Enterprise Strategy Group™

BRIEF

SEPTEMBER 2025

SentinelOne To Acquire Prompt Security To Take On Shadow AI

Gabe Knuth, Principal Analyst

Abstract: SentinelOne's acquisition of Prompt Security addresses the critical enterprise need to manage the security risks of widespread generative AI adoption. This move provides customers with granular visibility and real-time control over employee AI usage, tackling the pervasive problem of shadow AI and associated data leakage.

Key Highlights

- The disconnect between IT and employees is a major security risk; more than half (53%) of end users admitted to using "Shadow AI" tools they know are disallowed by their organization.
- SentinelOne's acquisition of Prompt Security adds capabilities to discover, monitor, and apply policy to employee interactions with generative AI tools like ChatGPT across browsers and desktop apps.
- The acquisition means SentinelOne isn't just using AI for security, it's leading the way in securing customer AI usage.

The Dual Challenge of Embracing and Securing AI

Generative AI tools within the enterprise represent one of the most significant shifts in workforce productivity in a generation. However, this "gold rush" for efficiency has created a parallel and urgent challenge for IT and security leaders. The core tension is no longer if an organization will adopt AI, but how it can be done without exposing the business to unacceptable levels of risk, from sensitive data leakage to novel attack vectors. This dilemma is a top concern for leadership, which is why "AI model security and integrity" and "data privacy and compliance issues" are among the top concerns around end-user AI usage, along with "increased complexity in IT management and support" and "IT lacking knowledge in AI Ops features."[1]

This challenge is compounded by a significant disconnect between IT policies and employee behavior. While IT teams work to establish guardrails, a gap in awareness and perception persists. For instance, while nearly three-quarters of IT decision-makers reported having an AI policy, fewer than half of knowledge workers were aware of its existence. This leads to a pervasive culture of "shadow AI," where employees use unsanctioned and unmonitored tools to get their jobs done. The scale of this problem is quite large (and likely under-represented): more than half (53%) of end-users admitted they use AI tools they know are disallowed by their organization, and 45% of knowledge workers said they believe their coworkers are sharing private, privileged, or confidential data with AI tools at least occasionally.

[1] Source: Enterprise Strategy Group Research Report, *AI at the Endpoint: The Impact of AI on End Users and Endpoint Devices*, April 2025. All research references in this Brief have been taken from this report.

Market Insight

45% of corporate knowledge workers believed their coworkers are sharing privileged, private, or confidential data with AI tools.

To deal with this, organizations are realizing that legacy security approaches are insufficient. Simply blocking access to known AI tools is a blunt instrument that stifles the very innovation the business seeks to foster and creates more opportunity for shadow AI to develop. To address this, solutions are evolving that provide granular visibility and control directly at the point of interaction—the prompt itself. The goal is to move beyond simple block/allow lists and embrace a strategy that can understand context, redact sensitive information in real-time, guide users toward safer practices, and protect against AI-native threats, all without impeding the flow of work.

SentinelOne's Acquisition of Prompt Security Adds Modern End-user AI Security to XDR

SentinelOne's intended acquisition of Prompt Security is a direct response to this need for visibility and control over generative AI. The move signals a strategic expansion for SentinelOne, extending its security umbrella from its traditional strongholds around XDR to the new, fluid perimeter of AI application usage. By acquiring Prompt Security, SentinelOne aims to give its customers the ability to embrace AI-driven productivity without sacrificing security.

The Prompt Security technology is designed to provide deep visibility into how employees interact with both sanctioned and unsanctioned AI tools. Through a combination of lightweight agents and browser extensions, the solution discovers AI usage across browsers, desktop applications, and APIs. It captures prompts and responses to create an auditable log, while enforcing policies to redact or tokenize sensitive data on the fly. The platform also focuses on preventing AI-specific attacks such as prompt injection and malicious output manipulation and provides inline coaching to help educate users on safe AI practices in real time. This model-agnostic approach is intended to provide a unified security layer across all major large language models, whether they are third-party SaaS or self-hosted.

Analyst Insight

This acquisition is a timely and strategically sound move for SentinelOne. The company is wisely capitalizing on a top-of-mind concern for virtually every CISO and CIO today: how to manage the explosion of generative AI without stifling productivity or driving end users to more and more insecure shadow AI practices. The move adds to SentinelOne's value proposition and positions them as both a company that uses AI to deliver security and one that also provides security for its customers' own AI initiatives. This addresses an emerging concern and a tangible customer pain point, and positions SentinelOne to have a more strategic conversation with enterprise leadership. The technical approach—integrating at the browser and application level—is a key advantage, offering deeper context than network-level or CASB solutions that might lack visibility into the specific content of prompts and responses.

> "[SentinelOne] is wisely capitalizing on a top-of-mind concern for virtually every CISO and CIO today: how to manage the explosion of generative AI without stifling productivity or driving end users to more and more insecure shadow AI practices."
>
> 
>
> **- Gabe Knuth,** *Principal Analyst, Enterprise Strategy Group*

Other XDR vendors will no doubt look to add features like this, so it's important for SentinelOne to hit the ground running when the deal closes. Success requires tight integration into the existing platform and a clear value story for customers.

Conclusion

IT needs to carefully consider how it goes about AI security for end users. While existing technologies that are already deployed in-house are easy to use and might check the box for "AI security", they can stifle growth and adoption of AI in important ways, like driving users to use unauthorized tools, which increases the risk of sharing sensitive data. It's encouraging to see modern approaches emerging around end-user AI security that will help address today's problems with new, bespoke methods that can both improve productivity and security at the same time. Organizations concerned about shadow AI should look to these modern approaches as they develop their AI usage and security strategies.

About Enterprise Strategy Group
Enterprise Strategy Group, now part of Omdia, provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com