


Databricks Genie Ontology

The Living Context Layer for Enterprise AI

Executive Briefing & Implementation Guide

Based on the June 16, 2026 Data+AI Summit Launch



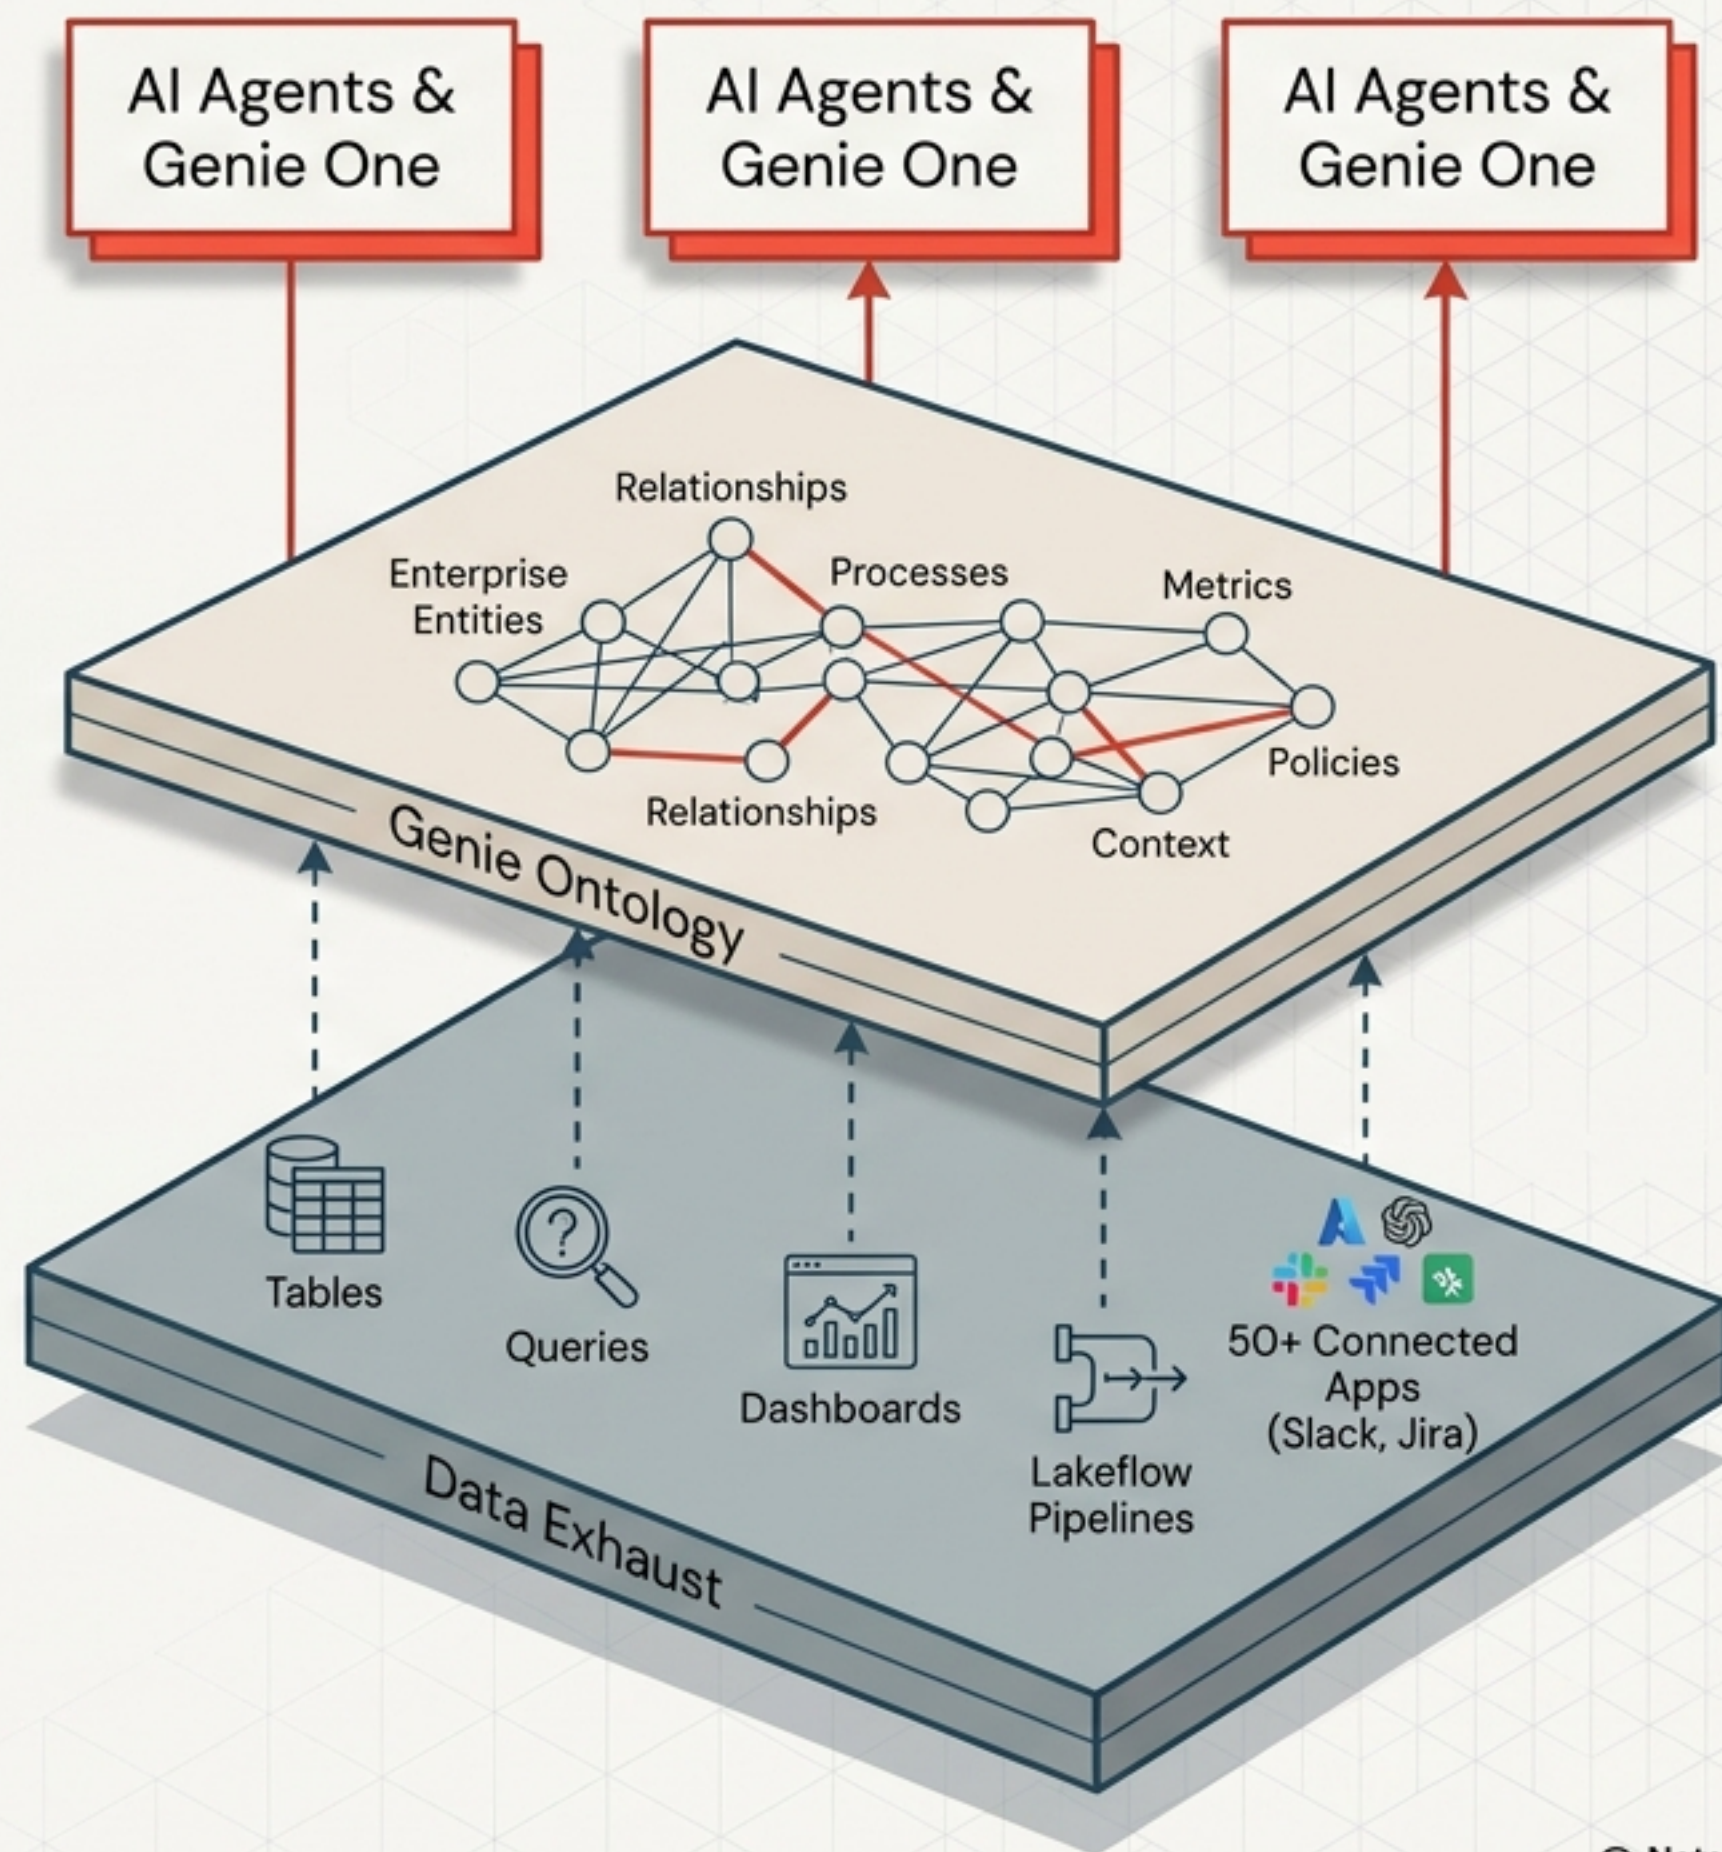
What is Genie Ontology?

The Paradigm Shift: From Static Schemas to Living Context

An automatic, self-improving knowledge graph built natively on Unity Catalog that extracts snippets of knowledge from enterprise data and organizes it into a living map of how a company works.

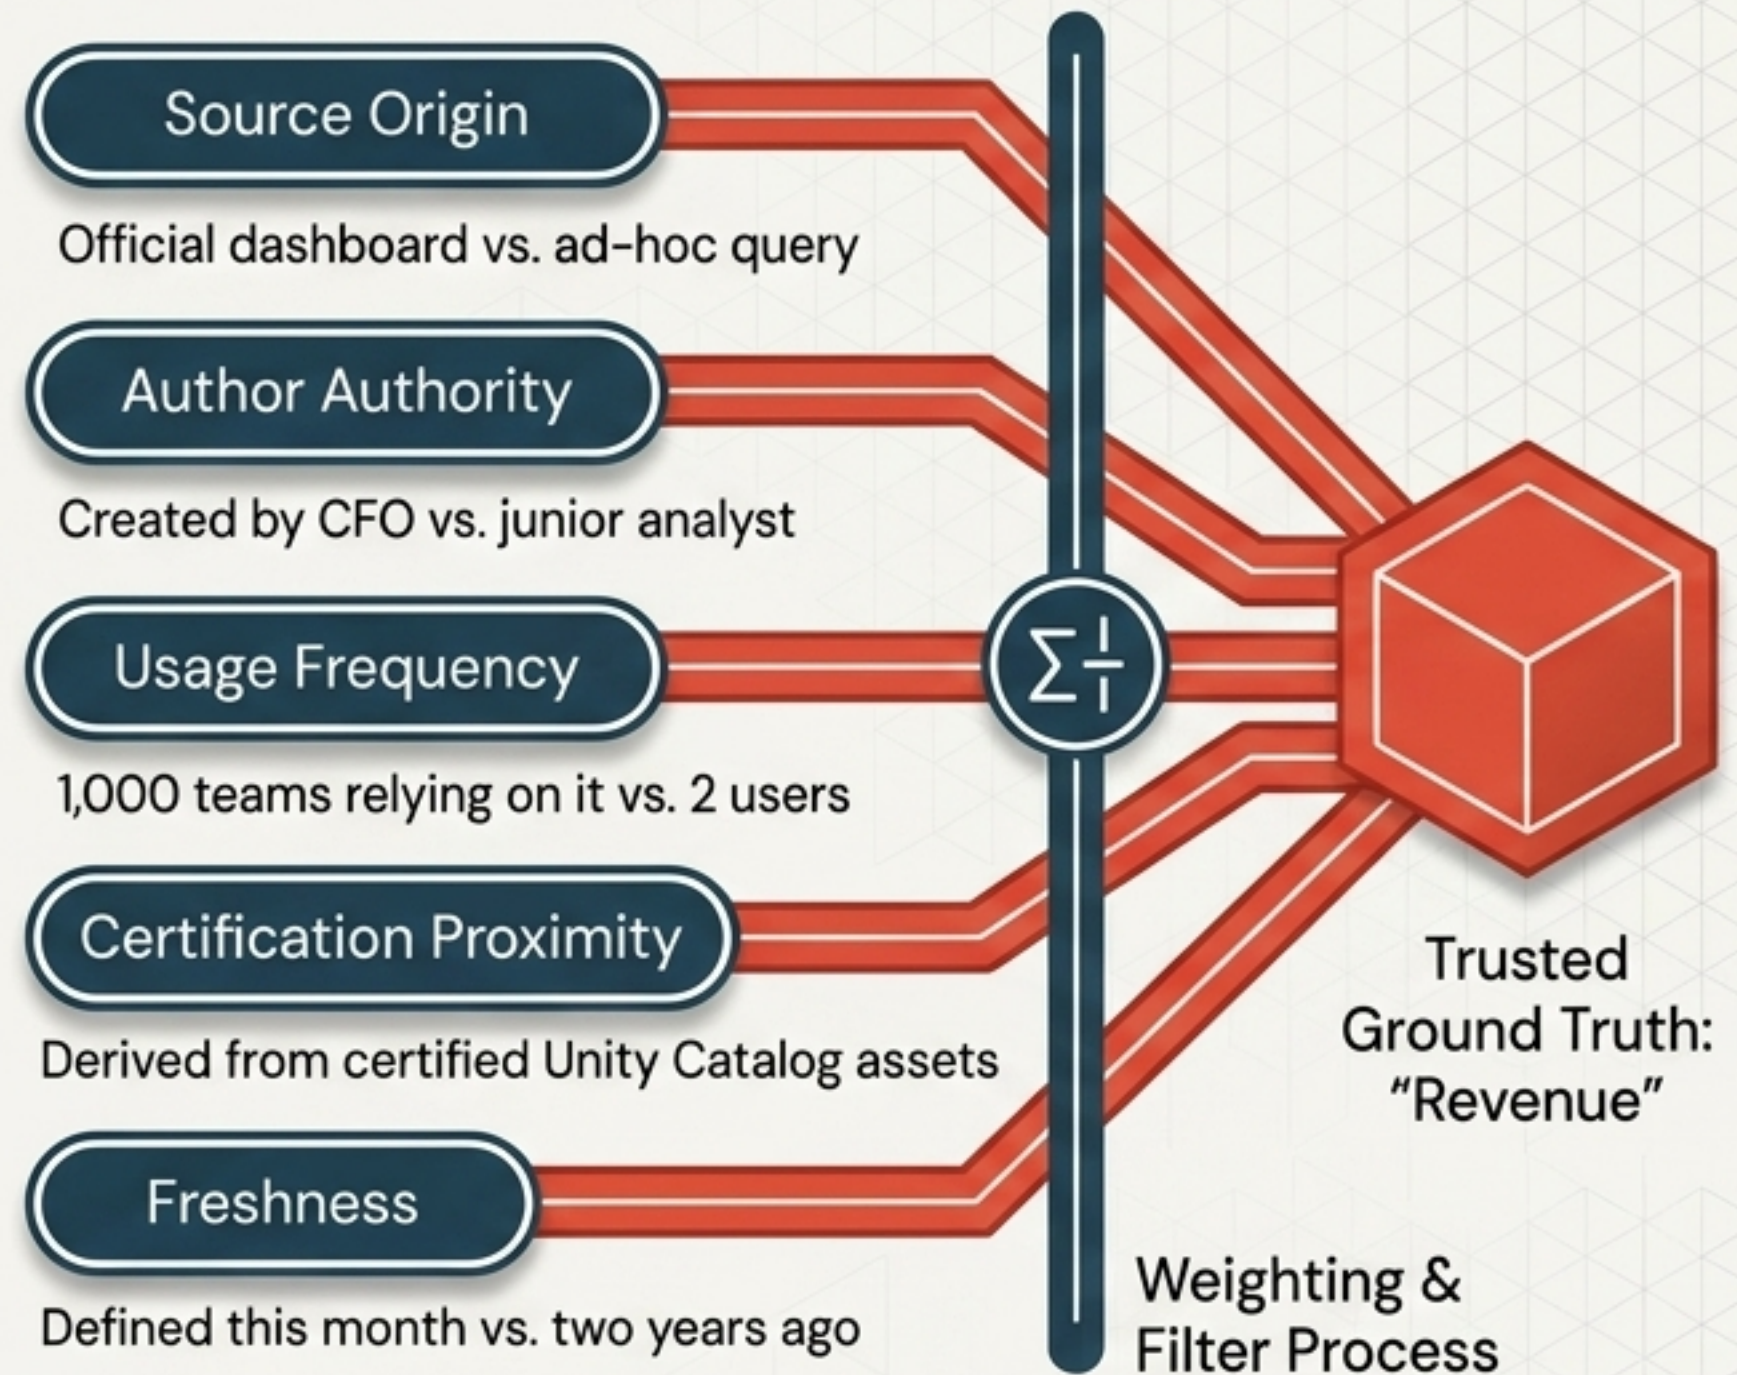
“Most enterprise AI today is just guessing with false confidence. Genie Ontology continuously learns context from data everywhere, so our answers are much faster and our agents are more accurate.”

— Ali Ghodsi, CEO



OntoRank: The Authority Engine

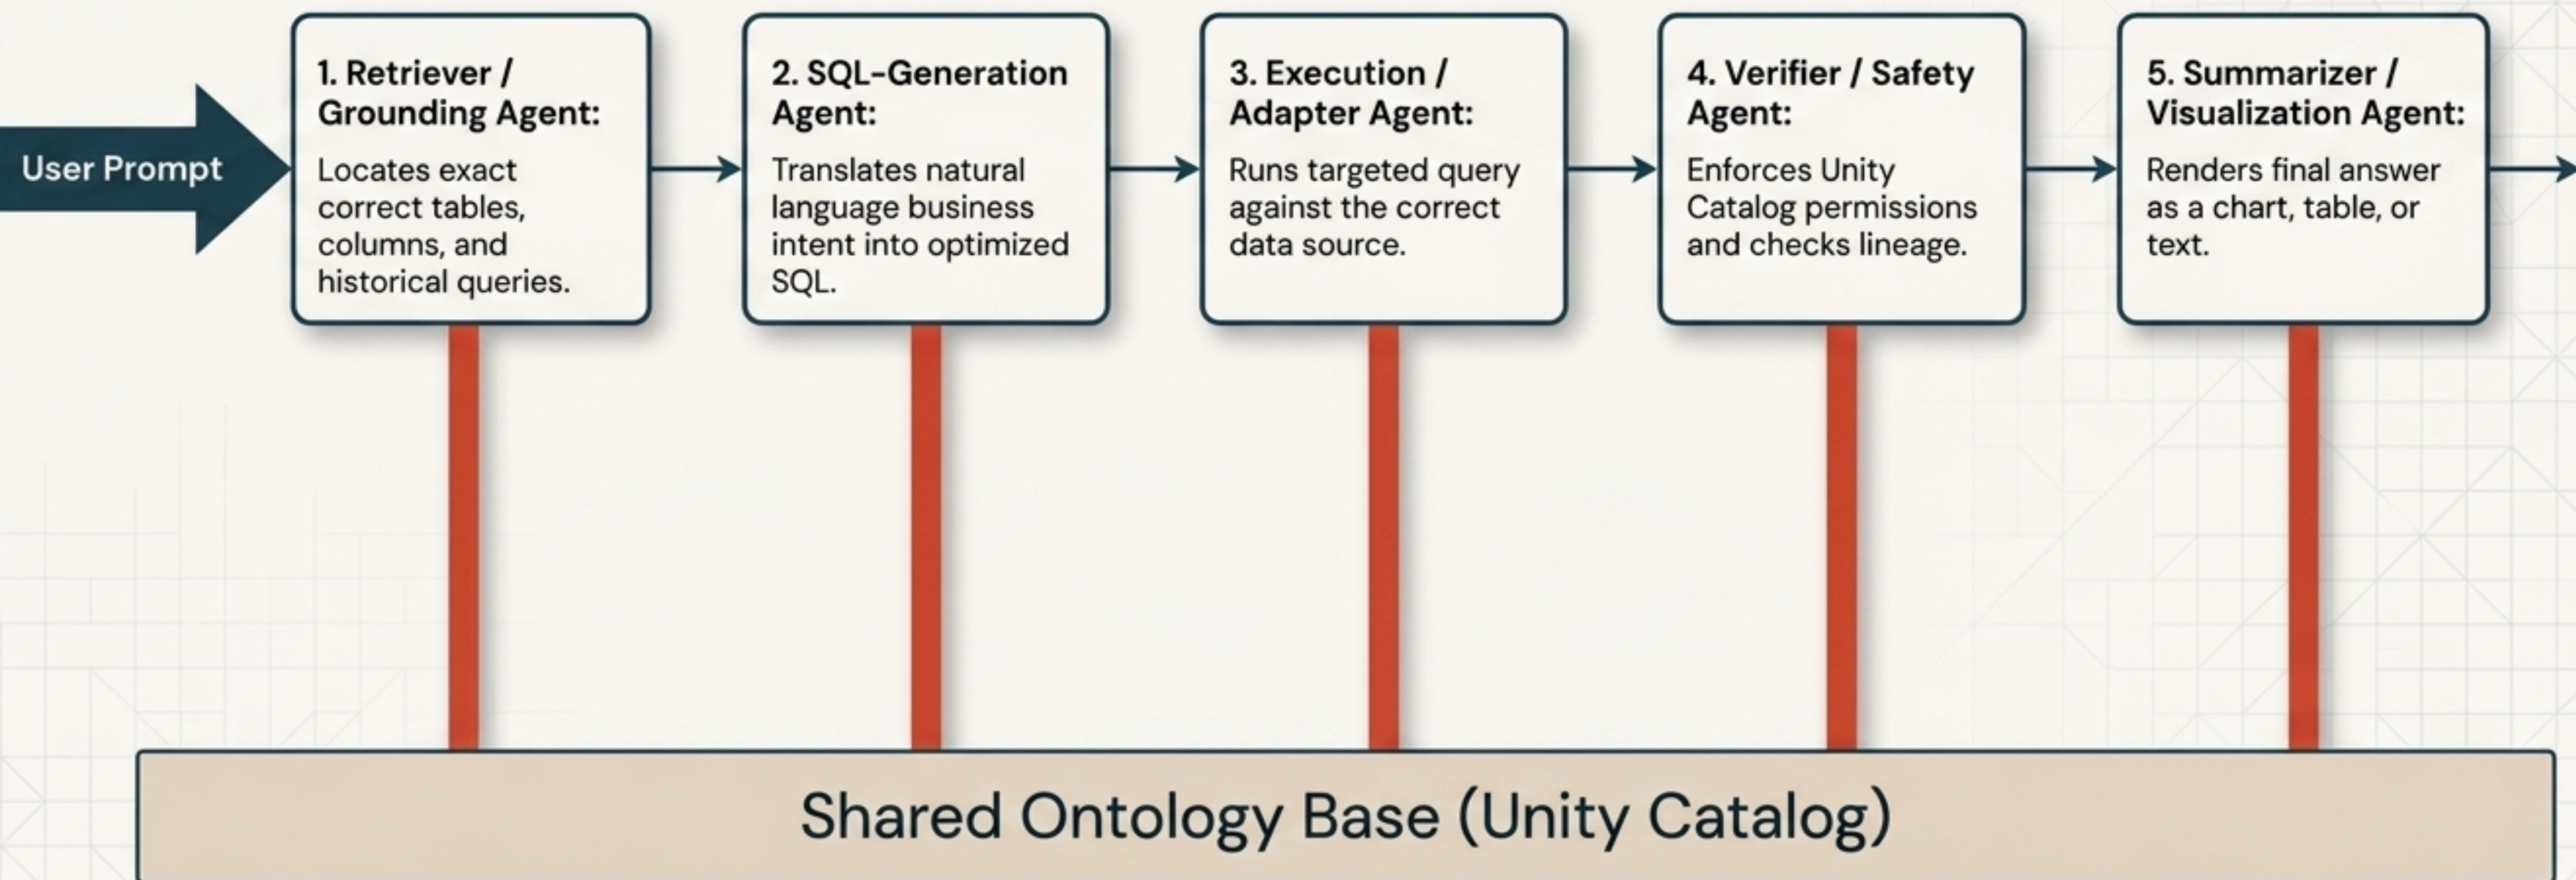
A PageRank-style weighting mechanism that filters enterprise popularity by certified authority to resolve business definitions.



OntoRank guarantees that when Sales and Finance both ask about Revenue, AI delivers the single most authoritative definition.

Multi-Agent Pipeline Architecture


Five specialized agents working in sequence, all reading from a single **shared ground truth**.



Performance Benchmarks: The Power of Context

Traditional RAG relies on proximity (guessing context). Genie Ontology provides meaning, replacing probabilistic hallucination with deterministic execution.

First-Attempt Text-to-SQL Accuracy



The Ontology Wars: Databricks vs. Palantir

	Legacy Incumbent (Palantir)	Modern Advantage (Databricks)
Creation Method	Human-curated, consultant-heavy	Auto-extracted, self-improving
Platform Philosophy	Closed, proprietary Foundry	Built on open standards (Unity Catalog)
Time-to-Value	6-18 months of setup	Days to weeks
Cost Structure	\$50M+ service contracts	Platform-native, included in license
Lock-in Risk	High	Low; open SQL, API, MCP standards

Palantir made ontology valuable by having humans build it. Databricks is betting it can make ontology scalable by having the system build and refresh it."

Implementation Playbook: Path to Value

Prerequisites Checklist

- Unity Catalog enabled
- Premium/Enterprise Tier
- SQL Warehouse active
- Ownership model defined

Phase 1

(Weeks 1-2)

Unity Catalog prep, table descriptions, Genie One workspace activation.

Phase 2

(Weeks 2-4)

Define business Domains, core Metrics, and initial Glossary drafts.

Phase 3

(Weeks 3-5)

Attach external databases via Federation; link 50+ connected apps.

Phase 4

(Weeks 4-6)

Establish pilot groups, deploy agents, set budget caps via AI Gateway.

Phase 5

(Week 6+)

Broad enterprise training, OntoRank quality monitoring, feedback loops.


Top Strategic Use Cases: Operationalizing Genie



1. Natural Language to SQL Business Self-Service

Before

After



Business users ask questions in Slack or Teams. Genie delivers 84.5% accurate answers instantly.

Impact: Search time reduced by 70%. Eliminates data engineering bottleneck.



2. Cross-Team Metrics Alignment Governance & Consistency

Before

After



Solving the 'What is Revenue?' debate. OntoRank enforces a single governed metric across all departments.


Impact: Total alignment via Unity Catalog Metrics; ends conflicting dashboards.



3. AI Agent Grounding RAG Modernization

Before

After



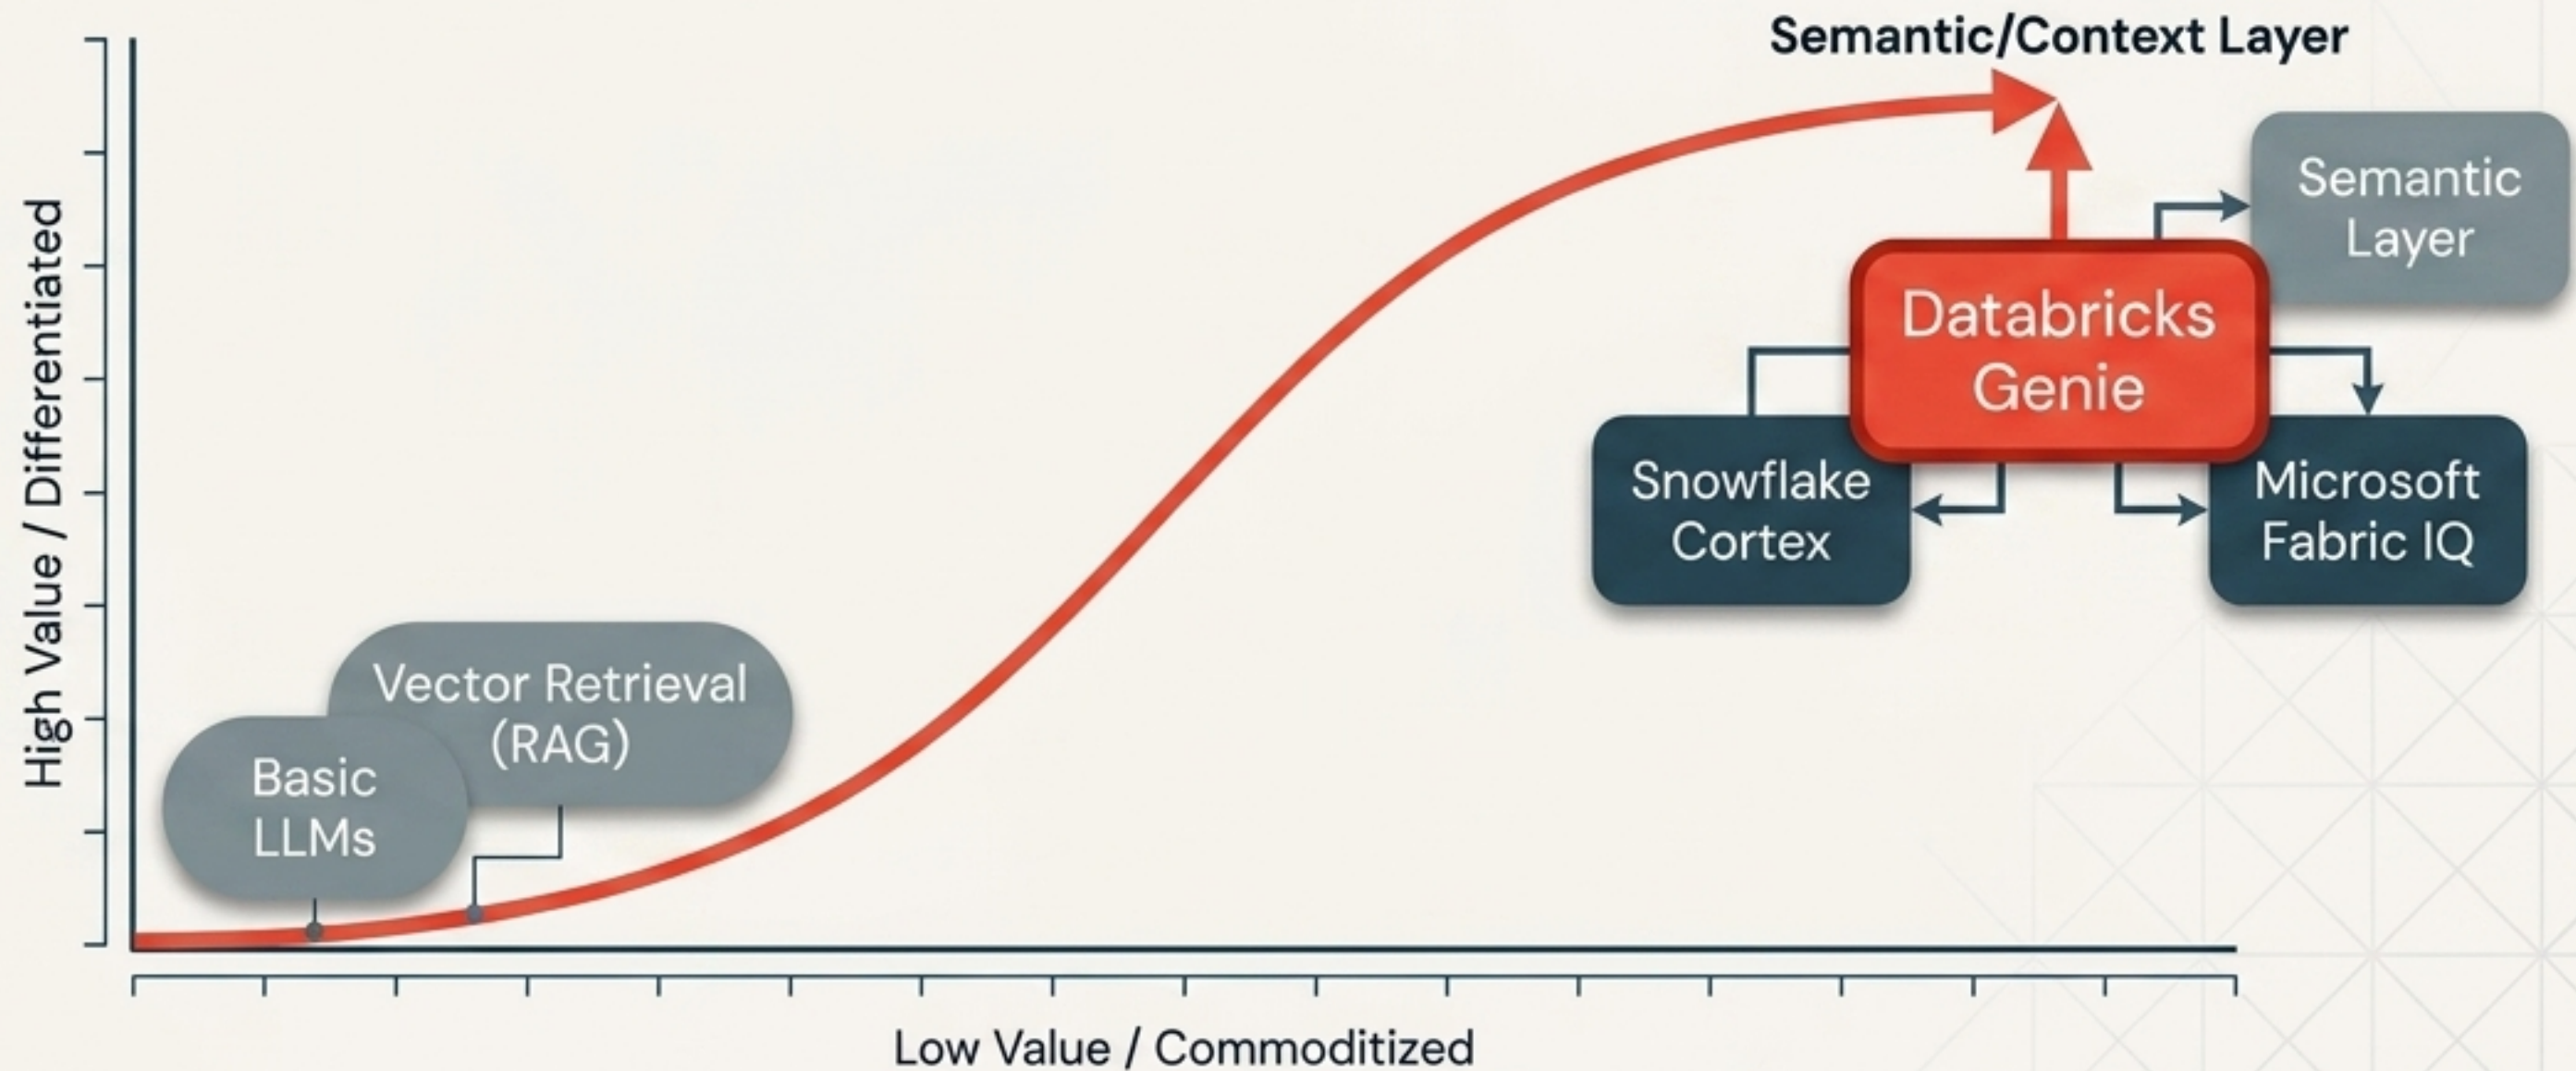
Ridding GenAI pipelines of hallucinations. Upgrading from blind Vector Search to relationship-aware GraphRAG.

Impact: 83% drop in AI pilot failures caused by missing business context.

The Ontology Renaissance

AI models are commoditizing; competitive advantage now lives entirely in governed, relationship-rich context.

The industry is shifting away from manual, static semantic layers toward auto-extracted, dynamic enterprise graphs.



"The model is swappable, the ontology compounds."

The Business Case: ROI & Pricing Architecture

Value Realization

5x Accuracy Improvement

Moving from 16% to 84.5% precision in business queries.

70% Faster Discovery

Eliminating manual data hunting across disparate systems.

Reduced Token Costs

Narrow, targeted SQL execution requires drastically less LLM reasoning.

83% Pilot Success Rescue

Fixing the context gap that kills enterprise AI initiatives.

Transparent Pricing

Component	Pricing Model
Genie Ontology	\$0 / Platform-native (Included in Unity Catalog)
Genie One (AI Coworker)	\$10 / user / month (Includes LLM token usage)
SQL Warehouse / Compute	Standard usage-based DBU consumption