

Governing at Machine Speed

Why the next two years of enterprise AI will be won or lost on governance, not models

TABLE OF CONTENTS

03 ▸ Section 1 . The symptom

04 ▸ Section 2 . Why this is different

05 ▸ Section 3 . Peeling back the layers

07 ▸ Section 4 . A way to think about it

08 ▸ Section 5 . The first surface, in depth

09 ▸ Section 6 . Where the layers meet

10 ▸ Section 7 . The shape of the answer

12 ▸ Section 8 . How Onix build it

13 ▸ Section 9 . The next two years

14 ▸ Section 10 . Industry recognition

16 ▸ Section 11 . What to do now



Section 1

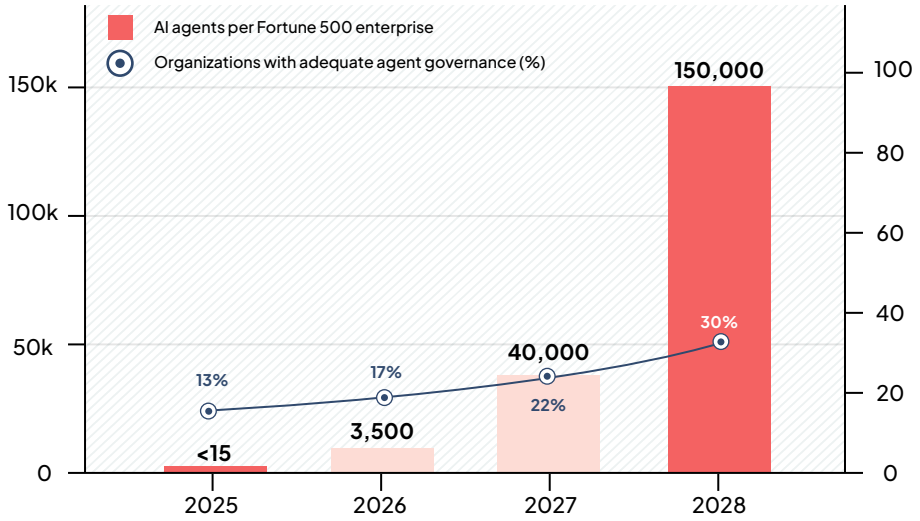
The symptom

Everyone is shipping agents. Almost no one can govern them.

Walk into any enterprise technology review in 2026, and you will hear two things that do not fit together. The first is ambition: teams are deploying autonomous AI agents into finance, operations, customer service, and security. The second is quiet unease. Ask who can say, with confidence, what every one of those agents did last week, on whose authority, and against which policy, and the room goes still.

That silence is the subject of this paper. Gartner projects that by 2028 the average Global Fortune 500 enterprise will run **more than 150,000 AI agents,** up from fewer than 15 in 2025. In the same research, **only 13 percent of organizations** believe they have the right governance in place for the agents they already have.

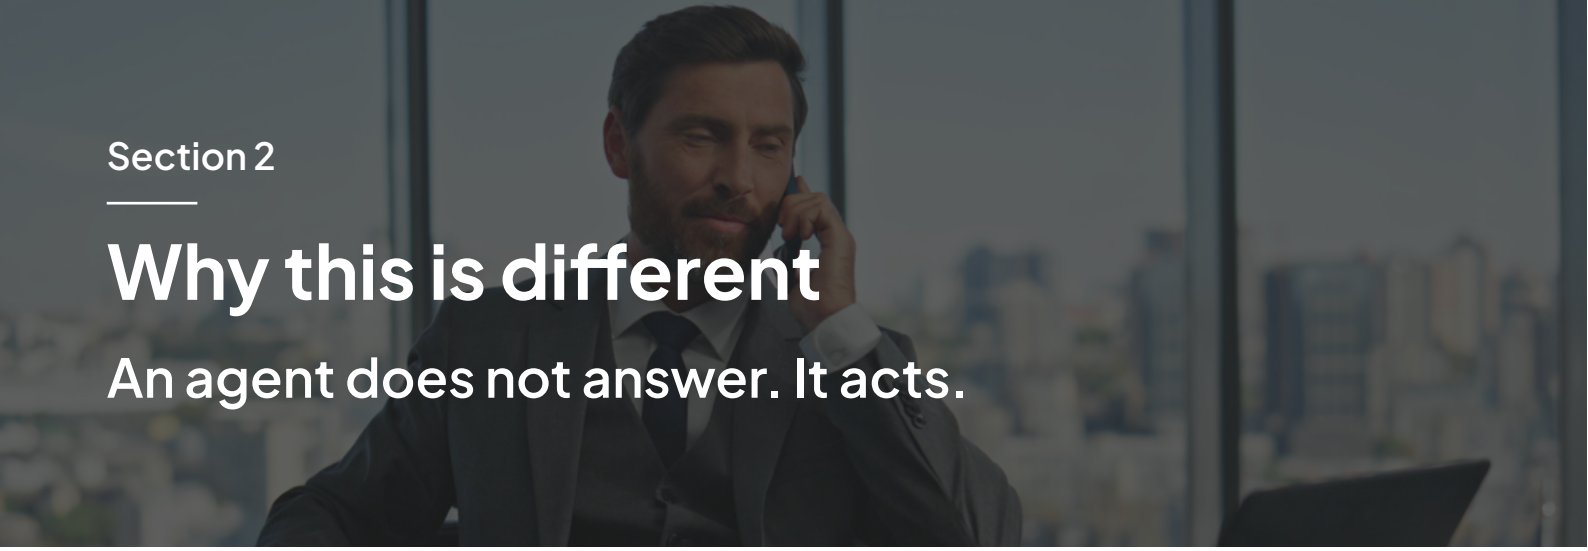
Figure 01 · Adoption is exponential. Governance is not.



The consequences are already showing up in the forecasts. Gartner expects more than 40 percent of agentic AI projects to be canceled by the end of 2027, citing escalating costs, unclear value, and inadequate risk controls.

> *"The last two years were about proving agents could do the work. The next two are about proving we can let them. That is a governance question, and most of the market has not caught up to it yet."*
> Mayur Jain, AVP Products, Onix

It would be easy to read this as a maturity problem that time will fix. It is not. The gap is structural, and the fix is architectural. The enterprises that put that architecture in place now will be the ones still standing when the regulators and the incidents arrive.



Section 2

Why this is different

An agent does not answer. It acts.

A chatbot that gives a wrong answer wastes a minute. An agent that takes a wrong action moves money, changes a record, grants an entitlement, or sends data somewhere it should not go. The generative AI wave of 2023 and 2024 taught enterprises to worry about accuracy. The agentic wave asks a harder question: what is this system allowed to do, and what happens when it does something it should not?

Three things change the moment an agent gains the ability to act.

It acts faster than anyone can review

A human approving each action is a control that works for ten actions a day and collapses at ten thousand. As agents multiply, the review burden multiplies with them, and human-in-the-loop stops being a safeguard and becomes a bottleneck that teams quietly remove to ship. The control has to move into the system itself.

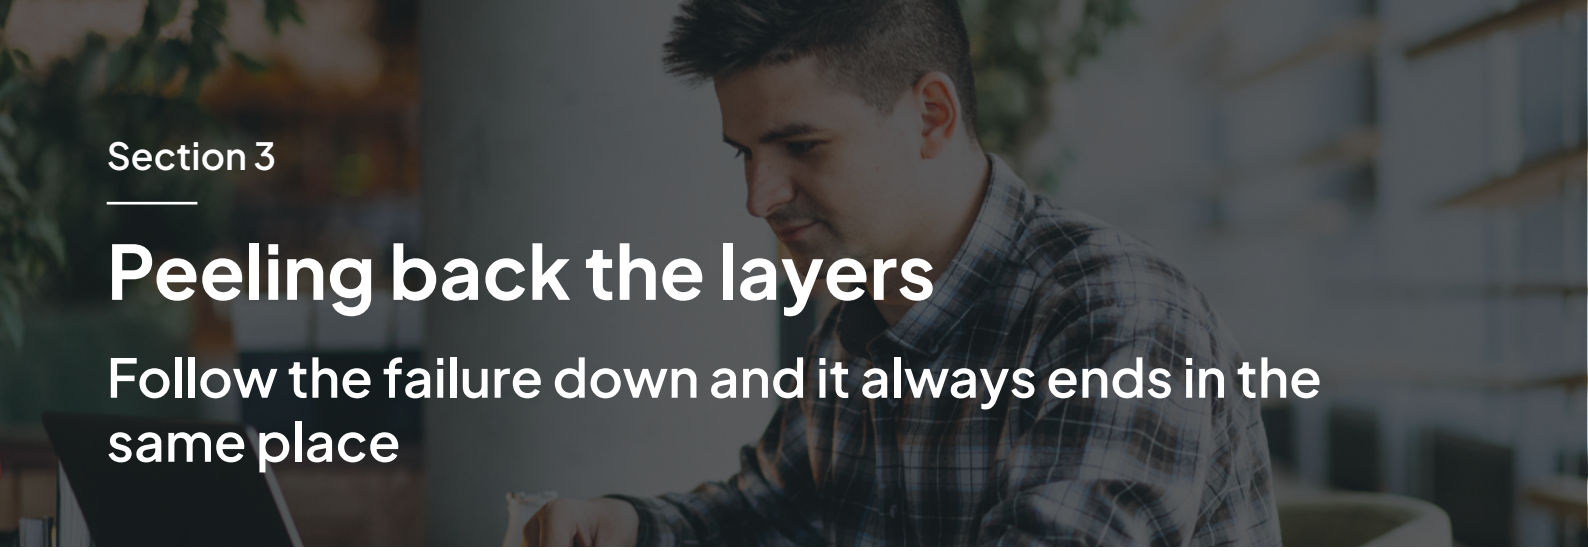
It is a new identity that authenticates and executes

Every agent is a non-human identity with credentials, permissions, and the ability to call tools and APIs. Legacy identity systems were built for humans who log in, and for static service accounts that do one predictable thing. An agent is neither. It reasons, chains tools, and acts on delegated authority, often inheriting far more access than the task in front of it requires.

It makes decisions no one can reconstruct

When an agent produces an outcome, the enterprise needs to answer a regulator-grade question: what data did it use, what policy governed it, and what path led from the source to the decision. Most agent deployments cannot answer this, because the reasoning, the tool calls, and the data lineage were never captured as one connected record.

> *"Hallucination is not a text problem; it is an interpretation problem. An agent that misreads what a column means is confidently wrong long before it writes a single sentence."*
>
> Mayur Jain, Product Head, Onix



Section 3

Peeling back the layers

Follow the failure down and it always ends in the same place

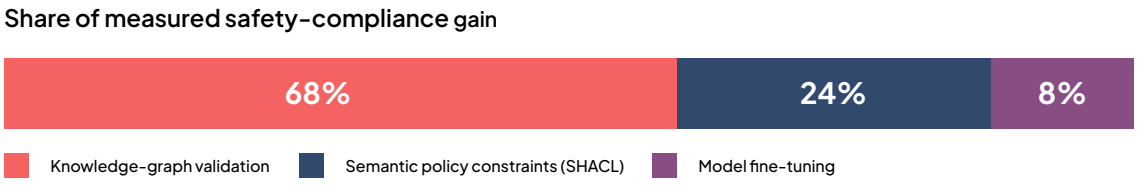
When an agent deployment fails in a way that matters, the incident report usually blames the model. Look closer and the model is rarely the cause. Peel back the layers, and the failure resolves to something more mundane and more fixable: the agent was acting on data whose meaning was never governed. Put plainly, AI does not fail as a model problem. It fails as a context problem.

This is now a measured finding, not a claim. In May 2026, Gartner stated it in language that reads almost like an engineering specification.

> *"Governed context is becoming a cost line and a trust line at the same time. The enterprises that treat it as optional will notice when both move the wrong way."*
> Kalpa Sharma, Senior Director, Onix

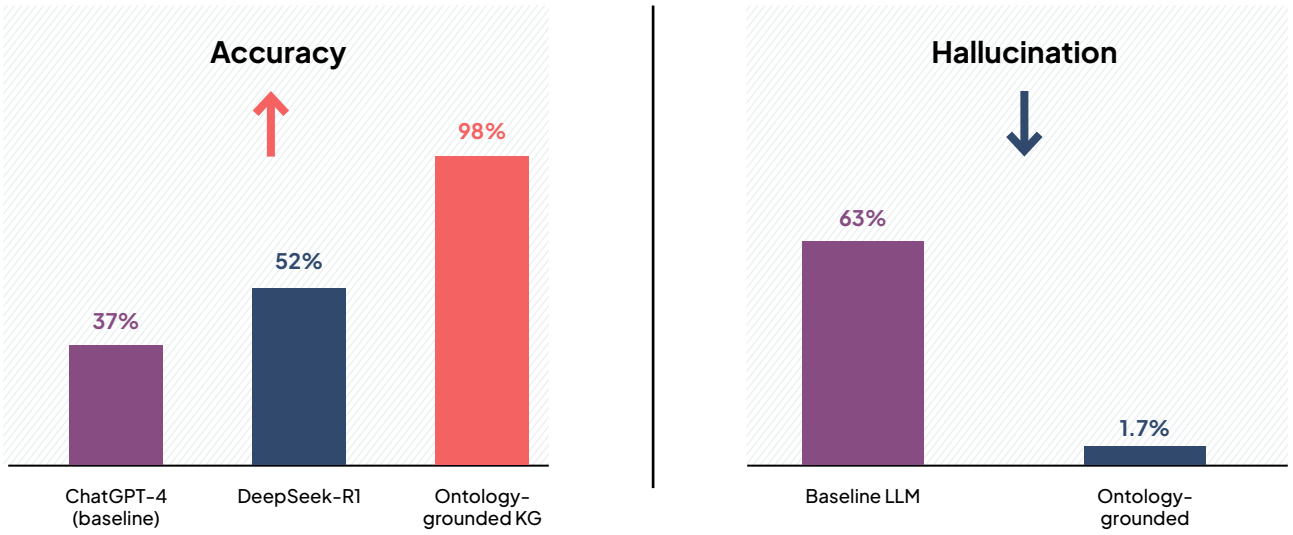
Gartner's prediction attached to that finding is the number to sit with: organizations that prioritize semantics in their AI-ready data will **increase agentic AI accuracy by up to 80 percent and reduce costs by up to 60 percent** by 2027. The academic literature of the last twelve months does not just agree. It quantifies the effect. Researchers building a neuro-symbolic governance framework, published in December 2025, decomposed exactly where an agent's safety and compliance actually came from.

Figure 02 · Where agent safety actually comes from.



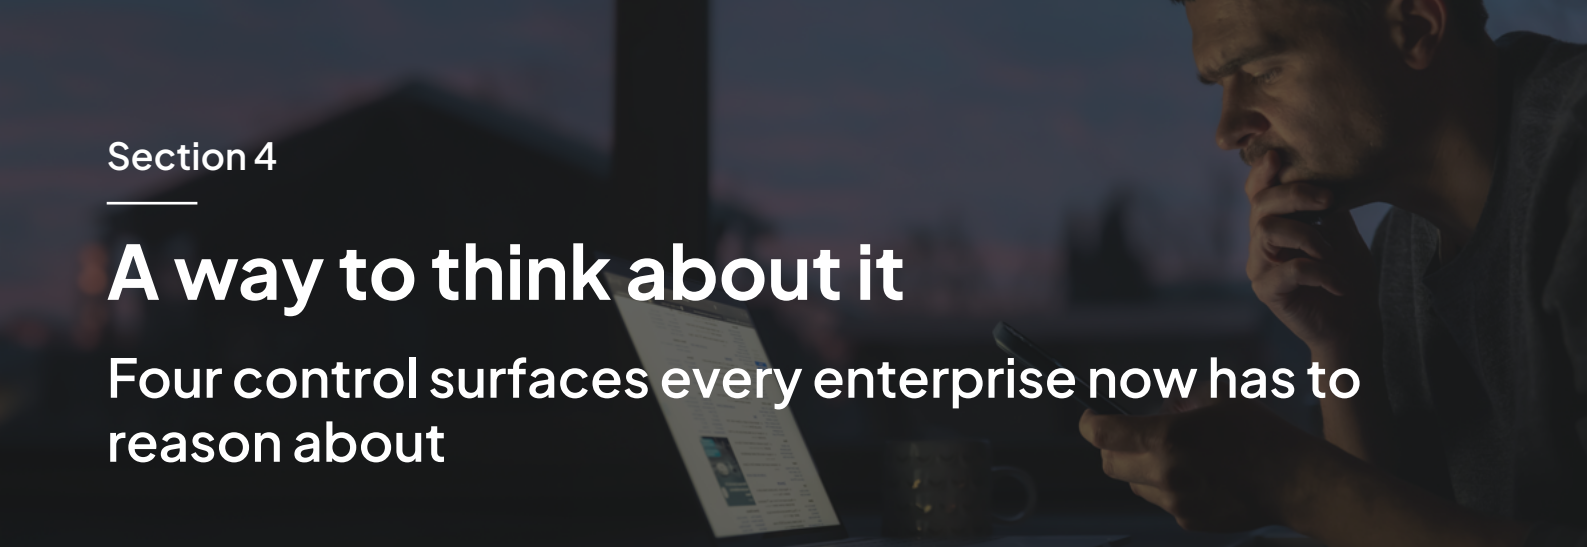
The same pattern holds for accuracy and hallucination. A 2026 study grounding a clinical agent in an RDF and OWL ontology reported accuracy rising from 37 percent to 98 percent, and hallucination falling from 63 percent to under 2 percent, against the same underlying model.

Figure 03 · Ontology grounding, against the same model.



> *"Most of what goes wrong with agents is not an attacker. It is an agent doing exactly what we allowed, on data nobody checked."*
>
> Binoy Samuel, Sr. Director, Product Engineering, Onix

That inverts the industry's instinct. The reflex when an agent misbehaves is to reach for a better model or more fine-tuning. The research says the leverage is upstream, in how the agent understands the enterprise's data before it ever generates a token. That is the layer almost no one has built, and it is what this paper is about.



Section 4

A way to think about it

Four control surfaces every enterprise now has to reason about

If governance is going to move from the human into the architecture, leaders need a structured way to think about what that architecture must control. The 2025 and 2026 research converges on four control surfaces. An enterprise that can answer for all four can govern its agents. One that cannot will be governed by its incidents.

Control surface	The question it answers	What it requires technically
Identity and delegation	Who is this agent, and on whose authority is it acting right now?	Agent-specific identity, scoped and time-bound delegation, on-behalf-of flows, just-in-time and just-enough access
Semantic grounding	Does the agent understand what this data actually means before it acts?	A governed ontology and knowledge graph the agent reasons from, not raw tables it guesses at
Policy enforcement	Was this specific action allowed, and was it stopped if not?	Policy expressed against business entities, evaluated at the moment of action, at machine speed
Provenance and observation	Can we reconstruct exactly what happened, from source data to decision?	End-to-end lineage plus runtime tracing, captured as the action happens, not reassembled later

These four are not a checklist to satisfy once. They are continuous properties of a running system, and they are connected: policy enforcement is meaningless without semantic grounding to define what the policy applies to; provenance is impossible without identity to attribute the action. Treating them as separate tools bolted together is how enterprises end up in the 13 percent. The insight of the last year of research is that they are facets of one architecture.



Section 5

The first surface, in depth

Identity was built for humans. Agents broke it.

Start with identity, because it is where the sprawl becomes visible. CyberArk's 2026 research, based on 2,930 global cybersecurity leaders, found that machine identities already outnumber human identities 109 to 1. And the numbers are moving fast: AI agents are expected to grow by 85% this year. As enterprises deploy more AI agents, each one adds another machine identity to an already expanding identity landscape. Unlike traditional service accounts, AI agents can reason, make decisions and act across systems, making their identities and access increasingly important to secure.

Figure 04 · The non-human identity majority.

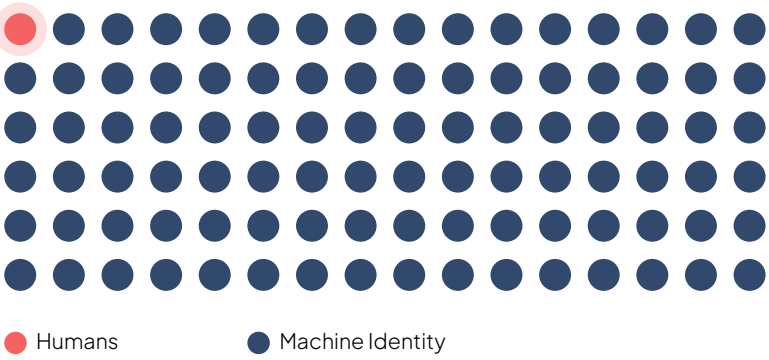
109:1

machine identities for every humans

+85%

projected growth in AI agent this year

Each new agent is another non-human identity to authenticate, scope and invoke.



The technical problem is subtle and deserves precise articulation. When an agent acts on a user's behalf, it needs authority, but it must not inherit the user's full authority. The research community calls this failure mode the confused deputy. An agent handed a broad token becomes a deputy that can be tricked into using powers the user never intended to delegate for this task.

TECHNICAL NOTE · How delegation is actually being solved

The emerging pattern, set out in "Authenticated Delegation and Authorized AI Agents" (South et al., ICML 2025), extends OpenID Connect with three tokens. A user ID token identifies the human. An Agent-ID token identifies the agent as a first-class principal, carrying its capabilities and limitations as verifiable metadata. A delegation token, signed by the human, explicitly authorizes this agent to act on their behalf for this scope, referencing both identities. Combined with OAuth 2.0 token exchange (the on-behalf-of flow, RFC 8693) and just-in-time issuance, the agent receives task-specific, short-lived authority rather than standing access. A policy engine evaluates the delegation at the moment of action, so that "read the calendar for the next ten minutes" is a grantable scope and "read everything, indefinitely" is not.

This is a zero-trust posture applied to agents: never trust the token, always verify the delegation against live policy. The direction is clear. Agent identity is becoming its own discipline, and the enterprises treating it as a variant of human identity management are building on sand.



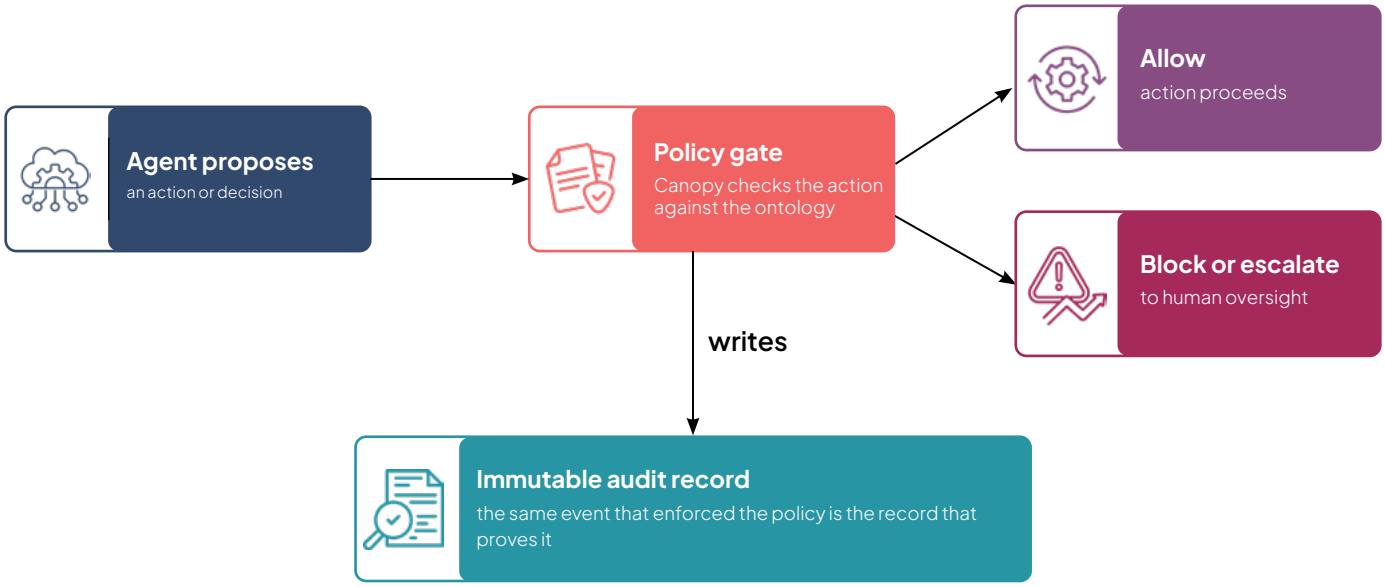
Section 6

Where the layers meet

Enforcement and audit are the same act

Here is the insight that ties the four control surfaces into one architecture, and it is the most important idea in this paper. In the agentic era, you cannot separate enforcing a policy from auditing it. If your governance layer is a separate observer that reads logs after the fact, the decision has already happened, the action has already been taken, and for a high-risk system you are already out of compliance.

Figure 05 · One event, not two.



This is not a philosophical point. It is close to a legal one. The EU AI Act reaches full enforcement for high-risk systems on **August 2, 2026.** Article 12 requires automatic logging across the system's lifecycle, with records that enable traceability. Article 14 requires human oversight, including the ability to intervene or interrupt. Penalties reach 35 million euros or 7 percent of global turnover. Recitals 99 and 100 extend the compliance boundary across every agent in a chain, so a multi-agent system does not escape the obligations by

> *"Every customer I meet has a pilot that worked and a rollout that did not. The difference is almost never the model. It is whether the agent knew what it was looking at."*
> Binoy Samuel, Sr. Director, Product Engineering, Onix

Gartner's warning matters for how the gate is designed. Uniform governance applied bluntly across every agent leads to failure, because it either strangles useful agents or waves through dangerous ones. The gate has to be context-aware, and context comes from the semantic layer that knows what each entity and action means to the business.



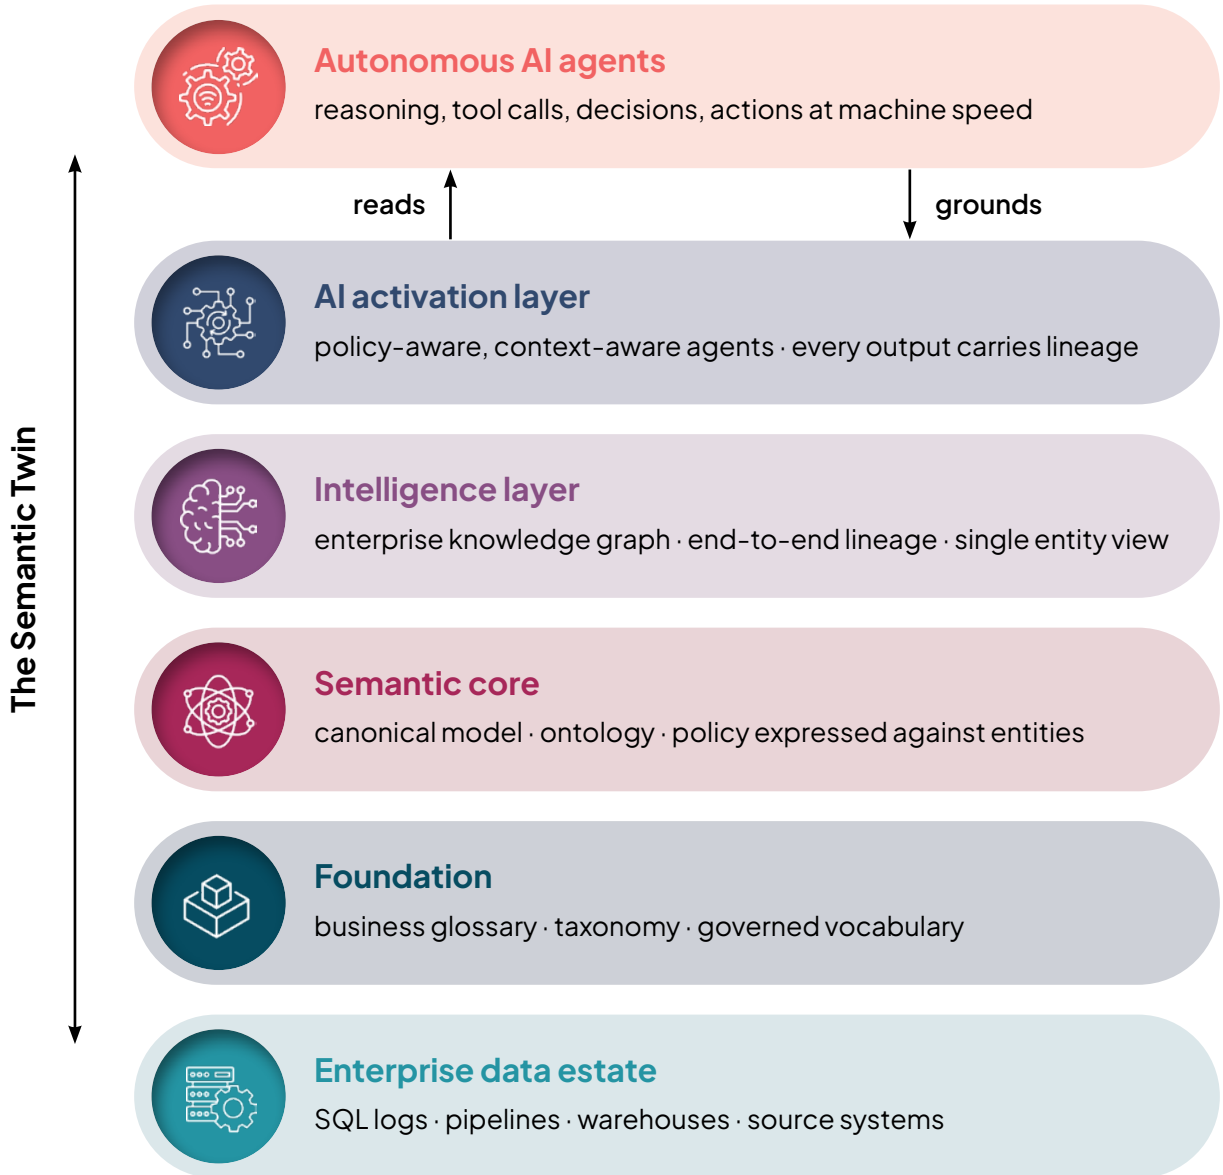
Section 7

The shape of the answer

Governance is not a layer on top. It is the ground the agent stands on

Put the pieces together and the architecture stops being a list of tools and becomes a single structure. The agent does not sit above a pile of governance products. It stands on a governed foundation, and every action it takes is grounded in, checked against, and recorded by that foundation. This is the **Onix Semantic Twin,** a governed context layer and the executable layer of business meaning, built once and continuously kept current, that every agent and every BI tool reasons from.

Figure 06 · The Semantic Twin, layer by layer.



Six parts, one layer

A useful test separates a real Semantic Twin from the things often mistaken for it. The twin is made of six parts. Remove any one and what is left is a catalog, a glossary, a graph, or a prompt, but not a twin.

The Meaning	**Ontology, KPIs and definitions**	What every term means, calculated one way, owned by one person.
The Bindings	**Sources, lineage and freshness**	Where each fact comes from, how it was derived, and how current it is.
The Guardrails	**Policy, purpose and entitlement**	What an agent is allowed to do with a given entity, enforced at query time.
The Engine	**A continuous refresh capability**	Sentinel processes that keep the twin current as schemas and definitions move.
The Memory	**Decisions, evidence and precedent**	What was decided before, on what evidence, so answers stay consistent over time.
The Surfaces	**MCP, API, graph and SQL**	How every agent and BI tool reaches the same governed context, through open standards.

The reason this shape works is that it puts policy where meaning lives. In most enterprises, a governance rule is written against a table column in one system, and has to be rewritten for every other system where the same concept appears. Define a policy against the semantic entity instead- the governed definition of "customer" or "protected health information", and it applies automatically everywhere that entity surfaces, across every source system. One definition, enforced consistently, is what makes governance survive contact with a sprawling data estate.

TECHNICAL NOTE · Open by default, so nothing has to be replaced

The twin is not a walled garden. It publishes through open standards, MCP for agent runtimes, plus A2A, OSI, OpenLineage and Iceberg, and exposes the same governed context through API, graph and SQL. Nothing in the estate gets migrated to use it: the catalog, the warehouse, and the BI tools stay where they are, and the twin grounds them. On the technical guarantee itself, the research is consistent. Deterministic validation against a governed graph, using standards like SHACL, provides assurances a probabilistic model cannot. In the G-SPEC study, graph validation and its SHACL policy corpus together accounted for the large majority of the safety gain. The pattern is that the model proposes and the governed graph disposes: the agent reasons, and a deterministic semantic layer validates that reasoning against the enterprise's actual rules before anything acts.



Section 8

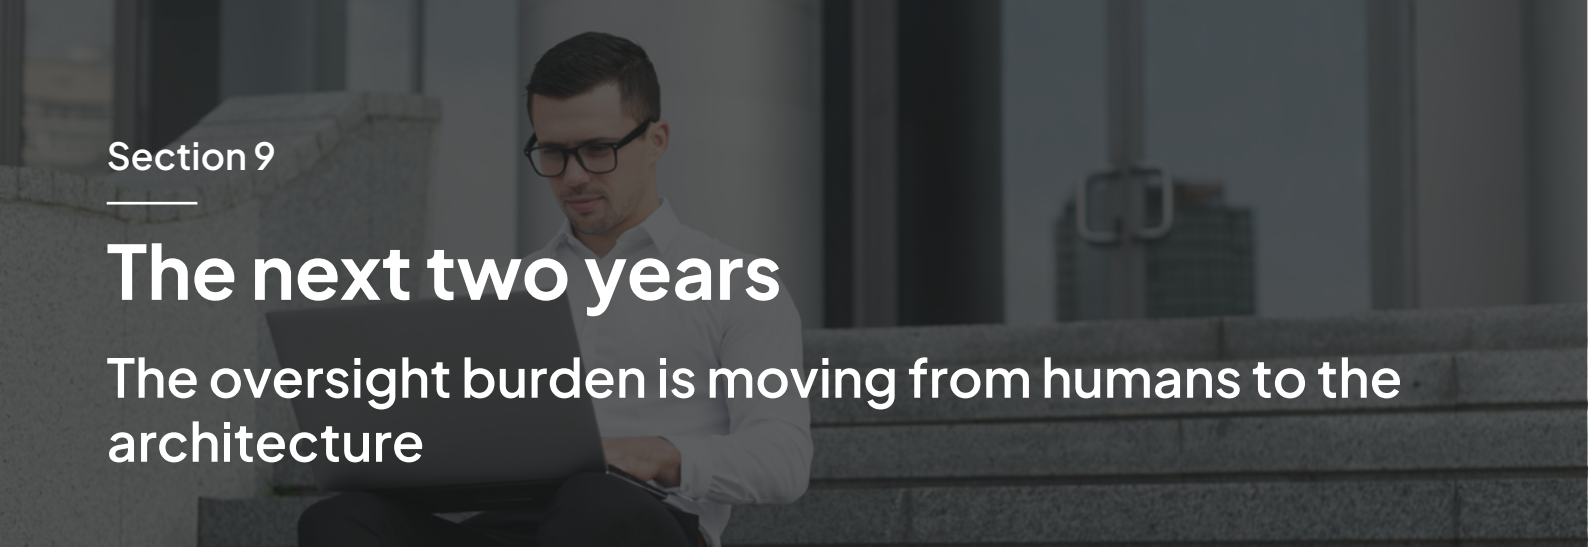
How Onix builds it

Wingspan, mapped to the four control surfaces

Onix Wingspan 2.0 is built on this architecture rather than around it. The **Onix Semantic Twin** is the governed foundation, reverse-engineered from the existing estate in weeks rather than quarters, and kept current on its own as data changes. Here is how the platform maps to the four surfaces the research says matter.

Control surface	What the research requires	How Wingspan delivers it
Semantic grounding	A governed ontology and knowledge graph the agent reasons from	The Semantic Twin, reverse-engineered from SQL logs and lineage. Agents reason from governed meaning, not raw tables. This is the surface that delivered the majority of measured safety gains in the research.
Policy enforcement	Context-aware policy, expressed against entities, enforced at machine speed	Policy is expressed against the twin's entities as the single source of policy truth. A rule defined against an entity applies everywhere that entity appears, avoiding the uniform-governance failure Gartner warns against.
Provenance and observation	Lineage plus runtime tracing, captured as the action happens	A single live map across monitoring, governance and security, at pipeline, job and table levels. Every AI output carries lineage from source through transformation to decision. Enforcement and audit, one act.
Identity and delegation	Scoped, auditable authority for each agent	Role-based access control, multi-tenancy and a dataset registry govern which agents reach which governed entities, with the twin providing the semantic scope that delegation is bound to.

The observability layer is where the abstract idea of runtime governance becomes something a team can operate. It monitors data platform health at both technical and functional levels, freshness, volume, schema, and service levels alongside the business impact of a failure. It delivers lower mean time to recovery through AI-based root cause analysis, runs continuous monitoring, and produces a configurable observability score across multiple dimensions rather than a single opaque number. An analyst can ask a question in natural language, get a diagnosis, and receive a suggested fix. The point is not the feature list. The point is that the same layer that watches the data is the layer that proves, after any incident, exactly what happened and why.



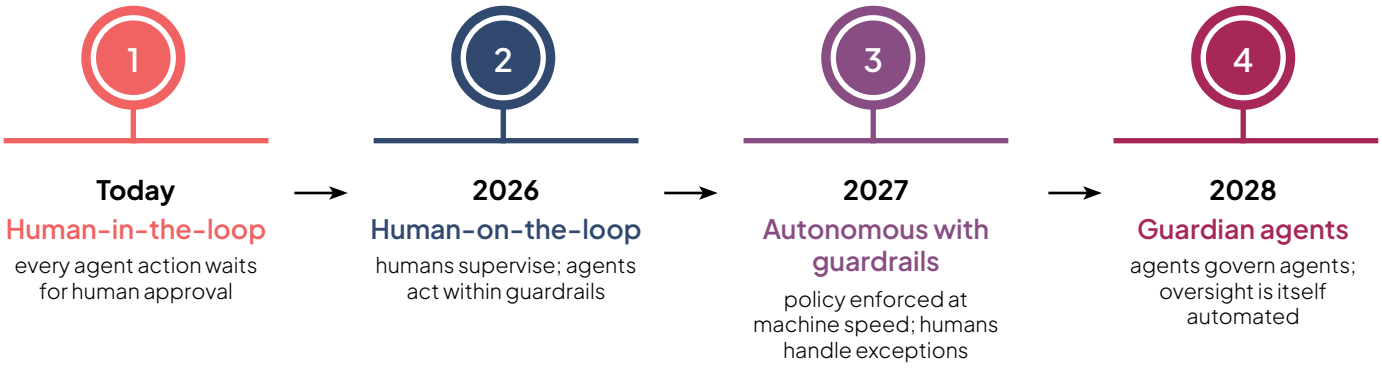
Section 9

The next two years

The oversight burden is moving from humans to the architecture

Where does this go by 2028? The trajectory in the research is consistent enough to plan against. The control that today sits with a human approving each action moves, step by step, into the system, because that is the only thing that scales as agents multiply.

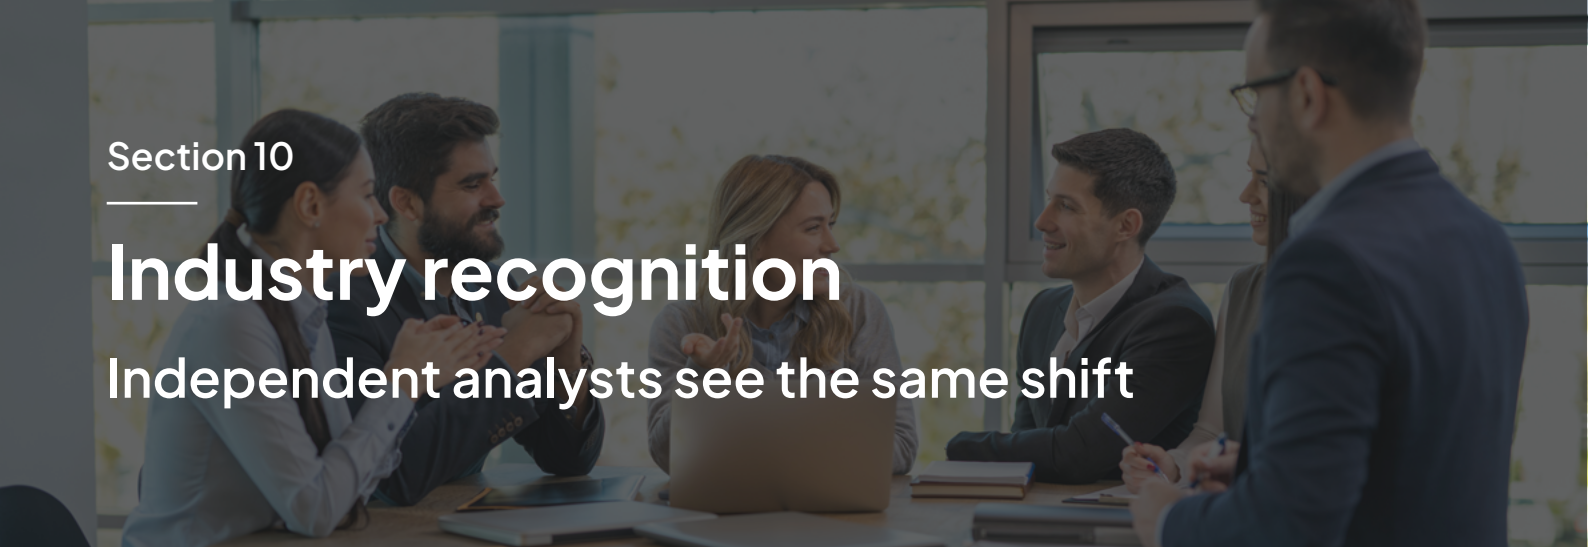
Figure 07 · The arc of the next two years.



Gartner has already named the endpoint. Guardian agents, AI that oversees other AI, are predicted to capture 10 to 15 percent of the agentic market by 2030, and 70 percent of AI applications will use multi-agent systems by 2028. When agents govern agents, the governing agents need something to govern against: a shared, machine-readable definition of what is allowed. That shared definition is the ontology. The enterprises that will be able to deploy guardian agents in 2028 are the ones building the semantic foundation in 2026.

A maturity model is emerging in the literature for exactly this, and the pattern it describes is the one this paper has traced: enterprises move from ad-hoc oversight, through tool-by-tool controls, to a unified semantic and policy foundation that makes governance a property of the architecture rather than a project. The move from the middle stage to the last is the hard one, and it is the one worth starting now, before the agent count crosses from hundreds into the tens of thousands.

The new governance parameters leaders will be asked about by 2028 are already taking shape: the identity lifecycle of every agent, the semantic coherence of the policies that govern them, the blast radius contained when one fails, the trust established between agents that interact, and provenance captured by default rather than reconstructed under subpoena. None of these can be retrofitted cheaply. All of them are easier if the semantic foundation is already in place.



Section 10

Industry recognition

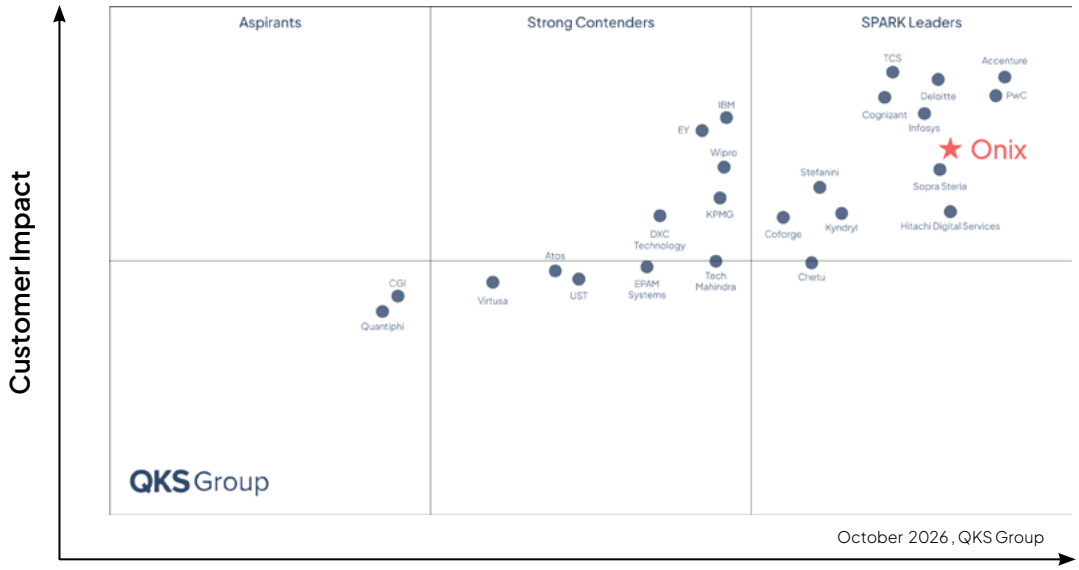
Independent analysts see the same shift

In its SPARK Matrix™ for AI Services, Q4 2026, the analyst firm QKS Group placed Onix in the leaders band, scored on service excellence and customer impact. The placement matters here for one reason. The direction QKS reads in the market is the direction this paper has argued for throughout.

SPARK Leader
AI Services, Q4 2026

QKS Group placed Onix in the leaders band of its SPARK Matrix™ for AI Services, scored on service excellence and customer impact.

SPARK Matrix™: AI Services, Q4 2026



QKS describes governance moving out of the post-deployment checklist and into the architecture itself, reaching past model outputs to the actions, permissions, and decision paths an agent takes. It names agent identity and authorization as a new enterprise control layer for coordinated, multi-agent work, and it lands on the same point this paper opened with. Trust in autonomous systems has to be engineered into the architecture, not assumed once an agent is live. That is governed autonomy, described independently.

What the analysts are seeing

- Governing an agent at the level of one model or application is no longer enough once agents act across connected systems.
- Agent identity, permission boundaries, and action-level auditability are becoming a distinct enterprise control layer.
- Governance is moving from a step performed after deployment to a property designed in from the start.

In its profile of Onix, QKS singled out the Semantic Twin as a persistent enterprise context layer, reused across AI applications rather than rebuilt for each one, connecting business definitions, KPIs, lineage, ontology, knowledge graphs, and governance into one foundation that holds context and explainability across applications. It also noted that Onix treats data readiness as an explicit engineering workstream, with lineage and ontology construction automated by Eagle and validation handled by Pelican. Both sit at the center of the architecture this paper describes.



Section 11

What to do now

Build the foundation before the sprawl arrives

The argument reduces to a single sentence. Enterprises are deploying agents faster than they can govern them; the missing foundation is a governed semantic context layer enforced and observed as one act, and the enterprises that build it now will be the ones able to scale agents safely when everyone else is canceling projects and answering to regulators.

The evidence bears the urgency out. More than 150,000 agents per large enterprise by 2028, against 13% governance readiness today. More than 40% of agentic projects canceled by 2027. Shadow-AI breaches cost an extra $670,000 each, in a year when 97% of AI-breached organizations lacked basic access controls. Full EU AI Act enforcement for high-risk systems in August 2026. Every one of these is a reason to treat governance as the first architectural decision, not the last.

> *"Context is turning into infrastructure. In two years, running agents without it will look like running a database with no backups."*
> Poorvi Sharma, Global head Product Marketing, Onix

The **Onix Semantic Twin** exists to build that foundation. It is reverse-engineered from the estate you already have, kept current on its own, and open by default, so every agent inherits the same governed context, and every answer can show the definition, source, lineage and policy it stands on. The result is agentic AI you can put into production and defend in an audit, because governance is not something you added on top. It is the ground your agents stand on.

Start the conversation: onixnet.com/wingspan

Govern the meaning, and the agents can be trusted to act.

The Semantic Twin is the governed foundation, reverse-engineered from the estate you already have, kept current on its own, and open by default

Start the conversation

onix®

onixnet.com

connect@onixnet.com

800.664.9638

Get in touch

Follow us:

Copyright © 2026 Onix . All rights reserved.