

Report

The State of AI Governance

Why AI Expansion Forces New Rules for Managing AI Systems Now




Table of contents

1.	Executive Summary	3
2.	AI Scale & AI Governance Urgency	7
3.	AI Governance Maturity & its Limits at Scale	15
4.	The Operational Burden of Manual Governance	19
5.	Barriers to Governing AI at Scale	25
6.	The Hardest Enterprise-Wide AI Governance Challenges	30
7.	Scaling AI Governance with Core Enterprise Capabilities	35
8.	AI Governance Ownership & Control	39
9.	Runtime Governance and Emerging Oversight Needs	42
10	Methodology	45
11.	About the Authors	47

Why AI Expansion Forces New Rules for Managing AI Systems Now

About This Report

This report examines the current state of AI Governance in large enterprises as AI and agents move from experimentation to production. Drawing on survey findings from 371 senior leaders, it explores the widening gap between AI adoption and governance, while pinpointing the specific operational and

organizational barriers driving it. Ultimately, the report is intended to empower CIOs, AI risk leaders, and AI Governance teams to better understand these challenges and identify the capabilities needed to operationalize and scale enterprise-wide AI Governance.

Market Momentum: A 15.8B Strategic Category

Reflecting this challenge, AI Governance is rapidly emerging as a critical enterprise category. Analysts project the AI Governance software market to reach

\$15.8B by 2030, underscoring a growing demand for systems that help enterprises scale AI safely and unlock its value across the enterprise.

The Adoption Gap: Innovation outpacing Governance

However, as this market matures, AI adoption is scaling across the industry at a pace that creates a renewed sense of urgency. While the need to address the growing number of AI systems and variety of risks is evident, organizations'

AI Governance practices are currently lagging behind the speed of AI adoption. This gap puts organizational AI safety at risk of legal, economic and reputational harm.

Operational Barriers: The Cost of Manual Processes

This vulnerability is largely driven by the fact that many organizations, including Fortune 500 companies, are constrained by a lack of AI Governance processes, and limited expertise in addressing emerging AI risks. Furthermore, the

majority of enterprises do not yet have the right tooling in place, relying primarily on manual governance or ad hoc solutions. This creates a mounting operational burden that slows AI builders and innovation in general.

Strategic Imperative: Scaling through Automated Governance

As AI adoption continues to accelerate, organizations are increasingly challenged by limited visibility and fragmented oversight, which are intensified by the rise of AI agents and Shadow AI.

Addressing these gaps requires a shift towards scalable, automated, governance that enables centralized oversight, effective risk management, and regulatory readiness essential for the modern enterprise.

By the numbers

60%+ deploy AI across multiple departments or company-wide.

AI adoption is rapidly moving beyond pilots into enterprise-wide deployment. As AI spreads across the enterprise, governance shifts from occasional oversight to a recurring operational need.

Once AI spreads beyond 1-2 departments, Governance urgency jumps from

61% to 88%



Only 12%

currently use a dedicated AI Governance platform.



45%


spend 11-20 hrs/week on governance workflows

AI Governance Infrastructure lags, resulting in an operational bottleneck

As AI expands, high-burden governance (21-40 hrs/week) rises from

5% to 23%

Core Capabilities for Operational Governance



As AI systems scale and operate continuously, governance must evolve into standardized, automated operational infrastructure that enables

automated risk management and lays the foundation for frontier runtime governance.

Takeaway

As AI Scales, AI Governance is shifting from isolated policy frameworks to operational enterprise-wide infrastructure. Organizations that standardize workflows, automate risk management, and build runtime governance capabilities will be best positioned to scale AI safely.

Implications for Enterprise AI Governance

AI Governance must evolve from frameworks to operational systems

- Manual AI Governance processes cannot keep up with the pace of enterprise-wide AI adoption and the growing use of AI agents.
- Organizations looking to safely scale AI will need scalable, operational systems that translate governance principles into repeatable workflows and oversight.

AI risk management and centralized AI visibility are becoming core AI Governance capabilities


- As enterprises scale their AI and deploy increasing numbers of AI agents, it becomes increasingly difficult to maintain visibility over all AI assets and manage their associated risks.
- Establishing centralized visibility over AI assets and operationalizing AI risk management are imperative for enterprises looking to scale their AI Governance processes.

AI Governance must extend to cover the entire AI lifecycle

- As organizations adopt low-code tools and autonomous agents, governance challenges quickly extend beyond initial approvals.
- AI Governance systems will need to incorporate and rationalize data from continuous monitoring and oversight as AI proliferates across the enterprise.

02 •

AI Scale & AI Governance Urgency



AI Governance urgency spikes when AI moves beyond pilots


As AI expands beyond pilots, governance shifts from defining principles to operating enterprise-wide processes.

Most enterprises are already scaling AI beyond pilots

60% of organizations are already deploying AI across multiple departments or company-wide, the point at which AI Governance becomes operationally complex. This surge in adoption renders early stage

governance frameworks obsolete, where the governance challenge shifts from defining frameworks to operationalizing and streamlining a AI Governance process that is scalable with AI adoption.


Current State of AI Adoption and Implementation Across Organizations



Urgency Spikes with Scale

Among organizations that already consider AI Governance relevant, perceived urgency for AI Governance increases sharply as AI adoption scales


from developing AI pilots and proof of concepts to deploying AI applications beyond one or two departments.



A step change in AI Governance urgency (61% → 89%) occurs when companies scale from deploying AI in 1-2 departments to multiple parts of the company, and increases to 92% urgency for organizations deploying AI applications company-wide.

When considering the AI Governance drivers behind this urgency, respondents rate Managing AI-related risks, accelerating and scaling safe AI adoption & meeting regulatory requirements as the top three AI Governance drivers.

Top Governance Drivers



The urgency step change is reflected in how top drivers shift from future-proofing needs like meeting regulatory compliance and building external trust for companies deploying AI in 1-2 departments to near-term


operational needs like managing AI risk and accelerating and scaling safe AI for companies deploying AI across multiple departments or company-wide.

Half the market is scaling faster than its governance capability

While most companies are already deploying AI across multiple departments or company-wide, almost half of these companies have

little to no confidence in their ability to effectively manage AI risk.

Confidence That Current AI Governance Practices Effectively Manage AI Risk



As AI expands, so does the complexity around managing risk, visibility and oversight, hence the need for more structured and scalable governance approaches.

“Existing governance structures were not designed to keep pace with the velocity of AI initiatives”

Survey Respondent

Dir. of Corporate Governance

November 2025

“The organization is rapidly expanding its use of AI across multiple business units and we face increasing regulatory scrutiny and emerging standards around AI risk, ethics, and compliance.”

Survey Respondent

VP of Procurement Excellence

November 2025

Governance Urgency Hits a Tipping Point at Scale

Patterns suggest that AI Governance becomes most pressing at the point where AI begins to extend beyond one or two departments. At this stage, AI shifts from isolated use cases to enterprise capability where risk becomes more distributed and harder to contain. In pilot phases governance can seem manageable. Lightweight mechanisms such as principles, informal review processes, and localized oversight often give the appearance of control when scale is limited and risks have not yet fully materialized. However, the tipping point associated with scaling is not just a function of early-stage governance gaps. Rather AI Governance becomes more complex as AI scales with risk exposure, accountability, and coordination demands increasing nonlinearly.

AI Governance deficits can become embedded as AI scales, and early practice in each phase of adoption often sets norms for how AI is built and used across organizations. This makes standardization and control more difficult as initiatives progress, particularly as AI expands across teams, tools, and use cases without a comparable improvement in governance practices. As AI moves beyond developing initiatives and pilots into cross-department deployment, the AI Governance challenge shifts from local oversight to enterprise coordination. The number of initiatives increases, teams adopt tools independently, and risk exposure becomes more widely distributed. At this stage, AI Governance moves from defining standards, to an operational requirement in maintaining consistency, accountability, and visibility across a growing AI footprint.

Among n=371 surveyed governance experts, comments highlighted surging AI adoption demand, and fragmented processes in association with a culture of duplication and inefficiency. This shifts the AI Governance challenge from whether to establish guardrails, to how to implement them in a scalable way across teams, tools, and use cases.

“We have over 100 requests across different functions, every team wants to use AI tools or products.”

Survey Respondent

Dir. of AI and Data Science

November 2025

At the same time, scaling increases the likelihood of unmanaged or untracked AI usage, and surveyed experts repeatedly reference loss of visibility and a need for strong mechanisms to detect and manage tools in navigating the “Wild West” of cumulative enterprise-wide exposure. Organizations can often tolerate ambiguity and informal controls during early experimentation, but as AI adoption spreads across departments so too does the complexity of governance.

“We feel like it’s the ‘Wild West’ with data going into AI tools.”

Survey Respondent

CTO

November 2025

Confidence Stagnates as AI Scales

While governance urgency increases sharply as AI scales, confidence remains flat. As companies scale their AI adoption, confidence in AI Governance practices does not improve, but rather stagnates. As shown in the following table, **58% of organizations still developing initiatives report confidence in current governance practices, with this percentage dropping to 47% for organizations deploying AI applications company-wide.**

Confidence that Current Governance Practices Effectively Manage Risk by Adoption Stage

AI Adoption Stage	% Confident or Very Confident
Developing AI initiatives, pilots, proofs of concepts, or applications, but none have been deployed	58%
Deploying AI applications in 1-2 departments	45%
Deploying AI applications in multiple parts of the company	54%
Deploying AI applications company-wide	47%

As AI scales, so too does the complexity of governance and risk challenges. Further limiting confidence, surveyed experts also frequently cited lack of existing AI Governance processes or personnel and uncertainty over whether governance tools meet current needs as barriers to scaling AI Governance. Taken together, patterns place many organizations at the convergence of weak confidence in governance effectiveness and limited capabilities and resources to address gaps.

“As company-wide use of AI becomes more widespread, we find existing basic data policies don’t cover all of the possible risk scenarios.”

Survey Respondent

Chief Administrative Officer

November 2025

03 •

AI Governance Maturity & its Limits at Scale



AI Governance maturity still lags behind AI Adoption

The AI Governance challenge changes once AI spans teams, use cases, and risk levels.

In this report, AI Governance maturity is the ability to operationalize AI Governance using defined processes, ownership structures, and tooling to enable enterprise-scale AI deployment.

Most organizations are still formalizing and optimizing their AI Governance practices

Despite the scale of AI adoption and the growing urgency for governing its use, AI Governance practices still lag behind AI adoption.


60%

Over 60% of organizations are already deploying AI across multiple departments or company-wide

4%

But only 4% are governing AI at scale, while the majority are still formalizing or optimizing their AI Governance programs.

While 60% of organizations have already scaled AI across multiple departments or company-wide, AI Governance maturity is failing to keep pace; currently, only 4% are governing these systems at scale.



Survey responses indicate a broad alignment on the understanding of AI governance as an enterprise-wide AI oversight – ensuring centralized visibility and coordination across AI initiatives (43%), with some more narrowly defining it as Regulatory and Compliance Alignment by meeting laws (22%), standards and audit requirements or data governance for AI, focusing on data quality, lineage, and privacy (18%). However, this conceptual clarity has not yet translated into operational measures characterized by consistent processes, workflows, tooling, or ownership models capable of operating at scale.

43% of respondents define AI Governance as enterprise-wide AI oversight, ensuring centralized visibility and coordination across AI initiatives, while smaller minorities narrowly define it as regulatory alignment or data governance for AI.

Adoption Outpaces Maturity

Mature AI Governance is characterized by an enterprise-wide process with continuous oversight and automation. While over 60% of organizations report deploying AI across multiple departments or company-wide, only 4% report full maturity in governing AI at scale. Though many organizations conceptually understand AI Governance as enterprise-wide oversight, most are operating in a transitional state—managing broad AI deployment with governance structures that were not designed to scale. As AI initiatives expand and governance becomes more operationally complex, a majority of organizations are still formalizing or optimizing governance practices rather than operating a fully scaled system. In practice, this is often reflected in AI Governance programs that are still building foundational elements such as clear ownership, standardized review processes, consistent documentation, and reporting mechanisms. This gap between AI deployment scale and AI Governance maturity suggests that many organizations are operating in a transitional state, forced to manage broad AI deployment while simultaneously trying to scale lagging governance processes.

“The push to adopt AI is moving faster than the governance approach can keep pace with.”

Survey Respondent


Cyber Audit Dir.

November 2025

Alongside the scarcity of mature AI governance, urgency remains high. 77.9% of surveyed organizations report urgency in establishing or strengthening governance. However, scaled maturity remains difficult to achieve as governance shifts from frameworks to operationalization. At enterprise scale, AI Governance requires more than principles or isolated reviews. Scaled AI Governance requires repeatable workflows, defined accountability, and visibility across the AI footprint. As organizations face challenges in manual workflows/processes and maintaining centralized visibility and control, scaled maturity can remain difficult to achieve even amid high urgency. Collectively, findings indicate a persistent gap between the pace of AI deployment, and the operational capacity required to govern AI effectively at scale.

04 •

The Operational Burden of Manual Governance




Manual governance does not scale, creating a growing operational burden

Despite readiness, most organizations are still relying on manual AI Governance or adhoc solutions

Despite the urgency and the scale of AI adoption, AI Governance tooling remains mixed and transitional. Nearly 60% of surveyed organizations rely on internal tooling or a hybrid model of

internal and third-party tools, while 27% still depend primarily on manual processes, indicating uneven tooling maturity across enterprises.

AI Tooling Approach Across Organizations



Only 12% of organizations operate a dedicated AI Governance platform to help them operationalize, automate, manage and scale AI adoption with AI Governance.

Operational burden rises as AI scales


45%

of organizations lose 11–20 hrs/week to manual governance intake, this scales with org size, where 27% of 100k+ orgs: lose 21–40+ hrs/week.

Across organizations, 45% report losing 11–20 hours/week to manual governance intake. This rises to high-burden AI Governance (21–40+ hrs/week) when AI expands beyond 1–2 departments. Across deployed stages, 11–20 hours per

week is the most common governance workload, while high-burden cases increase sharply as AI scales. High burden governance jumps from 5% to 23% when AI expands from 1–2 departments to multiple departments.

Weekly hours spent on intake and assessment by adoption stage



“We have overwhelming requests from multiple teams over 100. We were able to approve only 6”

Survey Respondent
Dir of AI and Data Science

November 2025

“Too many projects happening across the organisation. Everyone wants to do their own thing, creating duplicates, inefficiencies”


Survey Respondent
Chief Commercial Officer (CCO)

November 2025


High-risk use cases drive a disproportionate share of operational burden. Half of organizations report that low risk use cases are not an operational

burden with reviews completed within days, while 30% report high risk use cases that can take more than a month to complete a review.

High-risk governance reviews take longer




As AI scales, patterns suggest that governance increasingly becomes a capacity and operational challenge with manual workflows and uneven tooling maturity limiting organizations' ability to maintain consistent oversight.



52%

Of low-risk use cases complete in ≤ 7 days



30%

Of high-risk use cases take ≥ 1 month

Manual Governance Breaks at Scale

As AI adoption expands across teams and use cases, governance shifts from an occasional control function to a continuous operational workload. Manual intake, review, and documentation processes that were manageable during early experimentation become increasingly difficult to sustain as AI expands and the volume of initiatives accelerates. At scale, the volume and complexity of use cases grows, coordination demands increase, and manual processes begin to break forming governance bottlenecks. Notably, the share of organizations reporting high-burden governance intake and assessment (21-40+ hours per week) increases sharply as AI scales beyond 1-2 departments. Only 5% of organizations deploying AI in 1-2 departments spend 21-40+ hours per week on intake and approval, compared to 23% of organizations deploying AI across multiple departments. This high operational burden creates risk of delayed approvals, inconsistent oversight, or governance being bypassed altogether.

Tooling Remains Transitional


Surveyed experts frequently cite uncertainty over whether current tools meet AI Governance needs as a key barrier to scaling AI Governance. Consistent with this challenge, findings suggest that AI Governance tooling maturity remains mixed across enterprises. 39.1% of organizations rely on internal tooling or hybrid approaches, while 19.4% still depend on primarily manual workflows. While 12% report using third-party platforms for AI Governance, the transition to purpose-built governance platforms remains early, leaving many organizations without standardized structures and tooling to operationalize governance consistently at scale.

High-Risk Use Cases Drive Cycle Time

AI Governance workload is not evenly distributed across AI initiatives. While low-risk use cases may move through review in less than a week, high-risk applications often require deeper assessment, broader stakeholder involvement, and more extensive documentation. Surveyed experts report significantly longer review cycles for high-risk use cases, with 51% indicating that intake and review take up to one month and 30% reporting timelines of more than one month. As AI adoption scales, high-risk use cases can increasingly shape the overall AI Governance pipeline and can drive longer cycle times, particularly when governance relies heavily on manual workflows. This impact is amplified in barriers reported by surveyed experts. Namely, lack of existing AI Governance processes, uncertainty over whether tools meet needs, and limited specialized staff, can further exacerbate AI Governance bottlenecks and prolong cycle times for high risk use cases.

05 •

Barriers to Governing AI at Scale




Governance readiness is still constrained by operational and organizational barriers

Top barriers to scaling are operational readiness gaps, not technical limitations.

Many teams are constrained by limited capacity to operationalize governance at scale. Specifically, key challenges include:

- Lack of existing AI Governance process or dedicated personnel.
- Uncertainty over whether tools meet AI Governance needs.
- Lack of specialized staff or expertise.

Most Significant Barrier to Scaling AI Governance



Barriers Evolve by Adoption Stage

The highest concentration of barriers sit in process maturity, expertise, and tooling uncertainty. Findings suggest that even when teams have strong technical AI capabilities, they may struggle to operationalize governance across a growing AI portfolio. In other words, **the limiting factor in robust governance frameworks center on lack of automated repeatable scalable mechanisms.**

As AI Governance matures, governance pain shifts from foundational gaps (such as lack of existing governance process or functional ownership) to scalability constraints (such as manual or rigid governance workflows or uncertainty about evolving regulations).

While maintaining centralized visibility is a constant challenge across maturity levels, managing AI Governance with manual workflows tends to decrease as organizations mature. Conversely, difficulty tracking and interpreting regulatory changes tends to increase as organizations mature.


Early-stage organizations are primarily constrained by foundational gaps (such as lack of existing AI Governance process) while later-stage organizations are primarily constrained by operational and scalability challenges (such as manual or rigid AI Governance workflows).

Organizations not yet deploying AI or still in early development most often cite lack of executive support, clear ownership, or absence of AI Governance processes as their most significant barrier to scaling.

Once AI is deployed in multiple parts of a company, the most frequently cited barriers include:


- Incomplete inventory of AI systems.
- Manual or rigid governance workflows.
- Uncertainty about evolving regulations.
- Concerns that existing tools do not meet AI Governance needs.

Most significant barrier to scaling and current stage of adoption or implementation.



Operational Barriers Strongly Associate with High Effort


Most significant barrier to scaling and current stage of adoption or implementation.



While most respondents facing AI Governance barriers spend 11-20 hours per week on AI Governance activities, organizations citing manual or rigid workflows report disproportionately high effort spent on intake/review (21-40+ hours per week). Additionally, regulatory uncertainty and lack of functional ownership are also associated with high effort. Together, these patterns suggest that AI Governance effort is influenced not only by the scale of AI activity, but also by how AI Governance processes and ownership are structured. Conversely, those who cited lack of executive buy-in as their most significant barrier tend to spend fewer hours per week on AI use case intake and review, suggesting that the ability to perform AI Governance at all is conditioned on organizational prioritization.

06 •

The Hardest Enterprise-Wide AI Governance Challenges




Third-party risk and centralized visibility are the hardest AI Governance problems to solve

Third party risk and maintaining centralized visibility are the top AI Governance challenges for organizations


As organizations scale, they report challenges in assessing and managing the risk of 3rd party AI systems and having centralized visibility over all their AI assets. With scale, risks proliferate in number and type making their AI Governance more complex and

unmanageable with manual efforts. AI Governance professionals, challenged by lack of resources, buy-in and specialized expertise struggle to keep up with the growing volume, and speed of innovation.

How Challenging do Respondents' Rate Key Governance Activities



Among rated challenges, what do respondents find to be the most challenging aspect of AI Governance today?




AI Governance Maturity Reshapes Governance Challenges

AI Governance maturity does not eliminate governance challenges, but shapes them. Challenges managing third-party AI risk increase steadily as adoption broadens. Maintaining centralized visibility and control is frequently cited as “challenging” across every stage of AI Governance maturity, with roughly two thirds of organizations reporting this as challenging. Visibility is not just a hurdle to beginning AI Governance activities, but is a challenge that must be continuously managed as AI scales. Similarly, assessing and managing risks from third-party AI systems and vendors is widely cited as challenging across all stages.


Conversely, difficulty tracking and interpreting AI regulatory changes increases as organizations mature, indicating this challenge rising as organizations grow in maturity.

Managing AI Governance with manual workflows and processes decreases with AI Governance maturity, indicating that manual workflows and processes may be most acute during the middle stages of maturity. Together, these patterns suggest that advancing AI Governance maturity shifts governance effort from foundational coordination challenges toward managing scale, complexity, and external risk.




How AI Scale Reshapes AI Governance Challenges

Challenges managing third-party AI risk also increase steadily as adoption broadens. For organizations in early stage adoption, key challenges tend to focus on foundational AI processes such as establishing centralized visibility and control, and tracking and interpreting regulatory changes. However, the challenges of robust governance evolve as organizations scale their AI use, with activities such as assessing and managing third party risks spiking for organizations that are beginning to scale AI in more than one department. While some AI Governance challenges shift with maturity, manual workflows remain a consistent constraint across adoption stages as teams are continuously challenged to break free from rigid processes and move into automated, repeatable, AI Governance software frameworks.



07 •

Scaling AI Governance with Core Enterprise Capabilities




Core AI Governance capabilities are converging around centralized visibility and AI risk management

AI Risk Management and **Data Governance Visibility** are top governance needs for organizations scaling AI. Risk management and visibility also map most strongly to respondents' highest severity pain points and consistently show high

co-occurrence with primary barriers to scaling AI Governance such as lack of existing governance process, uncertainty of whether tools meet AI Governance needs and lack of specialized staff or expertise.

Most Significant Barriers to Scaling AI GRC



Once organizations have established foundational AI Governance, they will need an operational layer to report on their AI governance process through Reporting and Dashboards (68%) and AI Governance documentation and control (67%). These capabilities enable them to provide a system of record and prove compliance to executives, auditors and customers.

“We need to get a handle on the risks that AI adoption poses to our organization and our data.”

Survey Respondent

Security / IT Risk & Compliance

November 2025

“Visibility into what AI is in use, what it is being used for and what data it is consuming.”

Survey Respondent

Security / IT Risk & Compliance

November 2025

Core Capabilities Converge around a Baseline

The market is moving from “what is AI Governance?” to “what do we need to run it at scale?”. In particular, AI risk management and data governance visibility stand out as the most urgent needs for organizations scaling AI. These capabilities align strongly with high severity pain points and co-occur with common barriers to scaling AI Governance such as lack of governance processes. Practically, findings indicate that many organizations are not struggling to identify abstract risk, but to institutionalize risk identification, review, and decisioning in a way that is consistent, auditable, and sustainable as AI use expands.

Overall, patterns suggest that sustaining AI at scale requires strengthening the AI Governance operating model around a common set of foundational capabilities including clear risk management workflows, visibility into data and AI assets, and audit-ready documentation and controls. As organizations navigate pain points and tooling uncertainty, the differentiator will incredibly be the ability to operationalize AI Governance continuously.

Core Capabilities Persist across Roles

Analysis of capabilities by functional roles reveals strong alignment around core capabilities with meaningful divergences across stakeholder groups.

Across all roles, AI Risk Management is prioritized as a top desired capability reflecting its foundational role in AI Governance. This is followed closely by Data Governance visibility and Reporting and Dashboards which underscores the demand for transparency, traceability and measurable outcomes

Beyond shared priorities, roles diverge reflecting their operational responsibilities:

- AI Governance professionals prioritize reporting and dashboards as the 2nd most important capability following AI Risk Management reflecting their need to showcase the value of AI Governance and quantify their operational impact.
- AI/ML Management prioritize Data Governance and Visibility as the most important capability followed by AI Risk Management
- Data Engineering, Analytics and Procurement prioritize Dashboards and Reporting as the single most important capability in an AI governance platform.

Overall, while organizations converge on common governance capabilities, AI Governance platforms need to cater to the needs of the diverse stakeholders in the AI Governance process to ensure wide adoption and scalability.

08 •

AI Governance Ownership & Control



AI Governance teams drive requirements, and IT controls execution and spend

CIO is the economic buyer, while AI Governance drive requirements

Ownership of the AI governance budget is often treated as a proxy for the strength of executive sponsorship. In practice, more concentrated budget ownership, may enable faster AI governance purchasing decisions. Across organizations surveyed, AI Governance budget ownership is generally fragmented with no clear single decisive owner. However, the highest concentration of ownership tends to fall on the CIO/IT office. **26% say the CIO/IT ops office controls the budget.** In these organizations, they also control the final decision making and formal approval of AI Governance tools.

In some industry verticals, CIOs control is far more concentrated. In Manufacturing & Retail, CIOs control the budget in 50% of organizations while in Healthcare & Life Sciences, CIOs control the budget in 40% of orgs.

While AI Governance professionals drive the requirements and are an instrumental part of the decision making 59% of the time, they own the budget in only 5% of the organizations. Specialized AI Governance functions are growing but are still a minority compared to well established IT structures. Other roles such as security, legal/compliance, risk and procurement sit on the AI Governance committee acting as gatekeepers to the decision process.

59%

Of organizations, AI Governance professionals drive the final decision making in tool selection and purchase



5%

Of AI Governance professionals own the AI Governance budget.




Implications for Governance Professionals

The fragmentation of AI Governance budget ownership paired with relatively centralized execution authority in IT can create a misalignment between ownership and control. While AI Governance practitioners are frequently responsible for defining requirements and shaping tool selection, they often lack the formal authority needed to act on requirements. This structural gap can create an environment of slow decision cycles and high coordination overhead where AI Governance outcomes depend more on cross-functional negotiation than clear decision rights and streamlined decision making. It can also increase the likelihood that AI Governance programs are evaluated through an IT procurement lens rather than an AI Governance effectiveness lens. This dynamic can position AI Governance teams with accountability for outcomes, but low resources to deliver on them.

This concentration of budget control within CIO/IT offices signals that AI Governance is often still institutionalized as a technology operating model rather than an independent governance consideration. While this can accelerate execution, it also risks positioning AI Governance as just an IT element, rather than an enterprise accountability shared across risk and compliance. Consequently, AI Governance success increasingly depends on not only governance expertise, but stakeholder management, cross-team alignment, and translation of needs into IT-aligned propositions. Until budget ownership and decision rights align with AI Governance accountability, governance practitioners will need to have a strong presence in cross-functional alignment to drive progress in a fragmented environment.

09 •

Runtime Governance and Emerging Oversight Needs



Runtime AI Governance is an emerging frontier for AI Governance


Shadow AI and Agentic AI observability are two emerging desired capabilities for AI Governance platforms

While AI Risk Management and Data Governance Visibility top the desired capabilities for an AI Governance platform, respondents also rate Shadow AI visibility (65%) and Agentic AI observability (66%) relatively high.

Third-party AI risk (75%) and maintaining centralized visibility (69%) are top governance challenges for the segment who rate agentic AI observability as critical, both directly aligned with agentic AI observability needs. Nearly 3 out of 4 organizations viewing Shadow AI Discovery as critical also view Agentic AI Observability as critical which demonstrates that they are complementary capabilities.

As organizations navigate pain points and needs, clarity in runtime governance, risk management and audit readiness, and core operational activities become critical in sustaining AI scale.

How critical are the following capabilities in your governance program?



However, for many organizations, agentic observability and shadow AI are not Day 1 AI Governance capabilities, but rather aspirational needs. The agentic AI observability segment reported top AI Governance barriers that demonstrate their lack of readiness for agentic AI adoption, such as lack of existing governance processes or dedicated personnel, followed by tool uncertainty and skills gaps. Likewise, the Shadow AI segment's biggest barriers to AI Governance were lack of AI Governance process/personnel and lack of specialized staff/expertise. As it stands, many of these organizations lack the AI Governance structures, resources and capacity to handle the overwhelming deluge of ungoverned AI. **Organizations succeed by solving the operational execution problems first, then expanding into runtime AI Governance.**

From AI Governance Foundations to Ongoing Oversight

Runtime AI Governance is emerging as a new consideration in AI Governance maturity. While AI risk management and data governance visibility remain the most foundational needs for organizations scaling AI, the additional importance placed on shadow AI visibility and agentic observability indicate that AI Governance expectations are expanding beyond static frameworks. In the near term, many organizations are focused on accelerating AI adoption and gaining oversight on use cases. However, as the AI footprint grows, organizations will face a different operational reality: AI Governance that must persist beyond launch. They may encounter an “oversight wall” where AI initiatives outgrow governance capacity and become more difficult to monitor, control, and audit.

As AI Governance teams navigate barriers in lack of existing governance processes, limited dedicated personnel, and uncertainty around tooling, runtime governance can feel aspirational. Governance professionals recognize that unmanaged AI and more autonomous agentic systems introduce new categories of risk, but are not equipped with the operating model to govern them effectively. Taken together, patterns imply that scaling AI Governance requires a progression from risk frameworks, to operational execution, to runtime oversight. Foundational AI Governance capabilities such as documentation and controls, reporting and dashboards, and centralized visibility are not merely baseline features, but prerequisites for a scalable, built-to-mature, AI Governance system. In this sense, foundational capabilities are not just table stakes but the launch pad that will allow AI Governance to expand into runtime oversight without proportionally increasing AI Governance burden.

About Credo AI

Credo AI is the enterprise AI governance platform for organizations that need to scale AI with control, accountability, and proof. The platform helps enterprises discover, assess, govern, and monitor AI systems across the full lifecycle, including models, applications, agents, workflows, and third-party AI vendors. Credo AI gives security, risk, compliance, legal, and AI teams a centralized system to manage inventories, map policies to regulatory and internal requirements, assess risk, and generate audit-ready evidence for oversight and reporting.


Built for the operational realities of Generative and Agentic AI, Credo AI extends governance beyond static documentation into continuous oversight. Its platform includes capabilities such as GAIA for AI agent governance and assistance, AI Agent Registry for agent inventory and control, and Shadow AI Discovery to uncover unapproved AI usage and connect detection to governance workflows. Credo AI also offers Advisory Services to help enterprises design and operationalize governance programs that align with evolving policy, security, and business requirements.

Credo AI was founded in 2020 and has been recognized as a Leader in The Forrester Wave™: AI Governance Solutions, Q3 2025, a Gartner Cool Vendor, Fast Company 100 awardee, and a representative vendor in Gartner market guides covering AI governance and AI TRiSM.

[Learn More](#)

10 •

Methodology



Methodology

The findings in this report are based on an online survey that was in the field from November 22nd to December 22nd, 2025. The study garnered responses from 371 participants across the US, UK, France and Germany. These countries were selected specifically due to their active or impending AI regulations.

2. Data collection was conducted in partnership with Newton X.

Respondent Profile:

To ensure high quality insights, all participants were screened to ensure they hold Director-level positions or above at enterprise organizations with 1000+ employees. The sample was balanced to represent a diverse range of industries such as Financial Services, Healthcare/Life Sciences, Insurance, Technology/SaaS, HR, Energy, Government & Public Sector, Manufacturing.

Role & Influence:


The study looked beyond AI Governance professionals to capture the broader eco-system of decision-makers. Respondents were screened for roles with direct influence of the AI Governance process, including:

- Legal, Compliance, and Risk
- CISO, Security, and IT GRC
- AI, Data and ML Leadership (CDO/CIO/CTO)
- Procurement and 3rd party risk
- AI Governance and Responsible AI Leaders

Additionally, all organizations included in the sample were screened to verify they have an active or emerging need for AI governance.

11 •

About the Authors



About the Authors



Yomna Elsayed, Ph.D

Head of Research

Credo AI

Yomna leads user research and strategy at Credo AI, where she helps Fortune 500 organizations bridge the gap between technical AI implementations and human-centred AI governance. With over a decade of experience in mixed-methods research, she specializes in translating complex market and user signals into GTM and product strategies

An expert in Human-Computer Interaction (HCI), Yomna holds a PhD from the University of Southern California. Her academic work, which has been published in leading journals such as the Journal of Computer-Mediated Communication, focuses on the intersection of human behavior and algorithmic systems. In this report, she applied her expertise to uncover the organizational challenges shaping the current status quo of AI governance.



[View on LinkedIn](#)

About the Authors



Cecily K. Jones, MS

Senior User Researcher at Credo AI

Credo AI

Cecily is a senior user researcher at Credo AI where she supports the development of AI governance tooling and best practices through human-centered research. With eight years of research experience spanning digital enterprise environments, she specializes in technology adoption and personal productivity, drawing on expertise in behavioral analysis and industrial-organizational psychology.

An applied psychology specialist, Cecily holds a master's from Kent State University and bachelor's from NC State University, and has co-authored research published in peer-reviewed outlets such as Scientometrics and the National Library of Medicine.



[View on LinkedIn](#)

Contributors



Design and Layout

Lucy van der Meer

Brand and Design Manager

Credo AI

Lucy van der Meer leads brand and design at Credo AI, where she helps translate complex ideas around trusted AI governance into visual experiences that leave their impact in the hearts of people.

Her background spans traditional arts, digital design, and everything in between. But what's always driven her isn't the medium — it's the mark. The moment a brand stops being something people see and starts being something they remember. Lucy works at the intersection of human psychology and design strategy, building brands that don't just look good, but mean something to communities all around the world.

Before Credo AI, Lucy launched and relaunched brands internationally, working across user research, web design, and brand identity to create experiences that stuck around long after the first impression. Like graffiti on a wall, weathered, real, and impossible to ignore.



[View on LinkedIn](#)

Contributors



Content Development

Jerome J. Sanders

Senior Content Marketing Manager

Credo AI

Jerome leads content marketing at Credo AI, where he helps Fortune 500 organizations build trusted AI governance through data-driven strategic storytelling and thought leadership. With over a decade of experience in content strategy, product marketing, and product development at Cisco, he specializes in translating complex AI concepts into narratives that drive business outcomes, clarity, and industry influence for senior executives.

A 2024 Aspen Institute First Movers Fellow and United Nations Global Compact SDG Innovation Accelerator, Jerome focuses on expanding equitable access to foundational AI technologies. He holds a business degree from Santa Clara University's Leavey School of Business.



[View on LinkedIn](#)