Governing AI Inside the Enterprise: A Practical Framework for In-house Counsel

Information Governance

Technology, Privacy, and eCommerce

Banner artwork by Fahmi_Ruddin_Hidayat / *Shutterstock.com*

Cheat Sheet

- **Enterprise risk.** Treat AI as business and legal risk, not solely as a technology or innovation initiative.
- **Evaluate across dimensions.** Use risk tiering to align legal oversight with the real-world impact of AI use cases.
- **Accountability, not ambiguity.** Establish clear ownership and escalation pathways to ensure accountability for AI-driven decisions.
- **Well-calibrated records.** Create a defensible record by documenting AI governance decisions, risk assessments, and controls.

Artificial intelligence has moved from experimentation to infrastructure, now shaping how companies hire, price, forecast, negotiate, and serve customers. These systems embed themselves so deeply in everyday workflows that their presence dissolves into institutional muscle memory.

For in-house counsel, that invisibility creates the problem. AI systems advance faster than internal

legal oversight and far faster than external regulation, leaving legal teams asked to approve tools without a clear rulebook, consistent internal standards, or defined accountability.

When something goes wrong, whether through regulatory inquiry, litigation, or reputational crisis, the first question asks not how sophisticated the technology was, but who approved its use, why the organization deployed it, and what controls existed.

This moment demands that in-house counsel move beyond reactive review into a leadership role that shapes how enterprises adopt emerging technology responsibly. AI governance positions legal departments to do exactly that, not by mastering the technology itself but by applying the governance structures they already know how to run.

Find the latest AI resources, use cases, and more in the
ACC AI Center of Excellence for In-house Counsel

Disciplined reframing of AI

Effective AI governance requires neither that legal teams become technologists nor that they wait for comprehensive regulation that may arrive unevenly or too late. It requires instead a disciplined reframing of AI as enterprise risk, coupled with the rigorous application of oversight mechanisms legal departments have honed across cybersecurity, data privacy, and third-party vendor management.

Organizations succeed with AI governance when they stop treating AI as a novelty and begin treating it as a category of risk that demands the same rigor as any other material business exposure.

AI systems influence real people and real outcomes. They affect employment decisions, customer and vendor relationships, pricing integrity, intellectual property ownership, and compliance obligations across multiple jurisdictions.

The legal department stands uniquely positioned to lead AI governance precisely because it already operates at the intersection of uncertainty, accountability, and judgment. Legal teams routinely assess risk in environments where the law evolves, standards differ across borders, and perfect answers do not exist.

AI simply amplifies these conditions. The objective lies not in eliminating risk but in understanding it, allocating it thoughtfully, and ensuring that decision making remains defensible under scrutiny.

Most enterprises already possess the necessary architecture for AI governance. Enterprise risk management programs, compliance and ethics committees, data privacy councils, information security governance, and third-party risk processes all provide natural homes for AI oversight. Integrating AI into these existing frameworks reinforces consistency and signals that AI represents a managed risk rather than an exception requiring special treatment.

This integration also clarifies roles: IT and innovation teams remain responsible for implementation

and performance, business units retain ownership of use cases and outcomes, and legal defines risk thresholds, decision boundaries, and accountability. The alignment reduces friction and prevents AI governance from becoming an isolated exercise disconnected from business reality.

> Organizations rarely fail because they lacked rules; they fail because early warning signs did not reach the right decision makers in time.

Use cross-dimensional risk tiering

Risk tiering sits at the center of any practical AI governance framework. Not all AI tools pose the same exposure and treating them as though they do undermine both innovation and credibility. A generative tool that summarizes internal documents carries a fundamentally different risk profile than an AI system that screens job applicants, sets prices, or interacts directly with customers at scale.

Organizations must evaluate AI use cases across consistent dimensions: their impact on individuals, the degree of automation involved, the nature of the data used, and the potential legal or reputational consequences if the system fails or produces unintended results.

This tiered approach allows governance to scale intelligently. Lower-risk tools move forward with light oversight and standardized guardrails, while higher-risk deployments justify deeper review, clearer documentation, and senior-level approval. The framework directs legal resources where they add the most value and signals to the business that governance exists to enable responsible use rather than slow progress indiscriminately.

In-house counsel who helps their organizations implement risk tiering create a repeatable structure that adapts as both technology and legal landscapes shift.

ACC Members: Download the **Artificial Intelligence Toolkit for In-house Lawyers**

Establish clear accountability

Clear ownership forms the next critical element. AI systems typically sit at the intersection of multiple functions, which creates dangerous ambiguity when accountability matters most. Every material AI system must have a defined business owner who remains responsible for how the tool gets used, not merely how it performs technically.

Ownership means answering for outcomes, ensuring compliance with internal standards, and escalating issues when risk thresholds are exceeded. Without this clarity, organizations discover too late that no one felt empowered to pull the plug when early warning signs emerged.

Escalation pathways deserve equal rigor. Organizations rarely fail because they lacked rules; they fail because early warning signs did not reach the right decision makers in time.

Legal teams must define clear triggers that require senior management or board awareness: regulatory inquiries, systemic errors affecting multiple individuals, cross-border compliance concerns, or any circumstance where the AI system's operation conflicts with stated company values. These thresholds protect both the organization and the individuals involved by ensuring that material risks surface before they metastasize into crises that require crisis management rather than course correction.

> Regulators and courts rarely expect perfection from emerging technology; they do expect evidence of thoughtful decision making.

Ensure defensible records

Documentation often feels burdensome, but in the context of AI governance it provides one of the most effective defensive tools available. Regulators and courts rarely expect perfection from emerging technology; they do expect evidence of thoughtful decision making.

Legal teams should focus documentation efforts on capturing why an AI tool was adopted, what risks were identified, how those risks were mitigated, and who approved the deployment. The goal is clarity, not volume. Well-calibrated records demonstrate intent, discipline, and oversight, which matter far more in regulatory proceedings and litigation than flawless outcomes no organization can guarantee.

Many organizations rely on vendors for AI capabilities, which means legal exposure often originates outside the enterprise. Procurement and contracting processes must reflect the realities of AI systems.

Vendor agreements need to address how systems are trained, what data gets used, how outputs may be relied upon, what indemnification applies when systems fail, and what audit rights the organization retains. Contracts should align with internal risk tiering, with stronger protections and more detailed specifications reserved for higher-risk use cases.

In-house counsel who treat AI vendor agreements as standard technology procurement discover their mistake only when a vendor's model produces discriminatory results, hallucinates facts in customer-facing contexts, or violates data privacy commitments the vendor never actually made.

> AI governance represents a defining opportunity for in-house counsel.

Focus on credibility

Sustainable AI governance depends ultimately on credibility. If governance feels opaque or unpredictable, business teams will route around it, building shadow AI deployments that expose the organization to precisely the risks governance was designed to manage.

Legal departments prevent this by articulating clear expectations, training stakeholders on risk categories, and revisiting governance frameworks as technology and regulations evolve. Transparency builds trust, and trust allows governance to function without becoming a bottleneck that innovation teams learn to avoid.

AI governance represents a defining opportunity for in-house counsel. It offers the chance to shape how technology integrates into the enterprise responsibly, moving beyond the reactive posture that too often characterizes legal's relationship with innovation.

The organizations that succeed with AI will not be those that deploy the most advanced systems but those that can explain their choices clearly, defend them confidently, and adapt them thoughtfully as the landscape changes.

In-house counsel stands central to that outcome, not because they control the technology but because they shape the judgment behind it. That judgment, applied consistently and transparently, determines whether AI becomes a source of competitive advantage or a vector for institutional risk the organization discovers only when defense becomes necessary.

Join ACC for more AI guidance and insights!

Alexander Lima

Vice President, Associate General Counsel, and Regional Compliance Officer

Wesco International

Alexander Lima is Vice President, Associate General Counsel, and Regional Compliance Officer at Wesco International, a Fortune 200 company, advising on global transactions, M&A, governance, legal operations, and compliance across over 50 countries. His work spans multibillion-dollar business segments and sectors including infrastructure, AI, and high-growth international markets. Alexander also serves as Board Advisor and Corporate Secretary at Judy Security, a cybersecurity AI firm, and is an Adjunct Professor at the University of Miami School of Law.