INSIGHTS | TECH BRIEF

Securing AI Has Become a Board-Level Priority

Most cybersecurity programs were designed to govern human behavior. But AI introduces non-human actors into the enterprise as autonomous systems can access data, generate code, make decisions, and interact with infrastructure.

This changes the security equation entirely. Organizations are now responsible for governing AI agents, machine identities, model interactions, and autonomous workflows operating across cloud, SaaS, API, and endpoint environments.

For cybersecurity leaders, this creates a new governance challenge: *How can organizations securely adopt AI without introducing risk?*

Addressing that challenge requires a unified AI security architecture capable of governing identities, applications, agents, models, and runtime behavior across the enterprise.

As a Palo Alto Networks NextWave Diamond Innovator Partner, WEI helps organizations modernize cybersecurity strategies for the AI era through proactive, intelligence-driven security architectures focused on governance, operational resilience, and continuous risk reduction. Together, WEI and Palo Alto Networks are helping enterprises redefine AI-driven cybersecurity.

AI Has Expanded the Enterprise Attack Surface

AI has fundamentally changed how users interact with data, how software is developed, and how decisions are automated across the business. As organizations operationalize AI, the attack surface expands in several ways.

The challenge becomes even more severe as machine identities now significantly outnumber human identities across enterprise environments. According to Palo Alto Networks, machine identities exceed human identities by more than 80 to 1, while nearly 90% of organizations have already experienced an identity-related breach.[1]

AI Security Risk	Business Impact
Shadow AI usage	Sensitive data leakage and compliance exposure
AI-generated phishing and malware	Increased sophistication of cyberattacks
Autonomous AI agents	Excessive privileges and unauthorized actions
AI coding assistants	Introduction of insecure code and software vulnerabilities
AI plugins and browser extensions	Unmanaged supply chain risk
Machine identity sprawl	Expanded credential attack surface
Prompt injection attacks	Manipulation of AI behavior and outputs

Governing AI Requires a New Security Model

Most legacy security architectures rely on fragmented point products designed for static environments and human-driven workflows. AI changes those assumptions entirely as today's AI ecosystem operates across:

AI Operates Across:	Traditional Controls Lack:
Cloud platforms	Runtime AI visibility
APIs	AI-specific governance policies
SaaS applications	AI usage monitoring
Browsers	AI agent identity controls
Developer environments	Model-level protection
AI gateways	Prompt inspection capabilities
Autonomous agents	Centralized AI risk management
Machine-to-machine communications	Runtime AI visibility

An AI agent operating with privileged access can become the equivalent of a trusted insider operating at machine speed. Traditional endpoint detection and identity tools were not designed to monitor or govern this behavior.

At the same time, employees increasingly interact with unsanctioned AI services outside approved governance processes. Sensitive intellectual property, regulated data, and proprietary information may already be flowing into public AI platforms without organizational oversight.

AI Security Requires a "Left of Bang" Cybersecurity Strategy

Traditional cybersecurity models were built around detection and response. Organizations focused heavily on identifying threats after compromise indicators appeared inside the environment.

AI-powered attacks now move at machine speed. Threat actors are using automation to accelerate reconnaissance, scale phishing campaigns, weaponize vulnerabilities faster, and evade traditional detection models. By the time many organizations identify suspicious activity, operational damage may already be underway.

Borrowed from military threat prevention methodology, "Left of Bang" focuses on identifying indicators of risk before disruption occurs. In cybersecurity, this means proactively reducing attack exposure through offensive testing, threat intelligence, attack path analysis, runtime visibility, identity governance, and continuous validation. This philosophy aligns closely with Palo Alto Networks platform-based security strategy, particularly across Zero Trust architectures.

WEI helps organizations shift toward proactive cyber resilience strategies designed to:

- Reduce exploitable attack surfaces
- Identify weaknesses before adversaries do
- Improve readiness against AI-driven threats

Palo Alto Networks Vision for AI Security

Palo Alto Networks has positioned AI security as a foundational pillar of its long-term cybersecurity strategy. Rather than treating AI security as a standalone niche capability, Palo Alto Networks is building a platform-based architecture designed to secure three dimensions of enterprise AI risk:

Security Domain	Focus
Protection FROM AI	Defending against AI-generated attacks
Protection OF AI	Securing AI models, agents, and applications
Protection WITH AI	Using AI to improve cyber operations

This framework allows organizations to govern AI comprehensively across users, identities, applications, infrastructure, and runtime behavior.

Protection FROM AI: Defending Against AI-Generated Threats

Threat actors are already using AI to accelerate phishing campaigns, automate malware creation, improve reconnaissance, and evade detection. To address this challenge, Palo Alto Networks incorporates AI-powered threat prevention directly into its network security stack.

Advanced URL Filtering

Palo Alto Networks reports that its Advanced URL Filtering service processes more than 41 million AI-related URL queries.[2] This solution helps organizations defend against AI-generated threats while also governing employee access to AI services. Capabilities include:

- AI-powered malicious URL detection
- AI-specific web traffic categorization
- Governance policies for AI application access
- Protection against AI-generated phishing campaigns
- Inline enforcement through NGFW and Prisma Access

AI Access Security: Controlling Shadow AI

One of the fastest-growing concerns for cybersecurity leaders is the rise of Shadow AI — the unauthorized use of generative AI applications across the enterprise. Employees increasingly rely on public AI platforms to summarize documents, generate code, analyze data, and automate workflows. In many cases, these tools are adopted faster than governance teams can assess risk.

The challenge is that organizations often lack visibility into:

- Which AI applications are being used
- What data is being submitted
- Whether prompts contain regulated information
- How AI-generated outputs are being shared or retained

This creates significant governance concerns for organizations operating under regulatory mandates such as HIPAA, PCI-DSS, GDPR, and emerging AI governance frameworks.

Palo Alto Networks AI Access Security is designed to help organizations regain control over AI usage before these exposures become systemic risks. Rather than blocking innovation outright, the platform enables enterprises to discover, assess, and govern AI interactions through centralized policy enforcement and AI-specific risk analysis.

The platform continuously identifies sanctioned and unsanctioned AI applications across the environment while applying contextual controls around data movement, user access, and prompt activity. Integrated DLP capabilities help organizations inspect both prompts and responses for sensitive information exposure, giving security teams greater confidence that intellectual property, customer information, and regulated data remain protected.

Prisma AIRS: Securing AI Runtime Environments

As enterprises move beyond experimentation and begin operationalizing AI, security leaders face a new reality: AI risk does not stop at deployment.

Modern AI systems are dynamic. They interact with APIs, retrieve external data, access enterprise systems, invoke tools, and increasingly make autonomous decisions at runtime. This introduces an entirely different category of risk than traditional application security models were designed to address.

Palo Alto Networks Prisma AIRS was developed to govern this emerging runtime layer of enterprise AI.

Rather than focusing solely on static controls, Prisma AIRS is designed to continuously monitor and secure AI systems throughout their operational lifecycle. This includes visibility into AI agents, runtime interactions, prompt activity, model behavior, and AI-generated outputs.

The platform incorporates capabilities such as AI agent discovery, prompt injection prevention, AI model scanning, AI posture management, and autonomous AI red teaming.

The Emerging Importance of AI Gateway Security

Palo Alto Networks has announced its intent to acquire Portkey, an AI gateway platform focused on centralized governance and observability for large language model (LLM) interactions. While the acquisition remains pending, the announcement reflects a broader market realization: *Enterprises need governance at the AI interaction layer itself.*

As organizations deploy AI assistants, copilots, and autonomous agents, the volume of interactions between applications and language models is rapidly increasing. Each of these interactions may involve sensitive prompts, retrieved enterprise data, API requests, tool usage, or model-generated outputs.

Without centralized oversight, organizations risk creating an AI environment where model access is inconsistent, API keys proliferate without governance, and AI interactions become difficult to audit or secure.

A New Endpoint Security Challenge

Traditional endpoint security solutions were built to detect malicious files, suspicious processes, and known indicators of compromise. But AI is changing what actually operates on enterprise endpoints. Today's endpoint environment includes AI coding assistants, browser extensions, an AI-generated scripts operating with autonomy.

Many of these tools operate outside traditional security visibility models while maintaining access to sensitive systems, repositories, APIs, and enterprise data. This creates a new category of governance challenge often referred to as agentic endpoint security.

To address this emerging risk, Palo Alto Networks closed its acquisition of Koi in April 2026, a company focused on securing AI-driven endpoint activity. This move reflects a growing recognition that AI-enabled workflows are reshaping the software supply chain and expanding enterprise risk far beyond traditional malware prevention.

Why Does This Matter to the Modern SOC?

Security Operations Centers are already under pressure from alert fatigue, fragmented tooling, and expanding attack surfaces. AI adoption adds another layer of operational complexity as defenders must now monitor AI applications, machine identities, autonomous workflows, and AI-generated threats across increasingly distributed environments.

Traditional SIEM-centric operating models were not designed for this level of scale or speed. As a result, modern SOC strategies are shifting toward automation, AI-driven analytics, and platform consolidation to improve visibility and reduce operational noise.

Palo Alto Networks Cortex XSIAM reflects this evolution by combining AI-driven analytics, threat correlation, automation, and integrated response capabilities into a unified SOC platform to accelerate investigations.

Conclusion

As a Palo Alto Networks NextWave Diamond Innovator Partner, WEI helps organizations modernize SOC operations and operationalize cybersecurity strategies spanning AI security, Zero Trust, SASE, identity governance, cloud security, and platform consolidation initiatives.

Combined with WEI's proactive, engineering-led cybersecurity approach, organizations can strengthen operational resilience across the modern security stack.

Sources:

1. Palo Alto Networks. (2026, February 4). Palo Alto Networks completes acquisition of CyberArk to secure the AI era. https://www.paloaltonetworks.com/company/press/2026/palo-alto-networks-completes-acquisition-of-cyberark-to-secure-the-ai-era

2. Palo Alto Networks. (2025, June 17). Evolving AI traffic control: Why change requests are now limited. Palo Alto Networks LIVEcommunity. https://live.paloaltonetworks.com/t5/community-blogs/evolving-ai-traffic-control-why-change-requests-are-now-limited/ba-p/1242293

About WEI

WEI is an innovative, full service, customer centric IT solutions provider.

Why WEI? Because we care. We go further.

WEI is an expert in business technology improvement, helping clients optimize their technology environments and work efficiently. WEI works with clients to understand goals, integrate strategy with technology solutions, and leverage their current IT environment into one company-wide model to increase utilization and efficiencies around their unique business processes.