




SMARTER PERSPECTIVE: ARTIFICIAL INTELLIGENCE

AI Security Starts With Foundational Governance

By Mack Bodie

October 2026

Boards and executive teams are under pressure to have an AI strategy. Most already know that AI is active within their environment, from individual employee use to enterprise-wide deployments. What they lack is a structured methodology for asking the right questions. These include: How far has AI's use actually spread? What does it currently touch across the enterprise? How can and should it be governed to maintain proper controls and limit risk?

Generative AI has moved into the enterprise faster than any technology in recent memory. Employees are using AI tools to draft documents, analyze data, summarize meetings, and speed up decisions across nearly every department. Much of this adoption happens without formal review. An employee signs up for a free AI tool, connects it to a work or personal account, and starts feeding it company information within minutes. No one in IT or security ever sees it happen. Security teams have a name for this activity: Shadow AI. The tools themselves are rarely the problem. What expands faster than most governance programs can track is the attack surface and the data exposure it creates.

For middle market companies, many of which are already stretched thin on security staff and budget, this creates a specific kind of pressure. Leadership wants to capture the productivity gains AI offers. At the same time, they cannot fully explain or track how AI is being used across the business, what data it can access, or what happens if that data ends



up somewhere it should not. For regulated entities, the loss or exposure of sensitive data is more than an organizational issue and can become a significant regulatory and financial event. Regulators increasingly treat a data incident as a test of whether required safeguards were in place before the breach and whether the organization responded appropriately afterward. The consequences can be substantial. In 2025, HHS imposed a \$1.5 million civil money penalty on Warby Parker following a hacking investigation involving HIPAA Security Rule violations. Previously, New York's Department of Financial Services imposed a \$5 million penalty on Carnival for cybersecurity violations that contributed to the exposure of sensitive customer data. For organizations subject to the European Union's GDPR, serious violations can carry administrative fines of up to €20 million or 4% of worldwide annual turnover, whichever is higher. In other words, for regulated organizations, protecting

sensitive data is not simply a matter of limiting operational or reputational damage. Inadequate controls can translate directly into regulatory enforcement, corrective-action requirements, and material financial penalties.

The stakes rise further for private equity-backed businesses and portfolio companies operating across multiple entities. A single AI governance gap at one portfolio company can create exposure that ripples across an entire platform, particularly when shared systems, shared vendors, or shared customer data are involved. Investors evaluating add-on acquisitions or preparing a company for exit should expect AI governance to become part of that scrutiny, alongside the security questions that have long been standard in diligence

Cut Through the Noise

Much of the public conversation about AI security treats the technology as an

entirely new category of risk that requires an entirely new set of tools. That framing is not accurate, and it is not helpful to executives trying to make sound decisions with limited resources.

AI is a new technology, but it does not introduce a new category of business risk. The fundamentals of managing technology risk have not changed: know what is running in your environment, control who can access it, protect the data that flows through it, and put governance in place to keep all of that current as the business changes. AI simply raises the stakes on how quickly a company needs to do those things well.

Organizations that already have a functioning security program are closer to being AI-ready than they may realize. Identity management, endpoint monitoring, data loss prevention, and application oversight are not AI-specific tools. They are core security capabilities that most companies already have in some form. The work is less about building something new from scratch and more about extending what already exists to cover a new class of applications and use cases.

Five Questions Executives Should Be Asking

AI governance works best as a set of leadership questions rather than a technical checklist. Each one maps to a control area that security teams can build out, but the starting point belongs at the executive level.

- Start with visibility. Can the company see what AI tools are actually in use? Most organizations underestimate this number significantly. Formal visibility into approved and unapproved AI applications is the foundation for everything that follows. A company cannot govern what it cannot see.
- Access comes next. Is it managed with the same discipline applied to other business systems? Single sign-on, identity governance, and access policies that already protect financial systems, customer data, and internal networks should extend to AI platforms as well. AI should not become a side door around



controls that the company has already invested in.

- Detection matters just as much. Are existing endpoint and monitoring tools configured to recognize AI-related activity? Many companies have strong detection capabilities that were built before generative AI existed and have not been updated to flag AI-specific behavior. A tool that cannot see a risk cannot help a company respond to it.
- Data protection deserves the same scrutiny. Is sensitive data protected from exposure through AI platforms specifically? Proprietary information, customer records, and regulated data can move through AI tools in ways that bypass traditional data protection measures. Data loss prevention controls need to account for this new pathway, not just the familiar ones.
- Governance ties the rest together. Does the company have actual governance in place? Technology controls alone do not solve this problem. Policies, employee training, and clear ownership of AI-related decisions are what turn a set of tools into a durable program. Without governance, controls drift out of date the moment the business changes, and with AI, the business is changing constantly.

A leadership team that can answer these five questions with confidence is in a fundamentally stronger position than one relying on a single tool purchase or a policy memo that no one has read since it was written.

Cost-Effective Does Not Mean Complicated

Middle market companies often assume that a serious AI governance program requires the kind of budget and staffing available to much larger enterprises. That assumption is worth challenging. Because so much of the work involves extending existing security investments rather than replacing them, a practical AI governance program can be built incrementally and cost-effectively.

The starting point is usually an honest assessment of what a company already has, what it is missing, and where the biggest gaps sit relative to how the business actually uses AI today. From there, priorities can be sequenced in a way that matches available resources, rather than pursuing every control area at once. This sequencing matters as much as the controls themselves. A company that tries to address visibility, identity, endpoint monitoring, data protection, and governance simultaneously often ends up with a program that is broad but shallow, and difficult to sustain once the initial push fades. A company that sequences the work, starting with visibility and building outward, tends to end up with a program that actually holds up over time. Importantly, the goal is not about arriving at a finished, static state program. Instead, it should be viewed as a dynamic, disciplined process that establishes a habit of ongoing reviews and updates focused on keeping pace with how the business evolves as it adopts new tools.

This is also where the human element



matters most. Technology can flag unapproved applications and unusual data movement, but deciding what level of AI use is acceptable for a given role, or how a company should balance productivity gains against exposure, still falls to people. Those are governance decisions, and they require people asking the right questions on an ongoing basis, not a one-time project.

The Constant Underneath the Change

AI will keep evolving, and the pace of that change is unlikely to slow down. New tools, new use cases, and new ways of working will continue to appear faster than most companies can formally evaluate them. That reality can feel overwhelming from the outside.

What matters is that the discipline required to manage this risk is not new, even as AI's evolution and impact within a business remain hard to predict. Companies that have built strong security fundamentals, established clear ownership of technology risk, and a habit of ongoing governance, are already well positioned to adapt as

AI changes. Companies without those fundamentals will likely find that AI exposes these gaps faster than almost anything else has in the past.

We have seen that the organizations that get this right tend to be defined less by the number of AI tools they deploy or the pace of their adoption timelines, and more by leadership that asks the right questions early, builds governance around the answers, and keeps asking as the business and the technology keep moving.

For executives weighing where to start, the conversation is often more valuable than the checklist. A candid discussion about how AI is actually being used inside a business and measured against the controls already in place, tends to surface priorities faster than any generic framework. Companies evaluating their own exposure, or preparing for an acquisition, exit, or platform integration where AI governance will come under scrutiny, are well served by having that conversation with a proven team of expert cyber security advisors sooner rather than later.

About Hilco Global Cyber Advisors

Hilco Global Cyber Advisors is a dedicated cybersecurity, data governance, privacy, technology and AI risk advisory practice that brings decades of experience and a vast network to help growth, middle market, and nonprofit organizations strengthen security, drive resilience, and unlock growth. Our mission: equip clients with strategies, tools, and insights to protect their operations, customers, and reputations in an increasingly digital and high-risk environment.



MACK BODIE
DIRECTOR
GLOBAL CYBER ADVISORS
PROFESSIONAL SERVICES

Mack can be reached directly at
mbodie@hilcoglobal.com.