credo ai

Playbook

The ROI of AI Governance: A 2026 Executive Playbook

- **How to Move from Talk to Trust—and Scale AI with Confidence**

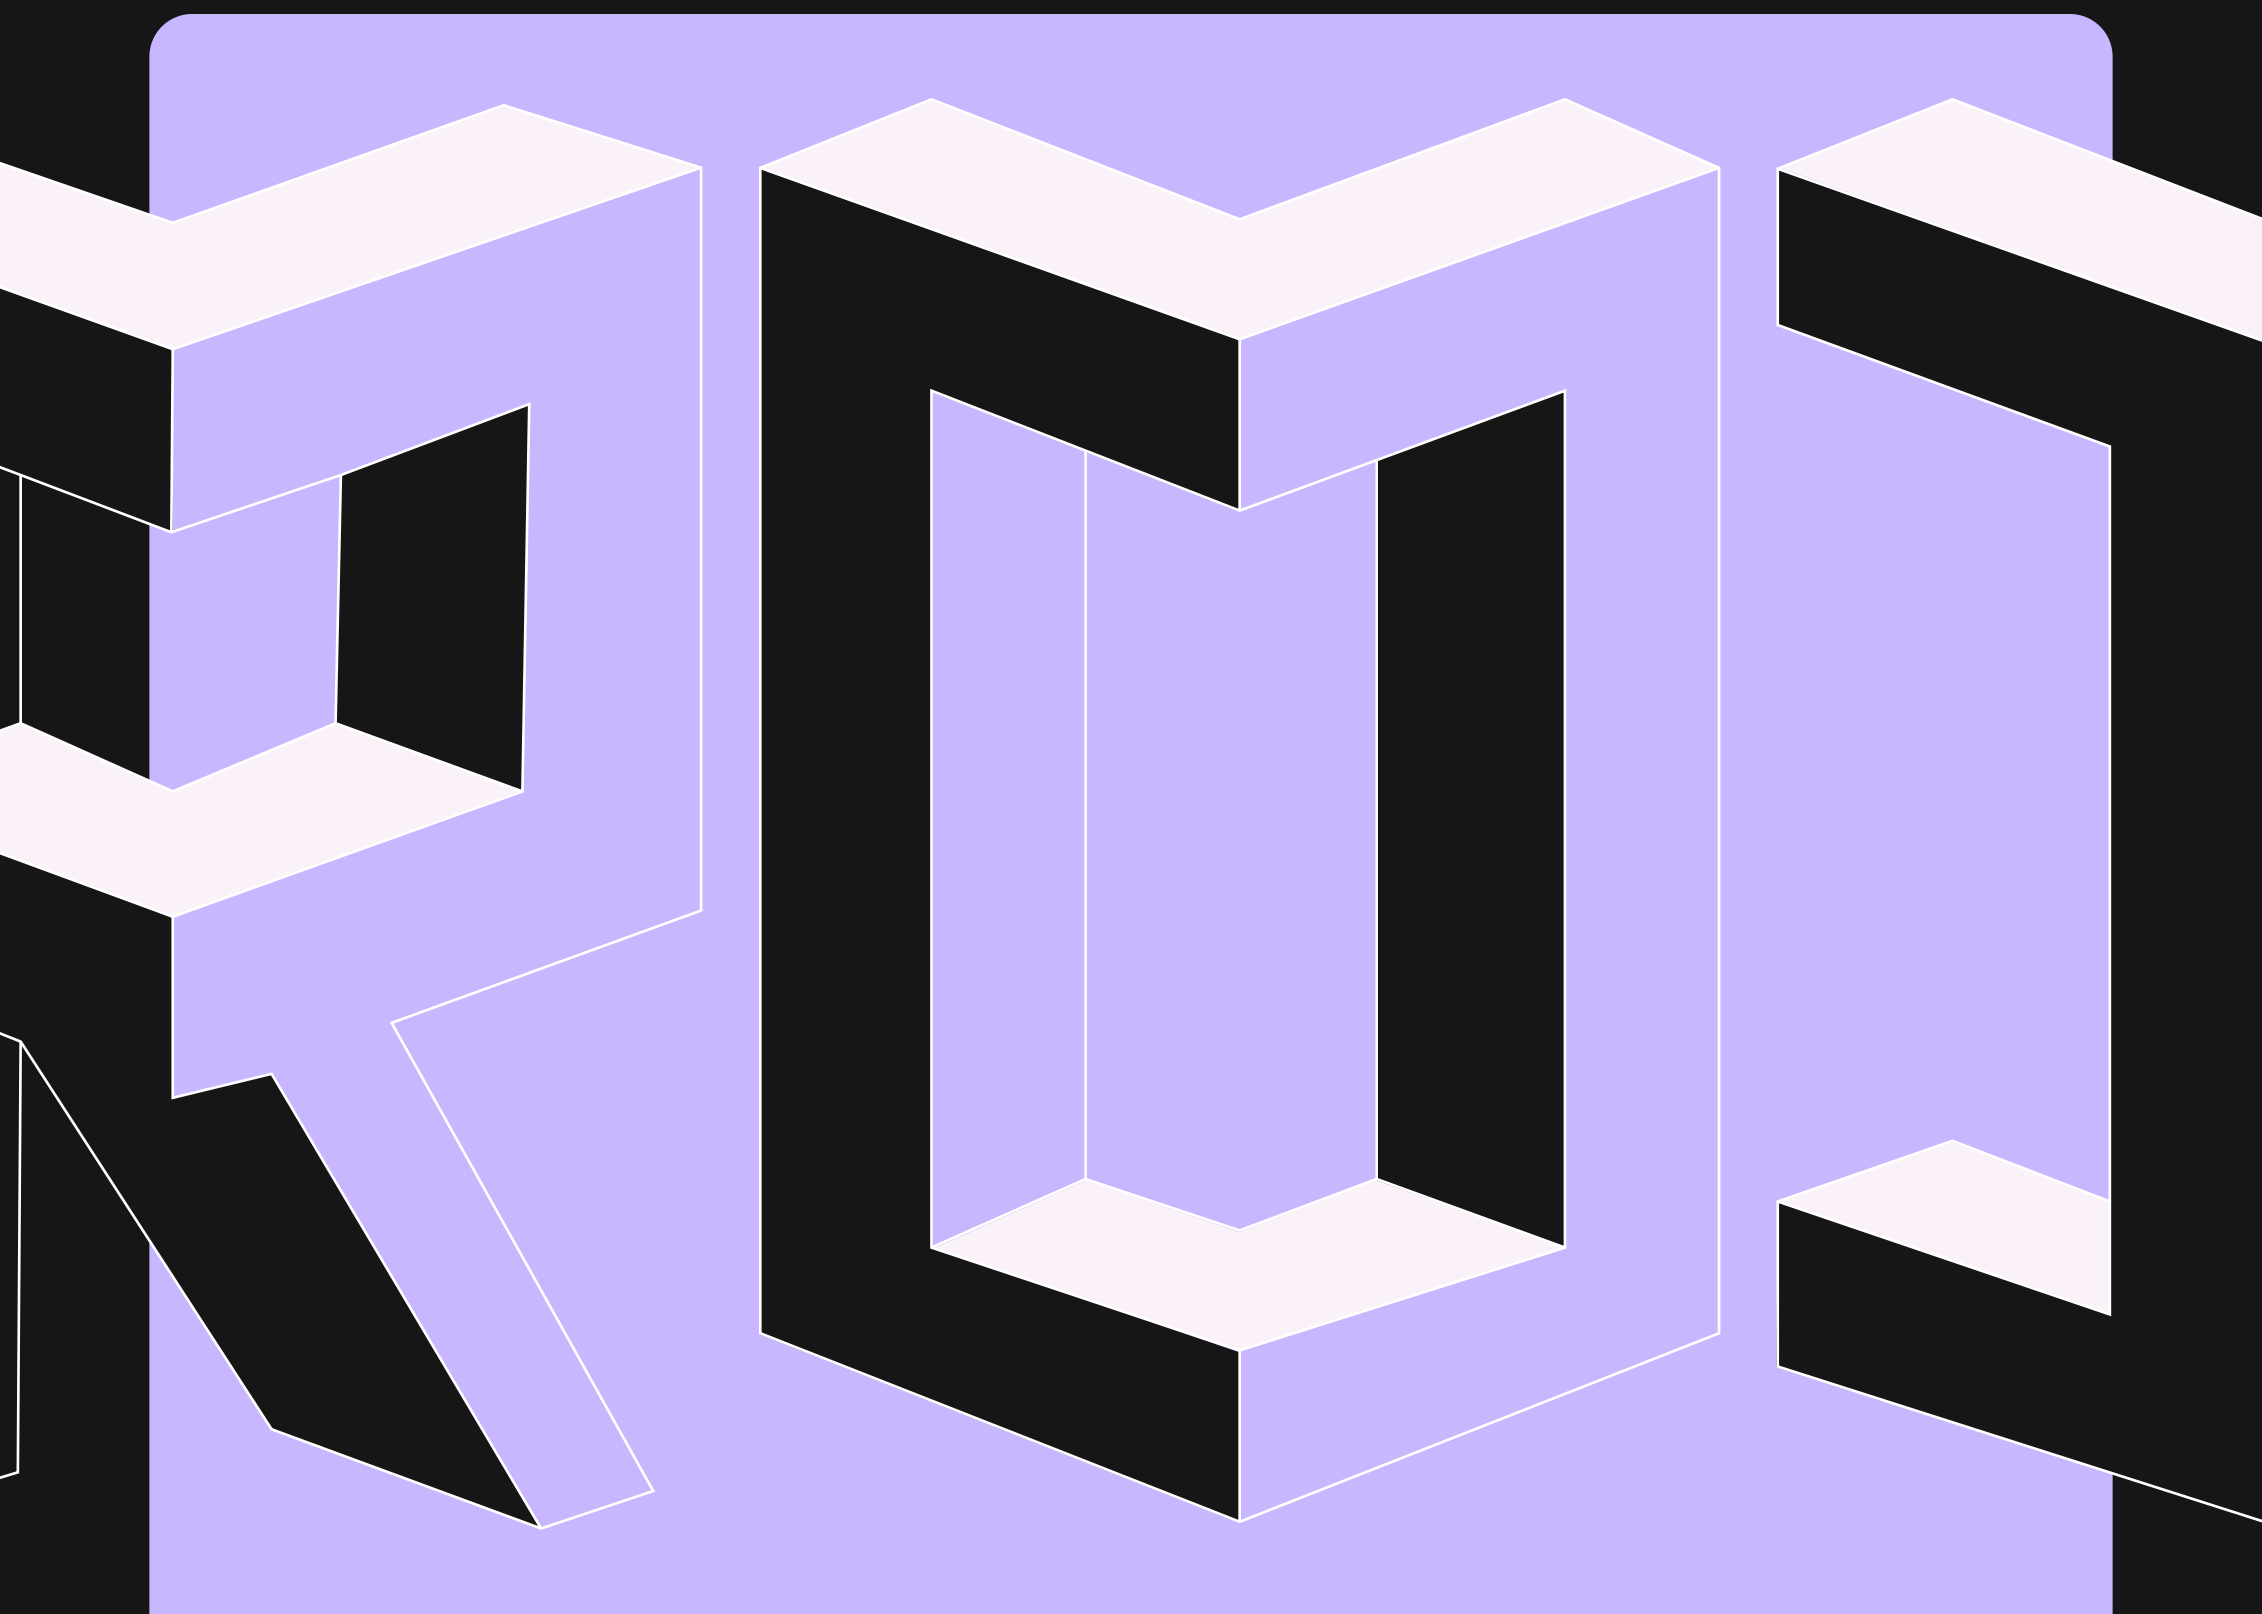

Table of contents

Table of contents

AI has become your biggest unmanaged super-user.

It is spinning up code, making credit suggestions, drafting legal language, talking to customers, and soon dispatching agents that chain actions across systems—often with more effective permissions than your human staff and far less scrutiny. Meanwhile, the “governance layer” for all this power is still mostly tribal knowledge, scattered spreadsheets, and workflow hacks.

The result is a widening governance implementation gap: AI capabilities are compounding exponentially while control surfaces grow linearly at best. That mismatch is already showing up as model-driven poor outputs, compliance near-misses, opaque agent behavior, and risk costs that don’t show on the balance sheet until something breaks. This playbook is about treating AI governance like production-grade infrastructure—not a slideware afterthought—before that gap becomes your next systemic incident.

> Navrina Singh, Founder and CEO of Credo AI, puts it plainly: "Trust and accountability should just be weaved into how we're building artificial intelligence. Governance is not a brake on innovation—it's the fuel that lets you scale it."

This playbook is for leaders who recognize that AI governance isn't about compliance theater. It's about competitive advantage. Organizations that embed governance early, move faster. They avoid costly incidents, build customer trust, and close more deals.

This guide will walk you through how to make the business case for AI governance, build internal buy-in, assess where you are today, design the core components of a governance program, and measure success in ways that matter to your board and your bottom line.

The 2026 Inflection Point: Why Now

AI is no longer experimental, it's business-critical. As per a Global AI Survey[1], AI adoption is broad and budgets are rising, but scaling remains limited and negative outcomes are mounting—from cyberattacks on unguided AI systems and biased outputs leading to lawsuits, to AI-generated content facing copyright and IP scrutiny—all underscoring the need for stronger governance.

On average, organizations now manage about four AI-related risks, up from two in 2022. And roughly half report they've already experienced at least one negative consequence[1].

At the same time, regulation is tightening. The EU AI Act is real. The US is hardening its approach to national AI security. Data sovereigns are multiplying. Board expectations are rising. And buyers—customers, partners, investors—are asking one simple question:

- **"How do you know your AI is trustworthy?"**

  For most organizations, the honest answer is:

- **"We're not sure."**

This playbook is designed to help AI governance professionals move from scattered, risk-only conversations about AI to a repeatable, value-driven governance program that unlocks trusted, scalable AI adoption. In 2026, organizations that have embedded governance will move faster, identify and reduce risk, and win more business. Those that haven't will face delays, incidents, and shrinking opportunities.

What AI Governance Is (and Isn't)

AI governance is:

- ✓ An operating model that embeds oversight, control, and accountability into the AI lifecycle.
- ✓ AI governance is the framework of rules, policies, standards, and practices that ensures artificial intelligence systems are procured , developed and used securely, reliably, transparently and in a compliant manner.
- ✓ The infrastructure that lets you scale AI without scaling risk.
- ✓ A competitive advantage that builds customer trust and investor confidence.
- ✓ A bridge between technologists, risk teams, compliance, and the business.

AI governance is not:

- ✗ A compliance checkbox
- ✗ A tool for slowing down innovation
- ✗ Something you outsource to a legal team and forget
- ✗ A list of principles nobody reads

In the analyst community, governance platforms are now recognized as a category. Gartner defines AI governance platforms as tools that give organizations central oversight of AI via inventory, risk management, policy enforcement, and continuous monitoring.

Forrester finds that the state of an organization's AI governance solution is one of the strongest indicators of its readiness to scale trusted AI beyond a handful of use cases.

The Market Noise: AI Operations, Security, and GRC

There's a lot of noise in the market. You'll hear about "AI operations" (MLOps), "AI security," "GRC for AI," and "responsible AI platforms." It's worth being clear about what's different.

AI Operations (MLOps) focuses on model training, deployment, and performance. It's about getting models to work.

AI Security focuses on threats: adversarial attacks, poisoned data, model theft. It's about protecting models from harm.

GRC (Governance, Risk, Compliance) tools are often adapted from privacy and security frameworks. They were built for static systems and at procedural level. It's hard to adapt them to dynamic AI systems requiring scientific measurements and technical controls across the value chain of datasets, models, agents , and applications.

AI Governance connects all of these. It defines what should happen, who decides, how we verify, and how we respond when there is an incident . It spans the full lifecycle of an AI system: from procuring , discovery, design , development , deployment through retirement.

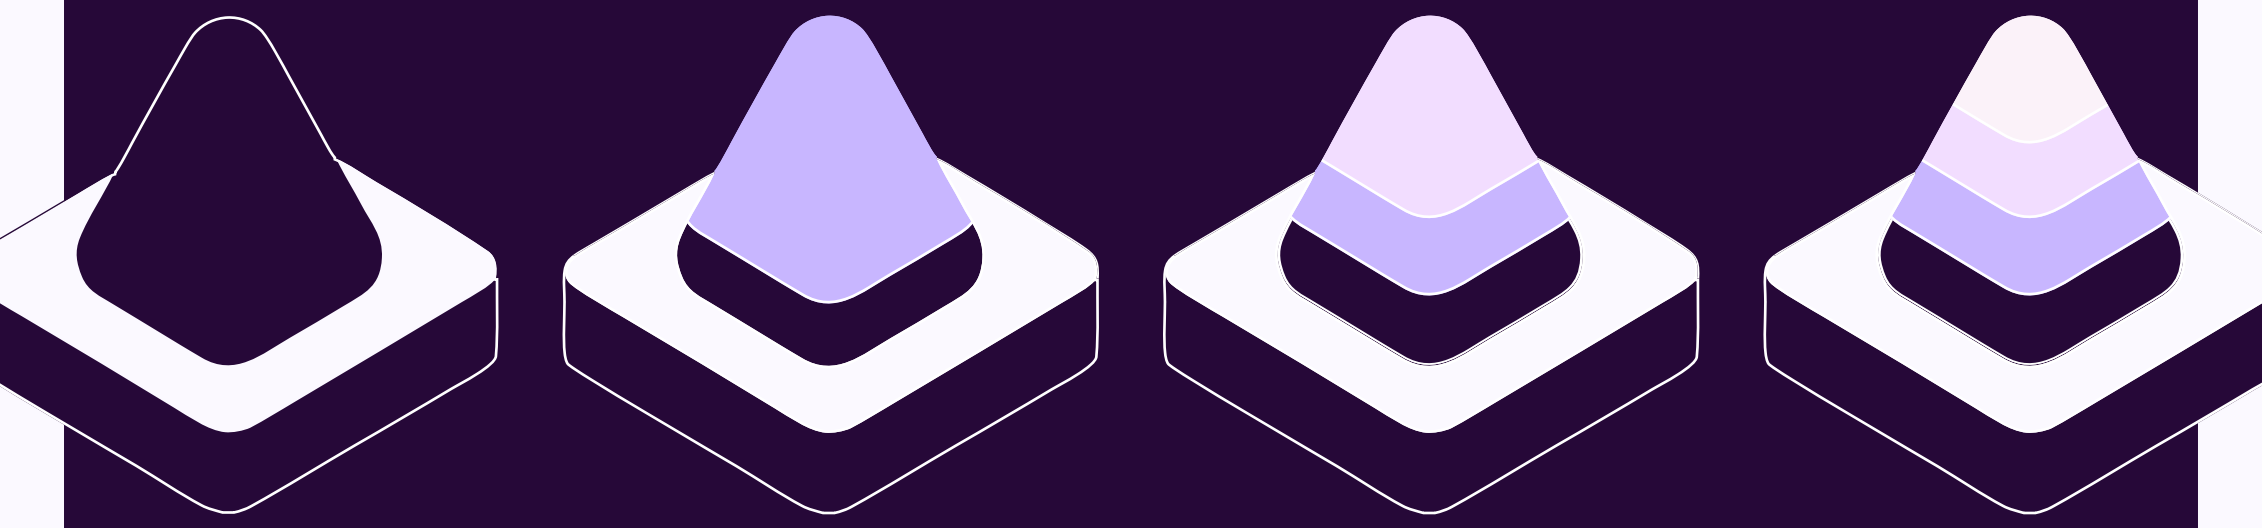

Two Pathways Through This Playbook

You may be reading this because:

1·

You're building the case for governance

You see the need but don't yet have a budget, executive alignment, or a clear roadmap. Start with Section 1 (The Business Case) and Section 2 (Building Buy-In).

2·

You're ready to design and scale governance

You have some mandate and now need to know: What does a mature program look like? How do we organize? Should we build or buy? Start with Section 3 (Maturity Assessment) and Section 4 (Program Design), then move to Section 5 (Build vs. Buy) and Section 6 (Measurement).

Both paths converge on the same truth: AI governance done right is an accelerator, not a brake.

What This Playbook Delivers

By the end, you'll have:

- A compelling business case backed by data on how governance correlates with faster AI adoption, better AI product quality, fewer incidents, and better outcomes
- A stakeholder map and objection-handler's guide to help you align legal, compliance, risk, product, engineering, the board, and the CFO
- A self-assessment framework to understand your governance maturity today and identify the highest-impact gaps
- A blueprint for a governance program including structure, roles, policies, risk frameworks, and the control plane across the AI lifecycle
- A build vs. buy decision framework to help you allocate resources wisely
- KPIs and measurement approaches that connect governance to financial and strategic outcomes your board cares about

03 •

Section 1: The Business Case for AI Governance

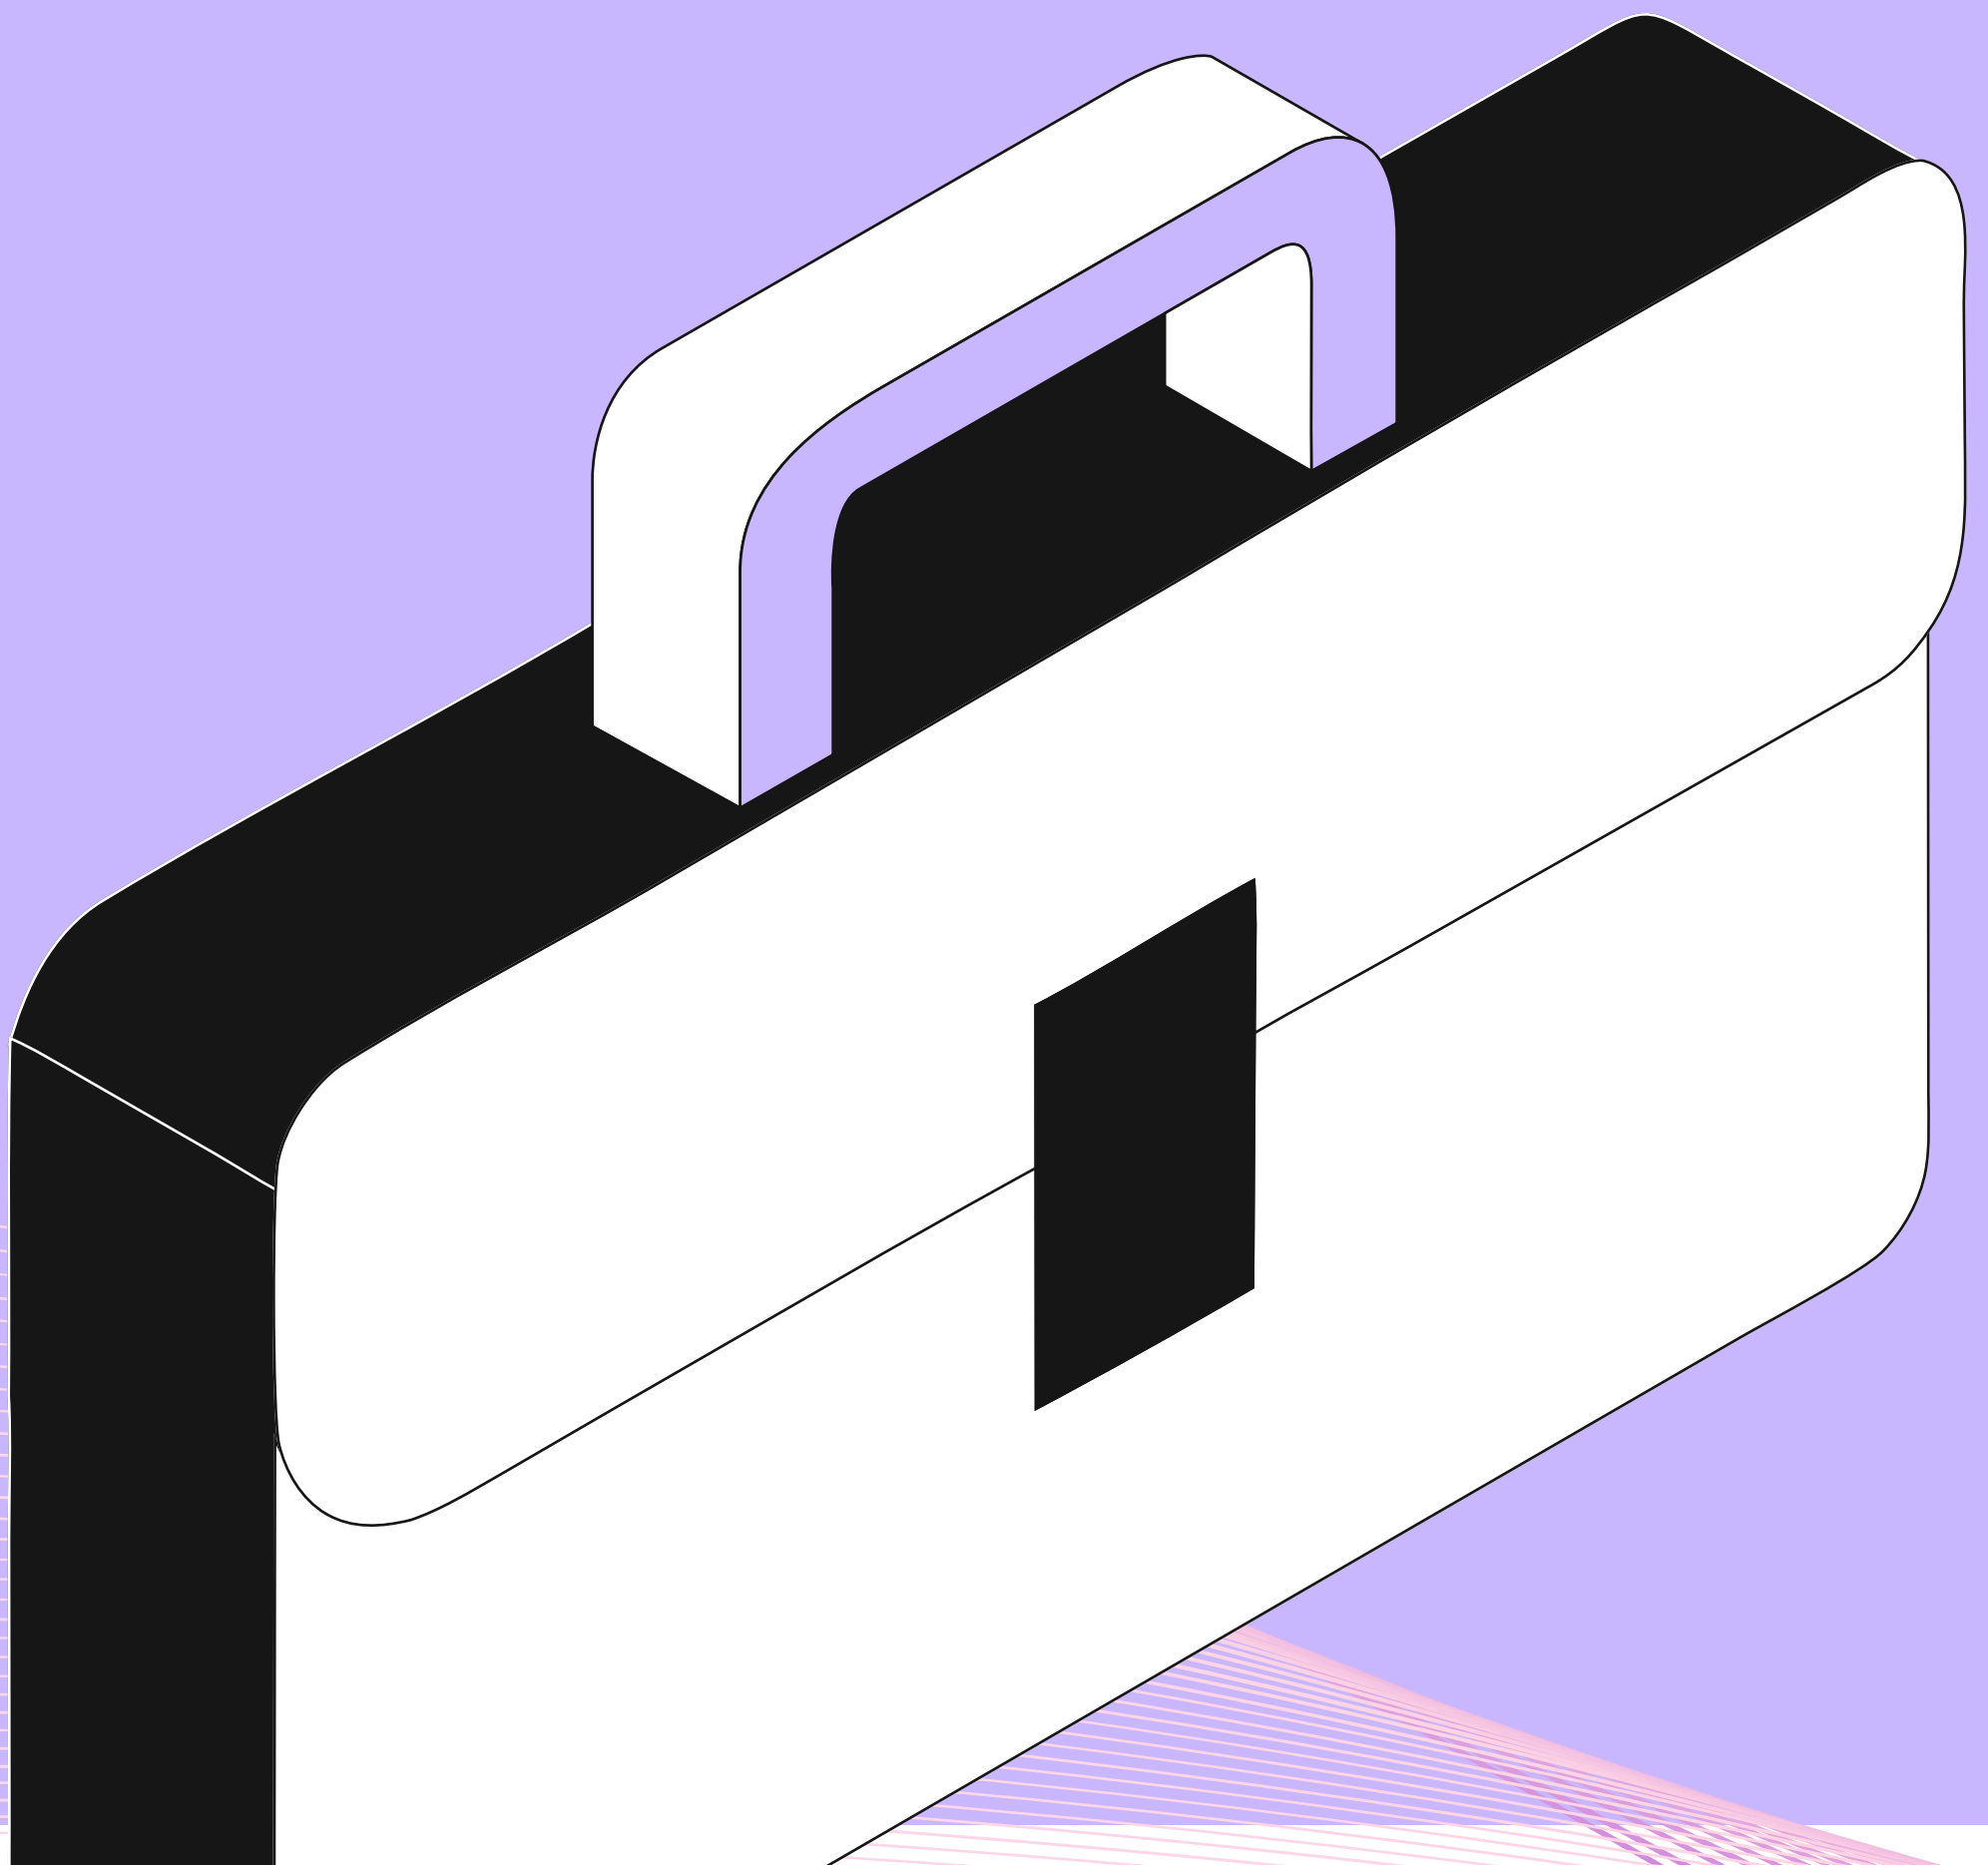

- 3.1

The Competitive Advantage: How Strong Governance Enables Trust at Scale

Here's what high performers do differently.

McKinsey's research shows that organizations with structured, proactive AI governance report significantly higher bottom-line impact from their AI investments.

They also report faster AI adoption, higher employee confidence, faster customer adoption of AI-driven products, and stronger vendor relationships.

Why? Because trust compounds.

When a company can say, "Every AI system in our portfolio has been assessed for security, robustness, bias, explainability, and regulatory alignment. We monitor it continuously. Here's the evidence," something shifts. Customers say yes faster. Partners ask fewer questions. Sales cycles shorten. And when a problem emerges, the company has the audit trail to act decisively and swiftly.

Contrast this with the majority of organizations. They run AI on reputation and hope. Governance lives in Slack messages and spreadsheets. When a model fails or a regulator asks questions, it’s a fire drill. There's scrambling, litigation risk, bad press and customer churn all compounding the crisis. Organizations that treat governance as an enabler rather than a constraint will own the next wave of AI scaling.

Real data supports this:

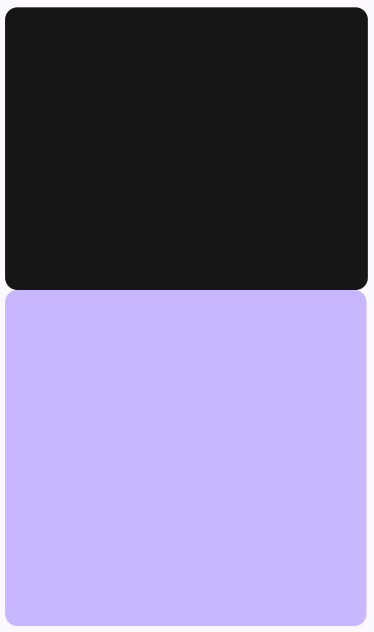

58%

of organizations now say AI is embedded in enterprise strategy (up from 26% a year ago).

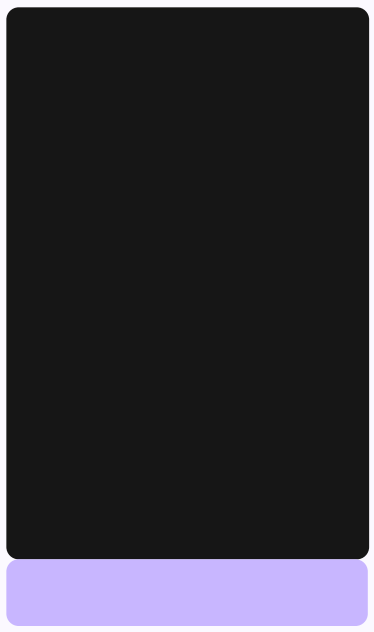

Only 19%

say they have fully implemented AI governance frameworks. That's a window of opportunity for fast-moving leaders.

- *The Competitive Advantage: How Strong Governance Enables Trust at Scale*

Companies with effective AI and data analytics risk management are significantly more advanced in technology adoption than those without.	Governance and scale are correlated, not opposed.	Operational efficiency: Organizations using platforms like Credo AI report that it takes them 10x less time to demonstrate compliance with regulations like the EU AI Act versus manual approaches.

Regulation is coming. Not someday—now.

The EU AI Act is enforced. The UK, Canada, Brazil, and others are following. The US is moving toward sectoral rules and potential federal AI regulations . And inside organizations, boards are demanding governance. The National Association of Corporate Directors now expects boards to have AI risk oversight on their agenda.

For most organizations, this feels like a threat. It's actually an opportunity.

Here's why: Governance requires cross-functional alignment.

Most organizations won't do it unless they have to. But organizations that build it early have a head start. They establish their governance baseline before the fire alarm goes off. They're ready for the audit. They're ready for the customer due diligence call. They're ready for the next regulatory shift and have future proofed their AI investments

- 3.2

The Risk Without Governance: Financial and Reputational

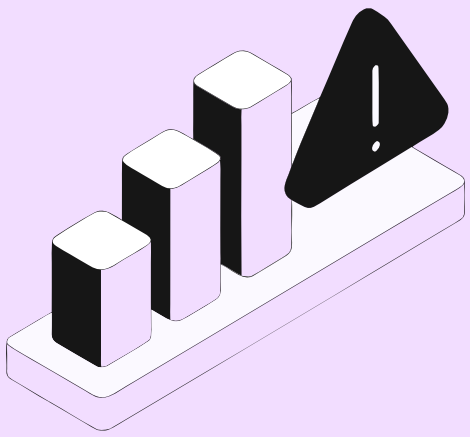

Poor product quality and Operational rework:

A model deployed without proper risk assessment lacks quality and might have to be pulled back, retrained, or reimplemented. That's time and money lost.

Compliance incidents and fines:

When an AI system breaks a regulation (EU AI Act, data privacy law, lending discrimination rules), the organization pays. Sometimes penalties can exceed a large percentage, e.g. 6%, of global revenue.

Customer erosion:

Customers increasingly require vendors to attest to AI governance. If you can't, you lose deals and that impacts your top line.
Litigation: If an AI system causes harm and there's no evidence of due

Litigation:

If an AI system causes harm and there's no evidence of due diligence, litigation exposure rises.

Capital costs:

Investors and insurers increasingly factor AI governance into their risk models. Weak governance means increased cost of capital or insurance.

Example:

A global financial services leader faced hundreds of generative AI use cases across the enterprise, putting pressure on manual approvals and risk processes. By centralizing AI use case intake in Credo AI, using an AI Registry and Vendor Portal to track internal and third-party applications, and automating risk categorization and workflow review, they reduced overall effort, accelerated time-to-governance, and gave executives clear visibility into where and how generative AI is used.

- 3.3

AI Is Becoming Agentic, Not Just Generative: Why Agentic Governance Matters

Generative AI afforded enterprises powerful content and prediction engines; agentic AI gives them actors. Agents don't just answer questions—they set sub-goals, call tools and APIs, chain actions together, and move money, data, and commitments on your behalf. For boards, that shifts the risk conversation from "bad advice" to "unintended actions at scale," which is why agentic AI is now a standing topic in many risk and audit committees.

From Credo AI's perspective, this is not a new problem—it's a higher-stakes version of the governance gap you already face. It's a shift "from AI that suggests to AI that does," which means your guardrails can't stop at model outputs anymore; they have to extend to what agents are allowed to decide, trigger, and touch. The same customers asking how you manage bias and explainability will now also ask very pointed questions about agent autonomy, agent permissions, chains of actions, and kill switches.

Boards and buyers are converging on five agentic risk questions:

Agent risks: What agents exist in the organization today, what can they do, and who owns them? Sanctioned and shadow agents are now a top concern in regulated industries.

Decision automation: Which decisions can agents make or materially influence without human sign-off, especially in credit, hiring, safety, and critical infrastructure?

Autonomy thresholds: How far can agents go on their own—what are the "red lines" for spend, access, and authority, and how are they enforced technically, not just on paper?

Chain of actions: Can you see the full chain of actions an agent takes across systems (APIs called, workflows triggered, data touched), and reconstruct it for an audit or incident review?

Exposure amplification: How could a single misaligned prompt, bug, or attack cascade across multiple tools, markets, or customer segments once an agent is live? Boards are already being briefed on how quickly agentic failures can propagate through supply chains and partner ecosystems.

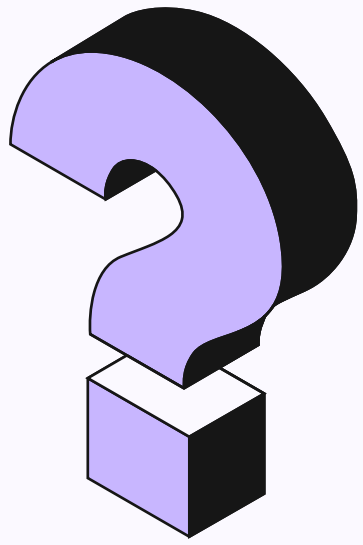

- *The Competitive Advantage: How Strong Governance Enables Trust at Scale*

This is where Agentic Governance comes in. In practice, it means extending your AI governance program and platform to:

> Maintain a live register of agents (internal and vendor-provided), with clear owners, privileges, and business purposes.

> Define autonomy levels and decision boundaries—for example, “observe,” “recommend,” “propose with human approval,” and “execute within limits”—and assign each agent a level tied to governance controls.

> Enforce least-privilege access and task-scoped permissions so agents can only act where they are explicitly allowed, for the smallest possible surface area.

> Require human-in-the-loop checkpoints for high-impact or regulated decisions, with clear logs of who approved what and when.

> Capture full action traces and evidence—not just model prompts and outputs, but the end-to-end chain of calls, tools, and changes an agent made—so you can answer regulators and customers when they ask, “What exactly happened here?”.

Analysts are already signaling that this is the next frontier. McKinsey urges enterprises to define autonomy levels, decision boundaries, and behavior monitoring specifically for agentic systems, not just traditional models. Forrester’s recent work on governance describes a shift toward “agentic governance systems” that actively enforce and remediate policies, not just document them. Credo AI’s own customers are now using the platform not only to inventory models and risks, but also to catalog agents, constrain privileges, and evidence how agentic actions stay within policy.

The takeaway for your board and your buyers is simple:

> AI is moving from answers to actions. Your governance has to move with it. Bringing Agentic Governance into your AI strategy and workflows —anchored in the same control plane you use for models—will be one of the clearest signals in 2026 that your organization is ready for the next wave of AI, not just the last one.

- 3.4

The ROI Framework: Value Levers That Connect to Your Bottom Line

1. Velocity and Scaling (Revenue and market opportunity)

- Faster time to market for new AI products
- Faster procurement of 3rd party AI for the business
- Ability to deploy more models with trust and accuracy (higher portfolio ROI)
- Faster sales cycles (customers trust you faster)
- New market access (you can now sell to regulated customers)

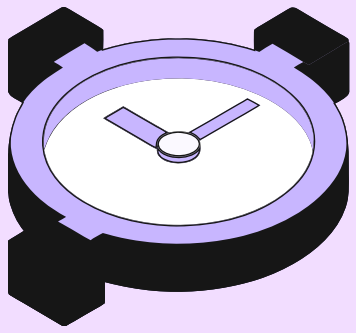

2. Trust and Brand (Indirect but real)

- Customer confidence and stickiness
- Employee attraction and retention
- Investor confidence
- Vendor and partner alignment
- Regulator and auditor confidence

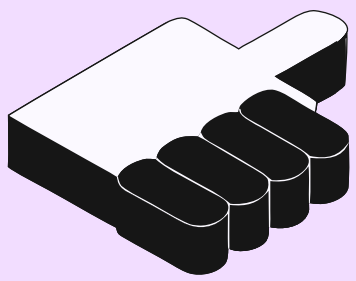

3. Operational Efficiency (Direct savings)

- Faster approval cycles (10x faster in regulated workflows)
- Reduced duplicate assessments and reviews
- Lower cost of evidence collection and audit preparation
- Reduced manual governance overhead

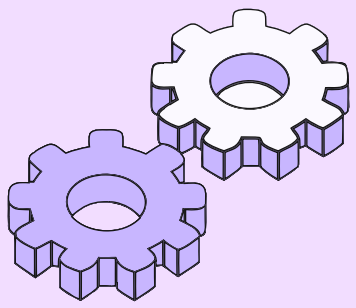

4. Risk Avoidance (Direct savings)

- Reduced compliance incidents and fines
- Avoided litigation costs
- Reduced customer churn from AI failures
- Avoided operational rework

- *The ROI Framework: Value Levers That Connect to Your Bottom Line*

The Model

Activating GenAI Governance in Financial Services:

A global financial technology company needed to manage Generative AI risks across thousands of models, vendors, and regions; without slowing innovation. Manual, spreadsheet-based reviews could not keep up with the volume and variety of new use cases emerging across functions from marketing to fraud and customer service.

What they implemented with Credo AI:

AI Registry and Automated Intake:

Centralized registry for all Generative AI projects (internal and vendor), with standard workflows to capture evidence, classify risk, and align to internal risk frameworks and global regulations such as the EU AI Act.

Cross-Functional Review and Approvals:

Automated routing to AI Governance, Security, Legal, Privacy, Brand, and Technology teams, enabling reviewers to collaborate and approve directly in the platform, with a persistent audit trail for every use case.

Vendor Portal for Third-Party AI:

A dedicated portal for external providers to submit documentation and evidence, giving the governance team a single source of truth for AI vendor risk and continuous validation of compliance.

The impact (within ~120 days)

- Operationalized Generative AI governance at scale across business units, regions, and vendors.
- Accelerated AI approvals by automating reviews and risk categorization.
- Unified oversight across Legal, Security, Privacy, and Technology through one platform.
- Strengthened compliance posture by aligning with global frameworks such as the EU AI Act and emerging U.S. standards.
- Reduced vendor assessment time via automated evidence collection and validation.
- Delivered executive-level visibility into all Generative AI use cases worldwide, enabling confident and compliant AI innovation.

- 3.5

Why 2026 Is the Tipping Point

Six forces converge:

AI adoption has normalized.

Pilots are moving to production. The question is no longer "Should we deploy AI?" but "How do we scale it with trust?"

Enterprises are proactive about trusted AI vendors

Enterprise buyers now audit AI governance as part of procurement.

Regulation is enforced, not proposed.

The costs of non-compliance are no longer theoretical.

Customers are demanding transparency.

Especially for high risk applications

Agentic AI is coming.

Autonomous AI agents make decisions on behalf of organizations. Governance becomes non-negotiable

AI capabilities are on the rise

and agentic AI is here demanding stronger governance.

Organizations that have governance in place today will scale confidently in 2026. Those that don't will face delays, constraints, and opportunity loss.

04 •

Section 2: Building Internal Buy-In

- 4.1

The Stakeholder Map: Who Needs to Align and Why

AI governance spans the entire organization. Each stakeholder has different priorities. Your job is to map them, understand their concerns, and give them a reason to care.

The Board

- What they care about: Fiduciary duty, regulatory exposure, strategic value, competitive advantage.
- What they hear: "Risk managed, confidence high, board-level reporting available."

Your message: Governance positions us for scale, enables new market access, and reduces regulatory and litigation risk.

The CFO

- What they care about: Cost, ROI, capital efficiency, operational expense.
- What they hear: "Faster ROI, lower TCO, reduced incident costs."

Your message: Governance accelerates time to value, by increasing top line, reducing rework, powering productivity, and enabling portfolio scaling.

General Counsel / Compliance

- What they care about: Cost, ROI, capital efficiency, operational expense.
- What they hear: "Faster ROI, lower TCO, reduced incident costs."

Your message: Governance accelerates time to value, by increasing top line, reducing rework, powering productivity, and enabling portfolio scaling.

CISO

- What they care about: Regulatory adherence, litigation defense, audit readiness, policy alignment.
- What they hear: "Evidence-ready, audit-prepared, policy-enforced, regulatory-aligned."

Your message: Governance embeds compliance into workflows, creates audit trails, and prepares you for regulatory inquiries.

- *The Stakeholder Map: Who Needs to Align and Why*

Chief Data Officer / Head of Analytics

- What they care about: Data quality, data lineage, responsible data use, talent retention.
- What they hear: "Data governance and AI governance are connected."

Your message: Governance ensures data quality, tracks data lineage, and builds team trust in data.

Chief AI Officer / Head of AI

- What they care about: Enabling fast innovation, scaling portfolios, avoiding constraints.
- What they hear: "Governance enables faster deployment, not slower."

Your message: Structured governance increases approval velocity, reduces incident downtime, and unlocks new use cases.

Product and Engineering

- What they care about: Speed, clarity, support, not being blocked.
- What they hear: "Governance gives us guardrails, not gatekeeping."

Your message: Governance provides templates, clear policies, and faster feedback loops—so you know that what you're building is trustworthy before you build it.

The Board's AI Committee (increasingly common)

- What they care about: Oversight, transparency, strategic alignment, risk posture.
- What they hear: "Monthly dashboards, clear metrics, incident visibility, governance maturity progress."

Your message: You'll have a single pane of glass for AI governance across the enterprise, with metrics the board can understand.

- 4.2

Common Objections and How to Respond

Objection 1

We already have GRC. Isn't that enough?

The reality:
Legacy GRC platforms were built for privacy and data security. They retrofit AI governance on top of frameworks designed for earlier risks.

Your response:
GRC gives us compliance infrastructure, which is valuable. But it's like using a database tool to manage machine learning pipelines—it can work, but it's not optimized. A purpose-built AI governance platform handles AI-specific risks: model drift, hallucination, adversarial attacks, emerging bias. It tracks the full lifecycle of models, agents, and datasets in ways GRC tools don't. Most mature organizations complement GRC with a dedicated AI governance layer.

Objection 2

Our AI is low-risk. Do we really need governance?

The reality:
Every organization underestimates AI risk. Models migrate. Use cases expand. Regulations tighten. Low-risk becomes high-risk overnight.

Your response:
Today, yes, your use cases may be lower-risk. But governance is about optionality and future-proofing. If we build governance now, when you want to deploy in high-stakes use cases- lending, hiring, healthcare- you're ready. If you wait, you'll have to retrofit. And regulation is moving faster than deployment. EU AI Act enforcement started this year. A governance foundation now means you're ready for the next wave without rebuilding.

Objection 3

Governance will slow us down.

The reality:
This is the most common objection and the most commonly disproven.

Your response:
The opposite is true with the right model. Manual governance—spreadsheets, email approvals, duplicate reviews—slows you down. Structured governance with a platform accelerates approval cycles, reduces rework, and avoids incidents that cause real downtime. Companies using AI governance platforms report 10x faster compliance evidence collection and faster deployment cycles for new models. The bottleneck isn't governance—it's chaos pretending to be flexibility.

Objection 4

We can't afford it. Budget is tight.

The reality:
Governance has upfront cost but saves money downstream. The conversation is about timing and ROI.

Your response:
Two ways to think about this. First, the ROI is real—faster approval cycles, fewer incidents, and reduced audit prep time typically pay for governance platforms within 18 months, often faster. Second, the alternative isn't free. Manual governance is expensive in person-hours. And a single compliance incident or customer due diligence loss is often many times the cost of a governance platform. We can phase implementation—start with high-risk models and scale. This is an investment in competitive advantage, not a cost center.

Objection 5

This is just 'responsible AI' theater.

The reality:
There's a lot of governance theater. This isn't it.

Your response:
You're right to be skeptical. A lot of governance is principles without practices. Here's the difference: we're talking about operational governance—policies that are enforced, risk assessments that are automated, compliance that's continuous, not annual. This is governance that changes how the organization works, not a report that sits on a shelf.

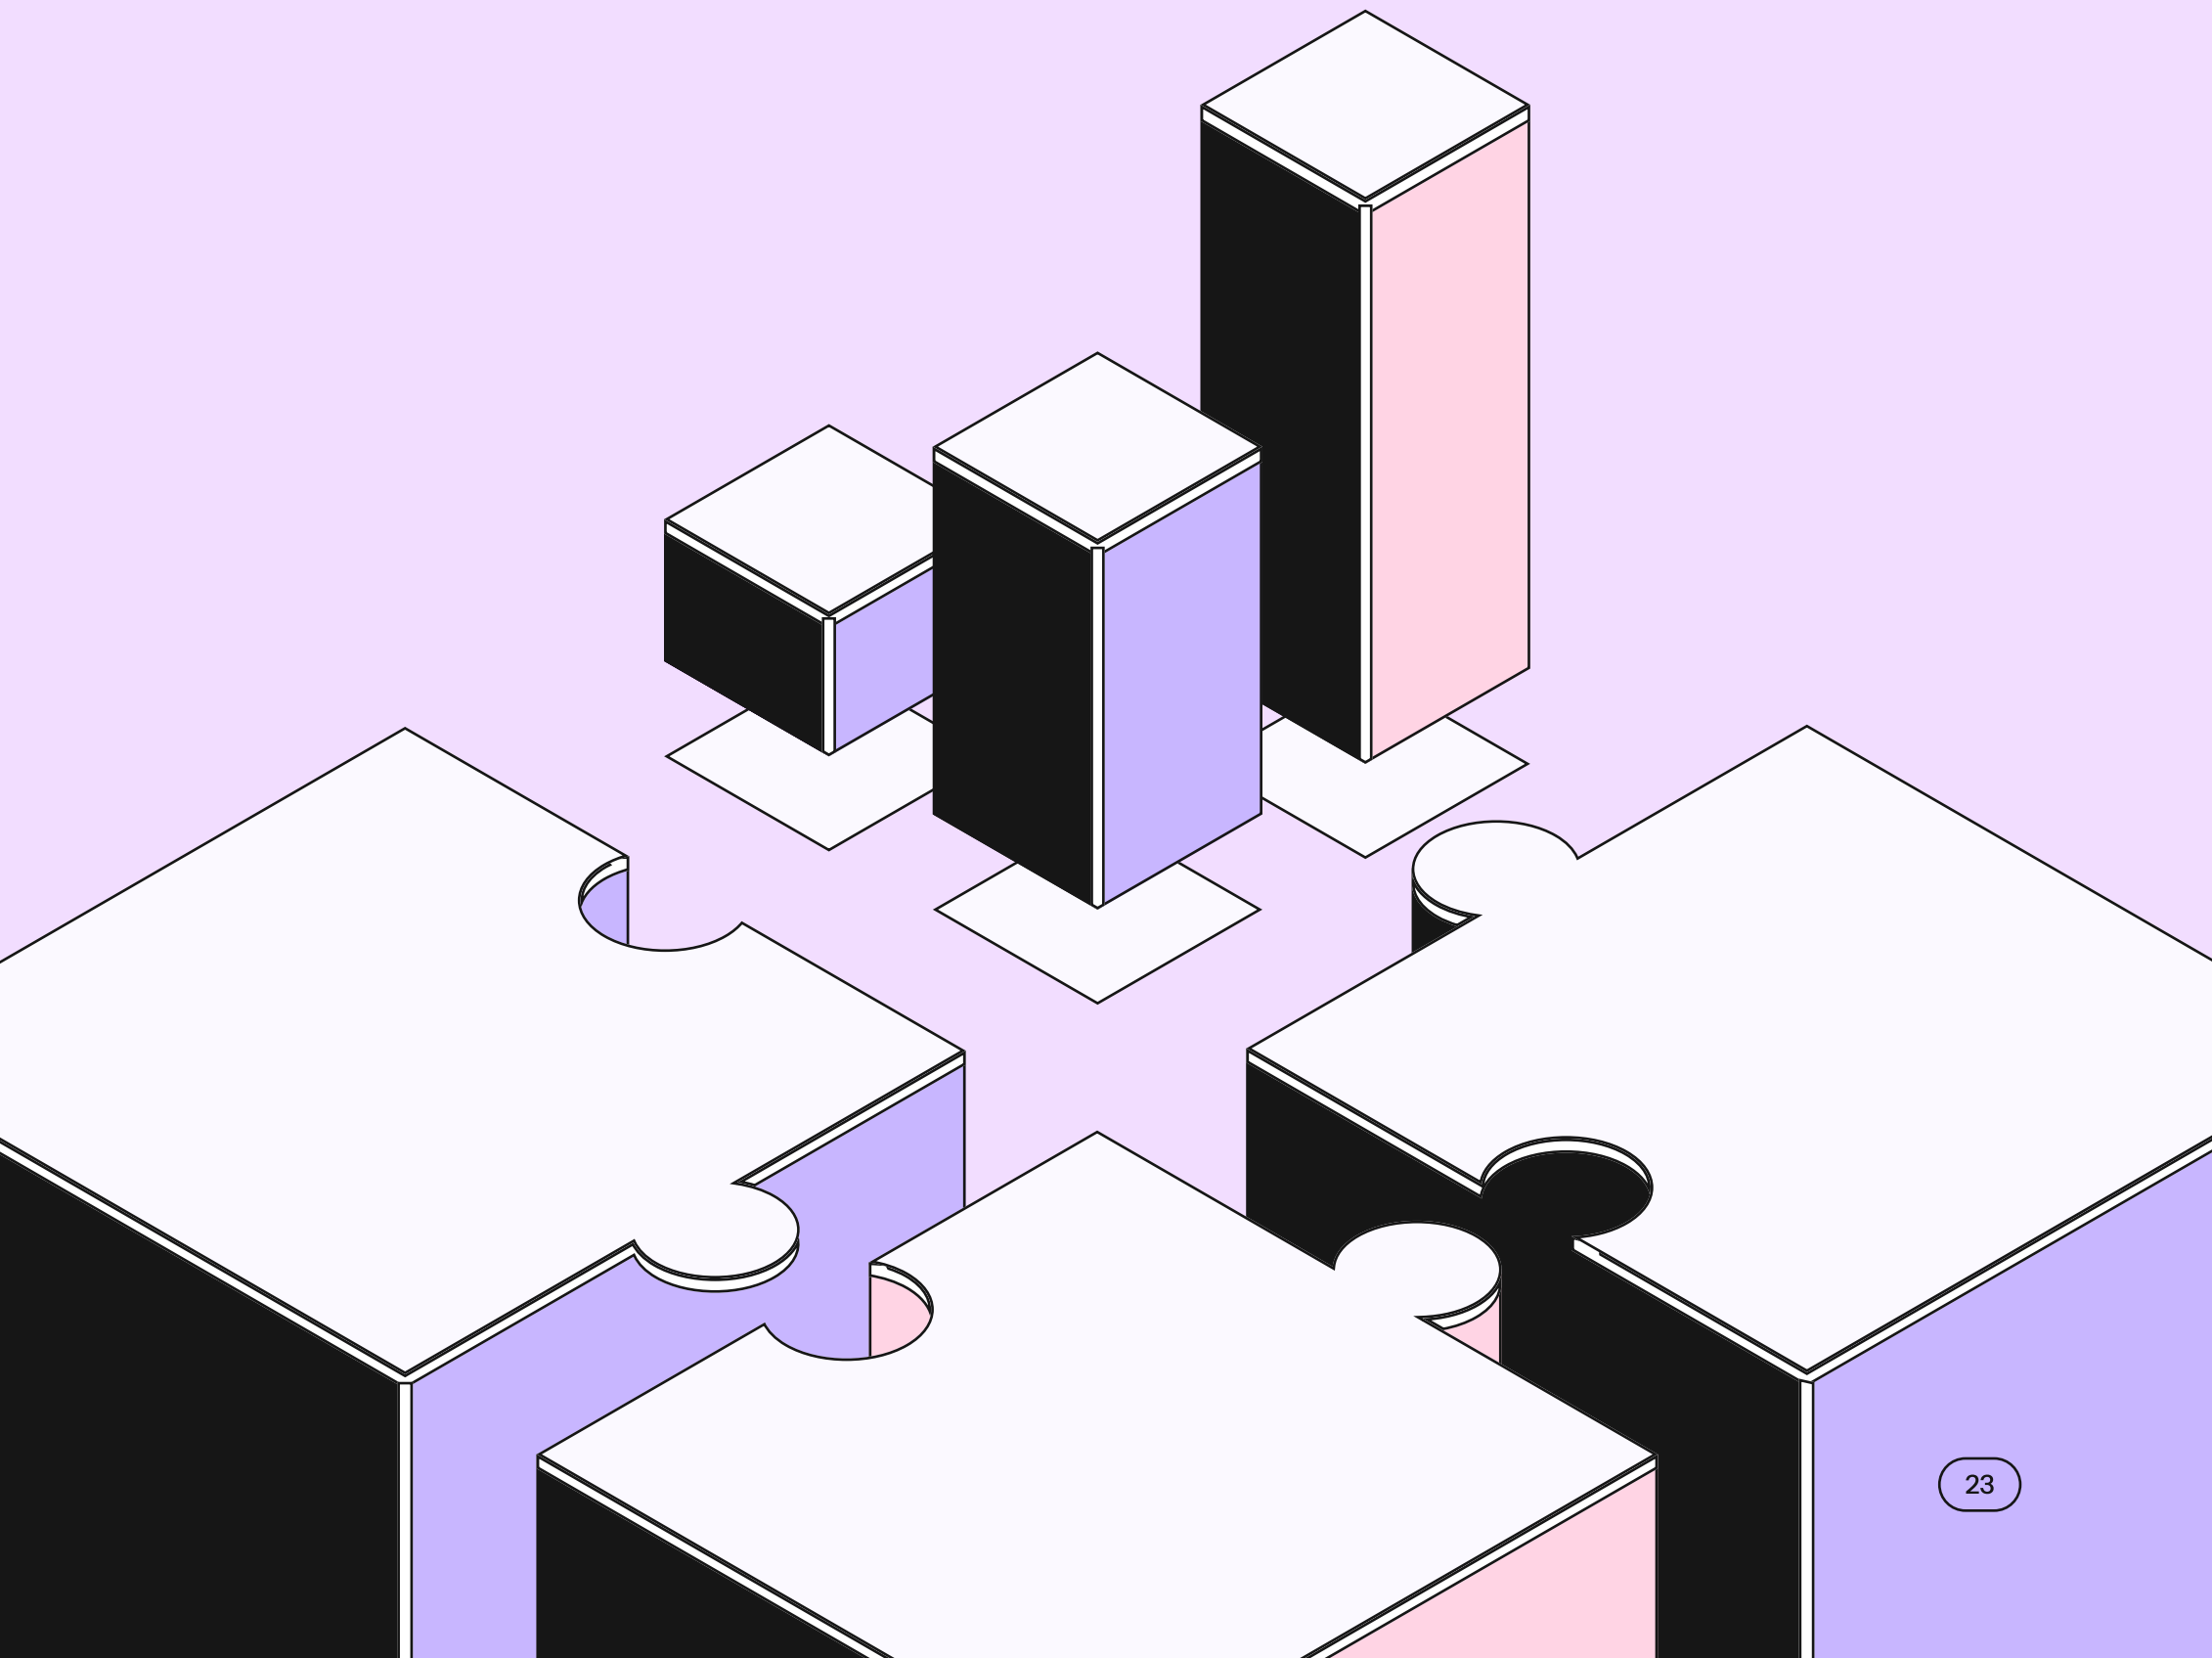

- 4.3

Top-down vs. Bottom-Up Empowerment: The Two-Sided Approach

Governance requires both top-down expectation and bottom-up enablement.

→ Top-down (Horizontal):

- Board and executive expectation that governance is a priority
- Policy requirements that are clear and consistent
- Resource allocation
- Accountability structures

→

↑ Bottom-up (Vertical):

- Easy-to-use tools and templates
- Clear training and guidance
- Examples and playbooks
- Support from governance teams

↑

The mistake most organizations make is choosing one or the other. Strong organizations do both. They set expectations from the top and enable teams at the bottom with the tools and training to meet them.

05 •

Section 3: AI Governance Maturity Assessment

- 5.1

Map Where You Are in the Journey

- Credo AI's Enterprise AI Governance Maturity Model describes how organizations evolve from early exploration to governing generative and agentic AI at speed.

  It applies across your entire AI estate—technology, people, and processes—and is the backbone for Credo AI's platform and Advisory Services.

  Most enterprises are uneven: they may be "Formalizing" in policy, but only "Exploring" in monitoring or inventory. That's normal. The goal of this model is to help you locate where you are today, clarify what "better" looks like, and prioritize your next set of moves.

- 5.2

The Six Maturity Levels and Outcomes

Level 1

Exploring

Organizations want to learn about AI governance—building awareness, baselining, and gap analysis.

- **Typical state:** Shadow AI, no common language, no inventory, unclear responsibilities.
- **Outcome of this stage:** A shared understanding of AI risks, a first map of AI activity, and agreement that governance is a strategic priority—not an afterthought.

Level 2

Aligning

Organizations want to define high-level enterprise AI governance principles.

- **Typical state:** Fragmented efforts across legal, risk, security, and data; ad hoc decision-making; limited visibility.
- **Outcome of this stage:** Clear principles, named owners, and an initial governance blueprint that aligns leadership on “how we will govern AI here.”

Level 3

Formalizing

Organizations want to adopt standardized workflows for AI governance.

- **Typical state:** Policies on paper, but inconsistent, reactive application; every new AI initiative feels bespoke; approvals are slow.
- **Outcome of this stage:** Repeatable intake, assessment, and approval workflows that reduce friction, remove guesswork for teams, and start generating reusable evidence.

Level 4

Optimizing

Organizations want to optimize their standardizing governance, risk management, and compliance initiatives.

- **Typical state:** Structures exist, but tools and processes are fragmented; multiple spreadsheets and point solutions; governance feels like overhead.
- **Outcome of this stage:** A unified control plane—often via a platform like Credo AI—that connects policies to practice, cuts manual work, and makes governance a tailwind for AI deployment rather than a speed bump.

Level 5

Governing at Scale

Organizations want to scale AI governance, risk management, and compliance.

- **Typical state:** Strong foundations, but hard to keep up with the volume of use cases, evolving regulations, and third-party/vendor AI; monitoring is periodic rather than continuous.
- **Outcome of this stage:** Enterprise-wide visibility across hundreds or thousands of AI systems, dynamic risk management, and regulatory traceability that allows you to respond quickly to new rules, audits, and board questions.

Level 5+

Governing at Speed (Agentic-Ready)

Organizations are using AI-augmented AI governance with real-time visibility into agentic and autonomous AI.

- **Typical state:** Mature governance program, but pressure to anticipate—not just react to—emerging risks from agents, autonomy, and multi-jurisdictional oversight.
- **Outcome of this stage:** Governance that learns and adapts in near real time, with clear autonomy thresholds, action traces for agents, and automated controls that let you innovate quickly without losing control.

- 5.3

The Self-Assessment Framework

So, where does your organization stand today?

Across these levels, you can score your organization on six practical dimensions that show up in every Credo AI deployment and maturity assessment:

Item	Description	Score 1-5
Oversight & Ownership	Executive sponsorship, governance council, clear RACI for AI systems.	
Principles & Policies	Documented principles translated into enforceable, AI-specific policies and standards.	
Visibility & Inventory	A live registry of models, agents, and AI use cases—including third-party and shadow AI.	
Risk & Prioritization	Consistent tiering and treatment of AI risks, with heightened scrutiny for high-impact uses.	
Monitoring & Control	Ongoing technical and process monitoring, especially for high-risk and agentic AI.	
Documentation & Traceability	Fast access to who approved what, which controls apply, and how systems behave over time.	

For each dimension, pick the level (1–5+) that best describes your current state. The pattern that emerges is your current maturity profile.

- 5.4

How Credo AI Can Accelerate the Journey to Trusted AI Adoption

This model shared on the last page is not just theoretical. It underpins how the Credo AI platform and Advisory services help customers move from Exploring to Governing at Scale and Speed—by baselining your current state, designing phased roadmaps, and implementing workflows and controls that match your ambition

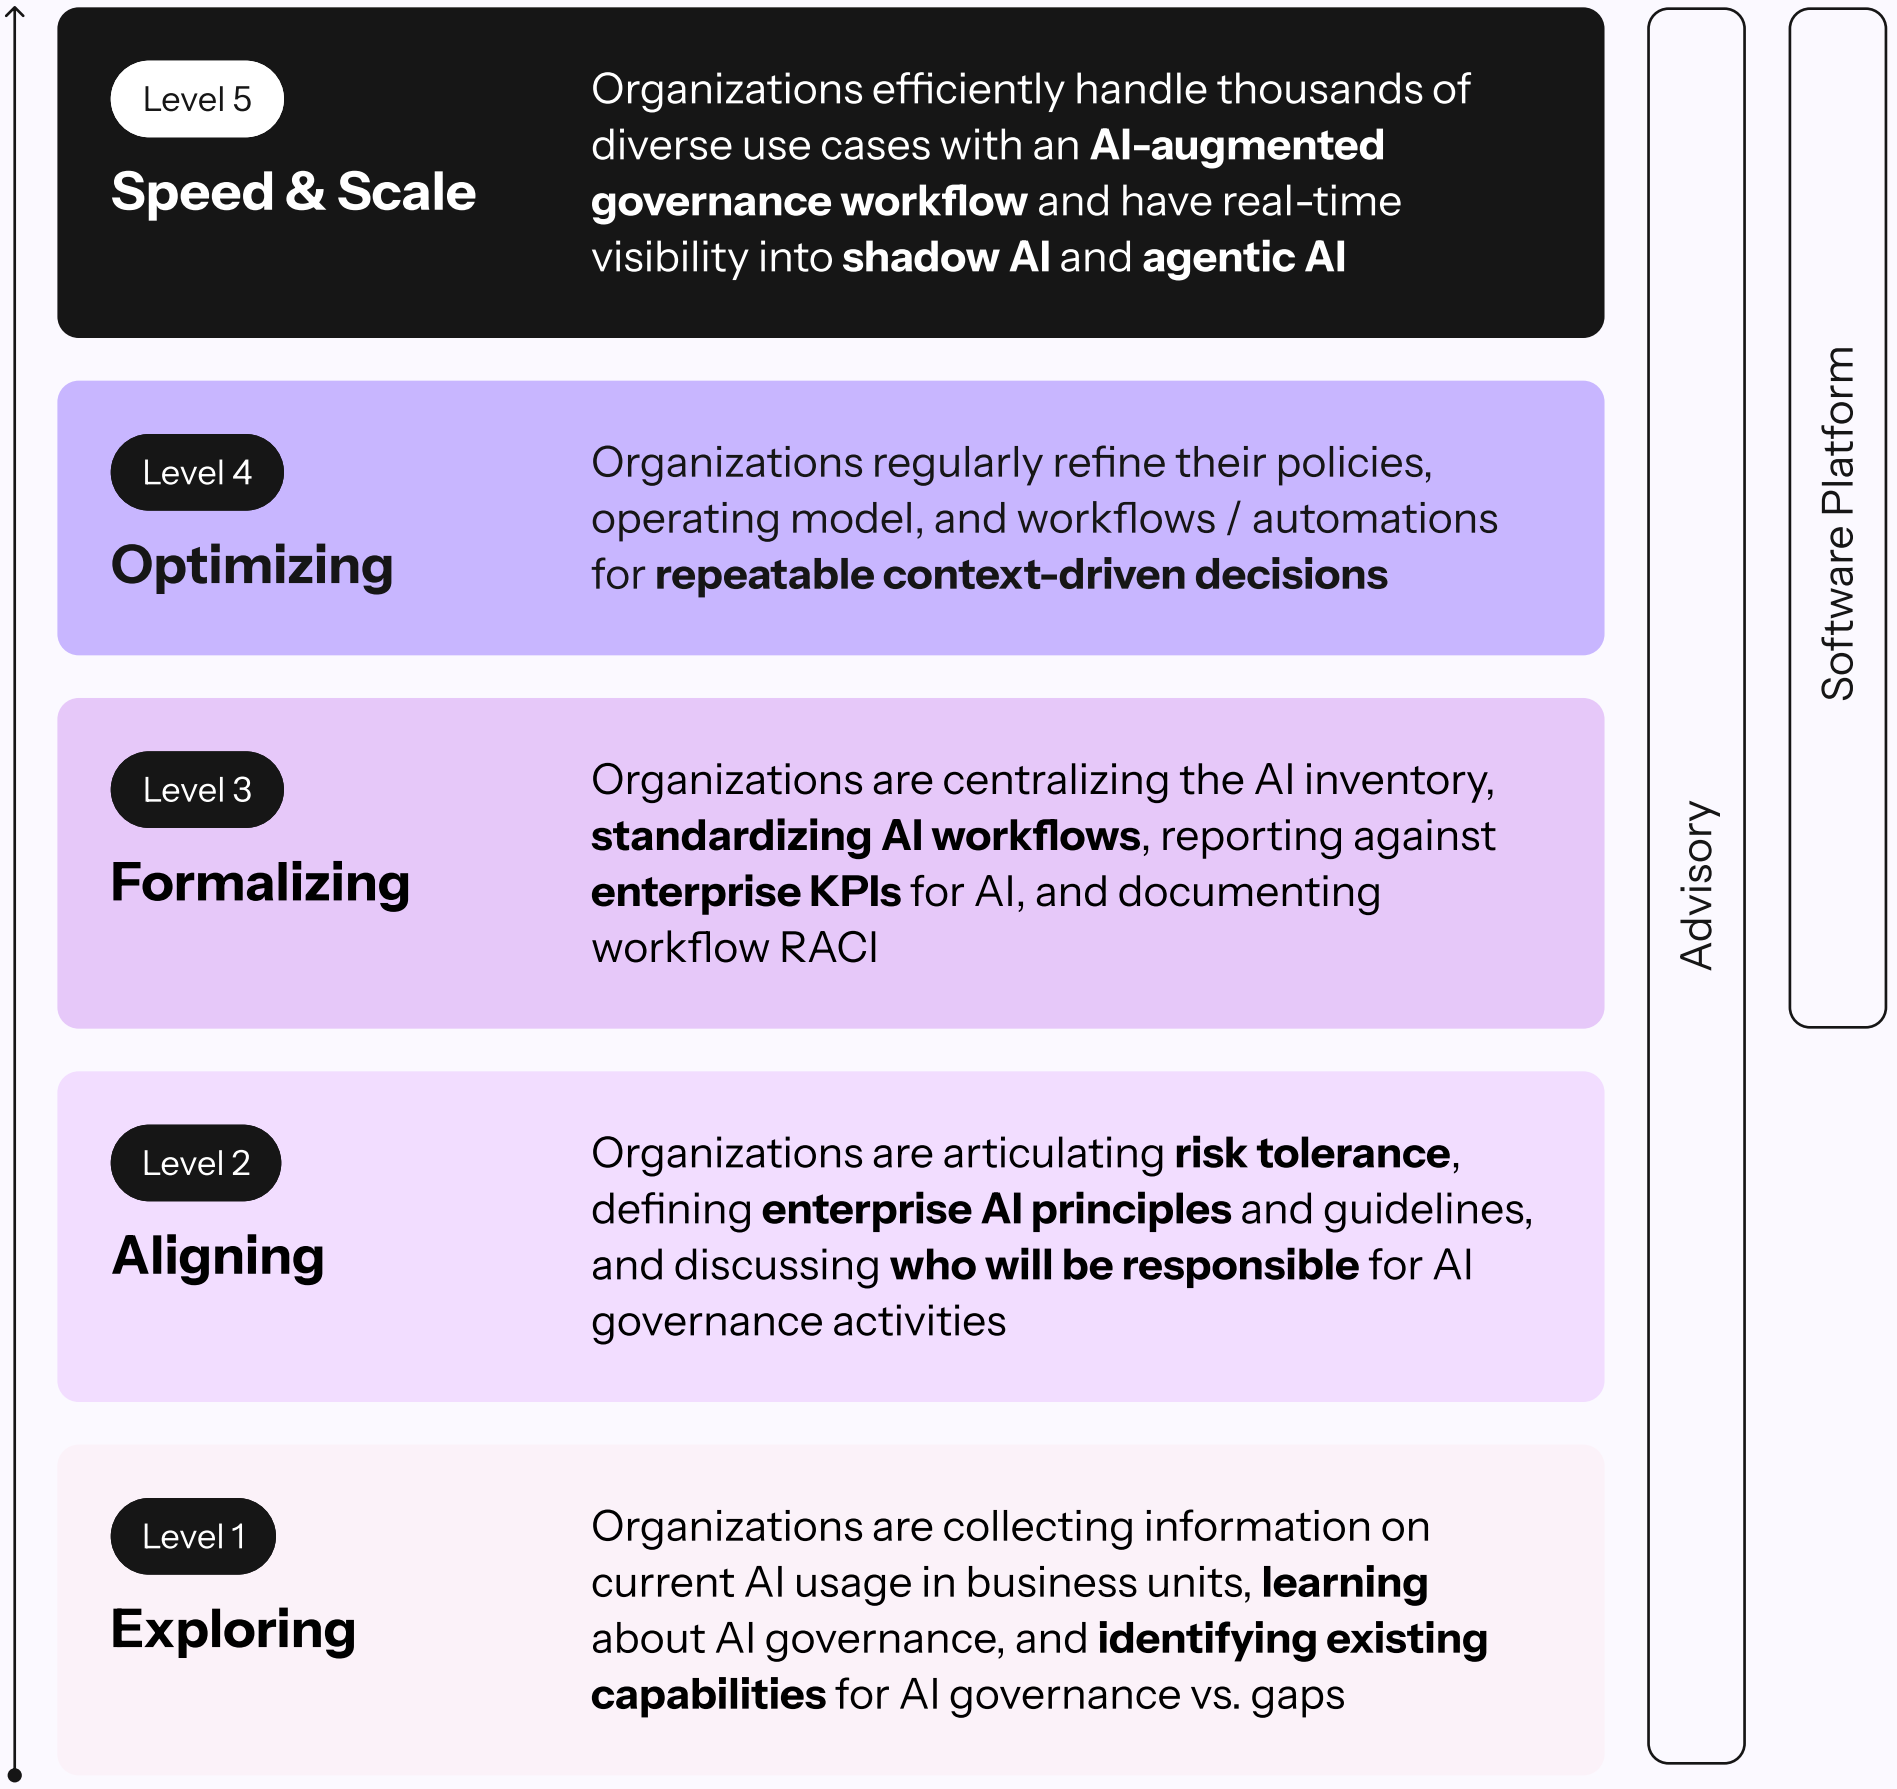

- 5.5

Benchmark and Prioritization

Once you've scored yourself across the six dimensions and mapped to the Exploring → Aligning → Formalizing → Optimizing → Governing at Scale → Governing at Speed (5+) levels, the next step is to decide where to act first.

After the self-assessment, focus on three questions:

1

Biggest gaps

Where are you sitting at Level 1–2 while the business is already behaving as if you were at Level 3–4?

2

Highest-impact gaps

Which gaps create the most regulatory, customer, or operational exposure if left unaddressed?

3

Quickest wins

Which gaps could you materially improve in the next 90 days to build momentum and credibility

Across Credo AI customers and broader market research, a familiar pattern shows up: many organizations are Exploring or Aligning in visibility and monitoring, Aligning or Formalizing in principles and policies, and only starting to Formalize or Optimize in risk workflows and documentation. This means early wins often look like:

Standing up an initial AI inventory/ registry, even if it starts as a simple but structured list.

Defining a clear risk tiering framework (e.g., low/medium/high, with examples) and using it for all new AI use cases.

Establishing a governance council with a clear charter and RACI, so there is a visible decision-making body.

Capturing and standardizing what you already do today into simple, shareable workflows and templates.

- *Benchmark and Prioritization*

Realistic Progression Timelines

Progress is not overnight, but it also doesn't need to take years to see value. Based on enterprise experience and advisory work:

Level 1 → **Level 2** — Typically 2–3 months

Exploring → **Aligning**

Focused on baselining, principles, and ownership.

Level 2 → **Level 3** — Typically 4–6 months

Aligning → **Formalizing**

Focused on standardizing intake, assessment, and approval workflows.

Level 3 → **Level 4** — Typically 12+ months

Formalizing → **Optimizing**

As you consolidate tools, embed a platform, and tighten policy-to-practice links.

Level 5/5+ — a multi-year journey

Reaching Governing at Scale and Speed

With seamlessly integrated governance architecture, automation, and agentic-ready controls.

Benchmarking Language You Can Use With Your Peers

When you talk to executives or the board, you can position your maturity relative to peers without naming competitors directly. Market studies and responsible AI surveys show that most organizations are still in the early to mid stages of trustworthy AI implementation, with the majority clustered in the first two levels and fewer than one in five having fully operationalized trustworthy AI practices.

- You might frame it this way:

> Across industries, most enterprises are still in the Exploring to Formalizing range—roughly Levels 1–3 in our model—with only a smaller set of leaders truly Optimizing or Governing at Scale. The gap between where most organizations are and where regulation, customers, and agentic AI will push them is significant. That gap is our opportunity: moving from Level 2–3 to Level 4+ over the next 12–24 months positions us as a leader, not a follower.

06 •

Section 4: Core Components of an AI Governance Program

- 6.1

Governance Structure and Roles

A mature AI governance program needs structure. Here's what it looks like:

Executive Sponsor (CEO, COO, or Board Member)

- **Role:** Sets tone, allocates resources, ensures accountability at the top.
- **Key responsibilities:** Board reporting, cross-functional alignment, breakthrough issues.

Chief AI Governance Officer, Responsible AI Officer, or Head of AI Governance (New role in many orgs)

- **Role:** Owns governance program design, implementation, continuous improvement.
- **Key responsibilities:** Policy development, platform management, stakeholder alignment, metrics.
- Reports to: CFO or Chief Risk Officer (not buried in IT).

AI Governance Council (Cross-functional)

- **Members:** Heads of AI/data, CISO, compliance, legal, business units, external advisors.
- **Frequency:** Monthly or quarterly.
- **Responsibilities:** Policy review, escalations, strategic alignment, external advisory input.

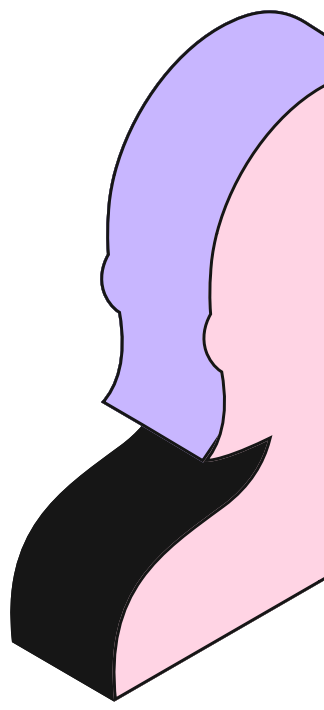

AI Risk and Compliance Team (Embedded in AI governance function)

- **Key responsibilities:** Risk assessments, policy enforcement, audit support, continuous monitoring.
- **Headcount:** Typically 2-4 FTE to start (assess, design, oversee platform).

Executive Sponsor (CEO, COO, or Board Member)

- **Role:** Sets tone, allocates resources, ensures accountability at the top.
- **Key responsibilities:** Board reporting, cross-functional alignment, breakthrough issues.

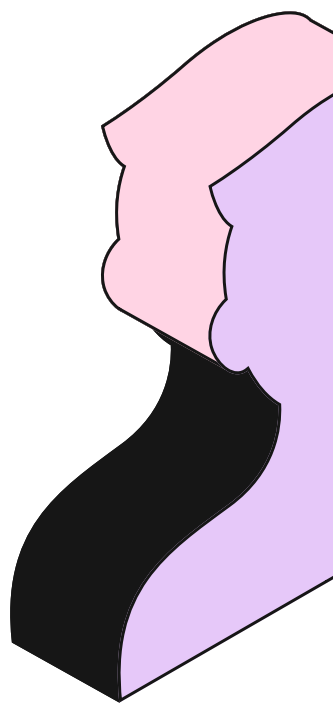

AI System Owners (Distributed across business units)

- **Key responsibilities:** Ensuring their systems comply with governance policies, responding to monitoring alerts, supporting audits.

Data Governance Team (Existing function, expanded)

- **Key responsibilities:** Data quality, data lineage, consent and privacy for AI datasets.

- 6.2

Policies and Standards

Core Policy Areas:

What should your policies cover?

1. AI Lifecycle Governance
   - Discovery and intake of new AI initiatives
   - Risk assessment criteria
   - Risk management process embedded in the entire lifecycle of an AI system, identifying foreseeable harms, mitigating them, and documenting why choices were made.
   - Approval gates
   - Deployment requirements
   - Iterative monitoring and review frequency
   - Retirement process

2. Model Governance
   - Model development standards (training data requirements, validation, testing)
   - Model documentation requirements (lineage, provenance, performance metrics)
   - Version management
   - Change management
   - Documentation, transparency and compliance support to downstream deployers.
   - Model safety and security

3. Risk Classification Framework
   - **High-risk:** High-risk AI systems are those used in especially sensitive, high-impact contexts where errors or bias could seriously affect people's safety, rights, or access to essential opportunities or services, so the law requires the strongest controls and proof.
   - **Limited-risk:** Systems that inform decisions or affect operations that require transparency mechanisms, but do not otherwise fall under high-risk or prohibited AI categories (e.g., customer segmentation, chatbots)
   - **Minimal risk:** Systems that are not deemed as impacting fundamental human rights/safety (e.g., autocomplete or spam filters)

- *Policies and Standards*

- General-Purpose AI Models:
  - Standard GPAI model risks: General-purpose models can be used in many ways, so their mistakes, biases, and opacity can spread widely into downstream products and decisions.
  - Systemic-risk GPAI model risks: At the frontier, these models can create large-scale, cross-sector harms through powerful capabilities, emergent behavior, and cascading impacts when many systems rely on the same model.
  - Different governance requirements for each tier - corresponding to their severity, probability and scale.

4. Bias, non-discrimination and fairness
   - Requirement for bias assessment of high-risk models
   - Documentation of known fairness limitations
   - Process for bias remediation and mitigation

5. Explainability and Transparency
   - Requirement to explain model decisions (especially for high-risk models) in order to not just inform ‘users know it’s AI’, but whether reasons are understandable enough to challenge or audit.
   - Customer transparency (when AI affects them, they deserve to know)

6. Data Governance for AI
   - Data source requirements (provenance, quality)
   - Consent and licensing alignment
   - **Control and document data-preparation steps** (annotation, labeling, cleaning, updating, enrichment, aggregation)
   - **Identify and analyse potential biases** that could harm safety, rights, or cause discrimination, including feedback-loop bias
   - Data retention and deletion
   - Process special-category (sensitive) data only if strictly necessary for bias detection/correction, and only with safeguards
   - Cross-border data transfer restrictions
   - **Ensure datasets reflect the real deployment context** (geographical, behavioural, functional, contextual fit)

- *Policies and Standards*

7. Human Oversight

- How autonomous is the system, and is there meaningful human oversight (ability + authority + time to intervene)?

8. Third-Party AI

- Vendor assessment requirements
- Contractual governance requirements
- Any GPAI/foundation model dependencies, vendor assurance, contractual evidence, and role-shifting risk
- Monitoring responsibilities
- Dependency mapping

9. Incident Response and Escalation

- How to identify and report AI incidents (and to whom)
- Continuous post-market monitoring
- Escalation procedures
- Remediation timelines
- Post-incident review process

Policy Development Approach:

Map definitions, risk-tier categorization, and mitigation strategies to external frameworks (e.g., NIST AI RMF, ISO 42001, EU AI Act) where applicable)

Compliance with global regulatory requirements based on operating jurisdictions, AI system uses/capabilities, and sectoral regulatory obligations

Make policies clear and prescriptive enough to drive consistent decisions and accountable behavior, yet modular and principles-anchored enough to adapt as risks, technologies, and regulations evolve

Write for your audience (i.e. technologists and product managers)

Review and update dynamically and iteratively (AI is moving fast)

- 6.3

Risk Assessment and Classification

A governance program is only as good as its risk discipline. Here's how to build it:

Step 1: Define High-Risk Categories

- Individual impact: Does the system make or heavily inform decisions affecting individuals? (hiring, lending, healthcare)
- Severity of safety/ rights impact: Not just whether individuals are affected, but how badly (loss of liberty, livelihood, safety, discrimination).
- Impact to Vulnerable Groups (including children, racial or ethnic groups).
- Probability of an impact to safety or rights occuring
- Dependency mapping
- Scale: How many people does it affect?
- Reversibility: Can decisions be appealed or overturned?
- Transparency: Do users know they're interacting with AI?
- Robustness & reliability: Performance under edge cases, distribution shift, adversarial inputs; error bounds tied to use-case tolerance.

Step 2: Build a Risk Assessment Questionnaire

Questions to ask about each AI system:

- What does the system do? What decisions does it inform or make?
- What data does it use? Is the data accurate? Is the data representative? Are there known biases?
- Who built it? What's their governance maturity?
- Who uses it? What is the intended purpose ?
- How accurate, robust and cyber-secure is it? How do you know?
- What could go wrong? What's the impact? And what is the risk management process in place?
- What are potential unintended impacts or consequences of the AI system from use (or misuse), both intentional and unintentional?
- Is there human review? What triggers it? How often?
- Is it explainable? Can you defend decisions?
- Does it affect protected groups? How do you ensure that there are additional augmented parameters corresponding to the needs of these groups?

- *Risk Assessment and Classification*

Step 3: Build a Risk Score

Create a simple scoring model: Impact × Likelihood × Manageability = Risk Score

- Impact: What happens if the model fails? (1-5 scale)
- Likelihood: How likely is failure? (1-5 scale)
- Manageability: Can you detect, respond quickly and mitigate substantially? (1-5 scale, inverted—lower is better)

This gives you a single risk number that determines approval gates, monitoring frequency, and escalation.

- 6.4

The Control Plane: Lifecycle Governance

The core of operational governance is a control plane that spans the AI lifecycle. Think of it as a command center.

Discover Phase

- [ ] Identify and catalog all AI systems (built, open-source, third-party)
- [ ] Assess the organization's AI estate
- [ ] Identify unknown systems

Tool: Centralized inventory/registry

Assess Phase

- [ ] Conduct risk assessment using the questionnaire above
- [ ] Document data sources, lineage, fairness considerations
- [ ] Identify gaps and compliance risks

Tool: Risk assessment platform with workflow and documentation

Approve Phase

- [ ] Route high-risk models to governance council
  - Medium-risk to risk team
  - Low-risk for self-certification
- [ ] Create approval audit trail

Tool: Workflow automation, approval gates

- *The Control Plane: Lifecycle Governance*

Deploy Phase

- [ ] Ensure deployment meets policy requirements (monitoring, documentation, etc.)
- [ ] Enable deployment only after approval
- [ ] Document deployment configuration

Tool: Deployment gates, configuration management

Monitor Phase

- [ ] Continuous monitoring for high-risk models (performance drift, bias drift, security issues)
- [ ] Alerts for anomalies
- [ ] Regular review for medium/low-risk models

Tool: Continuous monitoring, alerting, dashboards

Report Phase

- [ ] Executive dashboard for board and leadership
- [ ] Audit evidence collection
- [ ] Regulatory reporting
- [ ] Incident reporting

Tool: Reporting platform, audit trail

Retire Phase

- [ ] Decommission systems that are no longer needed or that have been replaced
- [ ] Archive documentation
- [ ] Delete data according to policy

Lifecycle management, data deletion workflow

- 6.5

Using a Governance Platform

At scale, this control plane needs to be automated and integrated. This is why organizations increasingly adopt AI governance platforms (like Credo AI, OneTrust, Alteryx, etc.).

A good governance platform gives you:

Centralized inventory: All AI systems in one registry

Integrated risk assessment: Workflows that guide teams through assessment

Policy automation: Policies embedded into workflows and enforced

Continuous monitoring: Integration with production systems to detect issues

Evidence collection: Immutable audit trails

Regulatory reporting: Pre-built templates for compliance frameworks (EU AI Act, NIST AI RMF, ISO 42001)

Escalation automation: Route issues to the right people

Executive dashboards: Board-ready visibility

07 •

Section 5: Build vs. Buy Decision Framework

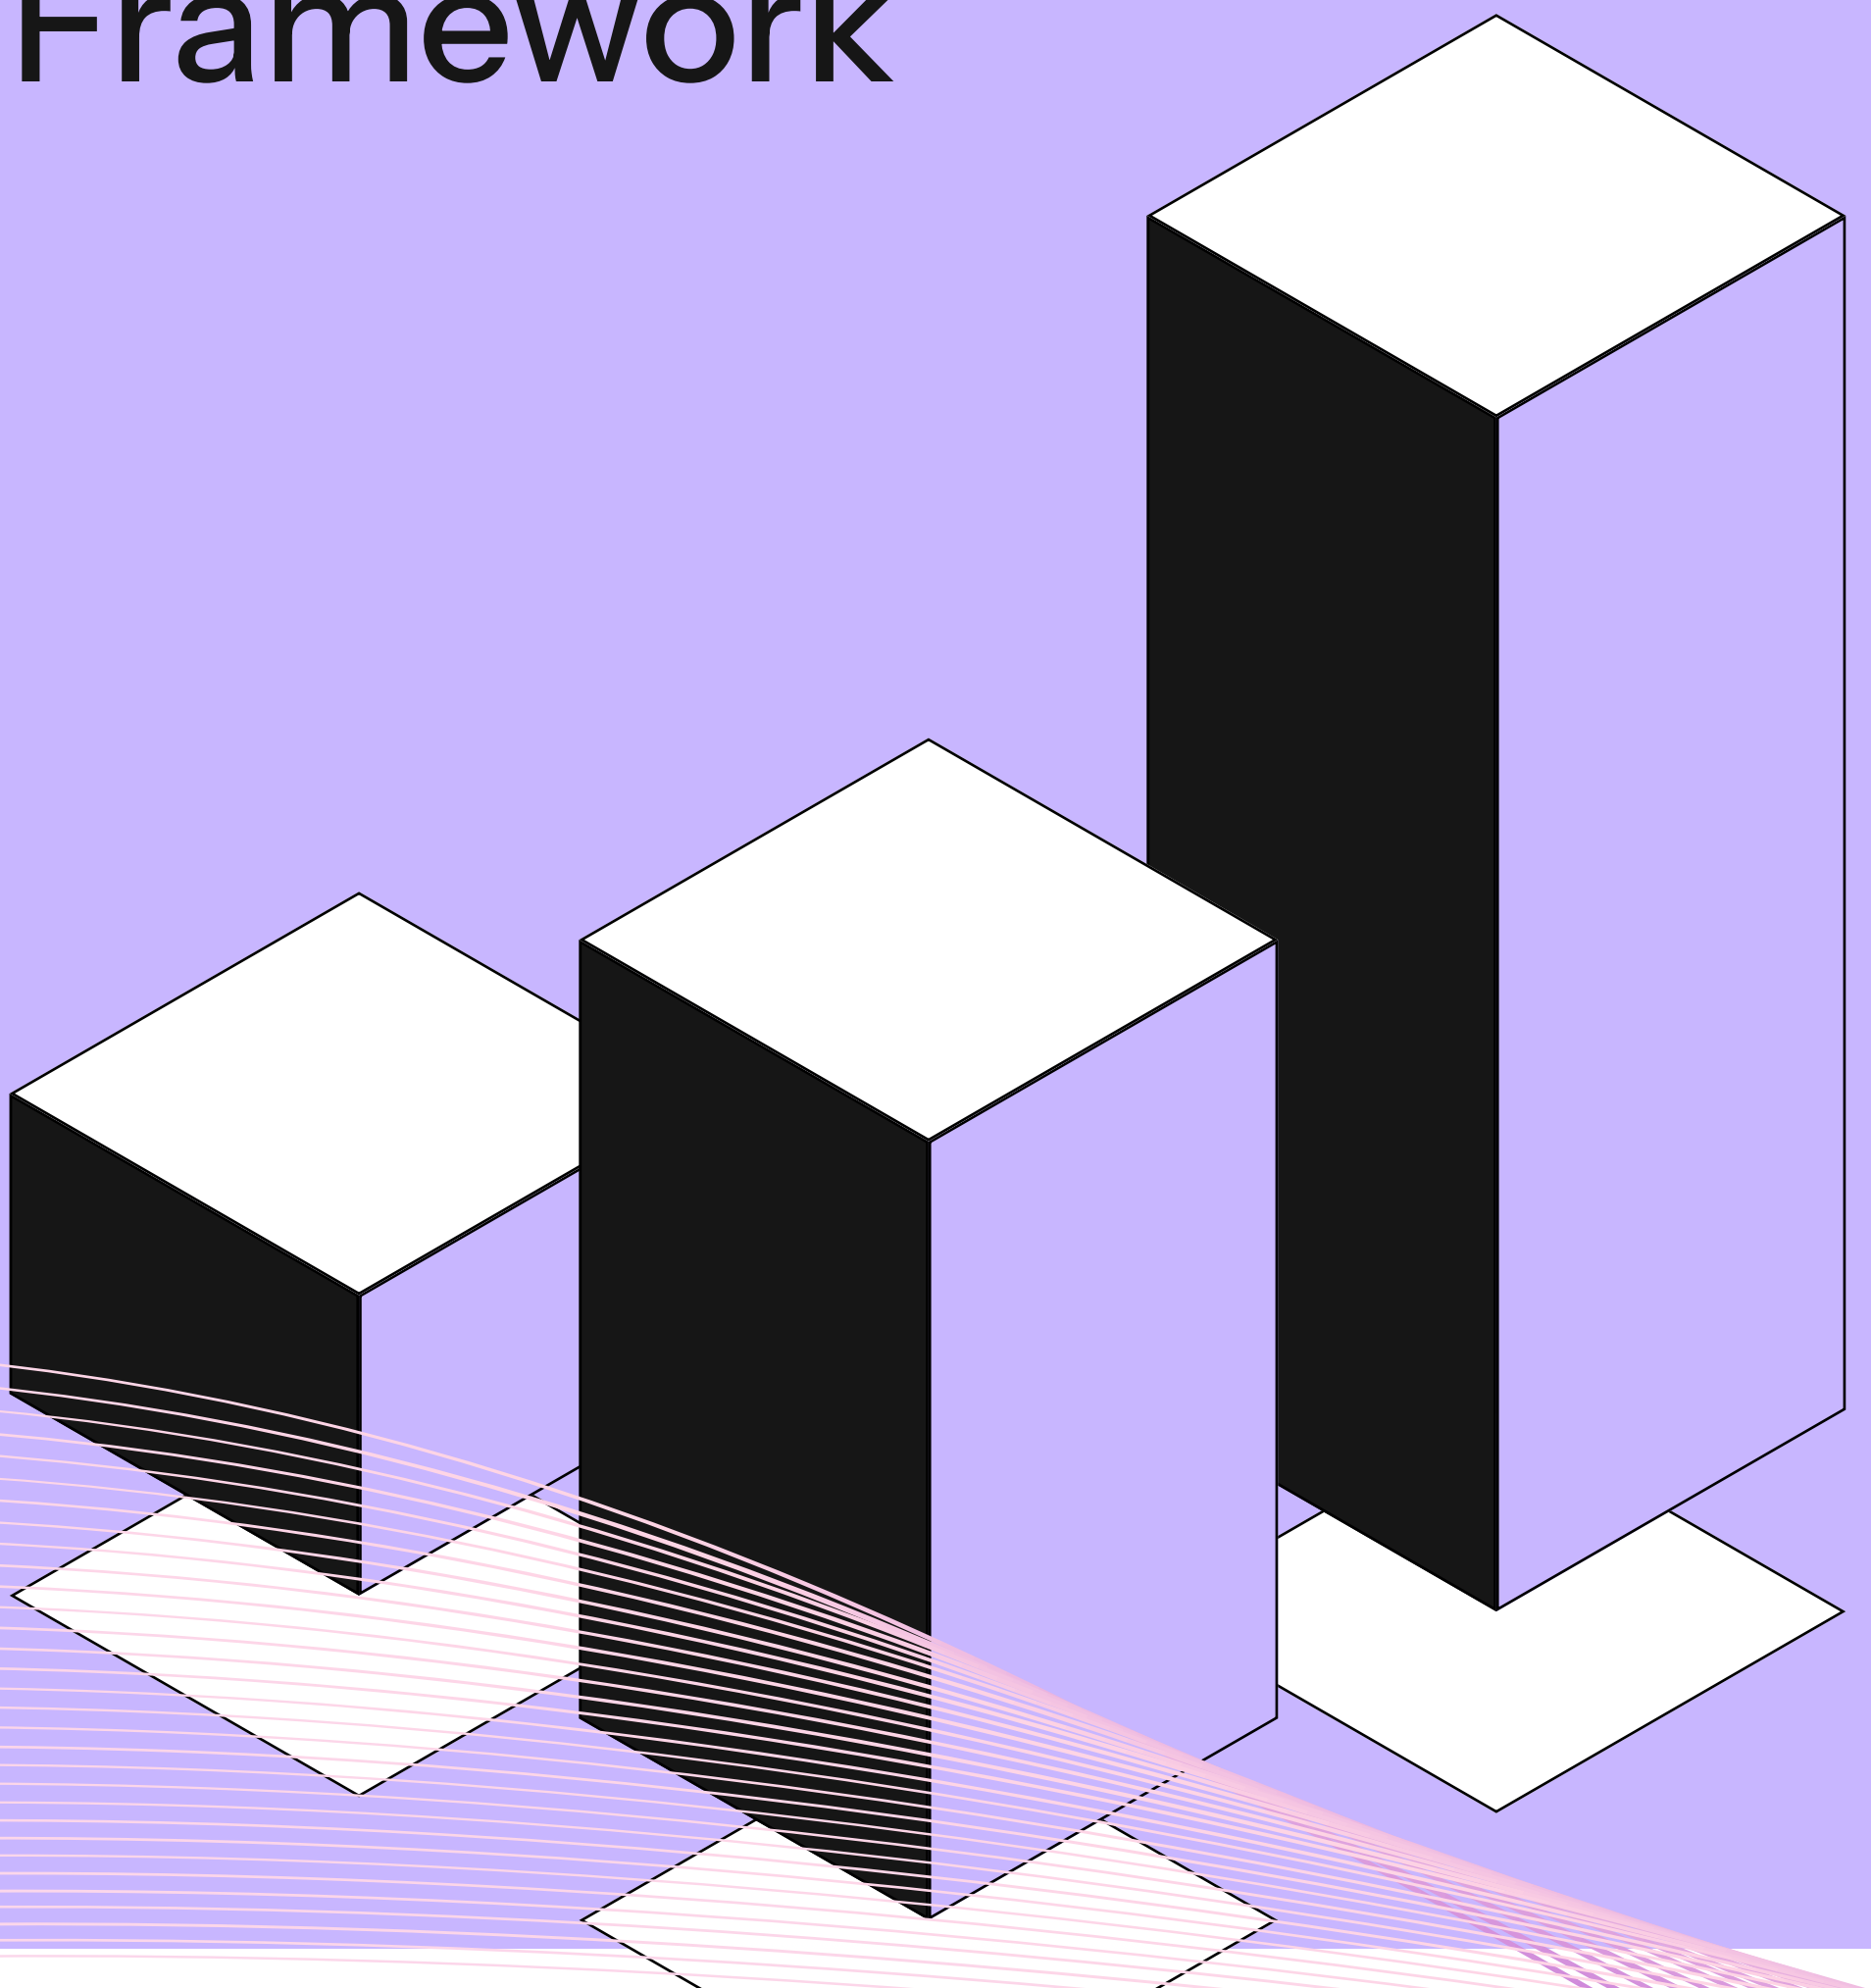

- 7.1

Internal Capability Assessment

Before you decide to build governance in-house, assess your organization's capability across four dimensions:

1. Governance Expertise

Question: Do you have people who understand AI risk, compliance, and governance frameworks?

- ✓ If yes: You can build (or at least configure a platform)
- ✗ If no: You'll need to hire or partner

2. Technical Infrastructure

Question: Do you have the data infrastructure and integration capabilities to build a governance platform?

- ✓ If yes: Consider building if you want deep customization
- ✗ If no: Buying makes sense

3. Resource Availability

Question: Do you have 2-4 FTE (initially, scaling to 4-6 over time) to dedicate to governance?

- ✓ If yes: You have the option
- ✗ If no: buying a platform and hiring a governance lead is more efficient

4. Budget Flexibility

Question: Can you invest in governance upfront, or do you need to see ROI within 12 months?

- ✓ If you can invest upfront: Building is an option
- ✗ If you need rapid ROI: Buying (and focusing on adoption) is better

- 7.2

Budget Justification Framework [1]

- # How to Frame the Investment Conversation

The Cost Structure

What you're investing in:

1. AI Governance Platform — $150K-$400K Year 1

SaaS option:
$150K-$400K annually for enterprise-grade platform with automated assessments, monitoring, and documentation

Custom build:
$500K-$1.2M upfront + $300K-$600K annual maintenance (only for large enterprises with unique requirements)

Hybrid approach:
$300K-$600K combining commercial platform with custom integrations

- **Platform includes: Risk assessment tools, model registry, policy management, audit trails, reporting dashboards**

2. Team Headcount (~2-4 FTEs to start) — $400K-$800K fully loaded

AI Governance Lead:
Owns strategy, stakeholder management, policy framework ($180K-$250K)

AI Risk Analysts (2-3):
Day-to-day assessments, documentation, compliance tracking ($120K-$160K each)

Technical Auditor:
Model evaluation, testing, technical validation ($140K-$180K)

- **Team scales to 4-8 FTEs by year two as AI portfolio grows**

- *Budget Justification Framework*

3. Advisory Services (as needed)

$150K-$500K

Strategic advisory:
Policy framework design, regulatory roadmap, governance operating model ($50K-$200K)

Implementation support:
Platform integration, workflow design, process optimization ($100K-$300K)

- **Advisory is front-loaded in Year 1, then drops to $50K-$200K annually for specialized support**

4. Training & Enablement

$50K-$150K

Organization-wide AI governance training for developers, product managers, legal teams

Certification programs for governance team members

Change management to embed governance into existing workflows

Annual refresher training and updates on emerging regulations

1. Estimates based on market research and representative data; see Appendix for methodology and limitations.

Typical investment range:

Mid-market (1K-5K employees):

$250K-$500K
Year 1

Enterprise (5K-20K employees):

$500K-$1M
Year 1

Large enterprise (20K+ employees):

$1M-$2M
Year 1

Note: The build vs. buy decision is the biggest cost driver. Most organizations should buy (SaaS) unless they have extraordinary requirements.

- *Budget Justification Framework*

The Returns: Three Value Pillars

1. Operational Efficiency: Faster Approvals, Reduced Rework

How to quantify YOUR savings:

Start with your current state:

- ☐ How many AI projects per year? _____
- ☐ Average approval/review time per project? _____ weeks
- ☐ How many need rework due to governance issues? _____%
- ☐ Average cost to fix governance problems post-deployment? $_____ K

Expected improvements with governance:

- ✓ **50-70% faster approvals** through automated risk assessments, standardized templates, and clear decision criteria
- ✓ **60-80% reduction in rework** by catching issues during development, not after deployment
- ✓ **30-50% less time** spent on legal/compliance reviews through pre-approved patterns and guardrails

Conservative estimate for your org: $____K annually (For mid-market: $500K-$800K; Enterprise: $800K-$1.5M)

Example calculation

Company deploying 20 AI projects annually:

Time savings:

- Current: 8 weeks × 3 FTEs × $2,500/week = $60K per project
- With governance: Reduce by 60% = $36K saved per project
- Annual savings: $36K × 20 projects = $720K

Rework avoidance:

- Current: 6 major rework incidents/year × $125K each = $750K
- With governance: Reduce by 75% = $562K saved

Total operational efficiency value: $1.28M/year

- *Budget Justification Framework*

2. Revenue Enablement: New Markets, Faster Sales, Customer Retention

How to quantify YOUR opportunity:

Assess your current constraints:

- ☐ How many deals require AI governance documentation? _____/year
- ☐ What's your win rate on those deals? _____%
- ☐ How many opportunities in regulated markets can't you pursue? _____
- ☐ Average deal size? $_____ K
- ☐ Current sales cycle length for AI products? _____ months

Expected improvements with governance:

- ✓ **30-40% faster sales cycles** when you can immediately provide governance documentation, certifications, and audit trails
- ✓ **25% higher win rate** on deals requiring governance proof (Gartner: 67% of enterprise buyers now require this)
- ✓ **Access to regulated markets** (financial services, healthcare, government) currently off-limits
- ✓ **15-25% better retention** on AI-enabled products through demonstrated responsible practices and transparency

Conservative estimate for your org: $____K annually (For mid-market: $800K-$1.5M; Enterprise: $1.5M-$3M)

Example calculation

Enterprise sales organization:

New market access:

- Regulated market opportunities: 8 deals/ year
- Historical win rate: 0% (couldn't compete)
- With governance win rate: 25%
- Average deal size: $500K
- New revenue: 8 × 25% × $500K = $1M

Faster sales cycles:

- Current pipeline value: $3M
- Accelerate close by 35% (from 9 months to 6 months)
- Time value of earlier revenue: ~$200K

Customer retention:

- AI product revenue base: $5M
- Retention improvement: 20%
- Additional retained revenue: $1M

Total revenue enablement value: $2.2M/year

- *Budget Justification Framework*

3. Risk Avoidance: Reduced Compliance Incidents, Litigation, Fines

How to quantify YOUR risk exposure:

Calculate your regulatory and operational risk:

- ☐ What regulations apply to your AI systems? (EU AI Act, sector-specific rules, state laws)
- ☐ What are the potential penalties? $_____ M
- ☐ Have you had AI-related incidents or near-misses? _____
- ☐ What's your model failure rate? _____%
- ☐ Average cost of AI system failure or bias incident? $_____ K

Expected risk reduction with governance:

- ✓ **70% reduction** in probability of major compliance violations through systematic risk assessment and controls
- ✓ **75% fewer production incidents** by catching issues before deployment
- ✓ **Avoid reputational damage** worth 3-10x direct costs through proactive governance

Regulatory landscape:

- **EU AI Act:** Fines up to €35M or 7% of global revenue for high-risk AI violations
- **US sector regulations:** Financial services (model risk management), healthcare (FDA guidance), fair lending laws
- **State AI laws:** Emerging requirements in CA, NY, CO, IL for AI transparency and impact assessments

Conservative estimate for your org: $____K annually (For mid-market: $400K-$800K; Enterprise: $800K-$2M)

- *Budget Justification Framework*

Example calculation

Mid-size enterprise with regulated AI systems:

Regulatory risk:

- Estimated penalty for high-risk violation: $10M
- Baseline probability over 3 years: 15%
- Expected loss without governance: $1.5M
- Risk reduction with governance: 70%
- Risk mitigation value: $1.05M over 3 years = $350K/year

Operational risk:

- Model failures: 3/year × $200K each = $600K
- Reduction with governance: 75% = $450K saved

Insurance benefit:

- Cyber insurance premium: $150K/year
- Reduction with demonstrated governance: 15% = $22K saved

Total risk avoidance value: $822K/year

Total Value Summary

Value Driver	Your Annual Benefit
• **Operational Efficiency**	$___K
• **Revenue Enablement**	$___K
• **Risk Avoidance**	$___K
• **Total Annual Value**	$__M

Typical ranges: • Mid-market: $1.7M-$3M annually • Enterprise: $3M-$6.5M annually

- 7.3

Strategic Positioning: Competitive Advantage, Not Cost Center

- Frame this as strategic capability-building, not compliance overhead:

Competitive advantages you're buying:

- ✓ **Speed to market:** Fast, confident AI deployment while competitors are stuck in legal review
- ✓ **Market access:** Ability to win deals and enter markets that competitors can't (yet)
- ✓ **Customer trust:** Transparent, trustworthy AI as a brand differentiator
- ✓ **Talent magnet:** Top AI talent wants to work where governance enables innovation, not blocks it
- ✓ **Regulatory resilience:** Future-proof against emerging AI laws and standards
- ✓ **Innovation enablement:** Clear risk boundaries let teams experiment with confidence

Board-level value story:

Risk management

AI is the new cyber—governance is how we protect and future-proof the business

Growth enabler

Governance unlocks $X M in revenue **we can't capture today**

Operational excellence

~40% improvement in AI development velocity

Strategic positioning

~12-24 month competitive advantage over reactive peers

- *Strategic Positioning: Competitive Advantage, Not Cost Center*

Language that resonates with executives:

"AI governance is to AI what quality management is to manufacturing—a competitive necessity"

"We're either leading with governance as an advantage, or scrambling after an incident"

"Every major enterprise will have AI governance; the question is whether we're proactive or reactive"

"This investment returns approximately $3-7 for every $1 we spend, with 12-18 month payback"

“Governance helps us deploy 2-3x more use cases with trust”→ translates to revenue opportunity

---

"Governance reduces audit prep time by ~80%" → translates to operational efficiency

---

"Governance enables us to enter regulated customer segments" → translates to revenue expansion

- 7.4

Call to Action

Present this as a three-part decision:

Do we believe AI is strategically important?	Do we lead or follow?	What level of investment?
Yes → governance is non-negotiable	Leader → invest proactively; Follower → pay more later	Match your AI ambition: minimal compliance vs. competitive advantage

Recommended ask:

Year 1 budget

$[insert your range] for platform, team, and setup

Governance model

Board oversight, quarterly value reviews, cross-functional steering committee

Success metrics

Track operational efficiency, revenue wins, and risk incidents

Decision timeline

60-90 days to avoid delays on current AI initiatives

- The alternative (doing nothing) costs $6.5M-$23M over three years in avoidable losses.

08 •

Section 6: Measuring Success

- 8.1

KPIs and Governance Metrics

“You cannot manage what you cannot measure”

A mature AI governance program must quantify adoption, efficiency, compliance, and business value using consistent measurement standards. Leading organizations increasingly evaluate governance using **operational metrics** and **business value metrics.**

This section outlines:

- Enablement and adoption
- Operational governance performance
- Risk & Compliance Coverage
- Stakeholder Buy-in & Engagement
- Business outcomes

It also incorporates **Credo AI’s 10 Gold Standard Metrics**—the industry’s only governance metrics framework purpose-built for AI adoption and enterprise accountability.

Trusted AI Adoption & Enablement Metrics

This section explains that AI governance is what makes AI adoption possible at scale, because trust and confidence in AI products come from clear governance policies and processes. It highlights that many organizations struggle to move generative AI experiments into production due to risk, regulation, and governance issues, and positions these metrics as a way to track how governance is enabling trustworthy deployment of AI across the enterprise

1.

Number of AI Use Cases Approved for Deployment per Month

Measures how many use cases are not just governed, but actively being deployed. This can be a strong indicator of how governance is enabling responsible innovation at scale.

2.

Growth in Number of AI Use Cases Registered Over Time

This metric can be especially helpful for companies that have just adopted governance as an initial indicator of increased AI adoption and the organizational-wide growing engagement with governance processes.

3.

Time to Production for Approved AI Use Cases

This metric captures how quickly AI use cases move from approval to deployment. Although full deployment may depend on factors beyond governance (like engineering resources or business priorities), companies should still track time to deployment as an indicator of how governance enables downstream responsible deployment of AI solutions.

Operational Efficiency Metrics

This section frames governance as a workflow that must be efficient, not just compliant, to sustain buy-in and accelerate innovation. It stresses that successful governance programs pay attention to how long each step in the process takes so they can remove bottlenecks, reduce friction, and keep reviews flowing smoothly.

4.

Average Time to Review an AI Use Case

Measures the average time it takes to review and approve a use case. This metric is a good indicator of streamlined workflows, stakeholder alignment, and ultimately AI governance maturity. It reflects how quickly governance teams can assess AI use cases, identify risks, and move them through the governance process. It's also valuable to track review time by risk level; for example, if low-risk use cases take disproportionately long review times, it may signal unnecessary friction, while unusually fast reviews of high-risk use cases could point to gaps in scrutiny.

5.

Use Case Approval Efficiency (Approved vs. Submitted Ratio)

In addition to average time to review a use case, throughput can be another useful measure of how effectively use cases move from submission to approval. As governance matures, this ratio should improve indicating better throughput, clear streamlined workflows and stakeholder alignment. Conversely, a declining approval ratio can signal friction and bottlenecks in the review process. A ratio of 70-80% can signal a healthy balance between approvals and rejections.

Risk & Compliance Oversight Metrics

This section positions governance as a way to maintain a strong regulatory and ethical posture while still enabling AI innovation. It notes that organizations face an expanding set of requirements (EU AI Act, GDPR, local and sector regulations), and that without cohesive ethical governance frameworks many will fail to realize the value of their AI use cases.

6.

Percent of AI Use Cases Compliant with Key Regulations

Measures alignment with regulations to ensure compliance with laws such as the EU AI Act, GDPR or US state specific laws such as mandatory bias audits and/or impact assessments under New York City's Local Law 144, Colorado's SB-169 and Texas TRAIGA.

7.

Percent of Mitigation Effectiveness by Risk Type

Average % reduction in risk (inherent vs. residual), tracked by risk category over a rolling time window. NIST AI RMF encourages organizations to track metrics for the effectiveness of risk treatments. This metric tracks impact by assessing whether risks were actually reduced and where governance is delivering results. Example: "We reduced fairness risk by 40% across use cases reviewed in Q2.".

8.

Time to Resolution for Risk/Compliance Issues

In addition to flagging risks, companies also need to track how well they respond to risk and compliance issues. Time to resolution closes the loop between risk identification and risk mitigation. As governance matures, time to resolution should decrease as an indication of both responsiveness and growing governance expertise.

Governance Adoption & Stakeholder Engagement Metrics

This section makes the case that governance has to be an organization-wide, deeply embedded activity involving many teams, not just a central function. It emphasizes the human side of governance: trust, collaboration, and engagement from product, legal, data science, and business stakeholders, all of whom must participate for oversight to be effective.

9.

Percent of Teams Participating in AI Governance

Captures adoption across business units and ensures governance is not siloed. As this percentage increases, it signals that governance practices are becoming increasingly embedded in business operations rather than being managed solely by a central governance team.

10.

Number of Monthly Active Users (Governance Platform MAU)

The more teams are engaged with governance, the more they will turn into governance champions promoting more user engagement. A continuous increase in the monthly active users is a good indication of growth, engagement and adoption of governance.

Bonus Metric

9.

Stakeholder Satisfaction with Governance Process

This metric measures perceived value and usability, in short buy-in, especially from non-governance teams that are critical to governance such as product, legal, data science, and others. Securing buy-in from these teams is a critical prerequisite for successful governance adoption. Because of its importance, governance teams dedicate significant time to fostering buy-in. They conduct awareness-building meetings, address concerns, and offer support, almost functioning like a customer success team. This ongoing engagement is essential to overcoming barriers and ensuring wide adoption across the organization.

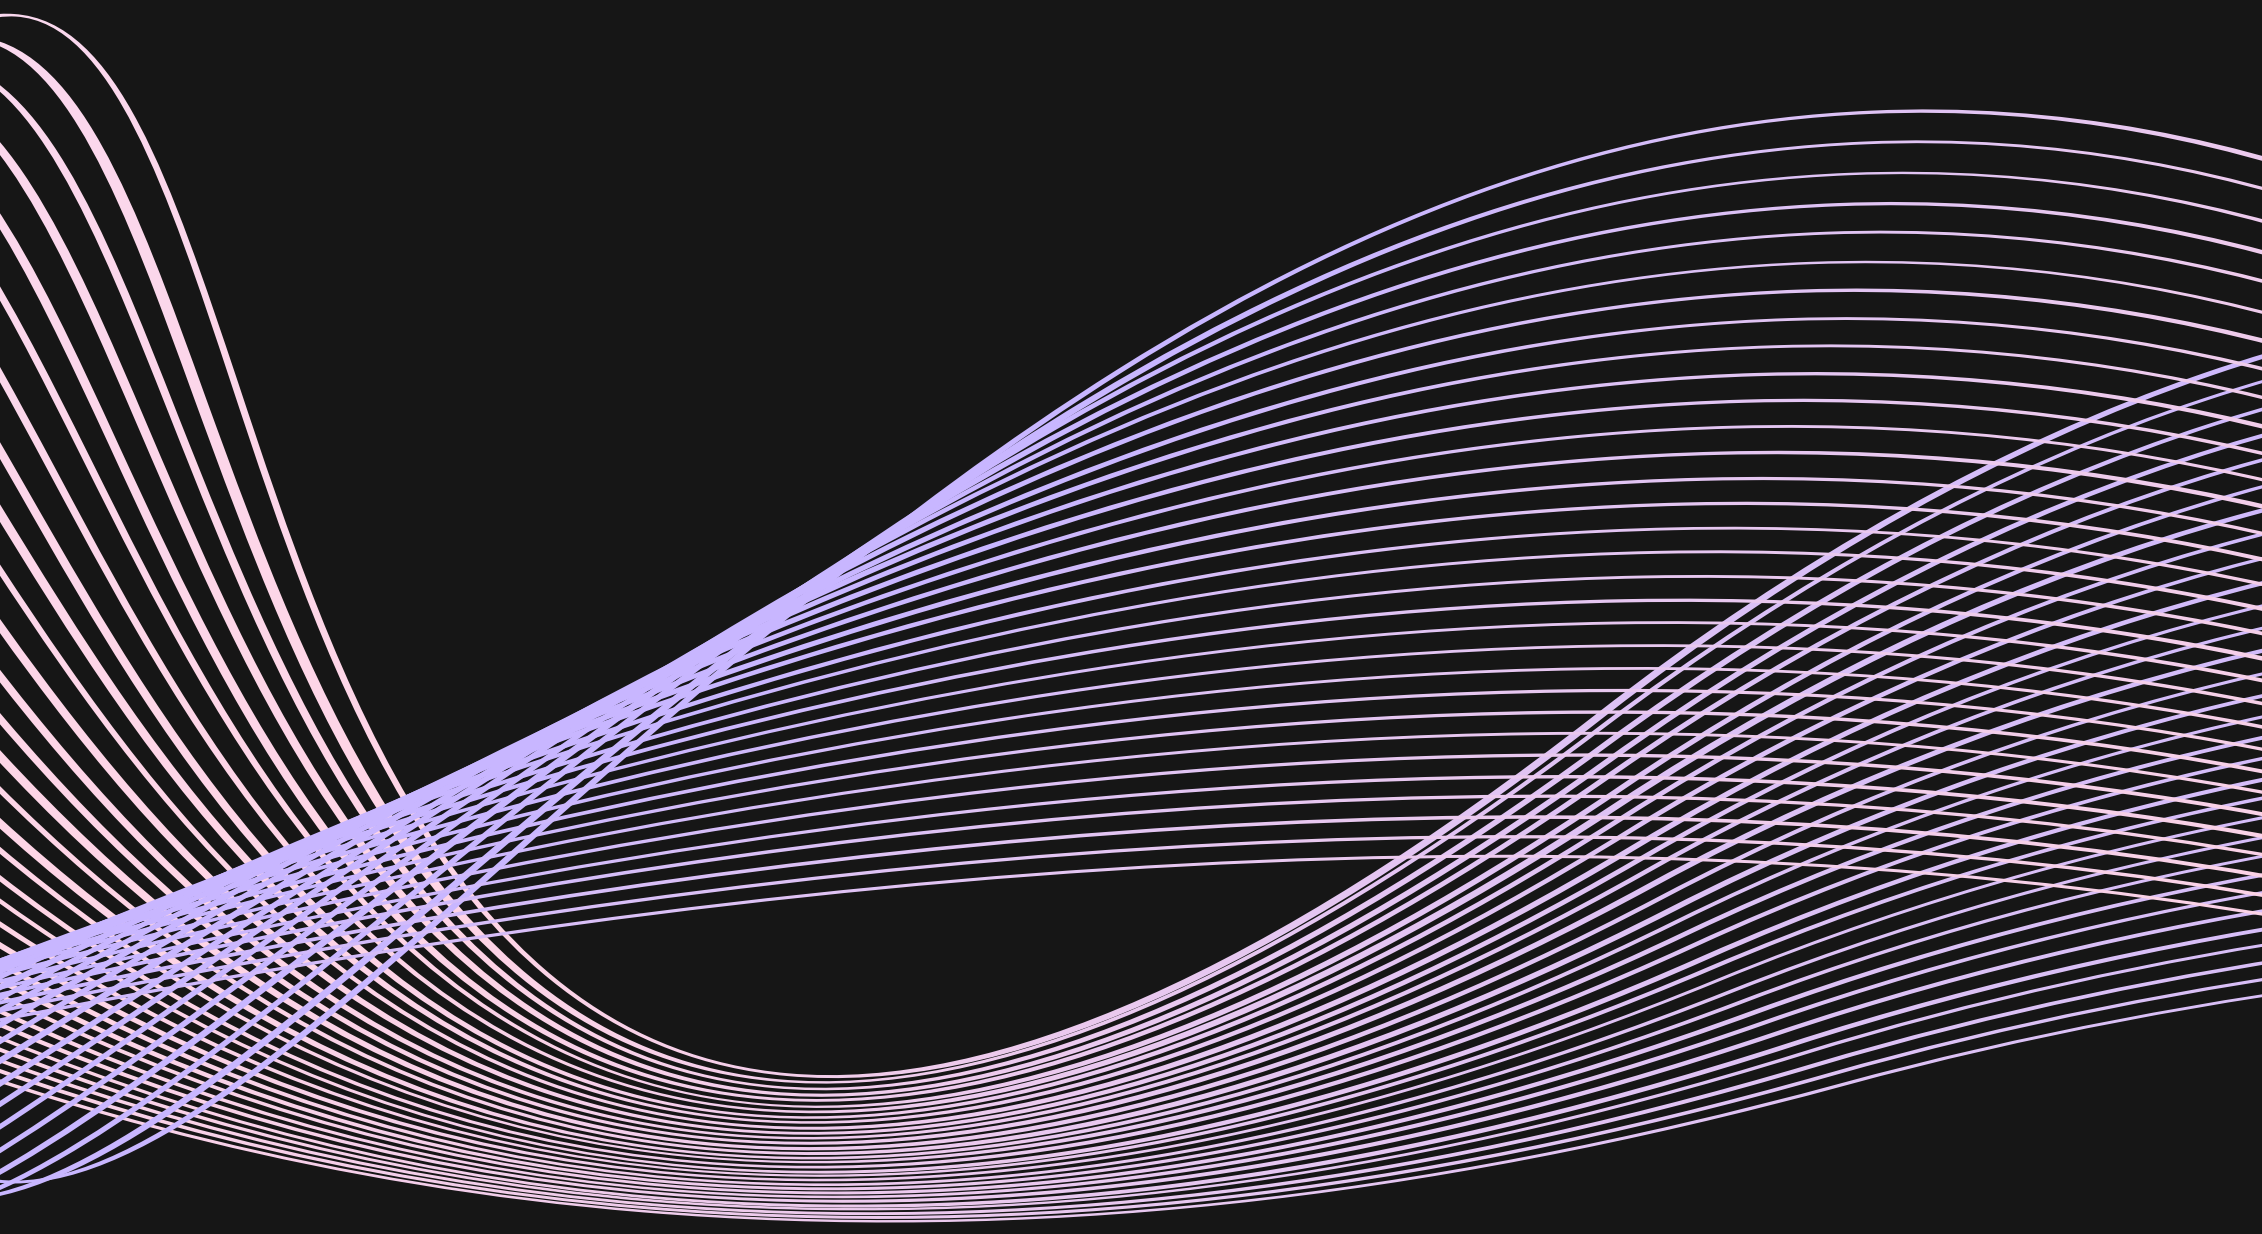

- 8.2

Reporting and Communication

For the Board:

Quarterly dashboard showing:

- [ ] AI portfolio size (number of models, use cases, investment)
- [ ] Governance maturity level (across six dimensions)
- [ ] Key risk indicators (high-risk model count, monitoring coverage, incidents)
- [ ] Regulatory readiness (alignment with key regulations)
- [ ] Strategic value (model deployment velocity, time to market)

For the Executive Team:

Monthly:

- [ ] Governance scorecard (KPIs above)
- [ ] Approval backlog and cycle time
- [ ] Incidents and responses
- [ ] Key risks or escalations

For the Governance Team:

Weekly:

- [ ] Approval queue and status
- [ ] Monitoring alerts
- [ ] Escalations
- [ ] Compliance gaps

For Business Units:

Monthly:

- [ ] Status of their AI initiatives
- [ ] Approval timeline
- [ ] Next steps
- [ ] Support needed

09 •

Conclusion: From Talk to Trust

9.1 • Conclusion: From Talk to Trust

AI governance isn't about compliance theater. It's about competitive advantage.

Organizations that embed governance early will:

- ✓ **Scale faster**
  Approval cycles are 4-6x faster; models go to production in weeks, not months
- ✓ **Avoid expensive incidents**
  Most high-risk models will have been assessed and guardrailed before they cause problems
- ✓ **Win more business**
  Customers and partners increasingly require governance audits; you'll pass where others struggle
- ✓ **Attract and retain talent**
  Technologists want to build AI they can be proud of
- ✓ **Prepare for the future**
  Agentic AI, new regulations, and more complex models are coming; you'll be ready

The window is open now. Regulation is becoming real. Customers are asking. Investors are watching. Boards are paying attention.

> **Navrina Singh, Founder and CEO of Credo AI, says:**
> "Don't think of AI governance as a constraint. Think of it as the fuel that lets you scale. Safety, security, and accountability should just be weaved into how we're building artificial intelligence."

This playbook is your roadmap. Use it to:

- ☐ Build the business case for governance in your organization
- ☐ Align your stakeholders
- ☐ Assess where you are today
- ☐ Design an AI governance program
- ☐ Decide on build vs. buy
- ☐ Measure and communicate success

The next six months will define the next three years. Organizations that move now will own the next wave of AI scaling and become true AI leaders. Those that wait will spend the next three years catching up, evaporating market share by the day.

Visit us: www.credo.ai

Playbook

Appendix A: Governance Framework Mapping

Alignment with NIST AI Risk Management Framework (AI RMF)

Credo AI's governance program aligns with NIST AI RMF across four functions:

- Govern: Leadership, oversight, controls, policies
- Map: Inventory, dependencies, risk classification
- Measure: Continuous monitoring, performance metrics, bias detection
- Manage: Incident response, escalation, remediation

Alignment with ISO/IEC 42001

ISO 42001 AI Management System covers:

- Leadership and governance
- AI risk management
- Human and organizational factors
- Operational controls
- Performance evaluation

Credo AI helps you operationalize each of these.

Alignment with EU AI Act

High-risk AI systems under the EU AI Act require:

- Risk assessment
- Data governance
- Documentation
- Testing and validation
- Human oversight
- Transparency

Credo AI has pre-built compliance templates for the EU AI Act.

Appendix B: Credo AI Budget Justification Framework Disclaimer

Methodology and Limitations

Source of Estimates and Projections

The financial figures, ranges, and projections presented in this AI Governance Investment Framework are derived from a combination of:

1. **Industry Research and Analysis:** Published reports and surveys from leading analyst firms including, but not limited to, Gartner, Forrester Research, IDC, and McKinsey & Company regarding AI governance market trends, enterprise AI deployment challenges, and regulatory compliance costs.

2. **Regulatory Framework Analysis:** Review of penalty structures and compliance requirements from the EU AI Act, U.S. sector-specific regulations (financial services, healthcare, government contracting), state-level AI legislation, and international AI governance standards including ISO/IEC 42001.

3. **Market Benchmarking:** Analysis of publicly available information regarding AI governance platform pricing, enterprise software implementation costs, professional services rates, and typical enterprise technology team compensation ranges.

4. **Industry Incident Data:** Published reports of AI-related compliance violations, model failures, algorithmic bias incidents, and associated remediation costs from media sources, regulatory enforcement actions, and corporate disclosure documents.

5. **Client Engagement Experience:** Aggregated, anonymized observations from Credo AI's advisory engagements with enterprise organizations across multiple industries, reflecting common patterns in AI governance challenges, implementation timelines, and resource requirements.

- *Appendix B: Credo AI Budget Justification Framework Disclaimer*

Representative Estimates, Not Guarantees

All cost estimates, return projections, payback periods, and value calculations presented in this framework are **illustrative examples based on representative market data** and should not be interpreted as guarantees, predictions, or promises of specific outcomes for any particular organization. Actual results will vary significantly based on numerous factors including but not limited to:

- Organization size, industry sector, and geographic footprint
- Complexity and maturity of existing AI portfolio
- Regulatory environment and compliance obligations specific to the organization
- Current state of governance processes and technical infrastructure
- Quality of implementation and organizational change management
- Market conditions and competitive dynamics
- Timing of regulatory enforcement actions
- Effectiveness of risk mitigation measures deployed

Individual Assessment Required

Organizations should conduct their own detailed financial analysis, risk assessment, and cost-benefit evaluation based on their specific circumstances, requirements, and risk tolerance. The ranges and examples provided are intended to reflect market variability and should be customized using organization-specific data.

No Professional Advice

This framework is provided for informational and educational purposes only and does not constitute financial advice, legal advice, accounting advice, investment advice, or professional consulting services. Organizations should consult with qualified experts before making investment decisions related to AI governance.

- *Appendix B: Credo AI Budget Justification Framework Disclaimer*

Forward-Looking Statements

Certain statements in this framework regarding regulatory trends, market evolution, competitive dynamics, and projected outcomes constitute forward-looking statements that involve risks and uncertainties. Actual regulatory developments, market conditions, and organizational outcomes may differ materially from those described. Regulatory frameworks such as the EU AI Act are subject to interpretation, implementation guidance, and enforcement practices that may evolve over time.

No Liability

Credo AI and its affiliates make no representations or warranties, express or implied, regarding the accuracy, completeness, or reliability of the information contained in this framework. Use of this framework and any decisions made based on its contents are at the user's sole discretion and risk. Credo AI shall not be liable for any direct, indirect, incidental, consequential, or special damages arising from the use of or reliance on this framework.

Copyright and Usage

For customized financial analysis specific to your organization's circumstances, contact Credo AI's advisory services team.

Last Updated: December 2025

- **References**

Footnotes

1. McKinsey Global Survey on AI, 2025.

https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

2. Navrina Singh, Founder & CEO Credo AI, Inside the ICE House, July 2025.

https://www.youtube.com/watch?v=pDCXcD-hR4s

3. McKinsey Global Survey on AI, 2025.

4. Gartner Market Guide for AI Governance Platforms, 2025.

https://www.credo.ai/blog/credo-ai-recognized-in-the-gartner-r-market-guide-for-ai-governance-platforms-2025

5. Forrester Wave: AI Governance Solutions, Q3 2025.

https://www.credo.ai/forrester-wave

6. Info-Tech Research Group, AI Trends 2026 Report.

https://www.morningstar.com/news/pr-newswire/20251117to26172

7. Gartner, cited in Credo AI case studies and market guidance.

8. Credo AI customer case study, AdeptID.

https://www.credo.ai/blog/credo-ai-named-a-leader-in-2025-ai-governance-solutions-report

9. Navrina Singh, Credo AI Blog, "Accelerating Enterprise AI Innovation with Governance into 2025," January 2025.

https://www.credo.ai/blog/from-pioneering-partnerships-to-groundbreaking-products-accelerating-enterprise-ai-innovation-with-gov

10. EU AI Act enforcement fines can exceed 6% of global revenue for high-risk violations.

• References

11. Credo AI case study, global CPG company.

https://www.credo.ai/casestudies/cpg

12. Accenture and Stanford Institute for Human-Centered AI, AI Index 2025 Annual Report.

https://aiindex.stanford.edu

13. Navrina Singh, TIME 100 AI, August 2025.

https://time.com/collections/time100-ai-2025/7305855/navrina-singh/