Futurum®

IBM

The Governance Gap:

Why Scaling AI Requires More Than Monitoring

JULY 2026

The AI Governance Problem Has Outgrown Its Tools

Enterprise AI deployment has moved well beyond discrete machine learning models. Organizations now operate interconnected systems of agents, tools, data pipelines, and use cases. The governance surface area has expanded faster than the frameworks designed to manage it. The spreadsheet-era approach to tracking AI assets was already straining under conventional ML workloads. Agentic AI, with its capacity to spawn sub-agents, call external tools, and act on live data, has made that approach inadequate.

Two compounding pressures define the current situation. The first is AI sprawl. Organizations are losing visibility into what AI assets exist, who owns them, what they connect to, and what decisions they are influencing, including assets deployed informally or without central oversight. The second is regulatory acceleration. The EU AI Act, US executive orders, and emerging frameworks in Singapore, Australia, and Canada have moved compliance from background concern to active operational requirement, with more jurisdictions expected to follow.

The result is a governance gap: a widening distance between the pace of AI deployment and the organizational capacity to manage risk, enforce policy, and demonstrate accountability (Figure 1). The question enterprises are now asking is not whether to govern AI, but how to do it at scale without impeding the programs, insights and capabilities – realized and still unrealized – that depend on it. Closing that gap requires governance infrastructure that is built into AI operations, rather than layered on after the fact.

Figure 1. The Governance Gap

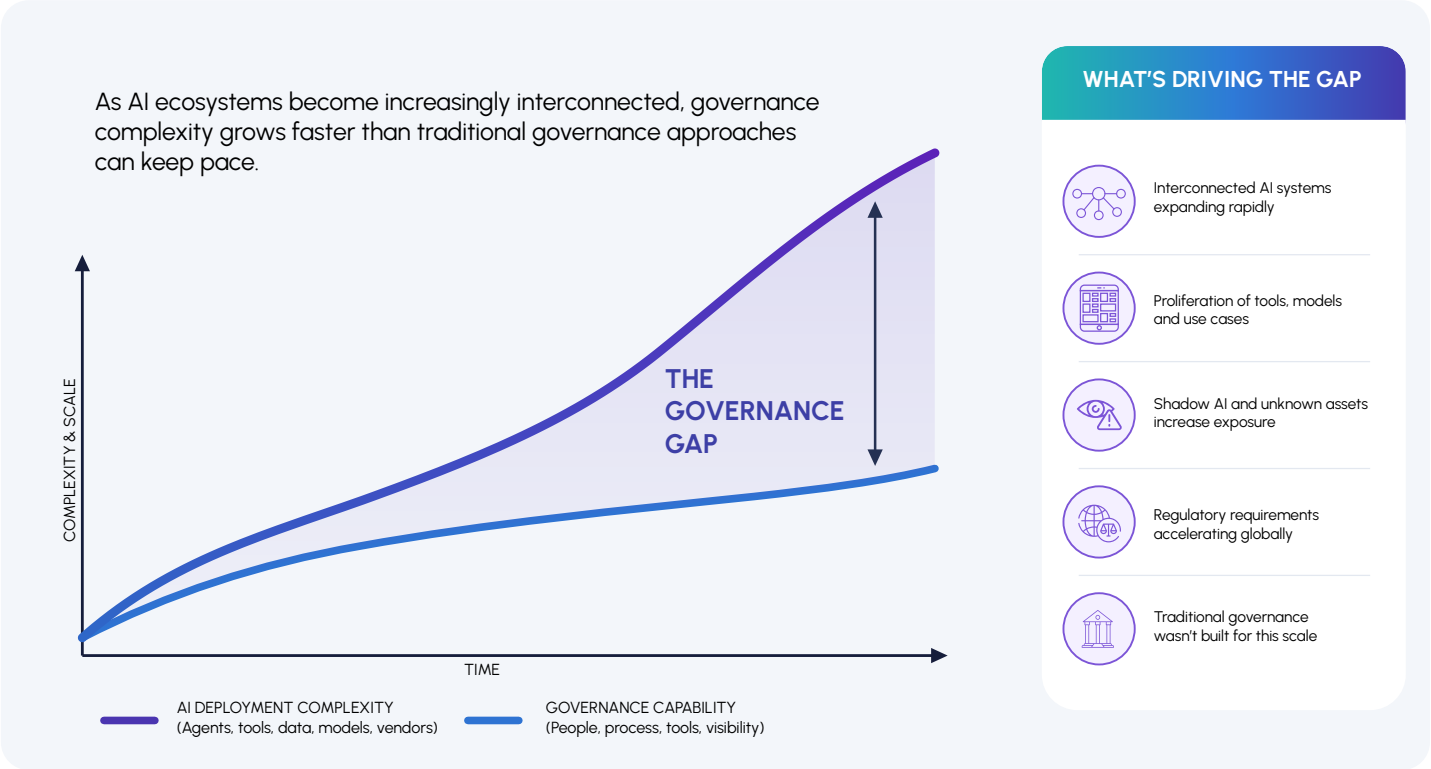

Source: Futurum Research

Visibility as the Foundation: You Cannot Govern What You Cannot See

The first-order governance problem is inventory. But static lists of AI assets are insufficient when those assets are interdependent, continuously changing and spread across multiple platforms, vendors and business units. What enterprises need is a relational view, showing which models power which agents, which agents call which tools, which data sets are in scope, and which use cases are affected if something changes or fails.

But an inventory only tells an organization what AI assets it has, not how those assets are being used, or how they relate to everything else running in the enterprise. Governance is, at its core, about relationships: a use case exposes a risk, that risk demands a control, that control is measured by a metric, and that metric is what proves a business outcome. Break any link in that chain and an organization is not governing its AI, it is guessing.

This matters in practice. If a foundation model is withdrawn, restricted by regulation, or found to underperform, understanding downstream exposure requires knowing the full dependency chain, not just that the model exists in a registry. IBM watsonx. governance addresses this through the Governance Graph – not a catalog, but a living, connected map of an organization's entire AI estate that captures relationships end-to-end, from AI assets through to the controls driven by corporate policy, enterprise risk, and regulatory requirement. With it, an organization can trace exactly which AI asset is being used, for what purpose, under what controls, and whether those controls are actually working, giving real-time visibility across every platform and production environment, including unmanaged AI.

Shadow AI compounds the visibility problem. Proof-of-concept projects left running, developer experiments on cloud platforms, AI features embedded in commercial SaaS products and models an employee brings into the organization without approval all represent governance exposure that conventional inventory approaches miss entirely. Organizations frequently discover these assets only when something goes wrong. IBM's AI asset discovery capability surfaces unknown assets and brings them into the

governance framework, enabling a deliberate decision about each one: assign it to a documented use case, decommission it, or accept it with recorded rationale. The capability is launching with coverage for AWS and Azure, with additional platforms planned later this year.

Vendor AI risk adds another dimension and it sits within a broader category that governance programs often treat separately from the rest of enterprise risk management, to their detriment. As AI becomes embedded in commercial off-the-shelf products, organizations inherit IT, operational, third-party, and business continuity risk from third-party vendors they may not be actively monitoring; the same categories long tracked in traditional GRC, now extended to AI. Tracking vendor AI posture – the models in use, the risk profile, the contractual obligations – is becoming a standard governance requirement, not an edge case. IBM supports this through dedicated vendor risk assessment partnerships with D&B, RiskRecon, SecurityScorecard, and Rapid Ratings, providing deeper vendor coverage and real-time incident monitoring rather than a point-in-time questionnaire.

Underpinning all of this is platform agnosticism, which matters for more than coverage alone. Governance coverage that stops at a single vendor's stack not only leaves material blind spots in multi-cloud, multi-model enterprise environments but also limits how much value an organization can get out of AI in the first place: a governance approach locked to one platform constrains an organization's ability to take advantage of what each platform does best. The governance layer must work across the full estate regardless of which platforms, models, or orchestration tools are in use, or the AI program ends up constrained by the very tooling meant to support it.

From Risk Identification to Control Enforcement

Enterprises face four parallel sources of control requirements and effective governance depends on identifying all of them before controls are built. The first is risk-driven: a mitigation required to address a specific risk, such as a rule that a customer service bot must detect and remove personal data with high confidence to prevent data leaks. The second is regulatory: obligations arising from frameworks such as the EU AI Act or sector-specific rules – for example, restricting production workloads for telecom and federal use cases in Canada to Protected B environments. The third is corporate policy: internally defined rules such as prohibiting open-source models in production. The fourth is business operations: constraints such as latency and payload limits that a line of business sets for a specific use case. The first three are typically defined and configured by risk, compliance, or security teams working in governance tooling; the fourth is typically defined by the line of business, closer to where the AI is actually running.

Custom risk taxonomies matter, and getting risk identification right is what determines whether the right control gets built in the first place. AI risk is not a fixed category set. Agentic AI in particular is generating new risk patterns rapidly and organizations with sector-specific or operationally-derived risk frameworks need governance tooling that can accommodate their own taxonomies alongside standardized ones. IBM watsonx.governance draws on built-in integration with established frameworks such as the IBM AI Risk Atlas and the NIST AI Risk Management Framework, on LLM-driven risk identification that flags patterns directly from use case data, and on dynamic risk scoring grounded in an organization's own taxonomy, so identification is not limited to whichever categories a vendor happened to pre-define. Tools that impose a rigid taxonomy on customers will find themselves misaligned with how risk is actually understood and managed within the enterprises they serve.

Identifying risk is necessary but not sufficient. The gap between risk assessment and control implementation is where many governance programs stall and where audit exposure accumulates. Effective governance requires a closed loop: obligations translated into controls, controls passed to the platforms where AI is built and deployed, evidence of implementation collected automatically, and ongoing monitoring of control effectiveness rather than binary compliance status (Figure 2).

Figure 2. The Governance Gap

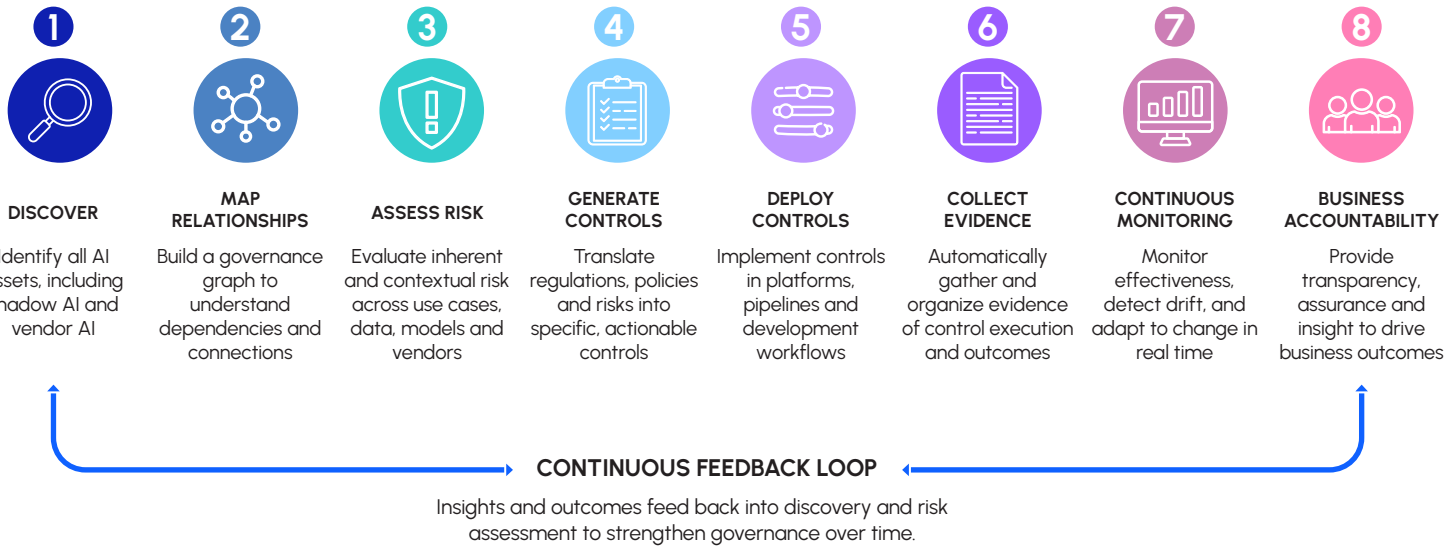

Source: Futurum Research

The manual approach to managing these requirements by translating regulatory text into controls, tracking implementation across development teams, and collecting evidence for auditors does not scale. IBM watsonx.governance approaches this through AI-driven control mapping that converts regulatory obligations and internal policies into what IBM refers to as AI Governance controls, defined and managed within watsonx.governance itself. Those controls are then passed as AI Operations controls to the control plane where the AI is actually running – such as watsonx Orchestrate – which implements and monitors them day-to-day. watsonx.governance pulls back evidence of implementation automatically and monitors effectiveness on an ongoing basis. This closes the loop between governance intent and operational execution without requiring manual coordination at each stage.

Regulatory horizon scanning is a related and increasingly important capability. Organizations that learn about new obligations at the enforcement stage are already behind. Earlier awareness allows time to assess applicability, define the necessary controls, and build implementation into development cycles before compliance becomes urgent. IBM has built this capability through a partnership with Cube and a broader partner ecosystem, which provide horizon scanning to surface potential incoming regulations and pull relevant obligations directly into watsonx.governance. The intent is to give compliance teams lead time rather than leaving them to react.

IBM's broader ambition in this area is to assemble the largest collection of AI-related regulatory content and frameworks available in an AI governance product – currently covering more than 200 regulatory frameworks – and to continue expanding that library through data partnerships rather than attempting to own regulatory content directly. The practical effect for customers is access to a continuously updated regulatory landscape without requiring an internal team to track it manually.

Accountability: Connecting AI Governance to Business Outcomes

Governance has historically been framed as risk prevention. The requirement that is now emerging particularly from boards and executive sponsors funding AI investment is to connect governance to value delivery. These are not competing objectives, but the tooling to support both simultaneously has been largely absent.

A pattern that has become common: AI programs proceed through build and deployment, and then several months later, the business discovers that costs exceeded projections and the expected return has not materialized, with no structured record explaining either outcome. This is a governance failure as much as a financial one. The absence of documented alignment between AI use cases, strategic objectives, anticipated outcomes, and actual costs leaves organizations unable to learn from the pattern or correct it.

IBM watsonx.governance addresses this through business value alignment, a capability that connects AI use cases to strategic initiatives at the point of inception rather than retrospectively. The platform enables organizations to capture what a use case is intended to achieve, which business objective it maps to, what outcomes and cost profiles are expected, and then to track actuals against those baselines by integrating with observability and financial tooling. The result is that governance becomes the mechanism through which AI programs demonstrate ROI – not just compliance – making it directly relevant to the executive stakeholders who are funding AI investment and need to prove its value. That shift matters because many organizations today are only accounting for what AI costs; the more useful question is what it is saving or generating in reduced operating cost or new revenue, and governance is what lets that case be made with evidence rather than assertion.

This capability also addresses a practical problem in larger organizations: the tendency to duplicate AI initiatives across teams without visibility into what already exists. Use case similarity detection, where the platform identifies overlapping submissions across a repository of active and proposed use cases, reduces redundant effort and focuses investment on net-new capability rather than rebuilding what is already in progress elsewhere.

The accountability layer does one further thing: it shifts the perception of governance from constraint to operational infrastructure. When governance is the system that makes AI programs legible to the business, auditable by regulators, and correctable when something goes wrong, it stops being the function that slows things down and becomes the function that makes sustained progress possible.

Conclusion: Governance as Operational Infrastructure

The governance challenge enterprises face is structural, not procedural. It cannot be resolved by adding headcount, tightening review cycles, or applying point-in-time assessments more frequently. The volume and complexity of AI assets in production and the pace at which that complexity is growing requires governance that is continuous, automated, and integrated into the platforms where AI is built and deployed.

Effective AI governance at scale requires three things working together (Figure 3). Visibility is the foundation: understanding what AI assets exist, including the ones an organization does not yet know it has. Control is what turns that understanding into action – enforcing policy, applying regulatory requirements, including the AI-related regulatory content needed to keep pace with a shifting legislative landscape, and setting guardrails across every environment where AI runs, not just one. Accountability is what proves the first two are actually working: tracking whether controls hold up in production, not just on paper, and connecting AI activity back to the business outcomes it was meant to deliver.

Figure 3. AI Governance at Scale

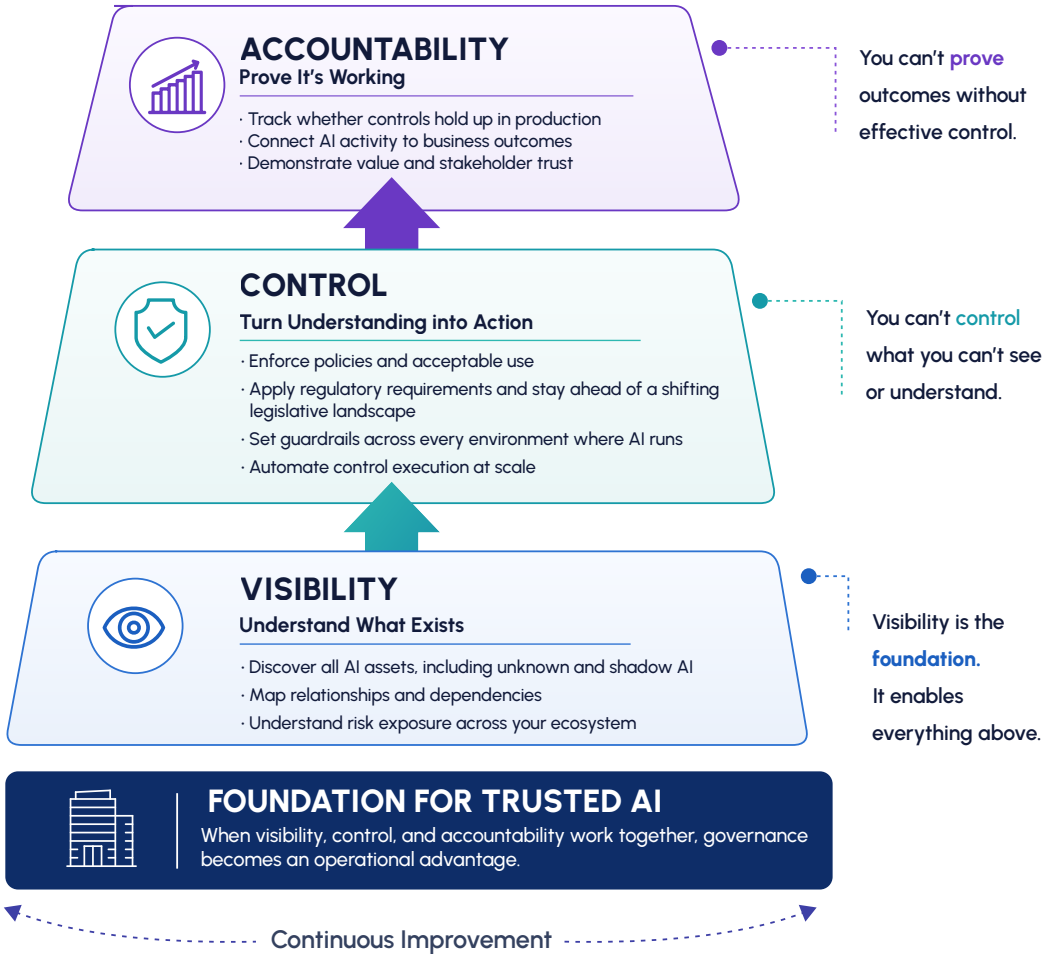

Source: Futurum Research

Organizations that treat governance as an afterthought will find it increasingly difficult to scale AI programs, satisfy regulators, or explain to the business what their AI is doing and what it costs. Those that build in governance from the start – as infrastructure rather than oversight – are better positioned to move at pace as the environment continues to change. IBM watsonx.governance is designed to serve that second approach and make governance an operational capability rather than a compliance exercise.

Important Information About This Report

AUTHORS

Nick Patience
Vice President & Practice Lead,
AI Platforms | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and The Futurum Group will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in context, displaying author's name, author's title, and "The Futurum Group." Non-press and non-analysts must receive prior written permission by The Futurum Group for any citations.

LICENSING

DISCLOSURES

The Futurum Group provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

ABOUT IBM

IBM is a leading provider of global hybrid cloud and AI, and consulting expertise. We help clients in more than 175 countries capitalize on insights from their data, streamline business processes, reduce costs and gain the competitive edge in their industries. Thousands of governments and corporate entities in critical infrastructure areas such as financial services, telecommunications and healthcare rely on IBM's hybrid cloud platform and Red Hat OpenShift to affect their digital transformations quickly, efficiently and securely. IBM's breakthrough innovations in AI, quantum computing, industry-specific cloud solutions and consulting deliver open and flexible options to our clients. All of this is backed by IBM's long-standing commitment to trust, transparency, responsibility, inclusivity and service. Visit https://www.ibm.com/us-en for more information & to learn more about our AI Governance solution see https://www.ibm.com/products/watsonx-governance.

ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

CONTACT INFORMATION: The Futurum Group LLC | futurumgroup.com | (833) 722-5337