

The Practical AI Governance Guide

Implement AI Governance to
Power Trustworthy AI at Scale

Governance should
not be seen as a
barrier to innovation.

If done right, it will
enable it.

Navrina Singh

Founder & CEO, Credo AI

CONTENTS

04 AI Governance in a nutshell

06 The necessity of AI Governance

09 Retrofitting tools for AI Governance

15 Establishing accountability for AI Governance

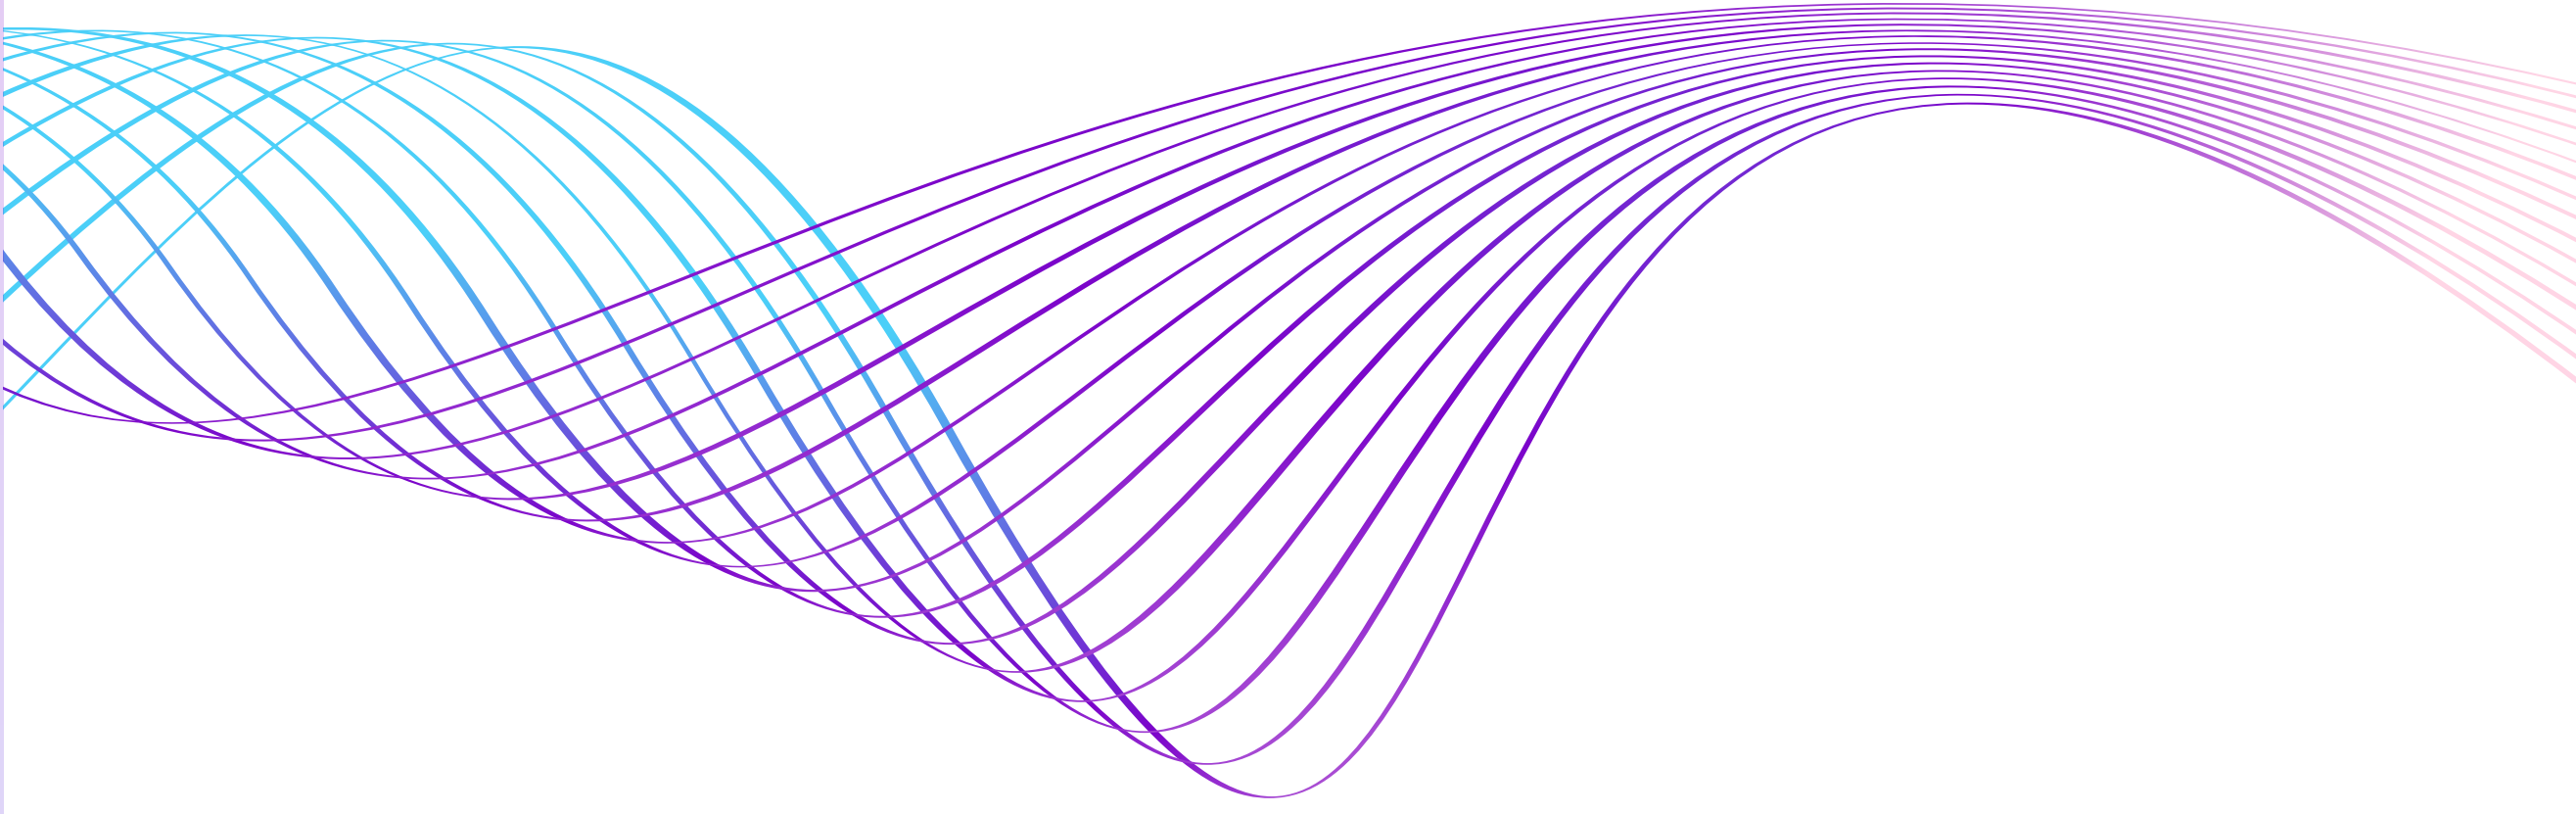
17 The three stages of AI Governance

22 Conclusion

AI Governance is the framework of policies and practices ensuring ethical, transparent, and accountable AI development and use. Explore the history of AI Governance and its importance to future businesses.

01

AI Governance in a nutshell



There's a lot of talk about 'AI governance' these days—and for good reason. From executives to boards to shareholders, adopting AI and Generative AI at scale is becoming a company-wide mandate. Meanwhile, enterprises are expected to keep pace with the emergence of AI-specific regulations like the EU AI Act and standards like ISO 42001. These standards are emerging in tandem (and often, in response to) headline-making AI pitfalls that can seriously impact company revenue and operations.

AI governance is the bridge to confidently adopting AI at scale while keeping AI within the guardrails of regulation, company best practices, and public safety.

However, the first practical steps to AI governance can get lost in the maelstrom of AI noise. New AI technologies emerge every day, and it's easy to get overwhelmed and confused by the sheer number of tools that offer similar functions in the market.

At Credo AI, we have extensive experience implementing AI governance at organizations, from SMEs to government agencies and Fortune 100 enterprises. Our goal with this guide is to cut through the noise and offer our practical AI governance knowledge so you can up-level your organization in AI maturity.

What is AI Governance?

AI Governance is the practice of ensuring that an organization's development and use of artificial intelligence systems is safe, secure, compliant, transparent, auditable, and human-centered. Responsible AI is the outcome.

The business benefit of AI governance is ensuring that you can adopt AI quickly and safely.

AI governance is needed to ensure AI systems are fair, safe, and respect people's privacy. It helps prevent misuse, biases, and errors that can harm individuals and society. Explore this chapter to understand each category where proper AI Governance makes a difference.

02

The necessity of AI Governance

There are multiple benefits to governing AI appropriately, but the ultimate reason is for control. Governance is often talked about in terms of reducing downside risk, but having an AI Governance approach is essential to steer AI development and use at an organization regardless of your purpose.

Maximizing efficiency and AI benefits is also achieved through good governance.

Below are listed some of the top reasons AI governance is helpful.



Risk Mitigation

- › Identifies and addresses potential risks associated with AI systems, such as bias, fairness, transparency, and privacy concerns
- › Reduces the likelihood of adverse events, unintended consequences, and reputational damage
- › Helps organizations avoid costly legal and regulatory penalties




Compliance and Trust

- › Ensures alignment with external requirements, standards, and regulations
- › Demonstrates commitment to responsible AI development and deployment
- › Enhances trust among stakeholders, including customers, employees, partners, and regulators
- › Improves public perception and strengthens brand reputation




Improved Decision-Making

- › Provides a framework for consistent and transparent AI decision-making
- › Ensures AI systems are aligned with organizational values, goals, and ethical principles
- › Enables organizations to make informed, data-driven decisions while considering potential impacts



Operational Efficiency

- › Streamlines AI development and deployment processes through clear policies, procedures, and standards
- › Reduces the risk of project delays or failures due to unforeseen ethical, legal, or technical challenges
- › Facilitates collaboration and communication among teams, breaking down silos and improving efficiency



Innovation and Competitive Advantage

- › Enables organizations to confidently explore and deploy innovative AI solutions
- › Provides a foundation for responsible AI experimentation and iteration
- › Allows organizations to stay ahead of the curve in adopting AI technologies while managing associated risks



Continuous Improvement

- › Establishes a framework for ongoing monitoring, evaluation, and refinement of AI systems
- › Supports data-driven optimization and updates based on real-world performance and feedback
- › Fosters a culture of continuous learning and improvement, ensuring AI systems remain effective and aligned with evolving needs and requirements

By implementing effective AI governance, organizations can unlock AI's full potential while mitigating risks, ensuring compliance, fostering trust, and driving innovation.

A robust governance framework provides the foundation for responsible, ethical, and successful AI adoption, enabling organizations to harness AI's power for competitive advantage and positive impact.

As AI quickly progresses and evolves, proper tools are a necessity for governing AI according to the newest laws and regulations. Why is managing AI so different from other processes? Explore this question and its answer in the following chapter.

03

Retrofitting tools for AI Governance

Can I retrofit a tool my organization already uses to address my AI governance needs?

The answer is simple: no. The reality is that if your organization wants to manage AI risk and comply with AI-specific laws and regulations, you will need a new process that deals with the specific challenges and requirements of AI governance.

What differentiates AI governance from:

data privacy,

cybersecurity,

and other

existing

governance

processes?

Let's discuss. →

The bigger picture: AI governance includes governance of datasets, models, and AI use cases.

AI governance is about establishing oversight and guardrails over all of the components of AI—not just models but also the datasets that are used to train and test those models. Perhaps the most critical component of an AI system for the purpose of governance, however, is the use case—the context of where, how, and why the model(s) will be deployed. AI governance is the “control center” that oversees everything!

Data privacy and model risk management in existing governance processes do not effectively account for all of the entities—use cases, models, and datasets—involved in AI governance. Data governance and privacy workflows focus on data, helping organizations manage data usage and access, consent, and sharing; not AI.

Model risk management and MLOps workflows manage risk related to models, with a focus on technical risks related to robustness and performance issues. However, neither of these two separate practices addresses use cases, nor are they set up to bring together these different entities to provide a “big picture” view of risk and compliance of an AI system.

For effective AI governance, you must have a system of record that tracks use cases, models, and datasets for every AI system or product in use across your organization, and the use case should define governance requirements for the models and datasets based on laws, regulations, and standards. This can only be done with specific AI governance processes.

Adaptability: Enterprise AI governance programs must support governance of highly flexible 3rd party AI models and applications.

Before the advent of large language models, many organizations were focused on building all or many of their own models in-house; these classical machine learning models were trained on proprietary data, purpose-built for the specific application into which they were going to be deployed, and organizations had a high level of control over these models from a governance perspective.

Today, however, most organizations are planning to leverage third-party models for generative AI use cases, and these third-party models are not purpose-built but instead designed to be highly flexible and deployed in many different contexts.

For example, GPT-4 could be used to power an internal document search tool for employees, and it could also be put inside an application designed to help doctors evaluate patient data and make diagnostic decisions; even though GPT-4 is the underlying model in both of these use cases, the governance requirements and appropriate guardrails are very different for these two systems.

Governing highly flexible third party models—ensuring that their risks are effectively managed, and that they are deployed in a compliant and safe manner—presents a new set of challenges for enterprises, which traditional model risk management and technical governance workflows are not equipped to address.

Model risk management and MLOps-driven governance processes are not designed to support defining different governance requirements for the same model depending on where and how it's being used. They are also not designed to support managing governance of models that are developed externally; these technical governance approaches often require that you have full control over model design, development, and training—which is not the case when you're working with a third-party LLM.

Existing governance processes that are focused only on managing the **risks of AI** that have been developed internally are not going to adapt well to the new needs of AI governance in a generative AI world.

Instead, you need a flexible governance process that enables you to assess and mitigate risks even when you are leveraging a model that was developed by somebody else; and that enables you to define governance requirements based not just on which model is being used, but also the context—where, how, and why that model will be deployed.

Tech Meets Business: AI governance must be integrated with technical and business infrastructure.

An effective AI governance process defines governance requirements for your AI use cases, models, and datasets; in order to understand whether those requirements are being met by an AI system, you need to evaluate both technical and documentation evidence.

A mature AI governance process is tightly integrated with the tools and workflows through which your AI teams are already generating evidence. This reduces the manual effort of providing evidence for governance, which in turn supports adoption of the AI governance process throughout the enterprise. For example, your AI governance tooling might be tightly integrated with Confluence and Jira, where teams document the AI systems they're building. It could also be connected to your AI deployment infrastructure, where model monitoring metrics are calculated. This setup allows the tooling to automatically extract relevant evidence from these systems. Technical teams don't have to manually upload documentation or metrics to governance teams for review – the information just flows, leading to a seamless oversight process.

Because AI governance touches such a wide set of entities, your AI governance program is going to need to be tightly integrated with a wide set of tools. This includes more than just your technical infrastructure, where model assessments and monitoring are taking place. It also encompasses existing governance tools like your data privacy management system, data governance tooling, and cybersecurity systems, as well as project management, where AI use cases are being defined and put on the roadmap.

Again, existing governance processes have their strengths and weaknesses when it comes to integrations—model risk management is often tightly integrated with technical model assessment tooling, while data governance is often integrated with your data infrastructure and systems of record. But none of your existing governance processes are going to have the breadth of integrations across all of the tools and evidence needed to bring insight into AI risk and compliance.

It is all about AI: AI governance requires AI-specific risk and compliance content—risk scenarios and controls.

Another thing that sets AI governance apart from your existing governance processes are the kinds of risks and compliance requirements that your AI governance program needs to address, which are—this may seem obvious—AI-specific.

Your existing security, data privacy, and enterprise risk management workflows are not going to effectively account for AI-specific risks like the risk of adversarial attacks on an AI system or the risk of generating copyright-infringing content through the use of an LLM. Your existing compliance processes are **not** going to enable you to evaluate **compliance against emerging AI-specific laws** like the EU AI Act, New York City LL144, or Colorado SB 21-169.

A new library of risk scenarios and controls is required to support your AI governance efforts—one that is tailored to the specific risks of AI systems. This library of risk scenarios and controls must cover a comprehensive set of possible risks, including fairness, robustness, security, privacy, and transparency. Narrowly focusing on one risk area—such as privacy or security—will only address one of the many different kinds of risks that your AI governance program must tackle.

This is again where your existing governance processes are going to fall short; your data privacy program cannot effectively support management of the risk of hallucinations in your LLM-based applications, and your cybersecurity program cannot help you evaluate the risk of harmful bias in your AI-powered applications. Bringing together these different risks into a single pane of glass is critical for ensuring comprehensive risk management throughout the AI lifecycle.

AI governance is quite different from the existing workflows that your organization already has in place for data governance, privacy, and cybersecurity. The tools that you're using for these existing governance processes each have their own strengths—and each have their relevant intersections into the AI governance process—but none is going to be able to fully support you in your AI governance journey.

Establishing accountability for AI governance is all about setting clear rules and responsibilities for those creating and using AI systems. It ensures accountability, but how do we clearly define this within an organization? Explore this topic in the following chapter.

04

Establishing accountability for AI Governance

The success of AI governance within an organization relies on the collaborative efforts of various stakeholders, but the **AI Governance Custodian** plays a central role. The AI Governance Custodian, often supported by an AI Governance Team, is responsible for the day-to-day management and execution of the governance framework.

This individual works closely with the AI Governance Board or Committee, which consists of senior executives and cross-functional leaders who oversee the implementation and enforcement of the AI governance framework, set strategic direction, and make critical decisions related to AI adoption and risk management.

AI Governance Custodian



To effectively implement AI governance, the AI Governance Custodian and their team require the support and collaboration of various departments and experts throughout the organization. Data scientists, engineers, and technical experts are responsible for designing, developing, and deploying AI systems in accordance with the governance framework, while legal, risk, and compliance professionals provide guidance on regulatory requirements and ensure adherence to applicable laws and industry standards. Ethics specialists contribute insights on the ethical implications of AI systems, helping to align AI development with organizational values and societal norms. Business leaders and domain experts from departments such as marketing, finance, and operations offer valuable knowledge to ensure AI systems support business objectives and deliver value.

Fostering a culture of responsible AI and engaging all employees through training and awareness programs is crucial for the success of the AI governance framework. As this is a continuous process, ensuring that directly responsible individuals are identified and empowered to forward AI governance is the most critical early step in maturing an AI Governance process.

Let's delve into the distinct stages of AI Governance and the complexities involved in managing AI systems. We discuss the evolving strategies needed to ensure responsible oversight as AI technology continues to advance and integrate into organizations.

05

The Three Stages of AI Governance

The ideal AI governance workflow consists of three primary stages. Each stage is critical in ensuring the responsible development, deployment, and monitoring of AI systems within the enterprise. The three stages are:

The Three Stages of AI Governance

STAGE 1

Organization-level Governance

During this stage general policies and processes are established to enable AI governance at an organization-level.

STAGE 2

Intake

During this stage AI systems are tracked and given an initial risk & compliance evaluation, and a “governance plan” is developed, establishing the requirements for development and deployment of the system.

STAGE 3

Governance

During this stage the governance plan is executed. Evidence is gathered to establish appropriate execution of the plan, which may include technical assessments or establishing oversight and accountability processes and the AI system is reviewed for adequacy. Post-deployment the system is monitored and relevant stakeholders are kept informed.

Let’s

discuss

in

detail →

Organization-level Governance

Organization-level Governance focuses on establishing the foundational principles, frameworks, team responsibilities, and tooling necessary to support AI governance across the organization. This stage sets the tone for responsible AI practices and creates a shared understanding of the AI lifecycle and its associated risks and benefits.

Key steps include:

- Establishing AI Governance Principles aligned with the organization's values and goals
- Creating a shared understanding of AI lifecycle stages
- Designing a comprehensive Governance Framework
- Assigning Team Responsibilities, including an AI Governance Committee and an AI Governance Team
- Developing or Procuring Governance Tooling (e.g., an AI-focused GRC workflow tool or platform) to support efficient governance processes
- Providing Training and Awareness to foster a culture of responsible AI

AI Intake

AI Intake is the stage where individual AI systems are tracked, evaluated for initial risk and compliance, and assigned a governance plan. This stage is crucial for identifying and prioritizing high-risk systems, ensuring compliance with regulations, and developing tailored governance plans based on the specific use case.

Key steps include:

- **Ingestion:** Tracking all AI systems and their lifecycle stages
- **Procurement:** Connecting the procurement process for external AI applications to internal use cases
- **Initial Risk Evaluation:** Estimating the inherent risk of the use case to inform prioritization and governance planning
- **Identifying Compliance Requirements:** Applying "common ground" requirements and defining operationalization of regulations
- **Defining the Governance Plan:** Creating a plan for oversight and safeguards based on the use case's risk and impact

Governance

Governance is the stage where the governance plan is executed, evidence is gathered, and the AI system is reviewed for adequacy. Post-deployment, this stage involves ongoing monitoring, auditing, and reporting to maintain oversight and ensure the system continues to operate as intended. This stage also includes processes for updating risk assessments and governance plans, as well as procedures for rolling back or shutting down problematic systems when necessary.

By understanding and implementing these three stages, organizations can establish a comprehensive and adaptable AI governance workflow that effectively manages risks, ensures compliance, and maximizes the benefits of AI systems.

Key steps include:

- **Gathering Evidence:** Integrating existing processes, documenting decisions, implementing risk-mitigating controls, ensuring human oversight, and conducting technical evaluations
- **Generating Reports & Dashboards:** Creating compliance reports, audits, and internal dashboards to reflect gathered evidence
- **Ongoing Monitoring and Audits:** Applying MLOps monitoring, periodic evaluations, red-teaming, incident reporting & response, and compliance reporting
- **Updating Risks and Governance Plan:** Incorporating new information and updating risk assessments and mitigation actions
- **Roll Back or Shut-Down:** Implementing processes to shut down or roll back problematic systems

The main goal of AI governance is to ensure ethical, fair, and safe AI use. For organizations, it's crucial as it builds trust, reduces risks, and ensures compliance with regulations. This leads to sustainable innovation, protects reputation, and enhances customer and stakeholder confidence in AI applications. Credo ai welcomes you to explore what this means for your organization.

06

Conclusion

Implementing an effective AI governance workflow is essential for organizations to successfully navigate the complexities and risks associated with AI systems while maximizing their potential benefits. By following the three-stage approach outlined in this playbook –

Organization-level Governance, AI Intake, and Governance – enterprises can establish a comprehensive framework that ensures responsible, ethical, and compliant AI development and deployment.

However, it is crucial to recognize that AI governance is not a one-time exercise, but rather an ongoing process that requires continuous adaptation and refinement. The rapid pace of AI technological advancements, coupled with evolving societal expectations and regulatory landscapes, creates a dynamic environment that demands agility and flexibility in governance practices.

To truly succeed in this ever-changing landscape, organizations must embrace a culture of adaptability and commit to iterative improvements in their AI governance approach. This means regularly reassessing governance principles, frameworks, and processes to ensure they remain relevant and effective as new challenges and opportunities emerge.

It also involves staying attuned to industry best practices, regulatory developments, and stakeholder expectations, and being willing to adjust course when necessary.

Furthermore, fostering a culture of responsible AI and empowering teams with the knowledge, tools, and resources they need to make ethical decisions is essential for embedding adaptability into the fabric of the organization. By encouraging open communication, collaboration, and continuous learning, enterprises can create an environment that is receptive to change and well-equipped to handle the evolving demands of AI governance.

The organizations that prioritize adaptability and view AI governance as a dynamic, iterative process will be best positioned to harness the transformative potential of AI while mitigating risks and maintaining the trust of their stakeholders. By embracing the need for adaptability and committing to ongoing refinement of their governance practices, enterprises can not only keep pace with the rapid evolution of AI but also set the standard for responsible and ethical AI adoption in their respective industries.

About Credo ai

Credo AI is on a mission to empower organizations to responsibly build, adopt, procure, and use AI at scale. Its pioneering SaaS AI governance platform helps organizations measure, monitor, and manage AI risks while ensuring compliance with emerging global regulations and standards, like the EU AI Act, NIST, and ISO.



What we do

Credo AI is the leading governance platform for building compliant, secure, safe, auditable, fair, and human-centered AI. Our mission is to empower organizations to deliver Responsible AI (RAI) at scale through context-driven, comprehensive, and continuous governance, reducing the risk of deploying AI systems while allowing companies to harness AI's benefits.

AI use case inventory and oversight

Credo AI simplifies and streamlines the AI governance process. This begins with AI Use Case intake, which allows you to automatically route use cases to the right governance workflows and oversight based on information about where, how, and why an AI system will be deployed.

AI risk identification and mitigation

Credo AI automatically identifies relevant risks associated with any use case across internally developed or third-party AI within an organization. Credo AI then maps your system to risk-mitigating controls designed to address the identified risks, including technical risks like bias, legal risks like copyright infringement, and operational risks like cybersecurity vulnerabilities.

AI policy intelligence and compliance





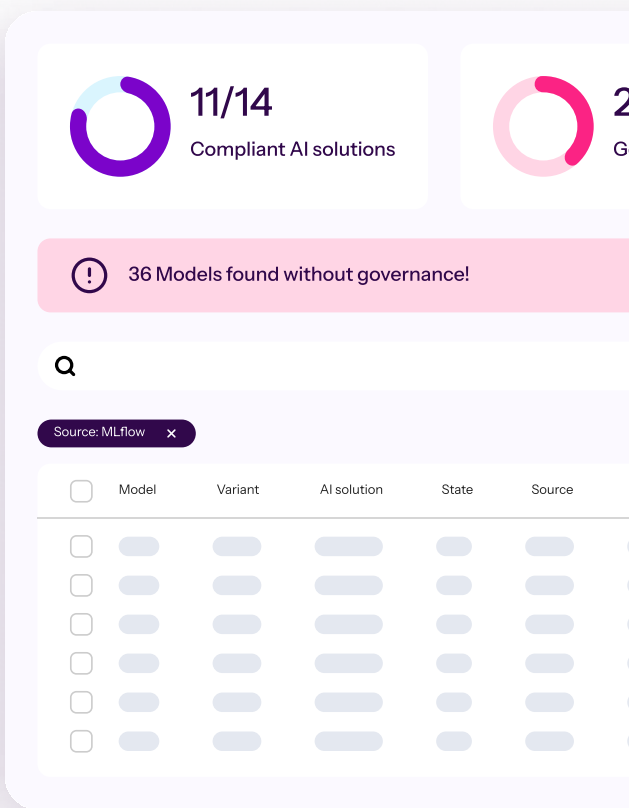
Credo AI's Policy Intelligence Engine translates AI policy into actionable requirements for teams that are developing and deploying AI. Credo AI Policy Packs help organizations comply with laws like NYC Local Law 144, or the EU AI Act, as well as cross-check alignment to industry standards, like the NIST AI Risk Management Framework or ISO 42001. Custom policy packs can also be created to operationalize a company's internal policies.

Unite AI stakeholders with streamlined governance workflows

The Credo AI platform unites AI stakeholders across Legal & Compliance, Privacy & Security, Risk, Data Science, Engineering, and Business teams using AI, to simplify decision making about risky AI systems. This includes customizable intake questionnaires, streamlined review capabilities, as well as aggregated dashboards of AI risk levels and progress to complete AI risk mitigation.

The orchestration layer for AI governance

AI governance requires gathering data and inputs from many different stakeholders, across many different tools and processes in your organization. Credo AI is designed to integrate with your technical AI infrastructure and your GRC tools and processes to serve as a centralized repository for governance-relevant data across the enterprise.





Ready to start your AI Governance journey?

Reach out to our team at:

demo@credo.ai