CYBER DEFENSE MAGAZINE

eMAGAZINE

JUNE 2026

CYBERSECURITY SOFTWARE IS ONLY AS GOOD AS THE PEOPLE CONFIGURING IT

Cyber Defense Magazine Editor's Choice Book Corner

The HIPAA Security Rule 2026 Problem Every Small Healthcare Practice Is Gambling on Right Now

How Not to Handle a Cyber Incident

From Compliance To Resilience: Why Modern CISOs Must Rethink Cybersecurity Strategy In 2026

...and much more...

MORE INSIDE

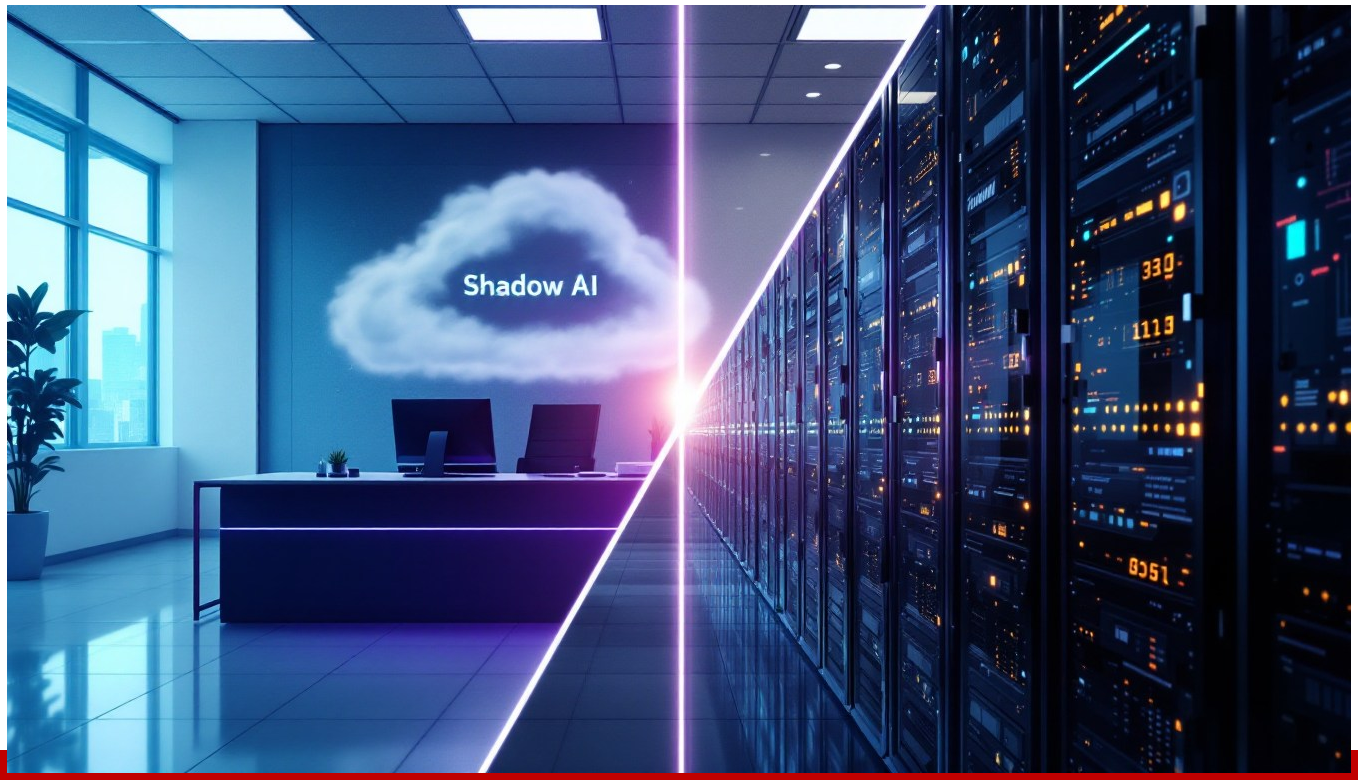

Shadow AI Is the New Shadow IT

By Almog Apirion, CEO and Co-Founder of Cyolo

For years, security leaders have battled shadow IT, including unsanctioned SaaS applications, unmanaged devices, and users bypassing policy in the name of productivity. The risks were significant, but they were also familiar: data leakage, expanded attack surfaces, and compliance gaps.

Now, a new challenge is emerging: shadow AI. And while the risks may look similar on the surface, they are fundamentally more complex and potentially more dangerous.

Unlike shadow IT, which typically operates outside established controls, shadow AI operates *within* them. It functions inside authenticated sessions, uses legitimate credentials, and increasingly interacts directly with operational systems.

Shadow AI as an Identity and Governance Challenge

Shadow AI occurs when AI tools, including copilots, scripting assistants, or automation platforms, interface with enterprise systems through legitimate access pathways but without formal oversight or governance.

This is not a theoretical concern. AI capabilities are rapidly being embedded into everyday workflows across IT and engineering teams. From generating code to optimizing configurations, these tools are becoming trusted assistants. And in many cases, they inherit the full permissions of the user operating them, without additional controls or constraints.

At the same time, organizations continue to struggle to gain visibility into access pathways. A 2026 Gartner report notes that cybersecurity leaders increasingly recognize “shadow access” as a critical gap, where undocumented remote connections bypass traditional controls and permeate cyber-physical environments.

Shadow AI builds on this existing challenge. It does not introduce new access paths but instead leverages existing ones to increase speed and scale and reduce accountability.

How Shadow AI Reshapes the Risk Model

The risk of shadow AI lies in how it amplifies actions within trusted sessions. Three characteristics are particularly important.

First, AI inherits user privileges and executes actions at machine speed. If a user has the ability to modify configurations or systems, AI can perform those same actions almost instantly. In operational environments, this can extend beyond IT systems to industrial controllers and safety mechanisms.

Second, AI lacks operational context. It is designed to optimize efficiency and outcomes. A recommendation that appears technically valid may conflict with physical safety constraints or process dependencies.

Third, AI enables rapid scaling of actions. A single command or recommendation can be replicated across systems, environments, and sites, turning isolated errors into systemic issues.

Why Traditional Security Controls Are Insufficient

Shadow AI changes where risk resides. AI tools gain access through existing enterprise mechanisms, including remote sessions, VPN connections, and third-party vendor interactions. As a result, the primary control point shifts to the authenticated user session, where all actions are executed under a verified identity and its associated permissions. In practice, this means AI systems operate through approved interfaces such as remote desktops or management consoles.

Most traditional security controls are designed to detect external threats or prevent unauthorized access at the point of entry. These controls can be effective when risk originates outside the environment. However, because shadow AI operates within authenticated sessions using valid credentials, its actions appear legitimate and are therefore unlikely to be flagged or blocked by traditional defenses.

This gap highlights the shortcomings of common security approaches.

- Endpoint detection tools are designed to identify malicious code, but not legitimate actions performed through valid credentials.
- Network segmentation restricts access between systems but does not govern behavior within authorized zones.
- Traditional privileged access management solutions focus on granting access, rather than controlling how that access is used in real time.

Even organizations with strong perimeter defenses may have limited visibility and control over in-session activity. Shadow AI operates directly within this blind spot.

Establishing Identity as the Primary Control Plane

Addressing shadow AI requires a shift in how access is governed. AI tools operate using the same identities as human users, and the scope of risk is therefore defined by the level of access granted to each identity.

In cyber-physical environments, where remote access is essential for maintenance and continuous operations, this becomes particularly important. Industrial organizations must move toward identity-centric models that enforce access at a more granular level. This includes limiting access to specific applications rather than entire networks, and separating control mechanisms from the systems they manage to reduce unnecessary exposure.

This approach aligns with zero trust principles by minimizing implicit trust and ensuring that access is continuously verified, tightly scoped, and context-aware.

Key Priorities for Security Leaders

Managing shadow AI risk does not require eliminating AI. It requires establishing clear control over how AI tools operate within authenticated environments.

Three priorities stand out:

- Reduce standing privileges so that users, and by extension AI tools, do not retain unnecessary access.
- Implement just-in-time access to ensure permissions are granted only when needed and then automatically revoked.
- Introduce stronger in-session visibility and control to govern actions after authentication.

These measures help limit the potential impact of both human and AI-driven activity while preserving operational efficiency.

Controlling Activity, Not Just Access

AI usage will continue to expand across both IT and operational environments. What began as a productivity tool is evolving into an execution layer, capable of acting with speed and scale that were not previously possible.

The challenge today is not controlling who can access a system but rather what can be done once access is granted. Actions that were once deliberate and contained can now be generated, executed, and replicated in seconds, often within fully legitimate sessions.

This is why traditional approaches to access control are no longer sufficient. Organizations must move beyond managing entry points and begin governing the activity within authorized sessions.

For environments where digital actions can influence physical outcomes, this distinction is critical. A single command, executed at scale, can have consequences that extend well beyond the system in which it originated.

In this new reality, security is not defined solely by who is allowed in. It is defined by how tightly actions are controlled once they are inside.

About the Author

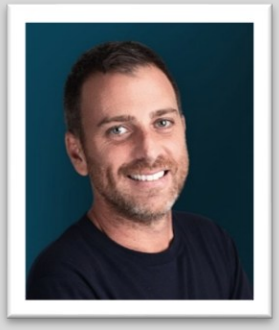

Almog Apirion is the CEO and Co-Founder of Cyolo. He is an experienced technology executive, a “recovering CISO,” and the founder of the Israeli Navy Cyber Unit. Almog has a long history of leading the cybersecurity and IT technologies domain, with a background that includes building and securing critical infrastructures at large organizations, and leading teams to success.