

Your AI is only as trusted as your **data.**

Aidan Reilly

Director of AI, Strategy

Investment is up. **Production is not.**

83%

of CEOs are increasing AI investment — it is the No. 1 stated growth priority across the enterprise.

SOURCE · INDUSTRY CEO SURVEY, 2025

59%

of AI initiatives never reach production. Of the ones that do, only 1 in 5 organizations reports significant value.

SOURCE · GENAI & AGENTIC SURVEY, 2025

The data underneath is fragmented, ungoverned, inconsistent.

01

Fragmented

"Revenue" means three things in your CRM, your warehouse, and your finance system. The model can't tell which one is right — so it guesses.

A faster model doesn't fix this. It just guesses faster.

02

Ungoverned

No audit trail. When AI makes a recommendation, most organizations cannot explain to a board member or regulator exactly what data it used.

If you can't trace it, you can't trust it. And you can't defend it.

03

Inconsistent

The same question, asked by two analysts in two markets, returns two different answers. That isn't AI. That's noise — at scale.

Two answers to one question is a governance failure, not an AI feature.



That is a very good point to make. Make a

A faster model doesn't fix a **broken GPS.**

87 **x**

more tokens without a governed context layer: 175,000 tokens doing the work of 2,000.

WITHOUT A GOVERNED CONTEXT LAYER

80 **%**

wrong-answer rate for enterprise text-to-SQL on complex schemas without proper business context.

INDUSTRY BENCHMARKS · 2025

10 **x**

prompt efficiency and the ability to use high-efficiency models.

PER AGENT · PRODUCTION WORKLOADS

What it looks like when you solve it.

Not a pilot. Not a proof of concept. A global pharmaceutical workforce running on one governed data foundation.

THE CHALLENGE

A global workforce that has to **trust** every answer.

Pfizer: one of the world's largest pharmaceutical companies. Their commercial field teams operate across 27 markets, each with distinct regulatory and data-residency rules. When a field rep acts on an AI insight in front of a physician, it needs to be right. In this industry, the consequences of a wrong AI answer aren't business costs. They're **patient outcomes**.

27

markets, each with distinct regulatory and data-residency rules.

15,000

Strategy users across Pfizer

THE CENTRAL QUESTION

How do you deploy AI that a global workforce can **actually** trust?

They built a **foundation** first.

PILLAR 01

Unified semantic layer.

One governed definition of every critical metric — revenue, patient segments, market share — shared across every AI tool, every team, every market.



PILLAR 02

Governed access & context.

Role-based, market-aware data access baked into the AI layer. The right analyst sees the right data automatically — no manual filtering, no leakage.



PILLAR 03

Provenance & audit.

Every AI-generated insight traceable to its source data, metric definition, and applied business rules. Regulatory confidence, built in from day one.



THE RESULTS

One foundation. **Every** data point consistent.

27

MARKETS, ONE MODEL

Consistent globally.

Every AI answer reconciles to the same governed definitions, regardless of market, language, or team.

15,000

USERS WITH TRUSTED ANSWERS

Decisions, staked.

Strategy users know where the answer came from and which definitions were applied, so they act on it.

Days → Hours

TIME-TO-INSIGHT

Velocity restored.

Clinical and commercial teams move from waiting on analysts to acting on the question they actually asked.

The maintainability story is the quiet one. When schemas changed, the semantic layer absorbed it. The AI tools didn't get rebuilt. The business logic stayed in one place — which is the only way this scales.

Three pillars. Start today.

Any CIO in this room can begin building toward this before they leave the building.

Give AI a **map**, not coordinates.

A semantic layer is not a product you buy. It's a capability you build: a machine-readable abstraction that turns your organization's **tacit knowledge** into something AI can use consistently.

START

One domain. The metrics that drive the business.

FEDERATE

Shared definitions across teams, not centralized ownership.

ITERATE

High-value use cases first. Don't try to boil the ocean.

REUSE

Define once. Every AI tool, every dashboard, every agent inherits.

AI-READY DATA

+80%

increase in agentic AI accuracy when organizations prioritize semantics in AI-ready data.

-60%

cost reduction across the AI portfolio over the same horizon.

Default behavior, **not** afterthought.

Most data flowing through your AI tools right now has no consistent access controls and no audit trail. Governed access bakes **role-based security, market residency, and compliance controls** into the layer, applied to every query by default.

ROLE-AWARE

Right analyst, right data, automatically.
No manual filtering.

RESIDENCY

Market-specific data rules enforced at the AI layer.

COMPRESSION

Send only what's relevant. 97% fewer tokens, fewer hallucinations.

DEFAULT

Governance applies to every query.
Bolted in, not bolted on.

THE GOVERNANCE GAP

90%

of organizations share data externally.

<1 in 3

do it in a governed way.

INDUSTRY DATA-SHARING BENCHMARKS, 2025

AI the board can stand behind.

When something goes wrong (and it will) you need to reconstruct exactly what the AI did, what data it used, and why it made the call. Provenance is the foundation of AI that can be **trusted at enterprise scale**.

LINEAGE

Every insight traces to source data, metric, and applied rules.

DEFENSIBLE

Built for the regulator, the auditor, the board.

RECONSTRUCTABLE

Reproduce any answer the AI gave, on demand.

OPERATIONALIZED

Governance tooling running with AI, not behind it.

FORECAST

2028

Half of all AI governance failures will trace to missing or inadequate semantic frameworks.

74%

of organizations say data governance tools help them operationalize AI governance.

Three things to do **before you leave the building.**

01

Audit your semantic debt.

Cost Zero. The audit reveals everything.

Map where your critical business definitions live. How many definitions of "revenue" exist across your systems? How many AI tools are rebuilding context independently?

02

Pick one use case.

Frame Land & expand. Measurable. Defensible.

Not your whole AI portfolio, one use case with clear ROI and measurable accuracy requirements. Build the governed context layer for that one. Prove the model. Then expand.

03

Treat semantics as infrastructure.

Horizon 2030. Compounding.

Same budget line as cloud platforms and cybersecurity. By 2030, that's exactly where it sits — and the organizations that got there first will have a compounding advantage.



The model is not the bottleneck. Governance scales your investment.

THE REPLICABLE TRUTH

27 markets. 15,000 users. Consistent, trusted AI. The playbook exists. The question is whether you build the foundation before your AI portfolio scales or after.

THE BEST TIME

The best time to build the foundation was before you needed it. The second-best time is today.

Thank you!

Happy to go deeper on any of the below and on the framework you'd apply in your own industry.

SEMANTIC-LAYER ARCHITECTURE

THE INTERNAL BUSINESS CASE

INDUSTRY-SPECIFIC GOVERNANCE

BUILD VS. BUY

MEASURING ROI