sonatype

MAY 2026

The CISO and CTO AI Governance Playbook Framework

Table of Contents

Executive Summary

AI innovation is moving faster than most enterprise governance models can manage. This means an effective AI governance framework is no longer optional for organizations adopting generative AI, machine learning, and autonomous software systems.

In modern enterprises, CISOs and CTOs face mounting pressure to accelerate AI-driven innovation while simultaneously managing escalating software supply chain threats, regulatory mandates, and operational risk.

CTOs are tasked with increasing developer velocity, integrating AI into products, and enabling teams to experiment rapidly with open source models and AI-powered tooling. CISOs, meanwhile, must contain risk associated with shadow AI, insecure model provenance, hallucinations, data leakage, malicious model payloads, and expanding compliance obligations under frameworks such as the European Union's Cyber Resilience Act (CRA), the EU AI Act, NIST AI RMF, and ISO/IEC 42001.

The organizations that succeed will not be those that slow AI adoption. They will be the organizations that operationalize governance at machine speed.

This AI governance playbook provides a practical strategy for CISOs and CTOs that bridges board-level policy with operational enforcement. It explains how organizations can:

- ▸ Build an enterprise AI governance strategy
- ▸ Operationalize AI risk management frameworks
- ▸ Secure AI software supply chains
- ▸ Automate software composition analysis (SCA)
- ▸ Simplify SBOM and AIBOM management
- ▸ Enforce policy-as-code inside DevSecOps pipelines
- ▸ Accelerate compliance readiness

More importantly, it demonstrates how governance can evolve from a perceived barrier into a competitive advantage.

Pillars of an Enterprise AI Governance Strategy

AI adoption requires more than policies written in PDFs or quarterly audit reviews. A modern enterprise AI governance strategy must function continuously across legal, architectural, and operational layers.

Without automation, governance becomes friction. And friction eventually drives teams toward unsanctioned tools, unmanaged models, and shadow AI environments.

One of the first major incidents in AI oversight occurred in 2023, when Samsung employees reportedly shared sensitive internal information with ChatGPT while using the tool to support engineering and debugging tasks. The incident raised concerns around exposing proprietary data to external AI services and ultimately led the company to tighten restrictions around the use of public generative AI platforms.

A scalable AI governance framework allows organizations to accelerate AI adoption while maintaining software supply chain security and regulatory compliance. In order to execute, your playbook should establish accountability across three critical dimensions.

What is AI Governance?

AI governance is the process of establishing policies, controls, and operational safeguards that ensure AI systems are secure, compliant, transparent, and aligned with organizational risk requirements.

LEGAL AND ETHICAL GOVERNANCE

At the highest level of AI governance, organizations must align AI initiatives with evolving regulatory frameworks and ethical expectations.

AI governance programs must support:

- Compliance with the Cyber Resilience Act (CRA)
- Preparation for obligations under the EU AI Act
- Alignment with NIST AI RMF guidance
- Conformance with ISO/IEC 42001 governance standards
- Data privacy requirements
- Intellectual property protections
- Bias and discrimination mitigation

The burden of proof increasingly rests with software manufacturers and platform operators. Organizations must demonstrate not only that controls exist, but that they are continuously enforced.

ARCHITECTURAL GOVERNANCE

The second layer of AI governance focuses on the technical infrastructure, underlying architecture, and software supply chain powering AI systems.

CISOs and CTOs must establish governance around:

- Model provenance
- Training dataset lineage
- Hosting environments
- Third-party AI APIs
- Open source AI dependencies
- Serialization formats
- Runtime execution behavior

This is especially important because AI artifacts behave differently from traditional software packages. Unlike conventional dependencies, AI models may execute arbitrary code during loading, deserialization, or runtime inference.

The research highlights a critical reality for enterprise security leaders: AI models must be treated as executable software components, not static datasets.

Unpickling PyTorch

Sonatype's Unpickling PyTorch research uncovered how malicious code can be embedded directly inside serialized AI models through insecure pickle files. The whitepaper documents real-world examples where PyTorch model files executed hidden malware in March of 2025 during deserialization, including DNS beaconing, remote code execution, and obfuscated payload delivery.

OPERATIONAL GOVERNANCE

Governance ultimately succeeds or fails operationally. A scalable AI governance framework requires automated enforcement embedded directly into development pipelines and software supply chain workflows.

This includes:

- Automated policy controls
- Developer feedback loops
- Continuous software composition analysis
- AI artifact scanning
- SBOM and AIBOM generation
- Runtime risk monitoring
- Automated remediation workflows

As Sonatype discussed in *The Last Mile Problem: AI Can Write Code, But Only Policy Can Ship It*, AI-generated velocity becomes dangerous when governance controls are manual, inconsistent, or disconnected from development pipelines. AI acceleration without automated governance creates exponential risk.

Assigning Accountability for AI Risks

One of the biggest failures in enterprise AI governance is ambiguity around ownership. When accountability is unclear, AI risks often fall between security, legal, engineering, and product teams, leaving critical controls unenforced. An effective AI governance strategy maps risks directly to accountable stakeholders.

AI Risk	Primary Owner	Operational Control
Hallucinations	Product leadership	Model evaluation pipelines
Data leakage	Security teams	Access controls and monitoring
Open source AI vulnerabilities	AppSec teams	AI-SCA and repository controls
Licensing violations	Legal and compliance	AIBOM validation
Model poisoning	Security operations	Model provenance verification
Shadow AI	IT and platform engineering	Policy enforcement and firewall controls

Organizations should also establish machine-readable governance thresholds, such as:

- Maximum acceptable model age
- Approved model providers
- Vulnerability severity limits
- Allowed licenses
- Data residency restrictions
- Approved serialization formats
- Required provenance metadata

By converting governance requirements into enforceable technical policy, enterprises move from static documentation and theoretical AI governance to operational AI risk management at scale.

Navigating the Regulatory Landscape: CRA and Beyond

The regulatory environment surrounding AI is evolving rapidly. For CISOs and CTOs, the most immediate concern is the European Union's Cyber Resilience Act (CRA), which fundamentally changes how organizations must manage software supply chain security.

WHY THE CRA MATTERS FOR AI

The CRA establishes mandatory cybersecurity and software transparency requirements for products with digital elements, including software applications, SaaS platforms, connected devices, and systems that integrate AI models or third-party APIs.

This means AI-enabled products may inherit compliance obligations related to:

- Vulnerability management
- Secure-by-design principles
- Continuous security updates
- SBOM transparency
- Risk disclosure
- Incident reporting

The regulation shifts the burden of proof onto manufacturers and software providers. Organizations can no longer rely on periodic audits or fragmented documentation. They must continuously demonstrate visibility into the components, models, and dependencies inside their software ecosystems.

That challenge becomes exponentially harder in AI environments where developers routinely import open source models from public registries.

ALIGNING GOVERNANCE WITH GLOBAL STANDARDS

A future-ready AI risk management framework should align with multiple emerging standards simultaneously. Rather than treating each regulation as a separate initiative, leading organizations are consolidating governance into a unified operational framework centered on software supply chain visibility and continuous enforcement.

A modern AI governance framework helps organizations comply with regulations such as:

- NIST AI Risk Management Framework (AI RMF)
- ISO/IEC 42001
- EU AI Act
- CRA requirements
- Industry-specific regulations

This unified approach reduces compliance overhead while improving software supply chain security and operational resilience. In this consolidated view, SBOMs, AIBOMs, and automated SCA become even more important.

Operational Controls: Software Composition Analysis for AI (AI-SCA)

Traditional software composition analysis (SCA) was built for open source libraries, packages, and application dependencies. AI introduces an entirely new attack surface that traditional SCA tools were never designed to analyze.

Unlike conventional software components, AI models can execute arbitrary code during loading, deserialization, and deployment. This creates new software supply chain risks across open source AI ecosystems.

As a result, modern AI governance requires AI-aware SCA capabilities capable of inspecting:

- Model tensors
- Metadata
- Serialization formats
- Embedded scripts
- Training artifacts
- Runtime behaviors

The need for AI-aware security scanning is not theoretical. Sonatype's research into unsafe PyTorch serialization behaviors demonstrated how attackers can exploit model-loading mechanisms to embed and execute malicious code directly within AI artifacts.

WHY TRADITIONAL SCA IS NO LONGER ENOUGH

Open source AI ecosystems create unique security risks because models can execute arbitrary code during loading and deployment.

Sonatype's Unpickling PyTorch research demonstrated how malicious actors exploited PyTorch serialization behaviors to embed malware directly into AI artifacts.

The research documented:

- Remote code execution payloads
- Hidden malicious pickle files
- Obfuscated malware techniques
- ZIP archive manipulation attacks
- DNS beaconing attacks
- Evasion techniques bypassing static analysis tools

This evolution requires organizations to expand software composition analysis into AI-aware security scanning.

SECURING THE MODEL SUPPLY CHAIN

A mature AI risk assessment strategy must include controls specifically designed for model supply chain security.

This includes detecting:

- Model poisoning
- Data exfiltration mechanisms
- Hidden backdoors
- License laundering in training datasets
- Malicious serialization behaviors
- Unauthorized model modifications
- Prompt injection risks

Organizations should also implement policy controls preventing developers from importing unverified models from public repositories. This requires organizations to move from reactive detection to preventative enforcement across the AI software supply chain.

One way to do this is with Sonatype Firewall, which can help organizations block malicious or unapproved AI artifacts before they enter development environments, reducing exposure to shadow AI and unvetted open source models.

Meanwhile, Sonatype AI and LLM Governance capabilities provide visibility into model provenance, usage, and risk posture across enterprise environments.

The Data-Centric Defense: SBOMs and AIBOMs

AI governance cannot succeed without visibility. For CISOs, one of the biggest challenges is obtaining a unified understanding of both traditional software dependencies and AI-specific components. This is where SBOMs and AIBOMs become essential.

WHY AIBOMS MATTER

Traditional SBOMs document software components and dependencies.

An AI Bill of Materials (AIBOMs) explains how an AI system was built, trained, modified, and operated by documenting:

- Base models
- Training datasets
- Third-party APIs
- Inference dependencies
- Fine-tuned models
- Embedding models
- Serialization formats
- Runtime environments

This creates what many organizations now describe as the "evidence spine" for enterprise AI governance.

CREATING THE EVIDENCE SPINE WITH AIBOMS

An effective AI governance framework depends on continuous traceability. AIBOMs help organizations:

Track Data Origination

Organizations increasingly face intellectual property litigation risks tied to training data provenance.

AIBOMs help document:

- Data sources
- Consent frameworks
- Licensing terms
- Geographic residency requirements
- Usage restrictions

Document Model Modifications

Most enterprise AI systems involve modifications to existing base models.

AIBOMs should record:

- Fine-tuning processes
- Internal retraining activities
- Reinforcement learning adjustments
- Safety alignment modifications
- Prompt engineering layers

Identify Third-Party Dependencies

Many AI systems rely on multiple external APIs and services.

AIBOMs provide visibility into:

- External inference providers
- Cloud AI services
- Embedded SDKs
- Vector databases
- Open source frameworks

This visibility becomes essential for both operational resilience and regulatory reporting.

AUTOMATED COMPLIANCE REPORTING

Manual compliance reporting cannot scale to modern AI ecosystems. Organizations need automated mechanisms capable of generating audit-ready evidence in real time.

Sonatype SBOM Manager enables organizations to generate and manage SBOM exports quickly while supporting evolving compliance requirements.

CISOs and AppSec teams can also leverage Vulnerability Exploitability eXchange (VEX) documents, which communicate whether known vulnerabilities are actually exploitable within a given environment. This enables faster audits, improved transparency, and reduced compliance overhead.

Policy-as-Code: The CTO's Engine for Scale

The most successful CTOs will operationalize governance directly inside development workflows. This is where policy-as-code becomes transformative.

Rather than relying on manual approvals, organizations can codify governance rules directly into CI/CD pipelines and software supply chain platforms.

GOLDEN MODELS AND APPROVED COMPONENTS

Developers move faster when security guidance is automated. Sonatype Lifecycle enables organizations to establish "golden models," or pre-vetted AI components and dependencies that meet organizational governance requirements.

This allows development teams to innovate rapidly while remaining inside approved security boundaries. Instead of blocking developers, governance becomes an accelerator.

AUTOMATED REMEDIATION WORKFLOWS

A mature AI governance playbook should include automated response mechanisms when models violate policy.

Examples include:

- Blocking prohibited model imports
- Flagging unacceptable licenses
- Detecting outdated model versions
- Preventing vulnerable AI dependencies
- Triggering remediation workflows
- Notifying security stakeholders automatically

This reduces governance bottlenecks while improving overall security posture.

RISK-BASED GOVERNANCE

Not every AI build requires the same level of scrutiny.

Leading organizations are implementing risk-tiered governance models that allow:

- Faster approvals for low-risk internal experimentation
- Stricter controls for customer-facing systems
- Enhanced oversight for regulated environments
- Automated waivers for sandboxed research projects

This balanced approach helps organizations maintain innovation velocity without sacrificing security.

Measuring Maturity: The AI Governance Dashboard

Governance programs fail when leadership cannot measure effectiveness. CISOs and CTOs need regular AI risk assessments for board-level visibility into operational risk.

An effective enterprise AI governance dashboard should provide visibility into both software supply chain security and AI operational risk. Key metrics to track include:

Metric	Why It Matters
Percentage of models with verified AIBOMs	Measures governance coverage
Average remediation time for AI vulnerabilities	Indicates operational responsiveness
Unapproved AI usage incidents	Reflects policy adoption
Percentage of approved vs. unapproved models	Measures enforcement effectiveness
AI policy violation frequency	Identifies governance gaps
AI dependency exposure trends	Supports proactive risk management

These metrics help leadership teams evaluate whether AI governance controls are operating effectively at scale while providing evidence for compliance, audit readiness, and enterprise risk management initiatives.

IMPORTANCE OF CONTINUOUS IMPROVEMENT

AI threats evolve rapidly. Governance strategies must evolve with them.

Attackers continue experimenting with:

- ▶ Prompt injection
- ▶ Training data poisoning
- ▶ Model inversion
- ▶ Supply chain compromise
- ▶ Adversarial inputs
- ▶ Serialization-based malware

Because the AI threat landscape changes continuously, enterprise AI governance frameworks cannot remain static.

An effective AI governance strategy must continuously evolve through:

- ▶ Threat intelligence integration
- ▶ Runtime telemetry
- ▶ Continuous scanning
- ▶ Behavioral analysis
- ▶ Updated policy enforcement
- ▶ Cross-functional governance reviews

Sonatype Enterprise Vulnerability Management helps organizations unify visibility across software dependencies, AI components, and operational risk signals.

Turning Governance Into a Competitive Advantage

The organizations that dominate the AI economy will not necessarily be the fastest adopters. They will be the ones that can scale AI responsibly and earn lasting trust from customers.

Customers, regulators, investors, and partners increasingly expect transparency into how AI systems are built, secured, governed, and monitored. This changes the role of governance entirely.

An effective AI governance framework is no longer just about reducing risk. It becomes a differentiator that enables organizations to:

- ▶ Accelerate secure AI adoption
- ▶ Strengthen software supply chain resilience
- ▶ Reduce compliance overhead
- ▶ Improve operational efficiency
- ▶ Increase customer confidence
- ▶ Enable safer innovation at scale

The "wait and see" approach is no longer viable.

The CRA, evolving AI regulations, and increasingly sophisticated supply chain attacks demand proactive governance today. Organizations that delay implementation will face higher remediation costs, greater compliance exposure, and increased operational risk tomorrow.

The Next Era of Governance

AI governance is ultimately a leadership challenge. CISOs and CTOs must align innovation velocity with operational resilience.

That requires:

- ▶ Continuous software supply chain visibility
- ▶ SBOM and AIBOM traceability
- ▶ Automated policy enforcement
- ▶ Runtime risk monitoring
- ▶ AI-aware software composition analysis
- ▶ Scalable compliance automation

The good news is that governance and innovation are no longer opposing forces. When implemented correctly, governance becomes the foundation that allows organizations to innovate confidently at enterprise scale.