Visibility and Control: The two things missing from every AI Governance Strategy

Artificial Intelligence, as everyone knows, has been the supercharger of innovation and productivity inside of businesses all over the world. The question, however, is who is governing the use of it?

This is a moment where businesses that do not adopt AI risk being left behind, but the security implications that come with the rapid implementation of AI also need to be considered. Palo Alto, one of the leaders in Cloud Security and AI governance, has produced a white paper on the proper frameworks and guidelines that they believe organizations should follow in order to maintain control of the AI in your organization. There is no universally accepted standard for AI governance yet, but business leaders still need to be aware of the consequences it could bring if nothing is done about it. If implementing AI into organizations can help increase productivity, imagine how fast it could leak private information as well.

Palo Alto's message for maintaining AI usage inside of businesses is simple, Visibility and Control. Most organizations already follow guidelines on how to maintain information security, data privacy, and operational trust; AI governance is no different. Keeping an eye on who has access and giving approval on what can and cannot be done is what is going to keep businesses operating securely and efficiently with AI. Palo Alto breaks this down into five areas of focus: Models, Data, Use Cases, Access, and Compliance. Each carries its own weight, and while not every category applies equally based on business needs, all five are worth understanding.

The starting point is Models, and knowing what AI tools are actively being used across the business. Without an inventory of what is deployed and who authorized it, there is no baseline to govern from. Documentation does not need to be complex, it simply needs to exist. Maintaining this visibility is what then provides the control to test and validate these models on a documented basis, consistency that also satisfies governing standards like SOC 2 Type II.

Data is the foundation of every AI output, what goes into a model determines what comes out of it. Organizations that cannot account for what information is feeding their AI systems cannot account for what those systems might expose. Classifying data by sensitivity, whether it’s for training, fine tuning, or production, and knowing where it lives is the minimum standard required before AI can be used responsibly. If the integrity of that data is compromised, the consequences extend far beyond a bad output. Palo Alto notes that a poisoned model cannot simply be corrected, and with machine unlearning still in its early stages, retraining a model is often the only option and will require hundreds of thousands of dollars.

Use Cases define where AI is appropriate and where it is not. Not every business question should be answered by feeding sensitive information into a public AI tool. Defining boundaries for appropriate use is not about limiting productivity, it is about ensuring that productivity does not come at the cost of security, or legal action. Organizations that can demonstrate controlled and documented AI governance have a competitive advantage; it is becoming a prerequisite for doing business, not just being a best practice.

Access should reflect the sensitivity of the work being done, and the roles and responsibilities of who is using it. Not every employee requires the same level of access as the other, and Palo Alto identifies guard rails that can be used like rate limiting to keep the usage in check. Access without defined boundaries is just an open door, and creating documentation on who has access to what, gives organizations the control to close that door if need be.

Compliance ensures that AI usage aligns with the regulatory frameworks already governing the business. For organizations in regulated industries this is not optional, it is a legal obligation to continue operating at that level. Even outside regulated industries, the absence of compliance around AI creates exposure that will only become more difficult to manage as regulation catches up to adoption. If policies need to be created, they should align with the risk tolerance the organization is willing to face.

This sounds easy enough to maintain, but the issue is dealing with a problem organizations cannot see. Everyone is using AI very quickly. If there is a problem or a question that doesn't have an answer, the instinct is to turn to AI for a solution. Sounds like no problem, but the information a user gave in order to produce that result is where the problem lies. That information used to produce the result doesn't disappear, and in most organizations nobody knows it was shared in the first place. Palo Alto calls this the Shadow AI problem, where employees within an organization use Artificial Intelligence without the approval, knowledge or oversight of the IT and security teams. The risk is not hypothetical, it’s already happening inside organizations whether leadership knows it or not. The issue is it’s only going to get worse as AI continues to advance and why leaders in the industry like Palo Alto want to identify it before it’s too late.

AI is not going away, and neither is the risk that comes with ungoverned use. The frameworks exist, the directions are clear, and the consequences of inaction are well documented. AI is meant to be a tool that anyone can and should use. The question remains, if no one is governing the use of it, who will?

References:

Palo Alto Networks. *AI Security and Governance: A Framework for Securing AI-Powered Applications.* Palo Alto Networks, 2024. https://www.paloaltonetworks.com/apps/pan/public/downloadResource?pagePath=/content/pan/en_US/resources/whitepapers/ai-governance