Technology

From AI Adoption to AI Control

by Dorene Rettas 7 min read (June 8, 2026)

From AI Adoption to AI Control 14:06

For the past two years, conversations about artificial intelligence have largely centered on opportunity. Boards have pushed executives to accelerate adoption. Business leaders have looked for productivity gains. Employees have embraced copilots, assistants, and AI-powered workflows that promise to eliminate repetitive work and speed decision making.

Today, however, the conversation has changed.

CISOs are no longer being asked whether their organizations should adopt AI. In many cases, that decision has already been made. Instead, security leaders are facing a new set of questions. What data can AI access? How are AI agents being governed? What happens when an AI system behaves unexpectedly? How can organizations defend against AI-powered attacks? And perhaps most importantly, how can security teams maintain control as AI becomes embedded across the enterprise?

During Gartner's Security and Risk Management summit I had discussions with ThreatLocker, Abnormal AI, AvePoint, Xona Systems, DeepKeep, and Menlo Security, and one theme surfaced repeatedly: the challenge is no longer AI adoption, the challenge is AI control.

Visibility Comes Before Governance

For many organizations, the first challenge is understanding what is happening inside the environment. Security leaders cannot govern what they cannot see, and the rapid growth of AI agents is making visibility more difficult than ever.

Abnormal AI believes traditional security approaches are struggling to keep pace with the speed and scale of modern attacks. According to Mick

defenses are becoming less effective. As he noted, "Every attack today can now be a zero-day attack."

Instead of focusing exclusively on known indicators of compromise, Abnormal AI advocates a behavioral approach. The goal is to understand what normal looks like across users, devices, locations, communications, and activities. Once a baseline is established, organizations can identify behavior that falls outside that norm.

That philosophy becomes particularly important as AI agents gain access to email systems, calendars, collaboration platforms, and business applications. Organizations increasingly view AI agents as digital workers, but Abnormal AI argues they should also be treated as non-human identities that require governance, monitoring, and oversight. If an AI agent suddenly begins accessing data, systems, or resources outside its expected role, security teams need visibility into those actions before damage occurs.

Leach sees behavioral AI as a way to help organizations adapt to a world where both attackers and defenders are leveraging artificial intelligence. In that environment, understanding behavior may ultimately become more valuable than relying on historical attack signatures.

Containing What AI Can Access

If visibility is the first challenge, containment is the second.

In a discussion with Danny Jenkins, CEO and Co-founder of ThreatLocker he argues that organizations should stop assuming software and AI systems will always behave as expected. Instead, security teams should focus on restricting what applications can access and what actions they can perform.

Jenkins believes many organizations continue to invest heavily in detection and response technologies while neglecting foundational controls that

fingerprinting. The concept is straightforward, whether an organization is deploying an AI agent, an accounting application, or another business tool, that technology should only have access to the resources required to perform its intended function. If an AI assistant needs access to specific files, databases, or workflows, it should not automatically receive broad access to everything else.

Jenkins noted that AI introduces a unique challenge because a system can make decisions that are not malicious but can still create significant risk. A model might delete data, access information it should not see, or interact with systems in unintended ways. Restricting permissions reduces the potential impact when something goes wrong.

ThreatLocker also views AI-powered vulnerability discovery as a major concern. As AI accelerates the process of identifying weaknesses in software, organizations must assume vulnerabilities will be discovered faster than ever before. In that environment, limiting the blast radius becomes increasingly important. They believe the future of cybersecurity may depend less on detecting attacks after they occur and more on preventing applications, users, and agents from exceeding their intended scope in the first place.

Governance Must Become Operational

Many organizations have already created AI governance frameworks. The problem is turning those frameworks into operational reality.

AvePoint approaches AI governance from the perspective of data and data protection. The company frequently reminds customers that AI may be the engine, but data remains the fuel. If organizations do not understand their data, permissions, classifications, and sharing practices, they cannot fully trust the AI systems built on top of them. That challenge becomes particularly significant in large enterprises where information is spread across collaboration platforms, file repositories, cloud environments, and years of accumulated business content. Employees may have access to

Timothy Boettcher, Sr VP, Product Marketing at AvePoint, shares that the arrival of AI changes the equation because information that was once difficult to find can now be surfaced almost instantly. An employee may only see data they technically have permission to access, but if those permissions were never appropriate in the first place, AI can expose governance problems that have existed for years. AvePoint focuses on helping organizations classify, protect, govern, and recover their information at scale. The company has also invested in AI-driven capabilities that help customers translate governance requirements into practical policies and controls.

One observation from the discussion stood out. Creating policies is relatively easy. Operationalizing those policies across hundreds of thousands of files, collaboration spaces, users, and workloads is where organizations struggle.

For CISOs, that distinction matters. Governance cannot remain a document stored in a compliance repository. It must become a living, enforceable framework that operates continuously across the business.

Securing AI Before It Goes Into Production

Another challenge emerging rapidly is the security of AI systems themselves.

Ben Stringer, VP Sales at DeepKeep compares the current moment in AI security to the evolution of application security. Years ago, organizations learned that security testing needed to move earlier in the software development lifecycle. The same lesson is now being applied to AI. According to Stringer, organizations should evaluate agentic applications before they are deployed rather than attempting to address security concerns after implementation. The company positions itself as an AI security platform, providing an AI firewall for runtime protection, execute

The rise of external and internal AI applications has increased the urgency of that work. Whether organizations are deploying chatbots, internal assistants, knowledge platforms, or automated decision-making systems, every deployment creates new opportunities for misuse, manipulation, and data leakage.

DeepKeep believes security leaders must ask difficult questions before deployment. Can the application be manipulated? Can it be tricked into exposing sensitive information? Does it stay within its intended purpose? Can users exploit it to gain access to information they should not receive? The company's emphasis on red teaming reflects a broader industry shift, from automated processes into a hybrid approach called vibe AI red teaming. Security teams are increasingly recognizing that AI systems must be tested the same way they test applications, networks, and infrastructure.

The objective is not to slow innovation. It is to ensure organizations understand the risks before AI becomes deeply embedded in critical business processes.

Protecting AI Agents in Real Time

While some organizations focus on securing AI models, others are focusing on the agents that use them.

Ramin Farassat, Chief Product Officer at Menlo Security states that AI agents are rapidly becoming a new class of user, "The next billion users are going to be AI agents."

That prediction is already influencing how organizations think about security. AI agents increasingly browse websites, collect information, access SaaS applications, automate tasks, and interact with business systems. In many cases, they are performing activities that previously required human involvement.

and malicious content, AI agents can be influenced through prompt injection, hidden instructions, manipulated content, and other emerging techniques. To address that challenge, Menlo Security has extended its browser security expertise into the AI space. The company focuses on providing visibility into agent activity, authenticating agents, enforcing policies, and protecting against runtime threats.

A key differentiator in Menlo's approach is its focus on real-time security. The company argues that organizations cannot wait until after an AI agent completes a task before determining whether something went wrong. Security controls must operate while the task is taking place. Farassat also highlighted evidence that attackers are already targeting AI agents. Campaigns designed to manipulate agent behavior are no longer theoretical, they are beginning to appear in real-world environments.

That reality reinforces a broader point, as AI becomes more autonomous, security must move closer to the moment decisions are being made.

Zero Trust Reaches Operational Technology

While many AI discussions focus on office productivity, cloud environments, and enterprise applications, operational technology presents a different set of challenges.

Xona Systems focuses on secure access for industrial control systems, critical infrastructure, utilities, manufacturers, transportation providers, and other operational environments. These organizations often depend on third-party vendors, remote technicians, and specialized equipment that require access to sensitive systems.

As Bill Moore, Founder and CEO of Xona stated, AI-driven attacks increase the importance of access control in these environments. If vulnerability discovery becomes faster and more automated, organizations must assume that critical systems will face greater scrutiny from attackers. Xona's

and provides detailed insight into how systems are accessed and used, creating a stronger foundation for both security and forensic investigations.

At the same time, Xona is exploring ways to use AI defensively. The company is developing approaches that analyze how users interact with operational technology systems and identify anomalies that could indicate risk. The importance of this work extends beyond cybersecurity, operational technology environments often support energy generation, manufacturing operations, transportation systems, and other critical services. A security incident in those environments can have real-world consequences that extend far beyond data loss.

As AI reshapes cybersecurity, organizations responsible for critical infrastructure may face some of the highest stakes.

The Future Belongs to Organizations That Can Maintain Control

The vendors interviewed for this article approach the AI challenge from different perspectives, some focus on visibility, others focus on containment, governance, model security, runtime protection, or operational resilience. Yet despite those differences, they all arrive at a remarkably similar conclusion.

AI is becoming embedded across nearly every aspect of business operations, employees are using it, customers are interacting with it, applications are integrating it and vendors are building it into products by default. For CISOs, the question is no longer whether AI should be adopted. That decision has largely been made by the business.

The new challenge is maintaining control once AI is everywhere. Organizations need visibility into what AI systems are doing, they need governance that can be enforced at scale, controls that limit access and reduce risk, confidence that AI models have been tested and secured,

The security leaders who succeed in the coming years will not necessarily be the ones who adopt AI the fastest. They will be the ones who can demonstrate that innovation and control can coexist. That may become the defining cybersecurity challenge of the AI era.

Previous story
← AI Is Reducing the Margin for Error in Cybersecurity Controls

You May Also Like

These Related Stories

Cybersecurity Careers and AI's Impact

(May 21, 2025) 4 min read

Data Security and Governance in the Age of AI

(June 3, 2024) 4 min read

Technology

Turning AI Risks into Strategic Board Conversations

(June 24, 2025) 5 min read

provide thought leadership and expertise from the world's most renowned cyber security professionals.