AI-powered drill-downs that helped analysts move from “what” to “why” in one place.



Himanshu Shekhar

VP of Planning & Analytics

“We implemented Strategy to enable our business leaders and analysts to become more self-serving, rooted in a single version of truth.”

These wins came without forcing a costly rebuild of downstream tools, thanks to a vendor-neutral layer that carries consistent logic wherever work happens.

Mosaic delivers outcomes that matter, quickly and repeatably.

Single truth, faster reports, deeper analysis—measurable impact without re-platforming.



Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Inside the semantic layer: Unified definitions

From raw tables to business-friendly objects and consistent metrics, automated relationships, in-memory calculations, and rich metadata make trustworthy answers travel across tools.

Mosaic's **semantic layer** turns technical schemas into business objects (e.g., Customer, Revenue) with consistent, reusable metrics that work everywhere: Power BI, Tableau, Excel, apps, and AI.

An **AI engine** auto-detects relationships and builds hierarchies, accelerating modeling while keeping structure transparent for data teams.

A high-performance **in-memory calculation engine** supports pushdown where possible and cross-source calculations where needed. Distributed data still yields one answer to questions like revenue or inventory.

Models carry metadata (including AI-generated descriptions) that improves discovery and fuels AI reliability by giving LLMs governed context rather than raw tables.

The net effect: a durable semantic backbone that unifies logic, speeds queries, and safeguards governance, without forcing centralization.

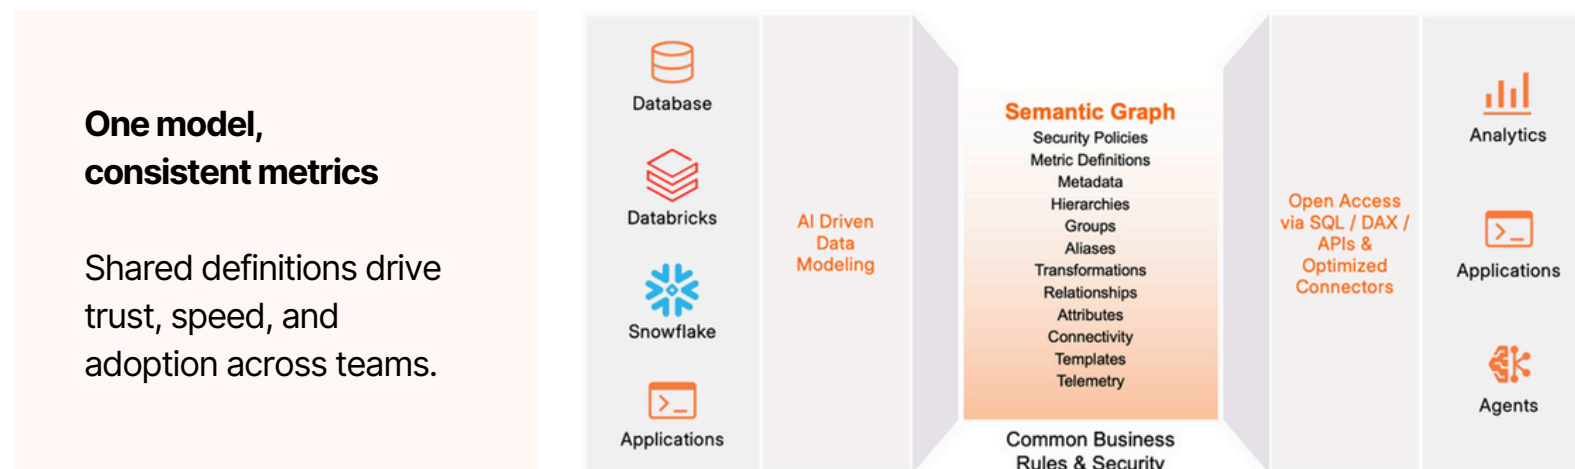


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Mosaic Studio: AI for data experts

Mosaic Studio speeds semantic modeling by coordinating dedicated **AI, Modeling, Data, and Compute** services.

The **AI engine** runs LLM-based operations via parameterized APIs, chaining steps and generating contextual prompts to refine models.

Automated **statistical profiling** inspects metadata and samples to identify entities, map semantics, and propose production-ready metrics.

A **relationship-detection system** evaluates column semantics, assigns confidence scores, and implements high-confidence joins while flagging others for review and validation against live data.

Meanwhile, **ETL-free preparation** surfaces and fixes common issues: missing values, duplicate rows, and inconsistent formats. Teams publish faster with fewer pipelines.

Customers report up to 87% faster time-to-first model and 10× throughput versus traditional methods, turning weeks of effort into minutes.

Foundation for universal Business Intelligence

AI, Modeling, Data, and Compute services collaborate to generate and validate semantic models.

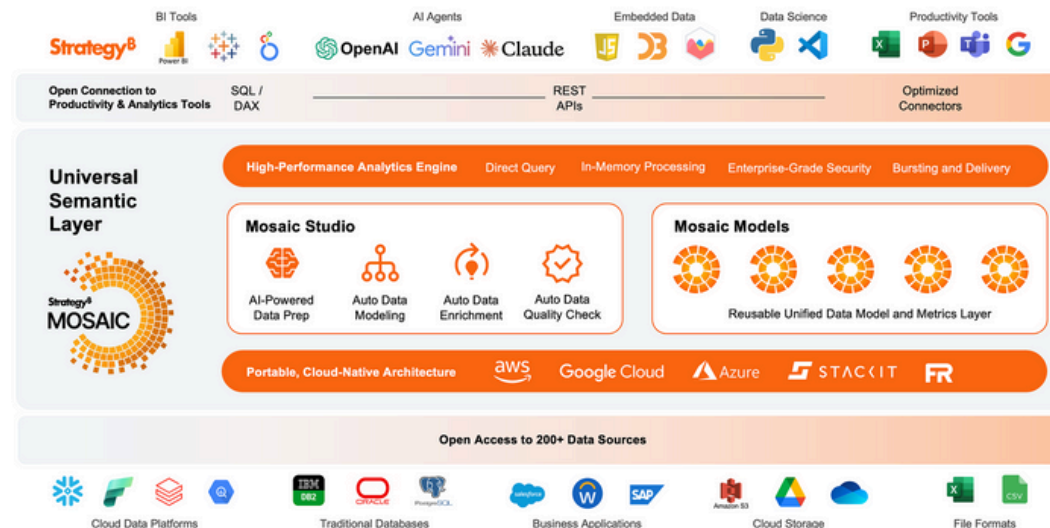


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Built for reliable AI applications

Direct LLM-to-database prompts miss business rules and create risk.

Mosaic places a semantic layer between prompts and compute: an abstraction of interconnected objects with unique definitions that enforces consistent rules across AI interactions and reduces hallucination risk.

Models are exposed through **standards-based interfaces** (SQL, REST, Python) while access controls prevent sensitive data from reaching LLMs.

Enterprise-grade governance includes automated metadata enrichment, relationship validation, and semantic consistency enforcement, with built-in lineage to trace how AI conclusions were derived, vital for compliance and trust.

The result: reliable, explainable AI that operates on governed semantics rather than raw tables, integrating safely with your existing applications and workflows.

Governed semantics reduce hallucinations and make AI outputs explainable, and auditable.

Built for AI applications: Semantic richness for AI

Strategy Mosaic is designed with AI at its core, both to accelerate model building and to power the next generation of AI applications by providing rich semantic context and trusted data.

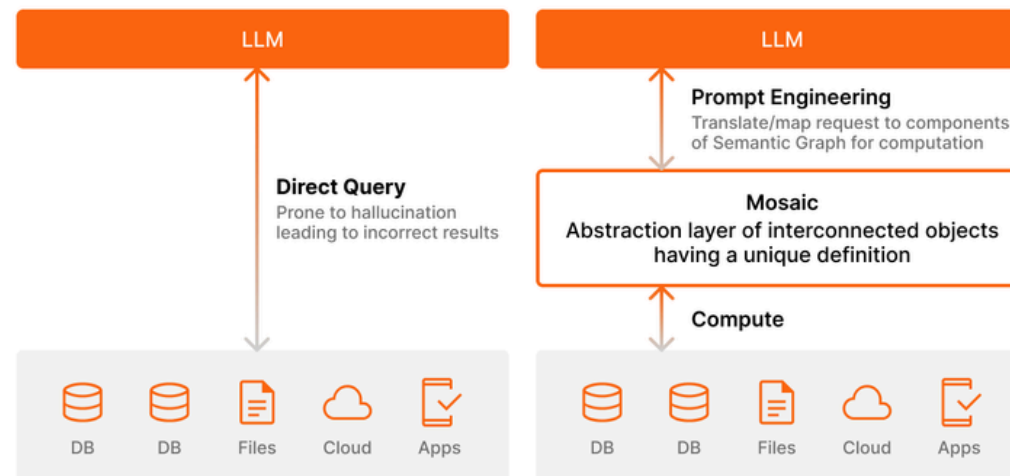


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Multi-cloud agility, zero vendor lock-in

Mosaic decouples business logic and security from storage and compute, so you can **switch clouds or databases** without refactoring models or rebuilding reports.

It's designed to run on AWS, Azure, GCP, StackIT, and hybrid environments, giving you true portability and future-proofing your data strategy.

Performance is accelerated across sources with an **in-memory engine** and optimizer, particularly when joining or querying older or slower systems.

Costs stay in check through **live-query vs. cached** execution, minimizing egress, and enabling **cost arbitrage** across platforms when prices or needs change.

Change clouds or databases without rebuilding analytics, models, or security policies.

One model, many environments

Port your definitions and security across clouds, databases, and tools. No re-architecture required.

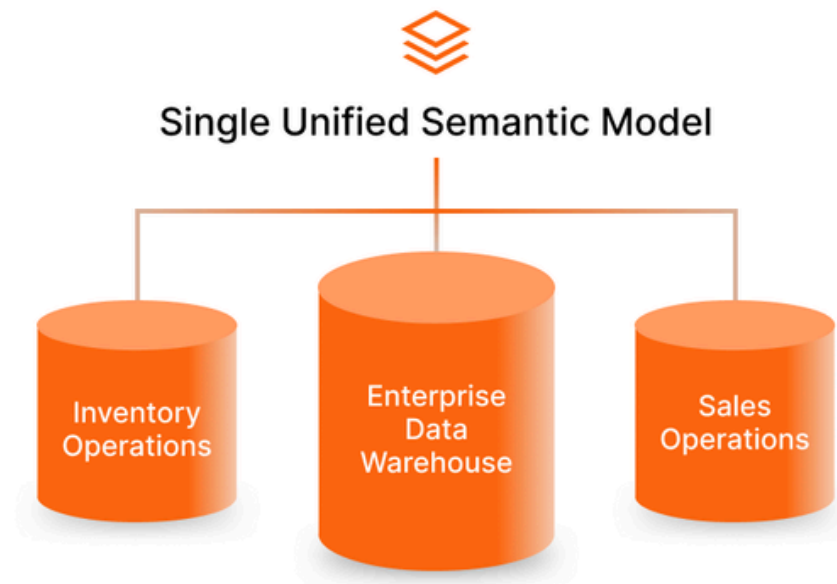


Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)

Connectivity, performance, and security

Mosaic exposes 200+ sources as unified, semantically enriched tables via its **Universal Access Service**, with intelligent routing for optimal performance across clouds and on-prem.

A **hybrid execution** engine chooses pushdown or in-memory paths; multi-level, semantic-aware caching adds freshness controls and versioned tables.

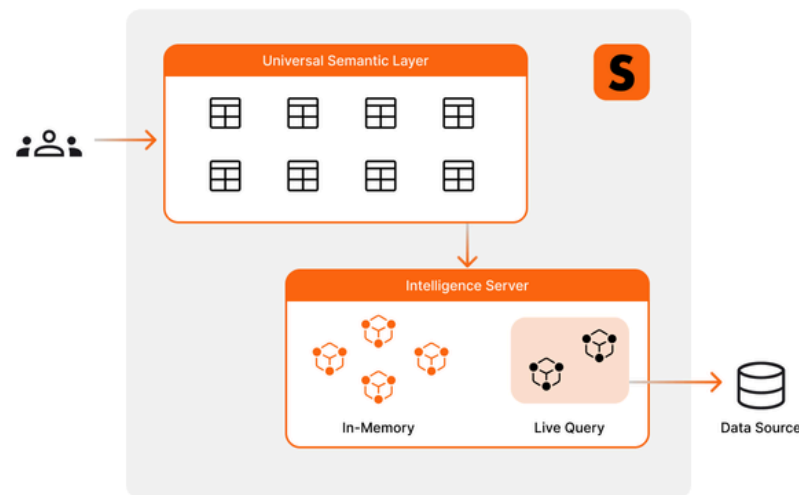
A **cost-based optimizer** (pushdown, join/aggregation re-ordering, materialized-view matching) accelerates responses.

Governance is centralized: row-level and object-level controls, masking, auditing, and guardrails that shield sensitive data from LLMs.

Delivered as a **fully managed cloud service**, Mosaic is instantly accessible as a modeling workspace and universal semantic endpoint—no infrastructure to stand up.

Universal connectivity

From tools to data, routed and governed through one semantic layer.



Strategy 

Copyright ©2025. All Rights Reserved.



Table of Contents

- 1 [The cost of fragmented data](#)
- 2 [Why traditional paths fall short](#)
- 3 [Universal Semantic Layer](#)
- 4 [Critical capabilities at a glance](#)
- 5 [Vuori: One source, faster insights](#)
- 6 [Inside the semantic layer: Unified definitions](#)
- 7 [Mosaic Studio: AI for data experts](#)
- 8 [Built for reliable AI applications](#)
- 9 [Multi-cloud agility, zero vendor lock-in](#)
- 10 [Connectivity, performance, and security](#)