


Compliance Today – May 2024



Marlyse Y. McQuillen (mcquillen.marlyse@gmail.com, [linkedin.com/in/marlyse-y-mcquillen/](https://www.linkedin.com/in/marlyse-y-mcquillen/)) is the Vice President, Regulatory Compliance and Privacy at Integra Connect LLC in Boca Raton, FL.

Implementing AI governance: Tips from the trenches

by Marlyse Y. McQuillen

Tired of the onslaught of promotional emails promising artificial intelligence (AI) tools that deliver increased efficiency while lowering operating costs? Wondering how many more of those emails your executive team is receiving? Up with night sweats at the thought that employees are entering confidential company data into ChatGPT? Or are you spurred on by the Biden administration's commitment to advancing the responsible use of AI in healthcare as seen in the October 30, 2023, Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (E.O. 14110)?^[1]

If the answer to these questions is yes, the time has come to put together an AI governance program.

So, where to start? Put together a list of guiding principles for the governance program before engaging other stakeholders. While some principles may differ based on your organization's products and services, key principles to keep in mind are management of uncertainty, flexibility, and transparency.

Guiding principles

Regardless of what sector of the healthcare ecosystem your organization occupies, regulatory activity following E.O. 14110 is bound to have an impact—if it has not already—given the January 2024 publication of Office of the National Coordinator (ONC)'s final rule on Health Data, Technology and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing (HTI-1 Final Rule).^[2] Since the scope of impending regulation remains unclear, your program needs to be able to manage that uncertainty. Moreover, the program should have built-in flexibility to account for new initiatives as well as oversight in the face of potential enforcement.

In E.O. 14110, the Biden administration simultaneously tasks the secretary of the U.S. Department of Health and Human Services (HHS) to advance responsible AI innovation in the healthcare sector, establish an HHS AI task force with a clear strategic plan by October 2024 involving policies as well as potential regulatory action on the use of AI and AI-enabled technologies and establish a mechanism for reporting and remediating unsafe healthcare practices involving AI. Given this scope, the governance infrastructure needs to be nimble and capable of overseeing new adoption of AI, AI use cases, and the implementation of wide-ranging regulation. It must also be flexible enough to respond to lessons learned from potential enforcement actions.

While certified health IT developers are busy operationalizing the HTI-1 Final Rule, other organizations await the release of guidance. What the HTI-1 Final Rule does tell us is the importance of transparency and disclosure in the use of AI. Establishing a new certification criterion for predictive algorithms used in healthcare IT, the HTI-1 Final Rule requires developers to disclose baseline information about algorithms to enable clinicians to assess

fairness, appropriateness, validity, and safety. (See 45 C.F.R. § 170.315(b)(11) Decision Support Intervention, which replaces 45 C.F.R. § 170.315(a)(9) Clinical Decision Support in the ONC certification standards.) More specifically, developers must describe the training data used for the algorithm/AI tool and provide information on whether it was tested and evaluated for fairness. For certified health IT developers, all this information is subject to online disclosure on the developer's registry on the Certified Health IT Product List. Even if you operate in another healthcare sector, appreciating the extent to which your organization may need to disclose its use of AI—particularly if automated decision-making results—should guide your approach to the oversight of AI tool development and AI vendor onboarding. Before greenlighting a development project or AI vendor, the governance committee should be able to describe the training data used for the proposed AI tool and have a sense of whether and how the vendor or internal user evaluated the tool for fairness, validity, and patient safety, as appropriate.

Privacy by design

Once you have established the guiding principles for an AI governance program, educate yourself and other stakeholders on privacy-by-design. In plain language, privacy-by-design is the concept that privacy safeguards and data protection should be incorporated at the earliest stage of the design process for a tool or technology. For stakeholders fluent in HIPAA, it enshrines the technical guardrails required by the HIPAA security rule and minimum necessary standards for data collection, processing, and access into the development process at the start rather than as an afterthought. (See the minimum necessary standard, a key protection of 45 C.F.R. §§ 164.502(b), 164.514(d), and 45 C.F.R. Parts 160, 164.A and 164.C.) In AI development, there is a natural tension between the use of diverse and representative training data and the collection restraint imposed by the minimum necessary standard. Understanding that tension and the potential issues that may arise with sweeping state privacy legislation granting data subjects rights to access their data and correct and delete it is critical to asking developers and proposed vendors the right questions. Ensuring that operational and risk management stakeholders are aware of those risks and speak the same language as developers and vendors is essential.

Stakeholders and structures

Aside from compliance, legal and information security serve as the usual suspects in any technology governance program involving technology but are crucial in the AI context given the heightened risk AI may pose to the organization's intellectual property (IP) assets and security infrastructure. Though the United States Patent and Trademark Office published in February guidance granting patents to AI-assisted inventions, the determination of IP rights and what constitutes fair use of generative AI outputs is unclear.^[3] Legal input is crucial to evaluating not only IP risk but also working with information security to determine if AI vendors have sufficient safeguards to prevent data leakage and resources to indemnify the organization should those safeguards fail. In addition to risk management stakeholders, consider adding operational stakeholders with the largest volume of AI usage requests to join in oversight.

Once you decide on stakeholders to include in AI oversight, evaluate existing governance structures to see if there is an overlap in participants and current proposed AI use cases. Creating an entirely new AI governance committee may not be necessary if an organization has a robust vendor management infrastructure and no plans to engage in the internal development of AI tools. Alternatively, if an organization plans to exclusively engage in internal AI development, consider whether the infrastructure or architecture review committee includes appropriate risk stakeholders. If done appropriately, revising existing committee charters, policies, and vendor and audit checklists—and getting those revisions approved—will save time and avoid unnecessary and unfamiliar administrative processes. That is particularly true if those committees satisfy requirements for flexibility and transparency and have proven dexterous in the face of fast-changing regulations.

The evaluation of existing governance structures may also pose an opportunity to mature an organization's vendor management program or cut down on silos between business divisions. Approaching stakeholders and achieving consensus on a common vision for overseeing the organization's adoption and use of AI by expanding an active governance body is preferable to sending kickoff meetings to yet another compliance subcommittee.

Documentation and release

Once you decide where the AI governance structure will reside, drafting or revising guidelines and policies—if not committee charters—will occupy a fair amount of time. Engaging stakeholders in the review process will speed up the initial approval and maintain the momentum achieved from your initial AI governance conversations. AI-related review requests will likely pile up during a protracted approval process. Also, evaluate the speed of granting approvals of policies versus procedures and desk references when structuring requirements that may change with the publication of new HHS guidance. For governance documents, consider whether all stakeholders should have a vote or merely participate in meetings and whether the governance committee's scope will include metrics oversight after AI tool integration or AI vendors onboarding.

Upon approval of governance documents and guidance, take an enterprise approach to publicizing the AI policies and governance structures. Though most inquiries regarding AI use may come from operations, shared services departments such as HR may already be using recruiting tools powered by AI without undergoing vetting for fairness and antidiscrimination, concepts also covered in E.O. 14110. Be prepared to have conversations about grandfathering AI tools or submitting use cases to ratify through the new governance procedures.

This document is only available to members. Please log in or become a member.

[Become a Member Login](#)