The ultimate enterprise AI governance report

How organizations will scale AI in 2026

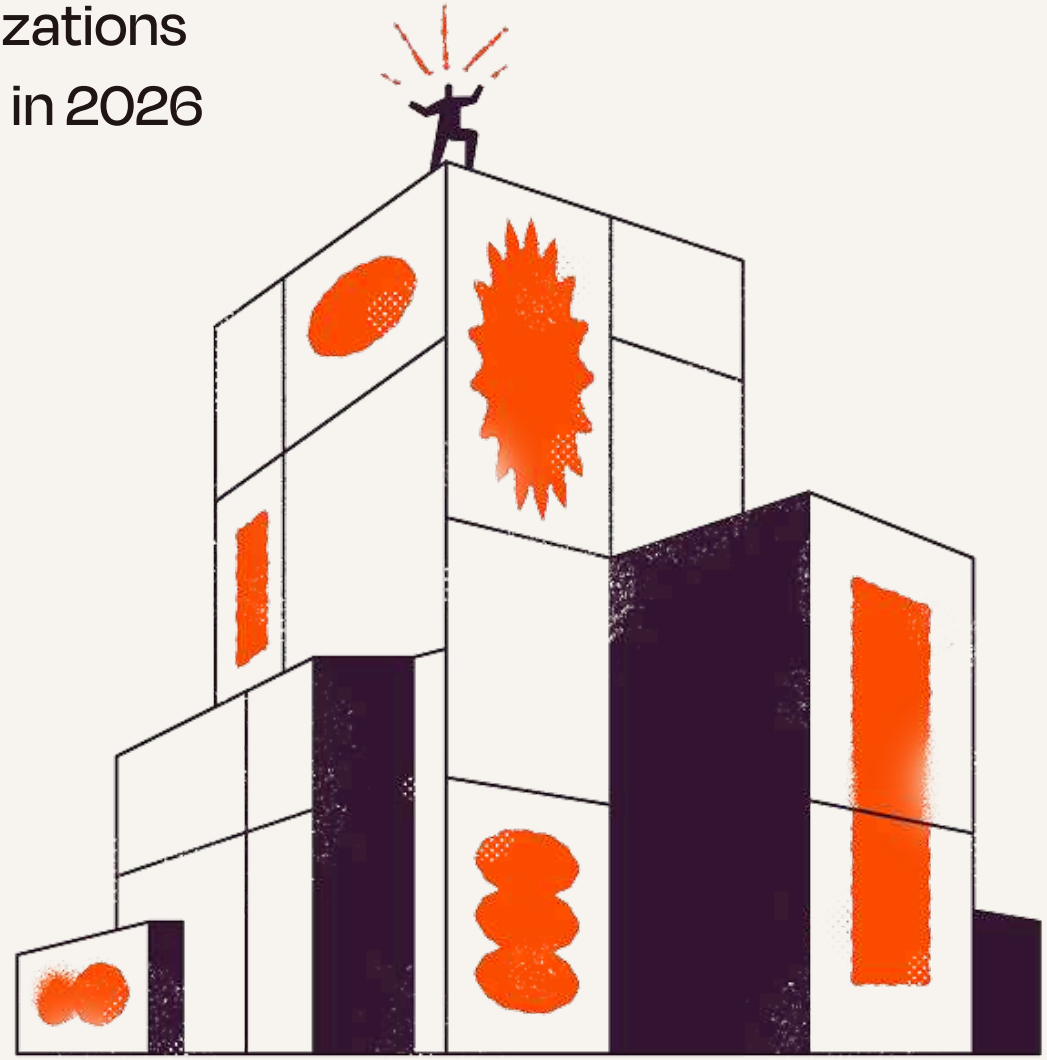

zapier

DATA-DRIVEN INSIGHTS

Introduction

Most enterprises think they have AI governance on lock. They don't—at least not the kind that works at scale.

As organizations push AI beyond pilots and into production workflows, they're discovering that existing governance models—manual reviews, centralized approval boards, fragmented policies—can't keep pace with real-world AI usage. The result: governance is now the biggest barrier to organization-wide AI transformation.

Based on a survey of 200 enterprise technical leaders and practitioners across IT, Engineering, Data, and Operations at companies with 1,000+ employees in the US, Canada, and Europe, this report examines how organizations govern AI today—and where the model is breaking down.

The AI governance reality gap

The AI governance reality gap

AI governance is present, documented, and acknowledged. It still doesn't work.

75% of practitioners surveyed agree that their organizations have clearly documented AI governance expectations. That looks like progress—until you consider a separate finding: 69% of practitioners agree that those expectations are difficult to translate into concrete implementation decisions.

Documentation alone isn't enough to close the governance gap. When a practitioner needs to determine whether specific data usage complies with policy, or when a team lead needs to decide whether to proceed with an AI deployment that sits in a governance grey zone, the documentation doesn't resolve the question. It describes the landscape without clearing the path.

The downstream effects show up in how practitioners describe their daily work.

When asked which governance challenge requires the most day-to-day judgment, no single issue dominates, but the top two are telling.

GOVERNANCE CHALLENGES DEMANDING MOST DAILY JUDGMENT

When asked which governance challenge requires the most day-to-day judgment, practitioners split evenly between data compliance and managing approvals.

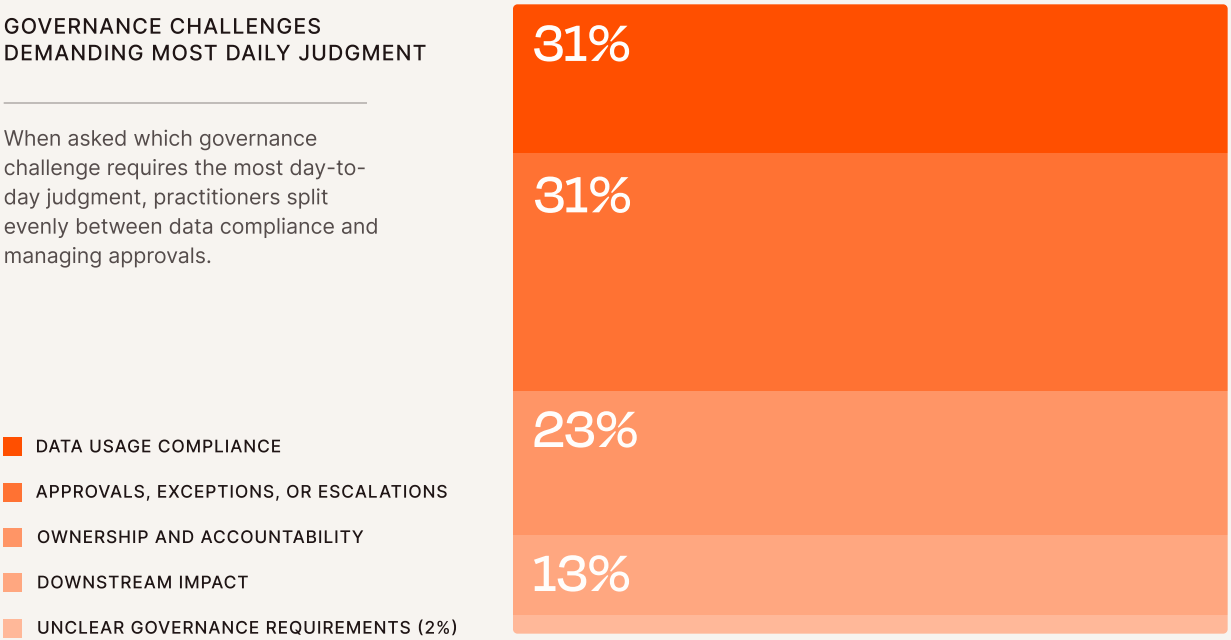

The governance challenges that demand the most daily judgment from practitioners:

- **31%** Determining whether specific data usage complies with policy
- **31%** Managing approvals, exceptions, or escalations
- **23%** Ensuring clear ownership and accountability for AI systems
- **13%** Preventing downstream business, security, or compliance impact
- **2%** Interpreting unclear or incomplete governance requirements

The top two answers are tied at 31%—and both are operational issues, not policy problems. The rules are known. Applying them is the hard part.

> “
>
> The [governance] gap is often a failure to translate high-level ethical and risk management princi-ples into the specific, day-to-day ‘how-to’ for an employee.

RAVI MALIK, SVP & GLOBAL CHIEF INFORMATION OFFICER, BOX

box

Recycled governance won't cut it

59% of leaders agree that their organization's AI governance is largely an extension of traditional IT governance models rather than a purpose-built model for AI at scale. That's the upstream root cause.

When organizations don't design governance based on how AI actually works—where decisions are distributed, models evolve, and data flows across tools and teams—the policies don't align with reality.

That structural misfit shows up in consistency. 56% of leaders agree they struggle to apply AI governance consistently across teams, tools, and workflows. The framework may exist in theory—enforcing it uniformly across a complex enterprise is another matter entirely.

The top governance challenges among leaders:

- **33%** Maintaining visibility into where AI is used and who owns it
- **30%** Enforcing data handling and usage policies consistently
- **19%** Managing approvals and exceptions without slowing delivery
- **13%** Meeting regulatory or audit requirements
- **5%** Preventing cost, usage, and performance surprises

Regulatory compliance—the concern that dominates public discourse around AI governance—ranks a distant fourth. The real friction is operational: visibility and enforcement.

Leaders already know the visibility problem is real. When asked about their confidence in having an accurate, real-time view of AI usage across their organization, 41% of leaders describe their confidence as slight or nonexistent. Compare that to the 45% that are moderately confident, and the paltry 14% of leaders who define themselves as very confident.

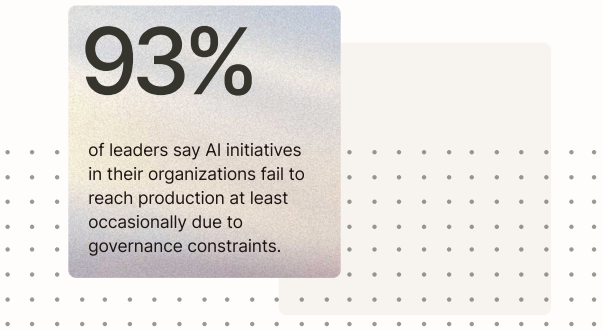

93% of leaders say AI initiatives in their organizations fail to reach production at least occasionally due to governance constraints. Moreover, practitioners are 3x more likely than leaders to say governance almost always blocks projects.

Governance by manual review

The majority of practitioners (75%) find AI systems move to production within defined governance processes.

People are manually enforcing governance—case by case, review by review, exception by exception. Only 25% of practitioners report that AI systems reach production either by working around governance constraints or because governance requirements slow or complicate delivery.

The AI governance reality gap stems from policies designed for a different era.

The question this raises isn't whether governance needs to exist—it's whether it can operate at the speed of the work it's governing. When 74% of practitioners describe governance as manual, the system is asking humans to scale a function that needs infrastructure.

Closing this gap means governance that's purpose-built for how AI actually operates: distributed, fast-moving, and embedded in the tools and workflows where practitioners do the work. ■

> “
>
> When governance is layered on top of how people actually work, instead of embedded into the systems they use, it inevitably breaks down at scale.

BRIAN ALVEY,
CEO, WORDPRESS VIP

Speed vs. safety in AI governance

The core tension in AI governance

The speed-versus-safety tension in AI governance plays out every day, in every sprint, across virtually every organization we surveyed. 99% of leaders say AI initiatives experience delays due to governance or risk-related reviews. Only 1% of leaders say governance "never" causes delays. Among practitioners, 97% report governance-related delays in their AI projects.

Governance friction isn't an edge case. It's the default, and that default exposes organizations to risk.

We asked leaders to indicate which governance challenge poses the greatest risk to their organization as it expands. Their answers were split across the board, with fragmented governance across functions ranking as the #1 risk to their organization.

The AI governance failure leaders believe poses the greatest risk to their organization:

- **31%** Fragmented governance across functions (IT, Security, Data, Product)
- **21%** Lack of visibility into where AI systems exist and how they're used
- **21%** Inconsistent enforcement of governance policies across tools and teams
- **17%** Reliance on manual reviews and approvals
- **10%** Unclear ownership and accountability for AI systems

> “
>
> AI moves fast, but governance models designed for quarterly review cycles don't.

BRIAN ALVEY,
CEO, WORDPRESS VIP

Proximity to governance work impacts perception

79% of practitioners agree that AI governance frequently forces tradeoffs between delivery speed and risk management. 53% of leaders agree that their organizations are forced to trade speed for safety when deploying AI systems.

That 26-percentage-point gap means the two groups are living fundamentally different realities. Leaders are split, while the majority of practitioners agree on what they're seeing in their day-to-day lives. The closer you are to the work, the more you feel the friction.

Leaders see a system that mostly balances

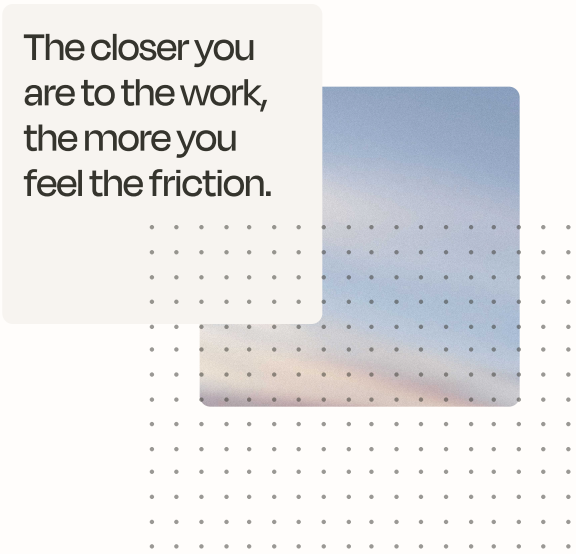

competing priorities. Practitioners see a system that routinely forces them to choose between moving fast and staying compliant.

At the frequency level, the experience is widespread but differently weighted. 49% of leaders say governance delays happen frequently or almost always. 45% of practitioners say the same, but practitioners are 2.2x more likely to say delays happen "almost always"—another signal that the people closest to the work absorb the sharpest friction.

Two realities: compliance vs. improvisation

When governance delays occur, leaders and practitioners describe starkly different responses—revealing not just a perception gap, but two different operating realities.

How leaders say teams navigate governance delays in AI projects:

- **44%** Teams wait for approval before proceeding
- **27%** Reduce scope to avoid governance complexity
- **27%** Exceptions or temporary workarounds are approved
- **2%** Projects are paused or abandoned

How practitioners respond when governance slows down their AI projects:

- **45%** Put exceptions or temporary controls in place
- **31%** Reduce scope to avoid governance complexity
- **19%** Pause work until it meets requirements
- **5%** Delivery continues with unresolved governance risk

Leaders see compliance—teams following the process, waiting their turn. Practitioners see improvisation—workarounds, scope cuts, risk absorption.

What leaders describe as "governance working" may be practitioners finding ways around it. When 45% of practitioners say exceptions and temporary controls are the default, governance has become an obstacle teams route around.

Ad hoc workarounds and temporary exceptions are how teams govern when the system doesn't. The improvisation isn't a failure of discipline—it's evidence governance lives in the wrong place. When controls live outside the tools where decisions happen, teams build their own informal controls inside them.

MONTHLY HOURS LOST
TO GOVERNANCE DELAYS

When asked how many hours governance delays add to their monthly workload, almost no practitioners said zero.

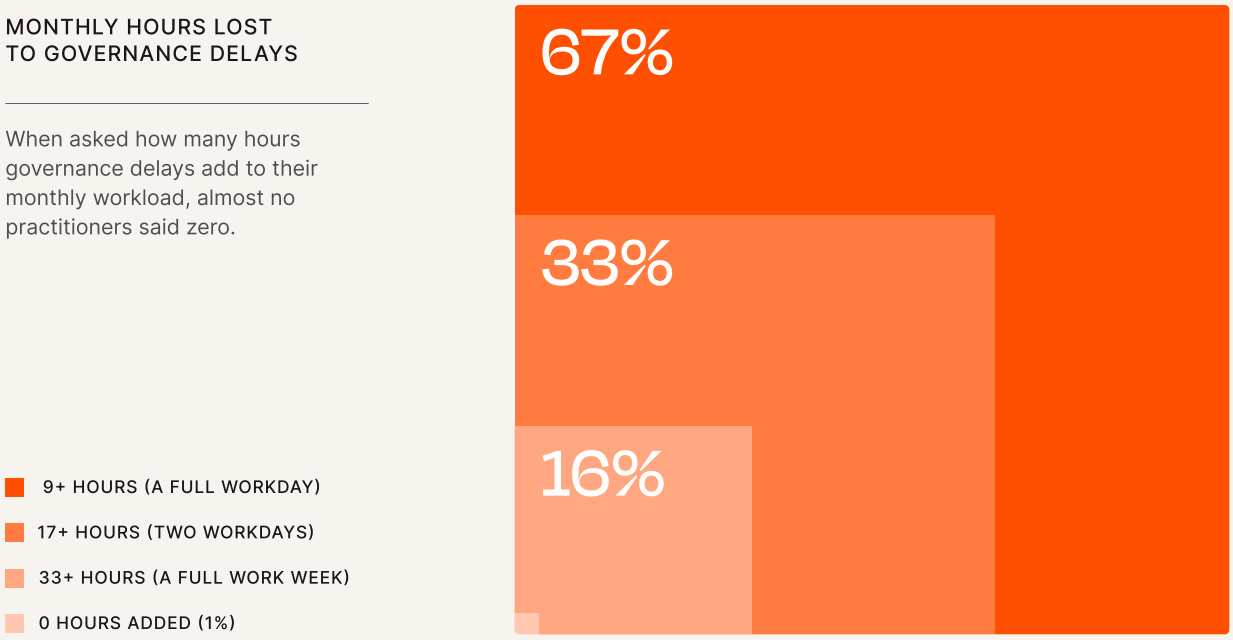

The time tax

The cost of governance friction isn't abstract—it shows up on practitioners' calendars.

When asked how many hours per month governance-related delays add to their work:

- **67%** 9+ hours per month (more than a full workday)
- **33%** 17+ hours per month (two full workdays)
- **16%** 33+ hours per month (a full work week)
- **1%** 0 hours added

67% of practitioners lose 9+ hours per month to governance delays—more than a full workday, every month.

Only 1% of practitioners report zero added hours. For the other 99%, governance delays are a measurable line item—a structural tax on delivery capacity that scales with every new AI initiative an organization launches.

Safety wins—but at what cost?

Despite the friction, safety wins when organizations force the choice—and it wins independently in both groups. 61% of leaders

say they'd prefer slower AI adoption with stronger governance controls over faster adoption with increased risk. 59% of practitioners say preventing governance and risk issues up front works better than moving fast and addressing risks as they arise.

That alignment holds when theory becomes practice. 81% of practitioners say that when governance and delivery priorities conflict, risk reduction wins—even at the cost of delivery speed. 71% of leaders say preventing worst-case risk determines the final call.

The tradeoff between speed and safety isn't inevitable. It's a symptom of governance that operates as a checkpoint rather than infrastructure. Both groups want governance—they just want a version that doesn't force a choice between doing the work and doing it safely.

What creates the tension

When we asked leaders which aspect of governance most often creates tension with delivery speed, no single factor dominated their answers.

What leaders think creates the most friction between AI governance and delivery speed:

- **28%** Difficulty enforcing policies consistently across tools and teams
- **24%** Manual approvals or reviews required before deployment
- **23%** Late discovery of data, security, or compliance issues
- **14%** Unclear ownership or escalation paths for governance decisions
- **11%** Unique complications given enterprise complexity

As you can see, leaders don't agree. The fragmentation is the finding—there's no single governance bottleneck to fix. The top three sources are within 5 percentage points of each other. The friction is structural and embedded across enforcement, approvals, discovery, and ownership.

That fragmentation gets worse as organizations add AI providers, surfaces, and tooling. When there's no interoperability between how you govern one AI surface and the next, every new addition multiplies the friction. ■

Why AI governance breaks at scale

Why AI governance breaks at scale and what starts to work

96% of leaders agree that embedding AI governance into tools and workflows helps teams move faster with confidence. But leaders also ranked embedded policy enforcement last among the capabilities that have helped them scale AI securely thus far—chosen by only 9% of leaders.

The consensus on where governance needs to go is near-universal. The progress toward getting there is concentrated in human-layer solutions—ownership, guidance, documentation—but the technical infrastructure everyone agrees on remains barely built.

This is unsurprising, 69% of leaders agree that governance in their organization depends too heavily on manual processes to scale effectively. That manual dependency shows up in how governance gaps get resolved in practice.

HOW PRACTITIONERS RESOLVE GOVERNANCE GAPS

When asked how governance gaps are most often resolved in practice, over half of practitioners pointed to ad hoc discussions or meetings.

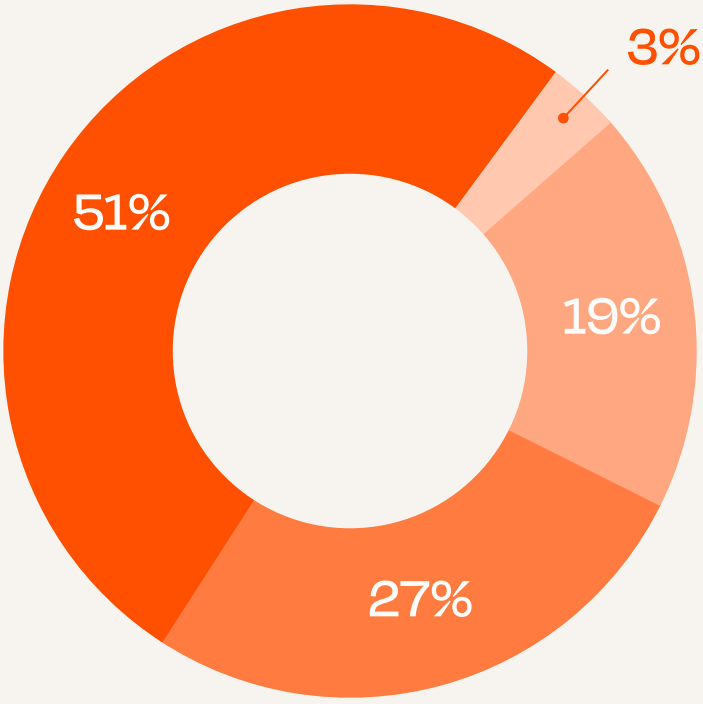

> “
>
> If teams are routing around your process, your process is too slow or too disconnected from reality.

BRIAN ALVEY,
CEO, WORDPRESS VIP

How practitioners say organizations most often resolve AI governance gaps:

- **51%** Through ad hoc discussions or meetings
- **27%** By applying temporary exceptions or overrides
- **19%** By deferring or limiting the AI use case
- **3%** By proceeding with unresolved governance risk

78% of practitioners say governance gaps get resolved through either ad hoc meetings or temporary workarounds.

19% defer the use case entirely. Improvisation becomes the de facto governance model—and manual processes create manual workarounds, which in turn require more manual oversight.

It's a loop that scales in exactly the wrong direction.

Why governance breaks down at scale

Leaders and practitioners were each asked to rank the top three structural reasons AI governance breaks down at scale. Their answers overlap—but the priority order reveals what's actually failing, and for whom.

For leaders, the bottleneck is mechanical. For practitioners, the bottleneck is structural. [See the full breakdown in TABLE A on page 14.]

The people writing the policies rank speed as the constraint. The people implementing them rank comprehensibility. Both groups flag manual processes, policy clarity, and ownership gaps. But leaders prioritize the approval queue, while practitioners prioritize the translation gap.

The data reveals where governance is least developed—and leaders and practitioners diverge sharply here as well. [See the full breakdown in TABLE B on page 15.]

Leaders are still writing the rules. Practitioners are still trying to figure out who enforces them.

Leaders rank ownership dead last at 8%. Practitioners rank it first at 30%. That's a 22-percentage-point gap on a question about what needs the most work—and it suggests leaders and practitioners are building governance from opposite ends.

TABLE A

Structural reasons AI governance breaks down at scale

Ranking	Technical leaders owning the strategy	Technical practitioners on the frontlines
1	Manual, approval-driven processes	Policies that are difficult to translate into actual controls
2	Policies are unclear, incomplete, or difficult to operationalize	Unclear ownership or escalation paths
3	Unclear ownership and accountability across teams	Manual approvals and exception handling

The governance challenge practitioners believe creates the most day-to-day friction:

- **36%** Waiting on reviews, approvals, or escalations
- **23%** Translating governance policies into technical controls
- **20%** Applying governance consistently across different tools
- **11%** Determining who owns an AI system or decision
- **10%** Maintaining auditability and documentation

Practitioners identify waiting on reviews and approvals as the single biggest friction point for AI governance—consistent with the manual dependency leaders acknowledge.

Despite these numbers, 69% of leaders agree their organization employs a well-positioned AI governance approach to scale over the next 12–24 months. They've isolated their pain points and, more importantly, their successes.

Leaders have a vision for what the future looks like, even if it's not the reality at present.

TABLE B

The least mature or operationalized aspects of governance programs today		
Ranking	Technical leaders owning the strategy	Technical practitioners on the frontlines
1	Definition of AI governance policies and standards	Ownership and accountability for AI systems
2	Monitoring, logging, and auditability of AI activity	Policy enforcement embedded in tools and workflows
3	Policy enforcement embedded in tools and workflows	Monitoring, logging, and auditability

Where governance works

Not everything is broken. Both groups identify specific governance capabilities that have delivered measurable value. [See TABLE C on following page.]

Practitioners are specific about what would close that gap. [See TABLE D on following page.]

What's working today is fundamentally human: ownership, guidance, and documentation tied to real workflows. The vast majority of leaders agree that embedded governance enables teams to move faster with confidence. Yet, only 9% of leaders say embedded enforcement has actually contributed to scaling safely so far.

The gap between consensus and implementation is the clearest signal in this data about where investment needs to go next.

TABLE C

The governance capabilities that have contributed most to scaling AI successfully and safely

Ranking	Technical leaders owning the strategy	Technical practitioners on the frontlines
1	Clear ownership and accountability	Well-documented policies tied to real workflows
2	Structured guidance and enablement	Governance controls embedded in tools
3	Centralized visibility	Clear ownership and escalation paths

TABLE D

Which capability would most reduce governance friction while keeping AI safe

Ranking	Technical practitioners on the frontlines
1	Guardrails and permissions embedded directly into workflows
2	Automated monitoring and assurance at build- and run-time
3	Real-time visibility into AI systems, data flows, and ownership

How AI governance must evolve

The 94% consensus

94% of leaders agree that AI governance must shift from a set of policies and approvals to a continuously operating system embedded in how AI is built and run.

Mirroring leaders, 94% of practitioners also agree that future AI governance should be enforced primarily through systems and guardrails, rather than through manual reviews and approvals.

The number happens to be identical, and the conviction runs independently in both groups. Practitioners go even further: 91% believe governance should guide how systems are designed, rather than reviewing them after deployment. The demand is for governance that shows up at design time—built into architecture decisions, not applied as a gate at the end.

Governance as a bridge, not a wall

86% of leaders say governance should act as a bridge—enabling AI use while dynamically

MOST CRITICAL STRATEGIC SHIFT FOR GOVERNANCE

When asked which shift is most critical for governance to scale, leaders ranked the move from centralized oversight to federated ownership first.

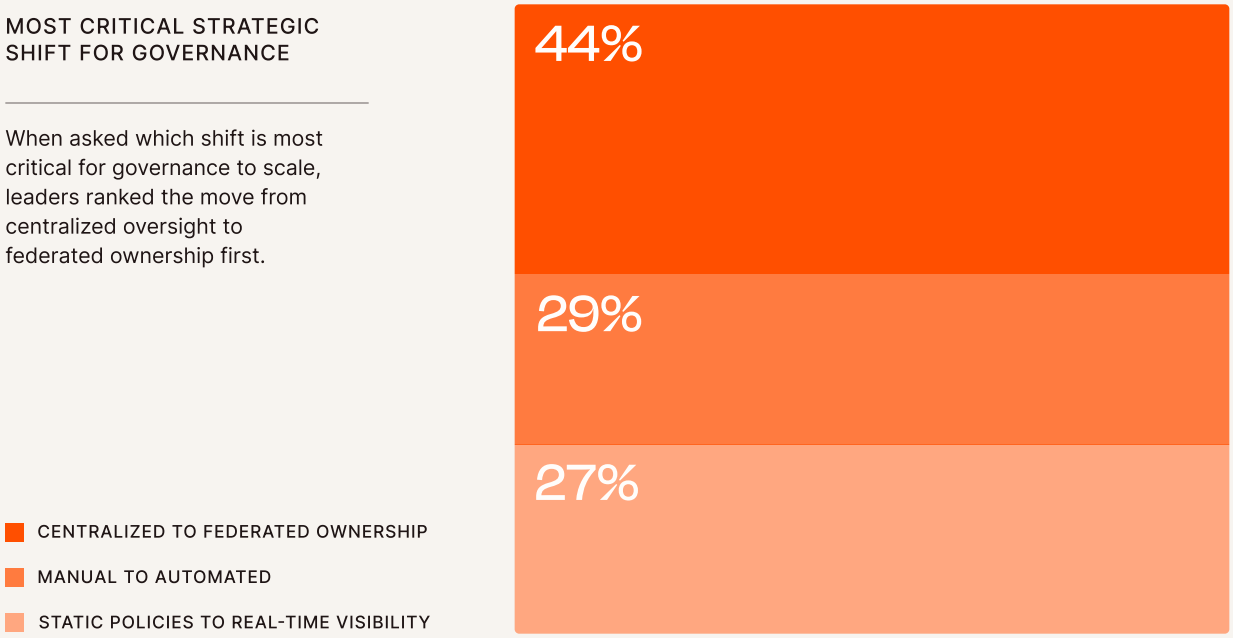

managing risk. 84% of practitioners agree, believing governance should guide safe actions rather than prevent unsafe ones.

83% of leaders say governance should primarily function as a strategic capability that accelerates AI-driven advantage. Only 17% of leaders see it as a control layer that restricts unsafe behavior. That 83% figure carries weight—it reframes how organizations approach investment, tooling, and operating model decisions.

> “
>
> In the AI era, trust and security is the currency which buys the agency to drive transformation and value.

RAVI MALIK, SVP & GLOBAL CHIEF INFORMATION OFFICER, BOX

box

The end of the centralized governance model

We asked leaders which shift is most critical for governance to scale over the next two to three years, and their top priority is the shift from centralized oversight to federated ownership.

What leaders think are the biggest strategic shifts organizations need to make when considering AI governance:

- **44%** From centralized oversight to federated ownership with guardrails
- **29%** From manual reviews to automated enforcement
- **27%** From static policies to real-time visibility and assurance

This mirrors what we see on the practitioner end. We asked practitioners which governance operating model will scale best in practice.

The operating model practitioners believe scales AI governance best:

- **39%** A hybrid model that adapts based on risk level
- **32%** Federated teams operating within shared guardrails
- **16%** Fully decentralized teams with minimal oversight
- **13%** Central teams governing all AI decisions

71% of practitioners favor federated or hybrid governance. Only 13% of practitioners want a single central team making every call.

WHICH GOVERNANCE MODEL SCALES BEST

When asked which operating model scales AI governance best in practice, practitioners overwhelmingly favored hybrid or federated approaches over centralized control.

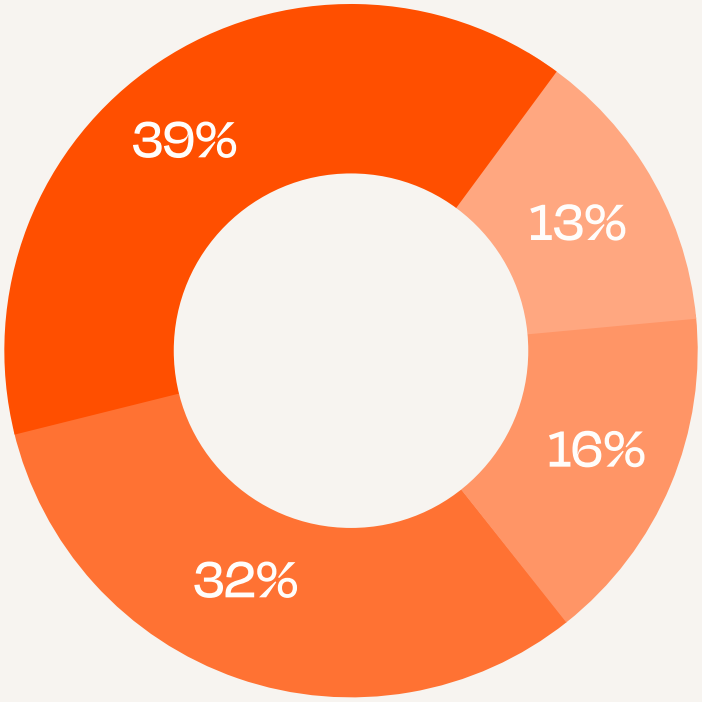

Practitioners want a model flexible enough to match the risk level of each use case. In both cases, the centralized governance model is losing ground.

In practice, hybrid governance means organizational rules that every team inherits—with the flexibility for team leads to tighten policies for their context. Org-level boundaries. Team-level ownership. The teams closest to the work make the decisions, within a framework they didn't have to build from scratch.

Where governance investment is going

The spending data reflects an industry that has already made this mental shift.

91% of leaders expect AI governance spending to increase over the next 24 months—43% significantly, 48% modestly. 0% of leaders expect a decrease. Not one.

91% of leaders say their organization currently spends $100K+ per year on AI governance—including tooling, infrastructure, staffing, and external services. 71% of leaders spend $500K+, and 59% of leaders say their organization currently spends $1M+.

The investment base is already substantial, and nearly every organization plans to grow it. But controls tied to a single AI provider lose value every time the model landscape shifts.

Interoperability—governance that holds across providers, tools, and AI surfaces—is what distinguishes a durable investment from a disposable one.

Why governance is the future of scaling AI

96% of leaders agree that effective governance builds trust with employees, leadership, customers, and regulators by making AI usage visible, controlled, and provably safe—without slowing teams down.

71% of practitioners expect their governance-related workload to increase over the next two to three years.

The teams closest to implementation are preparing for more governance work. And what they want from that work is consistent across every question in this survey: embedded guardrails, automated enforcement, and real-time visibility.

How enterprises will adapt their governance models in 2026

We asked leaders: "What is the single most important change your organization must make for AI governance to support AI scale rather than constrain it?"

We asked practitioners: "In your own words, what would meaningfully improve AI governance for the teams building and deploying AI systems?"

Both groups independently described the same model: governance embedded in the tools, clear ownership, and automated enforcement.

But building that model for a single AI provider or surface defeats the purpose. AI providers will continue to shift through pricing changes, capability leaps, and compliance requirements.

The controls, the policies, and the enforcement logic all need to work seamlessly, no matter which AI model you're using. →

What must your organization change for AI governance to enable scale, rather than constrain it?

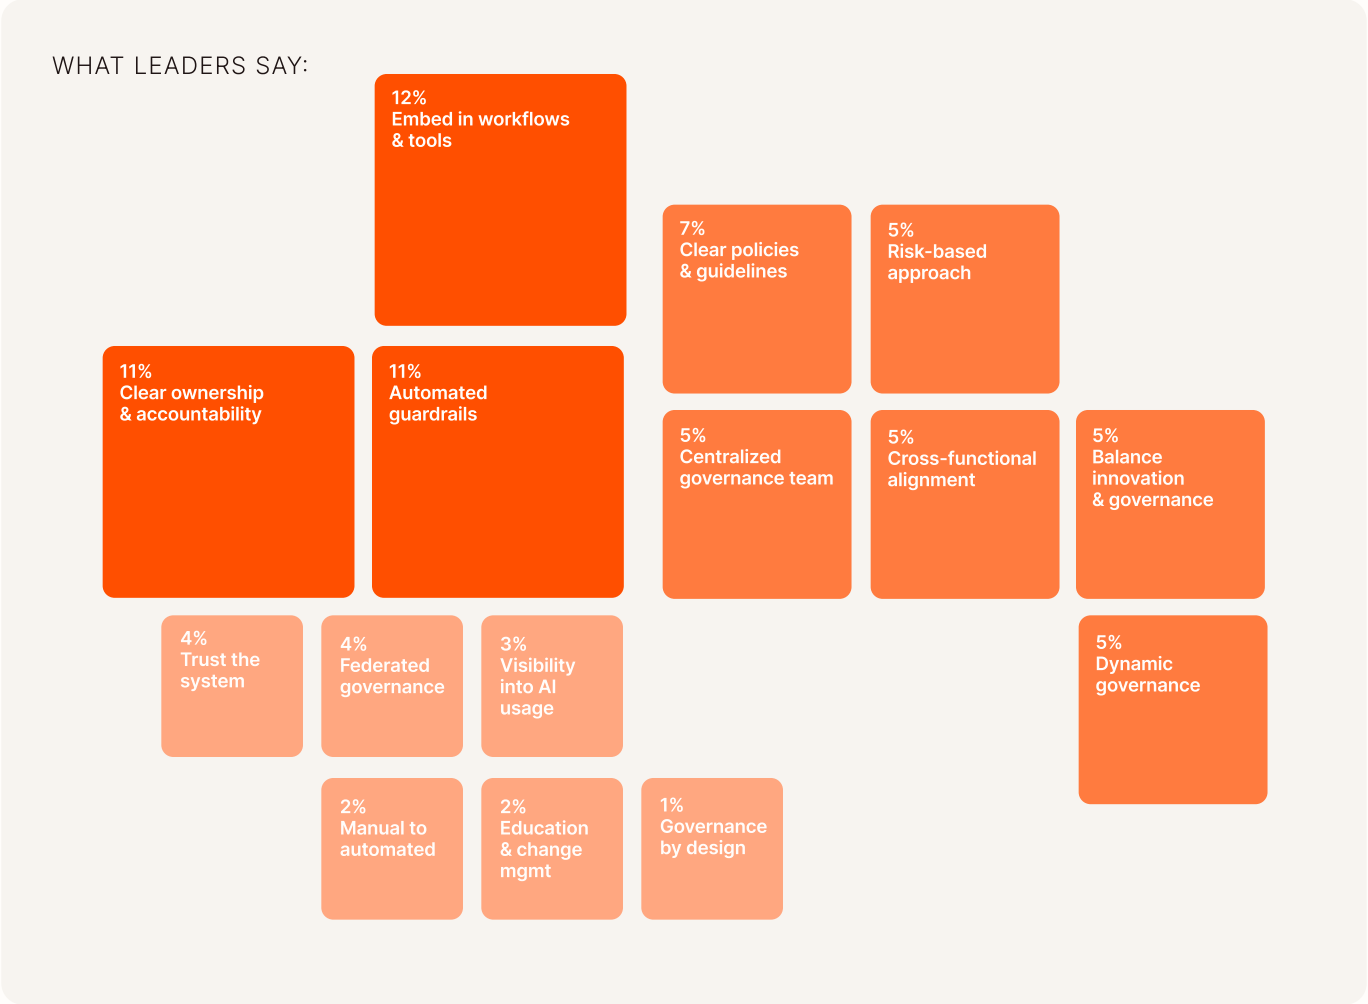

The top themes from leaders center on embedding governance into workflows and tools, automating guardrails, and establishing clear ownership and accountability. Secondary themes include clear policies, risk-based approaches, centralized governance functions, cross-functional alignment, and dynamic governance. Smaller clusters point to trusting the system over manual approval, federated models, visibility into AI usage, and governance from the design stage.

What would meaningfully improve AI governance for the teams building and deploying AI systems?

Practitioners' responses are even more concentrated. Clear guidelines and documentation dominate, followed by embedding governance into workflows and tools, and clear ownership and accountability. Real-time monitoring and auditing, automated compliance, transparency and trust, and centralized governance teams form a strong secondary tier. Smaller clusters include ethical guidelines, risk-based approaches, translating policies into technical controls, governance education, governance by design, human oversight, reducing friction, federated governance, and cross-functional collaboration.

4 principles for scaling AI governance

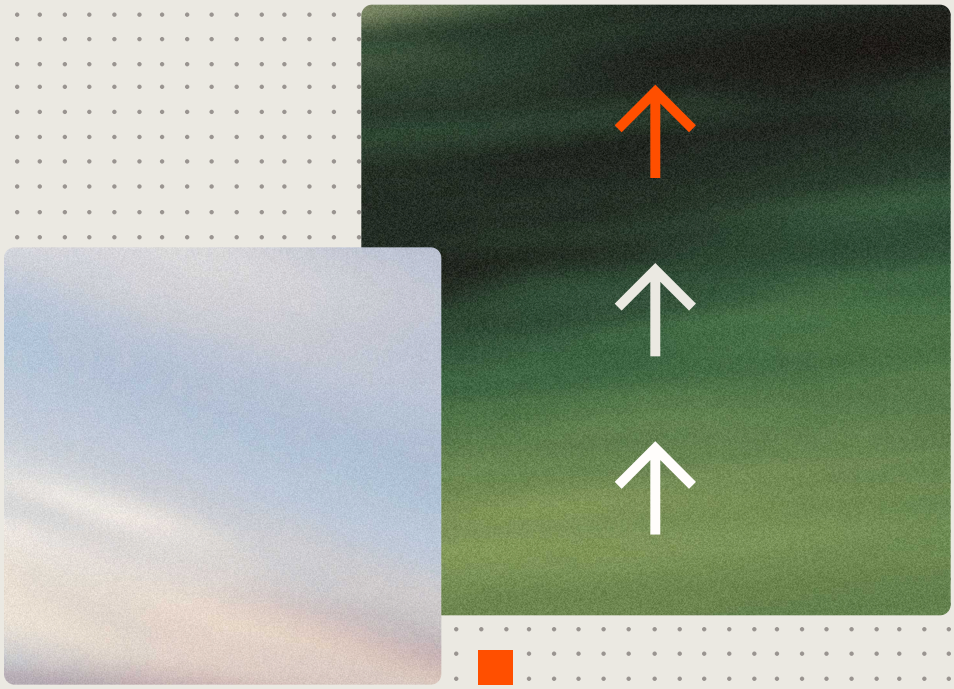

1. Embed governance where work happens

Leaders and practitioners both want the same thing: guardrails that live inside the tools where work actually happens, running automatically so teams don't have to stop and think about compliance. Today, 74% of practitioners say governance decisions are resolved through manual reviews. 69% of leaders agree governance depends too heavily on manual processes to scale.

The next generation of governance eliminates that handoff: guardrails and permission boundaries replace approval queues, starting with the highest-volume friction points—data handling policies, app-level access controls, and AI model usage boundaries.

2. Distribute ownership with guardrails

Both groups ranked clear ownership in the top three. Fortunately, federated governance allows for centralized control and team-level autonomy at the same time, resolving the ownership issue. The federated governance model defines boundaries at the organizational level—clarifying which apps, models, and data policies to use—then allows teams closest to the work to make their own decisions within those boundaries.

Set the rules once. Let everyone build within them.

> “
>
> The companies that win in the AI era will be the ones who find the right balance of confidence and speed, not the ones who move recklessly fast.
>
> Governance is an advantage when it reduces internal friction, speeds up safe deployment, builds trust, and enables experimentation at scale.

BRIAN ALVEY,
CEO, WORDPRESS VIP

3. Make governance visible

41% of leaders report low or no confidence in real-time AI visibility. Governance that can be observed earns trust—and governance that earns trust gets adopted. Audit trails, auto-generated documentation, and monitoring that streams into existing security tools make the system provable. When someone asks, “What’s running, who owns it, and where does the data go?" teams shouldn't need to spend a week pulling logs. If they can't immediately find an answer, there's still work to do around visibility.

4. Build for interoperability

Most enterprise AI today runs on a single provider family—and that provider's pricing, policies, and capabilities will change. Organizations that build governance tied to one model are building governance with an expiration date. Interoperability means your app connections, workflow logic, and governance controls travel with you across AI surfaces—regardless of which model sits underneath. The controls persist. The permissions persist. The logic stays. The model will always change, but the governance shouldn't have to.

The organizations that build on interoperability first won't just be more secure. They'll move faster because governance will compound their speed instead of constraining it. The conviction is universal. The architecture is clear. The organizations that act on it now will define how enterprise AI scales safely for the next decade.

Governance doesn't operate in isolation

It's one of four pillars—alongside leadership, talent and culture, and tools—that multiply to produce measurable AI impact. Organizations that advance all four together progress faster than those that try to fix governance alone.

For a framework that connects governance maturity to the broader AI transformation equation, read our AI Transformation Framework.

Ready to join the AI Leaders Lab?

AI transformation takes more than handing out ChatGPT licenses—it requires rethinking how people, processes, and technology work together.

That's why we created a monthly **AI Leaders Lab:** an invite-only session designed for leaders like you who are shaping the future of work.

This 60-minute conversation is small by design. We hand-picked a select group of AI leaders so every participant can contribute, exchange ideas, and walk away with new perspectives. If you can't make it, let us know—we'll open your seat to another leader.

What you'll experience:

- A roundtable discussion tailored to the group's priorities
- Peer-to-peer connection through interactive breakouts
- Fresh insights, data, and real-world patterns from Zapier's AI team
- Practical takeaways to accelerate AI adoption immediately

Spark viral AI adoption in your org, build habits and infrastructure that endure beyond the hype, and shift AI from a "project" into company-wide culture. Seats are limited—reserve yours now.

Talk to an expert today to assess your AI maturity and start your AI staffing journey.

Survey methodology

PUBLICATION DATE: APRIL 10, 2026 | TOTAL RESPONDENTS (N): 200

This report is based on a survey of 200 enterprise professionals directly responsible for AI governance, deployment, or operations. The survey was fielded by NewtonX in Q1 2026.

Respondent segments

Respondents were divided into two segments of equal size:

- Leaders (n=100): C-level executives (CIO, CTO, CDO, CISO), VPs, and Directors.
- Practitioners (n=100): Managers, Team Leads, and Senior Individual Contributors (technical).

Qualification criteria

All respondents met the following requirements:

Function: IT Operations/Infrastructure, Security/Compliance/Risk, Data/Analytics, or Engineering/Platform

Company size: 1,000+ employees

Geography: United States, Canada, or Europe

AI involvement: Directly builds, operates, or maintains AI systems; oversees teams that do; or implements or enforces governance controls on AI systems

AI maturity (personal): Strategic, Applied, or Basic usage of AI (respondents with limited or no AI usage were disqualified)

Organizational AI usage: Mission-critical and enterprise-wide, or limited but expanding rapidly (organizations with only experimental pilots or no AI usage were disqualified)

Governance responsibility: Primary decision-maker, shared responsibility, or advisor/influencer on governance decisions

TRUSTED BY 3.4 MILLION COMPANIES

NVIDIA airbnb Meta LOWE'S Disney SAMSUNG

mastercard. SIEMENS Dropbox hp CURSOR okta