


The Agentic Meaning Gap on Databricks

Turning Databricks-native governance into the business meaning layer for autonomous AI.

A standalone POV on why ontology-governed business meaning is becoming the foundation for trustworthy, autonomous AI. Part B adds optional field enablement.



Part A: Executive point of view

1. The problem, sharpened

The next AI failure in most enterprises won't be a hallucination. It will be an agent that makes a perfectly logical decision based on the wrong definition of the business.

The model may be accurate. The data may be governed. The answer may even be explainable. Yet the decision can still be wrong because "closed," "approved," or

"pending" meant something different to the agent than to the business.

We call this the agentic meaning gap: the distance between what an agent can compute and what the business actually means. It isn't a data problem. Data can be perfectly governed, yet the meaning gap can remain wide open.

What we mean by "ontology" in this document

A governed operational model that enables people and agents to interpret the business consistently: agreed-upon entities, metrics, relationships, decision rules, ownership, data contracts, and approval conditions. It is not a data model or a taxonomy project. It is what makes autonomous action safe.

Everything that shows is what that ontology looks like once it's built: what it prevents, what it costs to skip, how it

maps to Databricks, and how we'd prove it in a single domain before scaling.

Why now?

Ontology was helpful when AI answered questions. It becomes essential when AI starts making decisions. The more autonomous the agent, the more costly a meaning error becomes.

2. What it looks like

Without Governed Meaning Claims exception agent


“Pending” is interpreted from a status field alone. The agent may route the claim incorrectly or request documents that have already been submitted.

With Governed Meaning Claims exception agent

“Pending” is defined using claim stage, document status, reviewer assignment, and SLA rules. The agent routes the claim correctly or escalates it for approval when required.

That's the whole idea in one image. Nothing about the agent changed. What changed is whether “pending” means one thing or three.





3. What it costs to get wrong

Without governed definitions and execution controls, an agent can enter an incorrect interpretation into a system of record before it is reviewed. Fixing it later requires revalidating live agents, checking downstream workflows, and explaining the risk and compliance exposure.

Strip away the AI language, and it comes down to the same three things: grow revenue, reduce costs, and keep the organization safe. A meaning gap undermines all three at once. It slows the agent programs meant to grow revenue; it adds rework and review costs it was supposed to remove; and it creates the risk exposure that shows up in an audit.

From meaning to business outcomes

Governed meaning → Better decisions → Less rework → Reduced risk → Business outcomes

A question worth asking in the room:

"If one of your agents made a wrong call on a customer account last quarter, could you show us today which decisions it made, why, and what happened downstream?"

The answer to that question lies in the difference between "do you have GenAI" and "are you ready for autonomous AI."

4. The agentic trust model: What the ontology becomes in practice

The trust model is not a second framework. It is the ontology translated into operational questions: what data the agent can use, what meaning it applies, what policies constrain it, what workflows it can trigger, and how the decision is explained afterward. This is the governed operational model from Section 1, broken into five layers for use in a client conversation: a question at each layer, and who actually answers it.

The Agentic Trust Model

Data	Does the agent have access to the right information?	Databricks Unity Catalog
Meaning	Does it understand what that information means?	Databricks Genie Ontology
Policy	Does it know what it is allowed to do?	Databricks Unity AI Gateway
Action & Orchestration	Can it execute against the workflow?	Zensar, on Databricks AgentMesh for Databricks
Accountability	Can we explain what it did, why, and what happened after?	Databricks + Zensar Gateway logs + AgentMesh audit trail

Zensar

Implements this across every layer, in one business process, proven against a KPI the client already tracks.



Business Outcomes



Faster Decisions

Make the right decisions, faster



Automated Exception Handling

Detect, route and resolve exceptions



Context-Aware Approvals

Human-in-the-loop approvals with trust



Complete Auditability

Complete, immutable audit trail

Agent Orchestration Layer 5

Orchestrate governed agents to execute business workflows

ZenseAI.AgentMesh for Databricks

Orchestrate governed agents to execute business workflows and actions



Customer Service



Operations



AgentMesh (Orchestration)



Risk



Finance

Enterprise Semantic Services Layer 4

Governed business definitions, metrics, relationships & policies consumed by agents

Enterprise Semantic Services Layer Genie Ontology

The enterprise semantic meaning layer on Databricks that aligns people and agents on the same definitions, metrics, relationships and rules



Business Metrics



Business Context



Semantic Model



Enterprise Definitions

Governance & Control Plane Layer 3

Governs data and AI runtime traffic and agent actions

Unity AI Gateway

Governs AI runtime traffic, agent actions, and tool use



Model Calls



Agent Actions



Guardrails



Rate Limits



Budget Controls

Unity Catalog

Governs data and AI assets across the enterprise



Tables



Files



Models



Access Control



Lineage



Audit & Monitoring

Data Foundation Layer 1

Trusted, AI-ready data foundations on Databricks

ZenseAI.Data for Databricks

Databricks-native fabric for trusted, AI-ready, all-ready data foundations



Data Products



Data Quality



Metadata



Lineage



Business Definitions



Governance & Data Contracts



AI-Ready Data



Trusted Data → Business Meaning → Governed Agents → Business Action

Enterprise AI Trust & Compliance



Customer Data Never Used for Training
Databricks + Anthropic controls



Zero Data Retention
No request or response retention



Runtime Governance Enforcement
Policies, budgets & guardrails



Built for Regulated Industries
Auditability & traceability



5. Why Zensar

**Don't ontologize the enterprise.
Make one decision. Prove the outcome. Then scale.**

That's the core of our engagement philosophy. We aren't selling ontology as a new product; Unity Catalog and Unity AI Gateway are already Databricks-native. Our work closes the business-meaning gap on top of what the client already owns, one domain at a time, measured against a KPI they already track.

What's demonstrably in place today

A Databricks delivery framework already used across engagements, a reference architecture that connects Unity Catalog to ZenseAI.AgentMesh, and delivery accelerators from prior Databricks work. A named, quantified case study for this exact motion has not yet been added; once one is cleared for external use, flag that gap rather than papering over it.

6. What we can point to and how much weight to give it

The proof points below are tiered to prevent teams from overstating what is client-specific, platform-specific, or directional.

Use the hierarchy: client-specific proof carries the most weight; platform and market evidence should support the story, not replace it.

Published Platform Evidence Databricks Genie Ontology

2026 Data + AI Summit

Answer accuracy rose from roughly 50% to about 85% once agents were grounded in governed business context, in Databricks' own testing. Confirm methodology before quoting externally.

Zensar Delivery Evidence

Agentic pipeline creation, Databricks + Global oil & gas, Databricks Genie

First-hand, qualitative - not yet backed by a measured figure


Agents that query a governed entity model before touching a pipeline need materially fewer human corrections than agents working from raw table and column names. Separately, an anonymized oil & gas client reported a reduction in spend, time, and risk exposure.

Discovery Hypotheses — To Validate, Not Assert What we typically expect, not what we claim

Baselined and tested during the engagement

We typically expect governed meaning to reduce exception rework, review burden, and time-to-resolution in the chosen domain. That's a hypothesis to test, not a number to lead with in a pitch.

Snowflake published a similar pattern in its Cortex Agents (50% → 78% in the May 2026 benchmark), noted here only as independent, cross-platform confirmation that the mechanism isn't Databricks-specific. Not for use in Databricks-facing client material.



7. One domain, proven, measured honestly

A four- to six-week engagement on one domain, not a program.
What the client walks away with:

The goal is not to build an ontology. The goal is to improve a business KPI.

- Priority domain and use case
- Current-state meaning, governance, and platform gaps
- Minimum viable ontology and ownership model
- Governed agent/workflow design bound to the ontology
- KPI baseline, proof-of-value scorecard, and scale roadmap

No quantified case study for this exact motion exists yet, as stated in Section 6.
What we do have is a framework for measuring it honestly, baselined with the client rather than asserted in a deck:

KPI	What we'd baseline	How we'd measure it
Human review rate	Share of agent outputs currently requiring manual review	Workflow telemetry
Exception rate	Share of transactions routed to exception handling	Agent/workflow logs
Time to resolution	Average time from trigger to resolution	Operational system timestamps
Agent intervention rate	Share of agent actions overridden by a human	Agent telemetry
Decisions requiring escalation	Share of agent decisions escalated for approval	Approval workflow logs
Cost per exception handled	Current fully-loaded cost of a manual review or rework cycle in this domain	Finance/ops cost data, provided by the client

Targets are set once the baseline is measured, not before. The cost line is what turns the other four into a number a CFO will read. If it moves, the next domain sells itself. If it doesn't, we've learned that on one process instead of the whole account.

8. What this looks like, by vertical

Vertical	Why this matters
Banking and Financial Services	When “closed” means one thing to collections and another to compliance, agents can stop the right process for the wrong reason, creating risks of customer, regulatory, and revenue leakage.
Insurance	When “pending” means different things to claims, service, and risk, agents can trigger the wrong next step, increasing cycle time, rework, and customer frustration.
Life Sciences and Healthcare	When “adverse event” is treated as a keyword rather than a regulated definition, agents can miss escalation rules, turning an automation gain into a compliance exposure.
Retail and Consumer	When “in stock” blends shelf, warehouse, and in-transit inventory, agents can promise what operations cannot deliver, hurting margin, service, and trust.
Technology and Manufacturing	When “risk” ignores contracts, SLAs, and supplier obligations, agents can optimize for the wrong score, missing the business constraint that actually matters.

9. The agent readiness test

Five questions, asked in the room. Score one point for each the client can genuinely answer yes to.

Dimension	The question
Meaning	Would every team in your organization define this business term the same way?
Context	Can your agent understand the relationships among business entities?
Policy	Does the agent know which actions require approval?
Traceability	Can you reconstruct why an agent made a decision?
Outcome	Can you measure whether the agent actually improved the business KPI?

0-1
Experimental

2-3
Emerging

4
Production-capable


5
Agentic-ready

10. The question to leave them with

Which business decision would you trust an agent to make today, and which definition would it need to get right first?

That definition, when written down and enforced, is the ontology. Everything in this document is what happens once it's real.

The core POV ends here. The following field playbook is optional enablement material for account and solution teams preparing for client conversations.



Part B: Field playbook

For account teams: stakeholders, buying signals, platform adoption logic, and conversation openers.

11. Who, when, where

Who

Stakeholder	What they own in this conversation
Chief Data/AI Officer	Whether Unity Catalog becomes more than a governance checkbox
Head of Risk/Compliance	Whether agent decisions are explainable and auditable after the fact
VP Data Platform/Architecture	Unity Catalog and Genie deployment: the plumbing this connects to
Business domain owner (claims, lending, supply chain...)	Whether the definitions we encode are actually right

When: Signals worth listening for

- Genie or a Databricks assistant is deployed, but answers are inconsistent across teams asking the same question.
- More than one agent pilot is running, and none have graduated to production.
- An audit or regulator has already asked how a specific AI decision was reached.
- Unity Catalog is live, but no one outside the platform team can tell you what it actually governs.

Where

Full mini-stories, source systems, and open-with questions for each vertical are in Section 8 — not repeated here.


12. How ontology drives platform adoption

This isn't just a client pitch. Each domain we scope and validate adds measurable value and drives incremental platform consumption, with results ready for a joint conversation with a Databricks account team:

- More governed objects in Unity Catalog per onboarded domain, not just more tables: metrics, relationships, and policies that didn't exist there before
- Genie query volume grows once answers are trusted enough to route production decisions through it, not just dashboards
- Unity AI Gateway traffic grows directly with each agent action we bind to policy. That's the metering point
- Each proven domain is the case for the next, on the same account: a repeatable expansion motion, not a one-time project

The platform adoption case is simple: once a domain is trusted, the next domain is easier to fund, govern, and scale.





13. Where we're different, specifically

- We build on Databricks-native governance, not on a separate ontology product.
- ZenseAI.AgentMesh transforms governed meaning into workflows, approvals, exceptions, and audit evidence.
- Our differentiation is domain execution: resolving contested definitions and ensuring they stick.


This is the part that Databricks cannot productize on its own. The platform can provide governance, ontology services, and policy controls, but it cannot determine which definition of “approved,” “pending,” or “at risk” should govern a live business process. That judgment resides in the operating model: who owns the definition, which exceptions matter, which approvals are required, and which KPI proves the change worked. That is where Zensar's services add value.

14. Conversation starters, by audience

Audience	Best opener
CDO/Head of AI	How many agents are still in pilot because nobody trusts them to run unattended?
Risk/Compliance	Could you produce an audit trail today showing why an agent approved or denied something last month?
Platform/Architecture	Do your agents consult a shared definition of the business before acting, or does each one work it out on their own?
Business domain owner	If we fixed the definitions in a single process this quarter, which one would move the needle the fastest?



Sajeed Kazi
Principal Solution Architect
DE&A - Core



zensar
An  **RPG** Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145 leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the \$4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: info@zensar.com | www.zensar.com