Create an AI governance program from scratch

Optro

Table of contents

Who should use this guide

As businesses adopt AI in their operations and day-to-day processes, the need for strong AI governance has grown. The rapid pace of AI adoption across industries is driving the urgent need for robust internal policies and controls, with emerging frameworks like the NIST AI Risk Management Framework and regulations such as the EU AI Act providing valuable guidance to organizations as they begin to build effective AI governance processes. This guide provides best practices for establishing an effective AI governance program that helps to ensure the responsible and ethical use of AI in your organization. By following these structured steps, you'll build a governance foundation that not only mitigates risks but also positions your organization to harness AI's transformative potential while maintaining stakeholder trust and regulatory compliance.

Here's what you'll need to get started:

- ☑ Your in-scope AI models and use cases
- ☑ A list of AI frameworks and/or regulations that may be in-scope

The importance of a strong AI governance program

As teams benefit from implementing AI throughout their business processes, InfoSec, GRC, and Legal professionals must stay ahead of new AI-related risks through proactive risk management that drives accountability and transparency. A strong AI governance program is essential for helping businesses establish responsible and ethical AI practices.

Beyond risk mitigation, effective AI governance creates a competitive advantage by building stakeholder trust, ensuring regulatory compliance, and enabling faster innovation through clear guardrails. It's also rapidly becoming a critical evaluation criterion as InfoSec teams assess AI-related risks in potential vendor solutions. By establishing structured oversight without stifling creativity, organizations can confidently scale AI adoption while maintaining alignment with their core values and strategic objectives.

Creating an AI governance program from scratch

STEP 1

Identify relevant AI frameworks and regulations

Why this step is important: Your relevant AI frameworks and regulations will serve as the baseline for governance you will manage and the standards your AI governance program will adhere to.

Most organizations find it helpful to align with AI-specific InfoSec frameworks and/or requirements when starting to build out their AI governance program. Many may start with industry-recognized frameworks such as ISO 42001 and the NIST AI Framework, including its companion piece the NIST AI RMF Playbook. If your organization is building its own large language model (LLM), it may be helpful to consider the requirements of the EU AI Act because it establishes comprehensive, risk-based standards that will likely influence global AI regulation. Even for organizations without immediate EU exposure, these requirements provide a robust blueprint for responsible AI development, addressing critical areas such as model transparency, data governance, human oversight mechanisms, and systematic risk assessment protocols that can help future-proof your governance program against evolving regulatory landscapes worldwide.

STEP 2

Create an AI acceptable use policy

Why this step is important: Your AI acceptable use policy will shape your AI governance program. The drafting of this policy typically involves collaboration across several departments, including the Legal team or General Counsel, IT or InfoSec teams, Compliance, Risk Management, Data Privacy Officers, and the AI Governance team or Ethics Committee. The specific team that leads this effort may vary from organization to organization.

An AI Acceptable Use Policy (AUP) is a formal document that defines the approved ways employees and stakeholders can use artificial intelligence systems within an organization, outlines prohibited activities, and establishes consequences for non-compliance. An AUP is designed to ensure AI is used legally, ethically, and in alignment with organizational values.

When designing your AI AUP, the following are some fundamental elements to consider:

Purpose and scope
- Clear definition of which AI systems and applications are covered
- Explanation of why the policy exists and its organizational objectives
- Specification of who must follow the policy (employees, contractors, partners)

Authorized uses
- Specific business purposes for which AI can be used
- Required approval processes for new AI applications
- Conditions under which AI systems can access company data

Prohibited activities
- Explicit prohibition of discriminatory or biased AI applications
- Restrictions on using AI to circumvent human oversight in critical decisions
- Ban on using AI to process sensitive data without proper safeguards

Data governance requirements
- Guidelines for data quality and appropriateness for AI training
- Protocols for securing and anonymizing data used in AI systems
- Data retention and deletion requirements

User responsibilities
- Required training before using AI systems
- Documentation and transparency obligations
- Duty to report issues or unexpected behaviors

Oversight mechanisms
- Monitoring and auditing procedures
- Review cycles and compliance verification
- Roles responsible for policy enforcement

Consequences and accountability
- Clear escalation paths for violations
- Defined consequences for policy breaches
- Process for reporting concerns or seeking clarification

Special considerations
- Additional requirements for high-risk AI applications
- Industry-specific compliance considerations
- Procedures for handling AI-generated content

STEP 3

Create a centralized repository for your AI models

Why this step is important: Without a consolidated view of AI models in use within your organization, AI deployments can proliferate across departments with inconsistent oversight, creating significant compliance gaps and security vulnerabilities. Establishing a centralized repository transforms AI governance from **reactive to proactive** by providing comprehensive visibility into your AI ecosystem, revealing which models have undergone proper review, which require remediation, and which are still pending approval.

Implementation approach: Having established your AI governance foundations through identifying relevant frameworks (Step 1) and developing your Acceptable Use Policy (Step 2), creating a centralized AI model repository becomes a natural progression in your governance journey. While spreadsheets might suffice initially, purpose-built GRC or AI governance platforms offer significant advantages by:

- Automating the mapping between your AI framework requirements and specific controls
- Enabling efficient workflow management for model review and approval processes
- Providing real-time visibility into compliance status across your AI inventory
- Supporting dynamic risk scoring based on model characteristics and use contexts
- Facilitating systematic evidence collection for both internal and external audits

This repository should document each AI model's purpose, underlying training data, performance metrics, potential biases, deployment scope, and risk assessment results. By incorporating these elements into your existing GRC infrastructure, you can create a sustainable framework that evolves alongside your AI adoption, ensuring governance becomes an enabler rather than an obstacle to responsible innovation.

STEP 4

Create a centralized repository for your AI use cases

Why this step is important: It is important to keep track of all the different ways your organization is leveraging AI and the teams that are using it. Certain teams may have different privacy, security, and performance concerns. A centralized repository makes it easy to manage these different concerns and interact with your stakeholders if they need to add additional use cases.

Implementation approach: Make a central list of all your AI use cases within your organization and create a centralized place for people to input their use case requests. This is important for:

1. **Comprehensive oversight:** Without a complete inventory, AI applications may proliferate across departments without proper governance, creating potential regulatory blind spots and security vulnerabilities.
2. **Risk stratification:** Different AI applications carry varying levels of risk based on their purpose, data inputs, and deployment context. An inventory allows you to categorize use cases by risk level and allocate governance resources proportionally.
3. **Regulatory preparedness:** As AI regulations evolve globally, a centralized inventory enables rapid assessment of compliance impacts and facilitates appropriate response planning.
4. **Resource allocation:** Understanding the full scope of AI initiatives helps leadership make informed decisions about technology investments, talent acquisition, and control implementation.
5. **Use case optimization:** Identifying similar AI applications across departments creates opportunities for standardization, reducing redundancy and enabling best practice sharing.

Key components include:

- **Business function/department** owning the AI use case
- **Purpose and scope** of the AI application
- **Data types** being processed (especially sensitive categories)
- **Integration points** with other systems
- **Impact level** on business operations or stakeholder decisions
- **Deployment status** (planning, development, testing, production)
- **Review history** and approval documentation
- **Risk assessment** results and remediation plans

Implementation tips

- Leverage your GRC platform or dedicated AI governance tool to streamline inventory management
- Implement a user-friendly intake process for new AI initiatives to encourage compliance
- Connect your inventory to your change management process to ensure ongoing accuracy
- Establish clear ownership and review cycles to maintain inventory integrity
- Create dashboards for leadership visibility into AI adoption patterns and associated risks

This inventory will become the foundation for implementing proportionate controls and will serve as crucial evidence during regulatory examinations or audit reviews.

STEP 5

Perform a risk assessment

Why this step is important: Performing a risk assessment is important to help your team better understand the risks around your AI models and use cases. This helps to ensure your organization is properly implementing any relevant mitigating controls in your environment and tracking any necessary mitigation plans.

Implementation approach: Your risk assessment should align with your organization's risk tolerance and risk appetite around AI. Start by:

1. **Defining your risk criteria:** Establish clear parameters for evaluating AI risks based on your organization's risk appetite. Consider factors such as:
   - Potential financial impact
   - Regulatory compliance implications
   - Reputational consequences
   - Operational dependencies
   - Privacy and security considerations
   - Ethical concerns
2. **Applying a structured risk framework:** Leverage established risk methodologies tailored for AI applications:
   - Identify threats and vulnerabilities specific to each AI use case
   - Assess the likelihood and potential impact of adverse scenarios
   - Evaluate inherent risk levels before controls
   - Document existing controls and their effectiveness
   - Calculate residual risk after control implementation
3. **Prioritizing high-risk AI applications:** Focus initial governance efforts on AI systems that:
   - Process sensitive or personal data
   - Make or influence significant decisions affecting individuals
   - Operate with limited human oversight
   - Serve critical business functions
   - Face heightened regulatory scrutiny
4. **Developing targeted mitigation strategies:** For identified risks, create appropriate response plans:
   - Risk acceptance (for low-impact scenarios aligned with risk appetite)
   - Risk transfer (through insurance or third-party arrangements)
   - Risk reduction (through additional controls implementation)
   - Risk avoidance (by modifying use case parameters or approaches)
5. **Establishing continuous monitoring protocols:** Implement ongoing assessment processes:
   - Define key risk indicators (KRIs) for each significant AI risk
   - Set thresholds for escalation and intervention
   - Schedule regular reassessments as AI systems evolve
   - Integrate risk monitoring into your governance workflow

Documentation essentials:

- Risk assessment methodology and scoring criteria
- Comprehensive risk register with inherent and residual risk ratings
- Control mapping to specific identified risks
- Mitigation plans with clear ownership and timelines
- Evidence of executive review and acceptance of residual risks

---

Governance integration:

- Link risk assessment outcomes to your model inventory (Step 4)
- Incorporate risk levels into your approval workflow
- Use risk assessment results to inform control testing priorities
- Leverage findings to refine your Acceptable Use Policy (Step 2)

This systematic approach ensures your AI governance focuses resources where they matter most, providing appropriate oversight proportionate to actual risk levels while avoiding unnecessary constraints on lower-risk applications.

Conclusion

An effective AI governance program is essential for managing your organization's AI models and use cases by identifying, assessing, and mitigating potential AI risks. By following this best practice guide, teams can drive accountability, transparency, and confidence in building their AI governance program from scratch.

Learn how Optro can help by requesting a personalized walkthrough of our AI Governance solution here.

About Optro

Optro's mission is to be the category-defining global platform for connected risk, elevating our customers through innovation. More than 50% of the Fortune 500 trust Optro to transform their audit, risk, and compliance management. Optro is top-rated by customers on G2, Capterra, and Gartner Peer Insights, and was recently ranked for the sixth year in a row as one of the fastest-growing technology companies in North America by Deloitte. To learn more, visit Optro.ai.