

The General Counsel as Chief AI Governance Officer

Corporate, Securities, and Governance



Banner artwork by *Prostock-studio / Shutterstock.com*

Cheat Sheet

Widening governance gap. AI's autonomy, opacity, and constant evolution have outpaced the legal and compliance systems most companies still rely on.
Governance can't be siloed. Piecemeal ownership across departments creates conflicting fixes — AI risk needs one coordinating function, the same path data-privacy governance already took.
Natural fit. Legal can set guardrails, protect candid risk assessments under privilege, and translate fast-moving law into operational policy without owning AI strategy itself.
Raising the stakes. Recent officer-oversight rulings extend fiduciary duty to mission-critical risks, giving GCs a legal — not just practical — reason to lead AI governance.

Artificial intelligence (AI) is fundamentally transforming business operations and markets in virtually every company and industry. This rapid adoption of AI has already provided substantial benefits, enhancing productivity, efficiency, creativity and the capabilities of both organizations and their employees.

But AI adoption has not been without significant and growing concerns: biased or discriminatory algorithms, fabricated or fake information (including both “hallucinations” and “deep fakes”), data ownership conflicts, privacy and confidentiality transgressions, novel cybersecurity attacks, deceptive business practices, antitrust concerns, and unanticipated actions by AI chatbots and agents.

These AI risks and problems have created a growing governance gap, as many companies’ traditional legal and governance systems are unprepared for the unique capabilities, dimensions and risks of AI systems. The challenge is no longer merely how to govern a new technology, but how to govern a new way of operating. AI presents several distinctive governance challenges for organizations.

> The challenge is no longer merely how to govern a new technology, but how to govern a new way of operating.

What makes AI different

Perhaps most significantly, machine learning and generative AI systems — especially now in the form of AI agents — possess the capability to generate content or render decisions autonomously, without direct human supervision or direction. Whether the AI system is used to drive an autonomous vehicle, diagnose a medical condition, select a new employee, or process customer data, these systems can delegate consequential determinations to machine processes.

Moreover, humans often don’t fully understand how these complex AI systems make their decisions, and the machines are unable to explain their decisions in ways that humans can comprehend (the “black box” problem).

Another unique aspect of AI systems is that they are not static products, but rather evolve and change in use and in the field based on new data and new experiences they encounter, a phenomenon known as model or data drift, further undermining our ability to understand and predict their behavior.

A final unique dimension is that most AI systems are built on massive data sets and incorporate multiple pieces of software, some open-source and some proprietary, all of which can generate complex and unexpected downstream intellectual property, data privacy, data security, and product liability issues.

Who should own AI governance?

Traditional corporate legal, compliance, risk management and information technology (IT) practices are unprepared for the novel opportunities and risks of AI systems. Government enforcers and courts, along with corporate legal and compliance officers, confront substantial challenges in applying traditional legal frameworks to novel AI-generated harm.

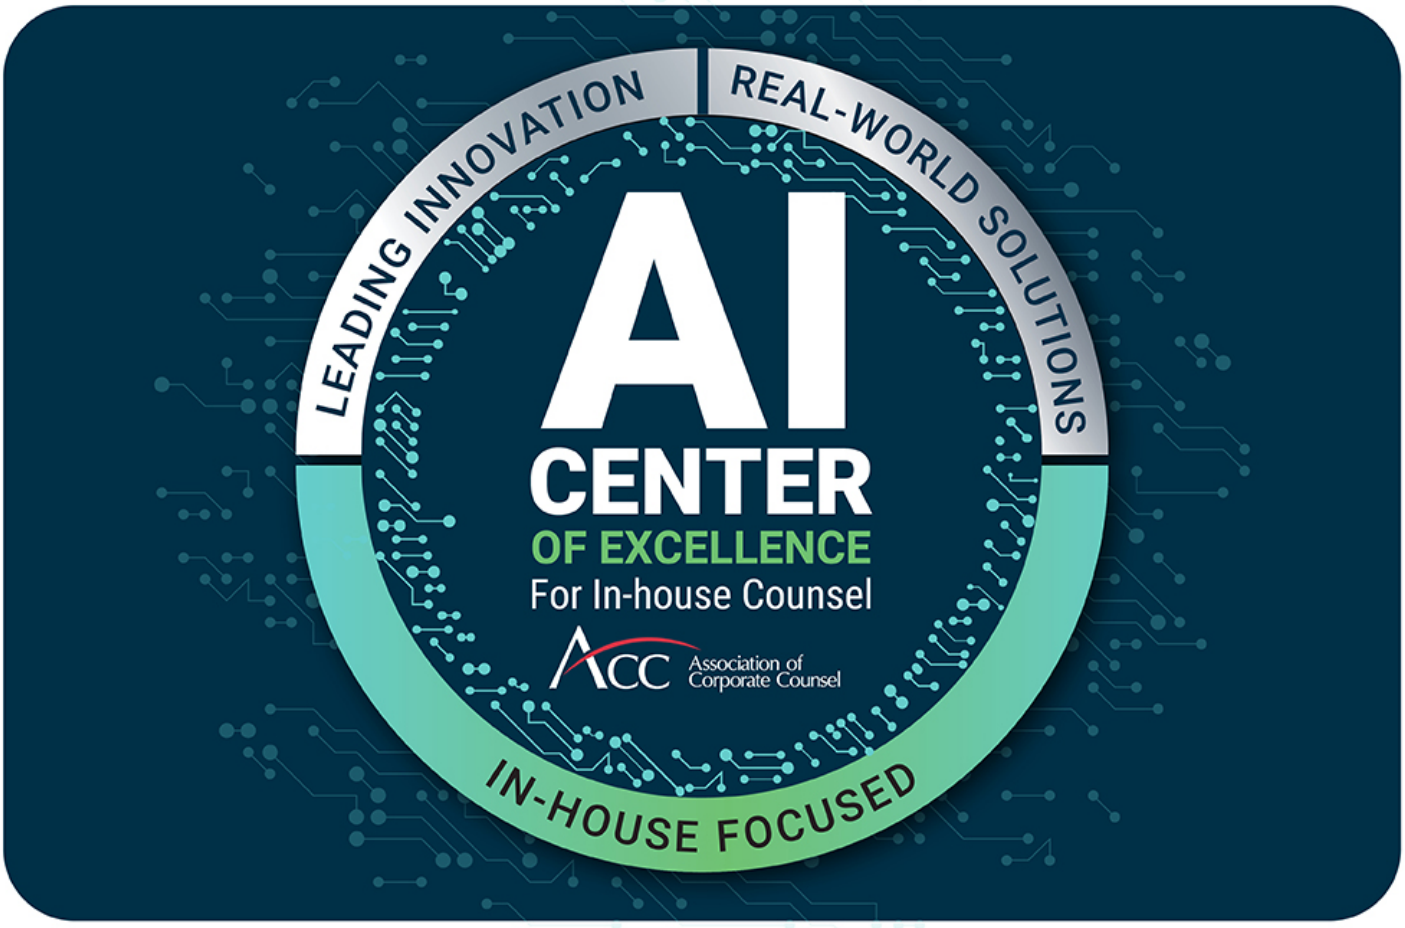
At the same time, governments across the United States and globally are responding with an expanding and often fragmented set of AI regulations. Corporate AI governance functions must therefore interpret this rapidly evolving regulatory landscape while integrating compliance with

enterprise risk management, liability mitigation, internal and external soft law frameworks, reputation management, and organizational strategies for deploying AI.

These efforts must be undertaken with a forward-looking perspective that anticipates continued technological and regulatory change.

Against this backdrop, organizations deploying AI confront a critical governance question. Because AI implicates multiple organizational functions simultaneously, effective oversight requires a coordinating governance entity capable of integrating legal, technical, ethical and operational considerations.

A threshold question is who within an organization should assume coordinating responsibility for enterprise AI governance. This article argues that the General Counsel (GC) or Chief Legal Officer (CLO) is uniquely positioned — institutionally, doctrinally, and functionally — to serve as a natural orchestrator of enterprise AI governance.



AI requires “all-of-enterprise” governance

Effective AI governance demands integrated expertise spanning regulatory compliance, liability exposure, contractual frameworks, intellectual property obligations, technology procurement and oversight, and organizational ethics. The GC/CLO is uniquely positioned to help lead and orchestrate corporate AI governance initiatives, ensuring board-level oversight that transcends traditional technology management to encompass comprehensive legal, ethical, and regulatory risk management.

Corporate governance of data privacy offers an illustrative precedent. Companies initially appointed IT specialists as their Chief Privacy Officers (CPO), driven both by internal risk management imperatives and external regulatory requirements like the European Union's General Data Protection Regulation (GDPR).

However, as privacy regulations scaled internationally and domestically, organizations increasingly hired lawyers for these roles. Lawyers bring essential expertise across key privacy domains, including regulation, liability, risk management and corporate ethics.

But technological fluency is equally critical. Attorneys serving as the CPO have had to develop sufficient technological sophistication and maintain strong collaboration with IT professionals. The lesson is not that lawyers should displace technologists, but that technology governance fails when either function operates alone. This evolution illustrates the broader organizational need for a central coordinating function as technological risk becomes more complex and cross-functional.

> The lesson is not that lawyers should displace technologists, but that technology governance fails when either function operates alone.

The many risks associated with AI, including bias, privacy, accuracy, and liability, cannot be managed piecemeal. They require integrated, enterprise-wide governance. Having multiple actors within the company attempting to regulate individual risks for the same AI system is a recipe for confusion and conflict.

Moreover, solutions to one problem may exacerbate other problems — for example, regulating against bias independently may result in privacy or accuracy problems. Greater transparency demanded by some stakeholders may result in greater security risks. An integrated, holistic and comprehensive "all-of-enterprise" governance approach is therefore necessary to consider all these issues through one overarching program or framework.

There are several available AI governance frameworks that take a "all-of-enterprise" approach, including ISO/IEC 42001, the NIST AI Risk Management Framework, and IEEE 2863. These AI frameworks all call for a central entity to coordinate AI governance across the enterprise, a role for which the corporate GC/CLO is a natural fit.

The GC/CLO as natural leader of AI governance

The GC/CLO should serve as a central orchestrator of the corporation's AI governance: a structure that discerns legal and fiduciary obligations and translates them into enforceable policies, board reporting, and accountability mechanisms.

This does not mean that Legal must own AI strategy, technology selection, or every use case. Rather, the GC/CLO is especially well positioned to provide three governance functions:

- **Guardrails without ownership.** The GC/CLO doesn't build AI models but helps set the guardrails within which AI development proceeds: defining acceptable use, approved data sources, testing requirements, and escalation thresholds.

- **Protected legal assessment.** Attorney-client privilege can create appropriately protected space for candid legal risk assessment when communications are made for the purpose of obtaining legal advice, allowing the organization to surface problems, assess algorithmic bias, and test compliance while preserving privilege where it properly applies.

- **Translation from law to operations.** The GC/CLO interprets fiduciary duties, statutes, regulations and private standards, then translates those mandates into operational directives. As AI evolves and new regulations emerge, the GC/CLO updates the governance framework to maintain alignment with foundational legal obligations in every jurisdiction in which the company operates.

The argument for GC/CLO leadership is not merely functional. It is doctrinal, rooted in how Delaware corporate law allocates oversight responsibilities among directors and officers. Recent Delaware decisions have clarified that fiduciary oversight duties extend to senior officers, reshaping the liability landscape for how “mission-critical” risks are supervised.

The 2023 decision in *In re McDonald's Corp. Stockholder Derivative Litigation* held that corporate officers may breach the duty of loyalty under Caremark where they consciously fail to implement reasonable information systems in their domains or ignore "red flags" indicating ongoing wrongdoing. The opinion makes explicit that oversight obligations attach to an officer with responsibility for a risk-laden area of the business. Where the GC/CLO has responsibility for privacy, ethics, regulatory strategy, or other dimensions of AI risk, these principles underscore the need for systems capable of surfacing and escalating AI-related red flags.

The Delaware Supreme Court’s decision in *Marchand v. Barnhill* held that the duty of loyalty requires directors to ensure they receive information about risks that go to “the core of the company's business.” Where a particular risk is mission-critical, a board’s failure to implement and monitor an appropriate reporting system can constitute bad-faith oversight.

As AI becomes the infrastructure powering pricing, underwriting, hiring, recommendation engines, and logistics, it increasingly may fit the *Marchand* profile of mission-critical infrastructure. The spectrum of potential risks and liabilities associated with AI reinforces its potential board-level significance. For organizations deploying AI at scale, treating it as a routine IT matter can understate the governance implications.

The GC/CLO is the officer structurally trained to interpret fiduciary obligations and convert them into board-satisfying reporting mechanisms. Fiduciary duties, enforcement priorities, regulatory guidance, and litigation risk must be decoded, prioritized, and embedded into the corporation's governance systems. Engineering, product, and data-science leaders may master the technical architecture of AI

systems, but they are not typically charged with mapping that architecture onto the duty of loyalty, disclosure rules, or the evolving contours of officer liability.

Legal leadership may also create a “privilege advantage” in appropriate circumstances: when AI impact assessments, model-risk reviews, and internal audits of algorithmic bias are conducted at the direction of counsel for the purpose of obtaining legal advice, elements of those processes may be protected by attorney-client privilege and work-product doctrines.

When a single AI system simultaneously implicates data privacy (GDPR), employment discrimination (Title VII), consumer protection (FTC Act), and securities disclosure obligations, candid internal legal assessment can require a protected space to evaluate exposure and whether to halt deployment. Legal involvement does not make ordinary business or technical discussions privileged, but it can help structure protected legal assessments where privilege properly applies.

Taken together, *In re McDonald's*, *Marchand*, and the broader *Caremark* jurisprudence provide an important governance framework for corporate responses to AI. These precedents push oversight duties both down (to officers) and deep (to mission-critical risks). Legal therefore cannot remain a peripheral reviewer of AI projects. The GC/CLO should help design, monitor, and update the governance systems that provide management and the board with meaningful visibility into material AI risks.

Governance as an accelerator, not the brakes

A modern company has many specialized units that perform sophisticated and specialized functions, such as finance, information technology, operations, human resources, and public relations. AI is rapidly becoming critical for all these functions, but the company needs an entity to coordinate and orchestrate AI governance across them.

As this article has argued, the GC/CLO is especially well positioned to perform this AI orchestration function. Corporate governance of AI must be centered on maximizing benefits while ensuring legal compliance with national and state regulations, conformance with international and internal standards, risk assessment and liability prevention, and protection against reputational harms to the organization.

> The lesson is not that lawyers should displace technologists, but that technology governance fails when either function operates alone.

That does not require Legal to own AI. It requires the GC/CLO to connect these domains, establish appropriate governance and escalation mechanisms, and help the business deploy AI with greater confidence. Good governance is not simply a brake on AI adoption; it is part of what allows an organization to accelerate responsibly.

Join ACC for more AI insights!

“Disclaimer: The information in this article should not be construed as legal advice or as a legal opinion on specific facts. The views expressed are those of the authors and do not represent the views of the authors’ employers, the article’s sponsors, or ACC. This article is not intended as a definitive statement on the subject addressed. Rather, it is intended to provide practical guidance and references for in-house practitioners and other readers.”

Whitney Stefko Dover



Director and Senior Counsel, Policy and Legal Operations & Adjunct Professor

Ford Motor Company | Sandra Day O’Connor College of Law at Arizona State University

Whitney Stefko is Director and Senior Counsel, Policy and Legal Operations at Ford Motor Company. She helps lead transformation across Ford’s global legal organization, with a focus on artificial intelligence strategy and governance, technology, operations, and the evolving delivery of legal services. Stefko is also an adjunct professor at the Sandra Day O’Connor College of Law at Arizona State University, where she teaches courses on AI law, ethics, policy, legal operations, and innovation. A frequent speaker and writer, she explores AI fluency, organizational change, and the skills legal professionals need to thrive as technology reshapes their work. She is particularly interested in the human side of transformation — how organizations redesign work, develop talent, codify expertise, and turn emerging technology into meaningful, lasting impact. She is based in Scottsdale, Arizona.

Gary Merchant



Regents' Professor

ASU Sandra Day O'Connor College of Law

Gary Marchant, Ph.D., J.D., M.P.P., is Regents' Professor and Faculty Director of the Center for Law, Science & Innovation at the Sandra Day O'Connor College of Law, Arizona State University. He is also the Clarissa Cerda Professor in Law and Technology at ASU. Professor Marchant's research interests include the governance of emerging technologies such as genomics, biotechnology, nanotechnology, artificial intelligence, neuroscience, quantum technology, and blockchain. Prior to joining ASU in 1999, Professor Marchant was a partner at the Washington, D.C. office of Kirkland & Ellis. He is the Chair of IEEE P2863, a Recommended Practice for the Governance of AI.