



It's Go-Time — Creating an AI Governance Program

Compliance and Ethics

Information Governance

Career Development

Technology, Privacy, and eCommerce



Original banner artwork by Darren Lin

Cheat Sheet

- **Stakeholder engagement.** It's better to cast a wide net for stakeholders early in the process than trying to develop a program with a small group.
- **Update your AI policy.** Detail how AI should be used, and state how it will be compliant, legal, transparent, and ethical.
- **Refine governance.** These processes — which include regulatory and privacy requirements, among others — will be needed in the initial system development as well as ongoing use.
- **Monitor and remediate.** Don't set it and forget it. An AI governance program will need constant updates and adjustments.

Pressure for companies to use generative artificial intelligence (AI) to gain a competitive advantage (or at least not fall behind competitors) is steadily rising. In 2024, companies will push their IT, legal, compliance, and privacy teams to deploy AI applications now, not later. While AI promises tremendous potential for innovation and productivity gains, the emerging compliance challenges and risks can feel overwhelming. New AI regulations emerge almost weekly, while courts are just starting to address copyright and IP issues, requiring companies to use AI both ethically and correctly. Saying “we can’t do this” is not a viable option. Companies need an AI governance program to help navigate the challenges of successfully and compliantly harnessing the power of AI.

The 10x productivity boon

The advent of generative AI offers the promise of tremendous leaps in productivity, new revenue, cost savings, and increased innovation. After going through the “AI Winter” in 1980s and 1990s, it slowly began to gain adoption in some niches in the ‘00s and the ‘10’s. In the past three years generative AI has created an explosion of interest, elevating it from the fringes to the mainstream.

Contrary to forecasts of doom from the digerati of AI replacing (and taking over) segments of society, it is the opinion of the authors that AI deployment will initially be piecemeal, with teams within finance, marketing, legal, and engineering creating AI-assisted applications. Companies are launching AI initiatives in legal, finance, marketing, product design, engineering, and nearly every single aspect of the organization. (Full disclosure: Contoural is launching an AI-based records management legal research initiative.) Without overstating, AI could be transformative.

Driving this rapid adoption is the promise of 10x productivity increase, (called the “10x engineer”), which argues that an employee leveraging AI can be 10 times more productive. For example, software engineers can use AI to develop code, finance employees can use AI to automate expense review, and HR use of AI can dramatically reduce the time it takes to create job descriptions or update policies. Much as the internet changed many business operations, generative AI promises to be equally impactful.

Think of generative AI as an intern

Perhaps the best analogy for generative AI’s capabilities and limitations is as an intern. Imagine hiring a bright, hardworking, knowledgeable but sometimes naïve intern in the legal department who works long hours and has nearly unlimited capacity to do work. As they have limited experience, this intern would not be given large or complex tasks. Rather, they would be given limited, specific assignments such as reviewing vendor agreements to ensure terms are in line with corporate standards.

Because the intern is new and somewhat inexperienced, all their work product would need to be reviewed by their (human) manager, then the intern would use this feedback to increase their work quality. Once one task is mastered, additional interns can be brought on board, with each trained in a slightly different task in adjacent areas. Interns could even be hired to coordinate and check the work of other interns, with final review still being completed by humans.


Compliance challenges, legal risks, and ethical concerns of AI

Generative AI’s explosive adoption has been met with a quick response from regulators. Every week governments across the world are proposing restrictions on how and where this new technology can be used. Wanting to become the global standard, [European regulators announced restrictions](#) on how AI can use information about individuals, as well as overall safeguards, especially around the use of personal information.


Creating an AI Governance Function. It's Go-Time.

Stacey Shaw, Assistant General Counsel, Careington
Mark Diamond, CEO, Contoural, Inc.
Tom Mighell, Esq., COO, Contoural Inc.

contoural



ACC
Association of
Corporate Counsel



[Log-in](#) to view the author's ACC webinar on this topic!

At the federal level, the Biden administration created a [new standard](#) for safety and security to protect privacy and civil rights. At the state level, a number of states have proposed AI rules, such as limiting how companies can use AI to make decisions affecting individuals, such as in the insurance and banking industries. (Note that at least one data protection authority in the EU has also created such limits.) These new regulations are just the beginning, as we expect to see many countries and states developing new rules limiting AI this year, creating a rushed and messy regulatory environment.

In addition to new AI regulations, there are significant copyright and intellectual property concerns around generative AI. In December 2023, the [New York Times sued](#) both OpenAI and Microsoft claiming that their generative AI products were based on and violated the *Times'* copyrighted information. The courts are just beginning to address some of these challenges, and it may take years for instructive case law to provide any guidance.

Companies deploying AI-assisted applications also need to ensure that these systems do not leverage or potentially expose proprietary, corporate confidential information, or trade secrets. The *Economist* reported last year that Samsung employees unintentionally leaked proprietary source code via ChatGPT. An unprotected disclosure of a trade secret to a third party – through an AI assisted application – generally vitiates the trade secret protections afforded to the information.

Another emerging issue is the ethical use of AI. AI systems are susceptible to biases from the “training data” used to build the system’s intelligence, and these biases can inadvertently discriminate against individuals or behave in other ways that do not reflect a company’s values. For example, if an HR application looking for candidates “teaches” an AI system to look for job candidates based on historical hiring profiles that do not reflect a company’s present diversity goals, its output will likely reflect an unintended bias.

Finally, “naïve” AI systems want to please and can produce inaccurate or unsafe information. Recently an eating disorder website added a chatbot to answer questions, only to find later that the chatbot was suggesting to possibly anorexic users that they should simply focus on eating less to lose weight. AI systems can produce correct answers most of the time, but detecting the incorrect outlier answers can be difficult.




Pressure to deploy AI

This year is setting up for a tremendous tug of war between the parts of the business that want to use AI and the chaotic compliance and risk environment of using it. Investors understand the tremendous competitive advantages AI offers, and CEOs realize that demonstrating that their companies are using AI has the potential to drive profitability. According to researchers Storable and WallStreetZen companies that mentioned AI saw an [average stock price increase](#) of 4.6 percent, almost double that of companies that did not. Worried about how companies may falsely be claiming to use AI, recently Gary Gensler, chair of the Securities and Exchange Commission, warned companies against “AI washing” by making false claims about their use of AI to drive up their stock price.

Despite these regulatory, legal, and safety concerns, IT and legal departments will face tremendous pressure in 2024 to deploy AI applications. Organizations may lose an advantage sitting on the sidelines. Waiting until the compliance and risk environment becomes better understood for many will not be a viable option.

Creating an AI governance program



A comprehensive AI governance program not only reduces risks and ensures compliance, but also speeds production deployment.

The dilemma of charging ahead with deploying a potentially risky AI application, or doing nothing while waiting for regulatory or legal clarity is a false one. Companies are developing AI governance programs today that can ensure compliance, evaluate legal risks, secure sensitive information, and drive responsible and safe use of AI. AI governance also allows AI applications to move into production more quickly.

4 steps to launch an AI governance program

Step 1: Stakeholder engagement and assessment

While it may be tempting to try and develop a program with a small group of stakeholders, this may slow down or even halt program development. As a first step, needs should be assessed and socialized with a larger group of stakeholders.

Assessment and roadmap

Assess current and targeted compliance and risk requirements, engaging key stakeholders in the process. Determining what needs to be done and at what level can speed up these types of complex projects.

Impact assessment

Some jurisdictions may specifically require an impact assessment focusing on people and organizations. This impact assessment can and should be borrowed from other elements in this step.

Steering committee

AI governance is complex, requiring the participation of various stakeholders, including compliance, risk, legal, privacy, information governance, data governance, IT, and business functions. This

committee should be formed early and a charter developed to ensure both that all risks and requirements are covered and that each group feels a sense of “buy in” to the process.

Roles and responsibilities

Many AI governance functions will require participation from a number of stakeholders, including the business units. Developing roles and responsibilities ensures that everyone knows their place in the organizational framework and how to execute on the other parts of the governance model.

Step 2: Create and update AI governance policies

Governance policies detail compliant and responsible use of AI.

Create an AI governance policy

The next step is developing and updating AI governance policies. An AI governance policy sets out the compliant, legal, transparent, and ethical use of AI for the organization. It details how AI should be used, and safeguards employees. This is an overall “guiding light” to demonstrate to regulators and others that the company is using AI responsibly.

Update data retention/records retention policy and schedule

Organizations may need to update their data retention policies or records retention schedules. The traditional practice of saving “all electronic information forever” creates risks that old, legacy data “polluted” with sensitive or incorrect information might be used in the development of AI systems. Good data hygiene — enforced through up-to-date retention policies — limits these risks.

Data security classification policies

A data security classification policy classifies information based on privacy, confidentiality, intellectual property, and other sensitivity factors. Organizations may also need to update this policy to ensure that appropriate controls are placed on sensitive information and used appropriately in AI.

Privacy policies

Many AI regulatory restrictions center on the use of personal information. Privacy policies need to be synchronized with AI governance policies and use.

Having up-to-date policies provides defensibility if an AI system faces review from a regulator. These policies demonstrate the organization is mindful in its use of AI and is diligent in its compliance efforts.

Step 3: Develop AI governance processes

Once the policies are in place, the next step is to develop governance processes. These processes will be needed both in the initial system development and ongoing use.

Regulatory review process

AI regulatory requirements are being announced seemingly every week. Organizations need to

develop a process to monitor regulatory changes to ensure their systems comply with any new rules.

Data provenance process

AI systems leverage both “training data” used by large language models and supplementary information used with an AI capability called retrieval augmented generation, in which the user can provide the AI large language model additional information. Companies need to undertake reasonable due diligence to confirm this input data is not copyrighted or, if it is, that they have the right to use this information. Furthermore, as this input data is often refreshed, ascertaining provenance must occur periodically.

Most AI applications leverage proprietary, closed, large language models from commercial vendors such as Open.ai and Anthropic. Concerns have been raised across the industry on whether these systems have been trained with copyrighted data. Some generative vendors have taken extra precautions to ensure that their products are based exclusively on fully licensed training data. The legal issues around AI and copyright are complex, and the case law is not settled. Furthermore, as these are closed systems, it is not possible to inspect their training data. Companies that use these products should seek assurances from AI vendors on the provenance of the training data. Some vendors go as far as offering to indemnify their users against copyright infringement claims when using their products. Ultimately, companies will have to determine their own level of risk tolerance.

Privacy and sensitive information review process

In addition to ensuring that input data is not copyrighted, organizations should develop a process to ensure the AI does not contain either personal or other types of sensitive information such as trade secrets or corporate confidential information. Training or other input data “polluted” with sensitive information may drive noncompliance or inadvertently disclose restricted information.

Ethical use review process

In addition to compliance, AI systems need to produce ethical results. For example, a visual generative AI application when asked to create a picture of “senior executives” should not consistently create historically stereotypical images. The AI output should be tested to ensure it is ethical and reflects an organization’s values.

AI accuracy, correctness, and safety review process

In addition to compliance, and legal assuredness, AI needs to be accurate, correct, and safe. AI’s polished output can lull a user to believe that all the information it produces is correct and accurate. However, these need to be tested both throughout development and on an ongoing basis. Additionally, AI also needs to be tested for safety to ensure it is not being misused.

All AI governance processes should be performed regularly, and documentation of this activity should be retained. Any issues, discrepancies, or problems should be noted, along with steps taken to remediate these issues. AI governance is similar to other types of compliance.

Step 4: Execute, monitor, and remediate

Once launched, AI systems need to be monitored. Any issues, discrepancies, or problems should be noted, along with steps taken to remediate these issues.

Initial and ongoing process execution

Ensure all processes are enacted during and after launch.

Ongoing monitoring

Monitor and review the results of the processes.

Updates and remediation

Update the system and/or approach as compliance and other rules change. Remediate any issues encountered, and document the actions taken.

In the event of a regulatory inquiry, being able to readily communicate what you intended to do (policies), how you intended to ensure you were doing it (processes), and how you addressed issues when they arose will demonstrate compliance and make the system more defensible.

Creating AI governance agility

With new AI regulations expected to be enacted this year, how can organizations create an AI governance function today that does not have to be updated every time a new rule is announced? Close inspection of many of the new global requirements shows that they address similar and common areas. Designing an AI governance function toward these common requirements will enable companies to become “AI Agile,” requiring minimal changes. This approach will be far easier in the long run than developing a program hard-wired for a single jurisdiction.

Early engagement will ensure compliant design

This new, complex technology faces a chaotic legal and regulatory environment. AI, from a legal and compliance perspective, can be both overwhelming and a bit scary. Fortunately, while the technology may be new, creating an AI governance program incorporating strategies, policies, and processes to address these challenges is achievable and leverages many existing in-house counsel skills. The key is getting involved early, engaging stakeholders, and ensuring the AI capabilities are designed compliantly from the start. This not only ensures smoother execution down the road, but also can prevent last-minute redesigns, allowing faster production implementation. Fear not AI. Instead, with a smart approach, embrace the technology and your role enabling your company to enjoy its benefits.

[Join ACC](#)

Disclaimer: The information in any resource in this website should not be construed as legal advice or as a legal opinion on specific facts, and should not be considered representing the views of its authors, its authors' employers, its sponsors, and/or ACC. These resources are not intended as a definitive statement on the subject addressed. Rather, they are intended to serve as a tool providing practical guidance and references for the busy in-house practitioner and other readers.

[Adam Shedd](#)



Vice President, Assistant General Counsel, Privacy and Marketing

Selective Insurance

Adam Shedd is vice president, assistant general counsel, privacy and marketing at Selective Insurance. He supports Model Governance, Data Protection, and Data Governance, IT, Contracts, Billing, Marketing,

Innovation, and Customer Experience, and oversees Records.

Shedd joined Selective in 2011 as Corporate Counsel, supporting Insurance Operations. Previously, he served as an insurance coverage attorney at Drinker, Biddle & Reath LLP (now Faegre Drinker Biddle & Reath LLP). He earned his bachelor's degree cum laude at Tufts University and his JD cum laude from Tulane Law School.

[Mark Diamond](#)



CEO and Founder

Contoural Inc.

Mark Diamond is the CEO and founder of Contoural Inc., an independent provider of information governance consulting services. His company works with more than 30 percent of the Fortune 500, plus many mid-sized and smaller companies.