



Mark Diamond (markdiamond@contoural.com) is the CEO of Contoural Inc. in Los Altos, California, USA.

Creating an AI governance function: Part 2

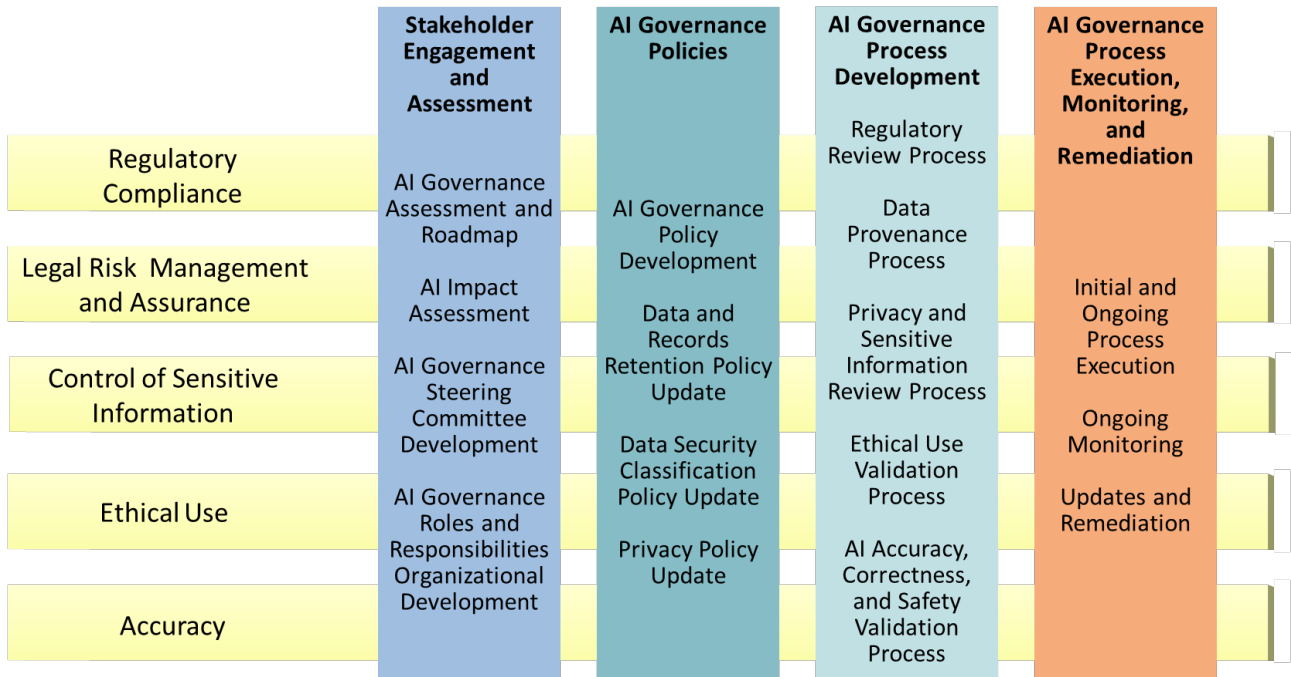
By Mark Diamond

This is Part 2 of a two-part series. Part 1 addressed the risks and restrictions organizations face in deploying artificial intelligence (AI) and the key elements of an AI strategy.^[1] This part details how to develop an AI governance function.

Steps in creating an AI governance function

The breadth of AI governance requirements can be overwhelming. Instead of addressing everything at once, break program development into a series of steps (see Figure 1).

Figure 1: AI governance function overview



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 early in the process.

Assessment and roadmap

Assess current and targeted compliance and risk requirements, engaging key stakeholders in the process. Develop a comprehensive AI governance roadmap reflecting the steps necessary to reach your targeted maturity level. Ideally, your AI governance roadmap will mirror AI application development steps, encouraging governance to be built into AI applications—not as an afterthought. 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 should be borrowed from other elements in this step.

Steering committee

AI governance is complex, requiring the participation of a variety of stakeholders, including compliance, risk, legal, privacy, information governance, data governance, IT, and even business functions. A cross-functional steering committee is the most effective model for managing governance on an ongoing basis. This committee should be formed early in the process 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 different groups. Establishing ongoing roles, and responsibilities as well as organizational design, ensures that AI governance will be an ongoing, continual process and not a one-time exercise.

Many AI initiatives start and are developed within IT and have little interaction with other stakeholders until the system is ready for deployment, often leading to delays when it needs to be vetted or redesigned. Compliance professionals and other stakeholders should engage with IT early, offering to work with them to help navigate these complex environments, explaining that designing governance *into* a system is much faster and easier than trying to retrofit it on the backend.

AI governance policies

The next step is developing and updating your governance policies. An AI governance policy sets out the organization’s compliant, transparent, and ethical use of AI. It details how AI should be used, safeguards employees, and ensures compliance with regulatory requirements. This is your overall “guiding light” to demonstrate to others that you are using AI responsibly.

Data retention/records retention policy and schedule

Organizations may need to update their data retention policies or records retention schedules. The traditional poor practice of saving “all electronic information forever” creates risks that old, legacy data “polluted” with sensitive or incorrect information might be used to develop 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 sensitive factors. Organizations may also need to update this policy to ensure appropriate controls are placed on sensitive information.

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 in the event an AI system faces review from a regulator. These policies demonstrate that the organization is mindful in its use of AI and is diligent of its compliance efforts.

AI governance process development

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 changing on a nearly weekly basis. Organizations need to develop a process for monitoring regulatory changes to ensure their systems comply with any new rules. While most of these rules provide a grace period for implementation, understanding upcoming requirements earlier enables better design decisions and compliance when enforcement begins.

As the saying goes, the best defense is a strong offense, and the same applies to AI governance. All AI governance processes should be completed regularly, and the results of these processes should be retained. Any issues, discrepancies, or problems should be noted, along with steps taken to remediate these issues. AI governance is like other types of compliance in that the courts and regulators realize AI is an imperfect system. 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 your system more defensible. Being proactive may stave off additional regulatory inspection.

Data provenance process

AI systems leverage both training data used by large language models and supplementary information used as for retrieval augmented generation. Companies initially need to undertake reasonable due diligence to ensure 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, provenance must be ascertained 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. The legal issues around AI and copyright are complex, and the case law is in its infancy. Additionally, as these are closed systems, inspecting what training data was used it is not possible. Deployers 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, when asked to create a picture of “senior executives,” a visual generative AI application should not consistently create an image exclusively consisting of older white males. 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, legal assuredness, and accuracy, AI needs to be accurate, correct, and safe. AI’s polished output can lull a user into believing that all the information it produces is correct and accurate. Correctness and accuracy 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.

What to do about AI hallucinations?

AI large language models can create hallucinations or generate inaccurate or incorrect output. Hallucinations include AI-generated legal research that contained fictional court decisions (for which the attorneys were sanctioned) or references to a nonexistent book, including a made-up ISBN number. Some are concerned that AI hallucinations are an inherent risk of the technology. Rather, AI engineers argue that hallucinations are borne from poor design, limitations, or biases in training data; poor testing; and finally poor review. Consistent use of AI governance processes along with thorough testing and review can largely eliminate these risks.

AI governance process execution and ongoing audit

The next step is execution. Once governance processes are developed, they need to be executed and applied both throughout system development and on a scheduled, repeated, ongoing basis after launch. Regulatory requirement changes should drive updates to policies and possibly a review of system design. Input data should continue to have its provenance established, and any issues detected should be resolved. Likewise, companies need to be vigilant that any supplementary data inputs do not contain sensitive information.

Creating a compliant and defensible program

How can an organization ensure it is using AI in a compliant way? In short, trust your processes. Develop a comprehensive approach with the appropriate governance processes. Regulators realize that the rapid growth of AI makes regulating the technology itself very difficult. Therefore, most regulators are focusing on how organizations are using AI. Companies need to demonstrate that they are trying to follow the rules, even if the rules themselves are general or not prescriptive.

Perhaps the biggest risk with AI governance is waiting for clear and well-prescribed regulatory and legal rules. The regulatory and legal uncertainty surrounding AI is not going away soon, and the wait for complete clarity could be a long one. Build out your policies and processes despite the gray areas. Be consistent in your execution. A well-executed, well-intentioned, but slightly imperfect approach is more compliant and defensible than waiting. Don’t let perfect be the enemy of good.

Final thoughts

AI and its required governance present a challenge for organizations. This new, complex technology faces a somewhat chaotic legal and regulatory environment. This challenge represents a leadership opportunity for in-house compliance professionals to change the tug-of-war conversation of “we have to do this” versus “this is what can be done” to let us work together and focus on what we can do. Compliance professionals to the rescue.

Takeaways

- Companies deploying artificial intelligence (AI) should create a comprehensive AI governance function.
- The first step is engaging key stakeholders and assessing needs.
- Next, organizations should create and update AI governance and related policies.
- The diverse requirements of AI require developing a series of governance processes.
- As systems develop and as they are brought into production, these processes need to be executed and monitored, and issues need to be remediated.

¹ Mark Diamond, “Creating an AI governance function: Part 1,” *CEP Magazine*, March 2024, <https://compliancecosmos.org/creating-ai-governance-function-part-1>.

This publication is only available to members. To view all documents, please log in or become a member.

[Become a Member Login](#)