

Now Is The Time For AI Governance

Thomas G. Stephens, Jr., CPA, CITP, CGMA

Artificial Intelligence (AI) is everywhere—and that much is clear. Given the advancements over the past four years, it is hard to imagine any business professional who is not aware of what AI can do. But how many are equally aware of the risks? And how many organizations still have not taken meaningful action on AI governance? Ignoring AI governance can lead to consequences ranging from embarrassing mistakes to the compromise of private information—and, in extreme cases, serious bodily injury or loss of life. Read on for practical actions you and your organization can take to address the growing need for AI governance.

It Starts With AI Governance Policies

Although most organizations have policies on data security, privacy, and even dress codes, many have not yet addressed AI governance. With the rapid growth in business use of AI, that gap can yield results ranging from embarrassing to catastrophic. For example, suppose a team member in a public accounting firm uploads a client's tax return to a popular AI platform to look for errors or inconsistencies. Does that create a risk? The answer is *maybe*. If the platform does not use uploaded data to train its language model, the action may be harmless. If it does, however, the firm may have created a data privacy issue that must be addressed immediately.

This example shows why AI governance policies should be a primary concern for business professionals today. In most cases, policies should clearly prohibit team members from including private or sensitive data in AI prompts or uploads. Just as important, employees should be trained on what counts as sensitive data, why these restrictions exist, and the potential consequences of violations—including disciplinary action up to and including dismissal.

A strong AI governance policy typically begins with scope and definitions. It should clarify what the organization means by "AI" (for example, generative AI, machine learning models used for scoring, and automated decision tools). It should also spell out what "use" includes—prompting, uploading files, connecting an AI tool to internal systems, or using AI-generated output in client deliverables. Defining these basics reduces the chance that someone violates the policy simply because they assumed it applied only to one type of tool or one department.

AI Platform Policies Are Also Necessary

In addition to governance policies, organizations should give team members clear guidance on which AI platforms are approved for organizational business. For example, an employee might prefer a specific AI tool for personal use—but that tool may not meet the employer's requirements for confidentiality, security, or administrative controls. Without realizing it, the team member could jeopardize sensitive information by using a platform that lacks adequate data protection features. Accordingly, employers should publish an approved list of platforms

(and account types) and make it clear that unapproved tools are not to be used for company work.

Creating an approved list is only half the work; the other half is performing basic due diligence on the tools that make the cut. At a minimum, the organization should understand how the vendor handles uploaded content, whether content is used for model training by default, and what settings exist to opt out. It should also assess where data is stored, how it is encrypted, and what the vendor's retention and deletion practices are. Many teams also evaluate whether the platform supports administrative controls, single sign-on, role-based access, and audit logs—features that make it easier to manage AI use at scale.

Require Accuracy and Intellectual Property Policies

No less important than the two policy types outlined above, organizations that allow AI use should also address accuracy and intellectual property issues. For example, suppose a team member uses AI to help solve a problem for a client. Now suppose the AI-generated solution is incorrect (yes, that happens) and no one catches the error. At best, the situation will be embarrassing. At worst, it could lead to the loss of a key client or, in extreme cases, serious injury or loss of life. For these reasons, firms should require AI-generated content to be verified for accuracy before it is used to inform decisions or client deliverables.

Verification should be defined in practical terms. For numerical work, this might mean recalculating results independently, tying amounts back to source documents, or running reasonableness checks. For research and writing, it can include checking primary sources, confirming quotes, and ensuring that any citations actually exist (AI tools sometimes generate “hallucinated” references). For code, it can include peer review, unit testing, and security scanning. The key point is that AI output should be treated as a draft or a suggestion—not as an authority—unless and until it has been validated.

This is also a professional judgment issue. Whether you work in accounting, consulting, healthcare, manufacturing, or government, the organization—not the algorithm—remains responsible for outcomes. Policies should therefore make it explicit that employees may not delegate accountability to an AI tool. In many cases, it is wise to document material uses of AI in workpapers or project files, including what tool was used, what inputs were provided at a high level, what output was received, and what steps were taken to validate the result. This documentation becomes invaluable if questions arise later.

Additionally, consider work products created with the assistance of AI platforms. A common question is: “Who owns the intellectual property (IP) created in part with help from a firm-paid AI subscription?” Does the IP belong to the employee(s) who developed the deliverable, or to the firm? Don't take the chance. Create a policy that addresses ownership explicitly and states that the firm owns AI-assisted deliverables created by employees in the course of their work.

Summary

Let's be clear—AI is here to stay, and it offers opportunities unlike anything we have seen before. But AI also injects real risk into business environments. Organizations that establish policies to manage these risks will be better positioned to keep them at a prudently acceptable level; those that do not may regret the oversight.

1. Publish an interim AI use policy that prohibits sharing sensitive data and requires human verification of AI output.
2. Identify and approve a short list of AI platforms and approved account types for business use.
3. Train team members using realistic, role-specific examples (including what to do when they are unsure).
4. Assign ownership for governance, periodic reviews, and an exception/approval process for higher-risk use cases.

Then, as usage grows, refine the policies to address higher-risk scenarios, including client-facing deliverables, automated decision support, and system-to-system integrations. Which choice will you make?

Tommy works with K2 Enterprises. At K2, Tommy focuses on creating and delivering content. You may reach him at tommy@k2e.com, and you may learn more about K2 Enterprises at www.k2e.com.