

The AI SOC Boom Is Real, But The Work Started Long Before The Buzz

By [Tony Bradley](#), Senior Contributor. © Tony Bradley covers the intersection o... ▼

Published Jan 12, 2026, 08:18am EST

[←](#) [🔖](#) [💬 0](#)



What's new about the AI SOC isn't the idea—it's how many companies are suddenly paying attention.
GETTY

Security operations is suddenly fashionable.

Over the past year, the “AI SOC” has been framed as an emerging category, complete with a wave of new startups, fresh terminology and plenty of claims about autonomy and agents. That surge of interest is real—and it isn't misplaced. But it also creates a misleading impression: that AI-driven security operations are a recent invention, still largely theoretical.

They aren't.

The SOC is one of the few enterprise environments where AI has moved beyond promise and into production, not because it was exciting, but because it was unavoidable. Alert volumes have been climbing for years. Attackers have automated faster than defenders can hire. And the core mechanics of SOC work—triage, investigation, correlation—haven't meaningfully changed in more than a decade.

That reality is what's driving the current influx of vendors. It's also why some platforms are entering a space others have already been operating in for years.

Why the SOC Looks Different From Other AI Use Cases

Most enterprise AI initiatives struggle because they attempt to optimize work that already functions, even if inefficiently. Security operations never had that cushion. It reached the limits of human scale long ago.

That operating model persists even as attackers automate reconnaissance, phishing and lateral movement at machine speed. [Ofer Smadari](#), co-founder and CEO of [Torq](#), told me, “AI is being used by the bad guys – they’re building new attack vectors every day – and the SOC is the only one that gets all those alerts from every detection.”

SOC teams aren’t overwhelmed because analysts lack skill. They’re overwhelmed because the number of decisions required every day exceeds what people can reasonably handle. In large environments, that can mean hundreds of millions—or even billions—of security events flowing through detection systems daily.

That dynamic isn’t lost on experts who work directly with CISOs and security professionals. “Security analysts aren’t struggling because they lack skill or training—they’re struggling because the volume of work is relentless and the margin for error is zero,” said [Den Jones](#), founder and CEO of [909Cyber](#). “That’s why the SOC is one of the few places where AI can deliver real value today. It’s not trying to be creative or insightful. It’s taking on the repetitive, high-volume tasks that exhaust people and obscure real risk, so humans can focus on decisions that actually require judgment.”

From Helping Analysts to Handling the Volume

Early security automation focused on assistance. Tools enriched alerts. Playbooks reduced some manual effort. Analysts still carried most of the burden.

What’s changed is scope. AI systems are now handling defined portions of SOC workflows on their own, particularly where the work is repetitive and rules-based. Travel anomalies. Identity checks. Phishing triage. Tasks that consumed analyst time without improving security posture.

Smadari described one common scenario of an executive or employee logging in from a different region or country than they normally do, and the steps an analyst has to go through to validate whether or not it's legitimate or a security event. "That's twenty minutes of work that can be fully automated with AI, and no one needs to do it anymore." The value isn't novelty. It's scale. Multiply those twenty minutes across thousands of alerts each day, and the operational impact becomes clear.

This is where the SOC diverges from many other AI efforts. The outcomes are visible. The math works. Response times improve, and analyst workload changes.

A Category That Feels New—But Isn't

As these capabilities become more visible, the market has responded by treating the AI SOC as a newly forming category. New vendors like Prophet Security and 7AI are launching with agentic messaging.

From a market perspective, that framing makes sense. From an operational one, it misses important context.

Some platforms have been applying AI-driven automation inside live SOC environments for years, well before the current surge of interest. Torq is one such example.

As new entrants crowd into the space with ambitious roadmaps and evolving terminology, Torq increasingly functions as the reference point others are measured against. It has already navigated sustained enterprise deployments, large-scale alert volumes and the operational expectations that come with running inside production SOCs rather than pilot environments.

In that sense, Torq is more or less the de facto leader of the AI SOC space. While the category is now being treated as emerging, Torq's position reflects something closer to incumbency—an established platform in a market that is only just catching up to what it represents.

Why Enterprises Care About Stability

Large enterprises don't buy security tools the same way startups do. They care more about the balance between managing risk and embracing innovation.

Earlier in my career, I saw promising technologies stall not because they didn't work, but because buyers couldn't justify placing critical operations in the hands of vendors without a track record. Features matter. So does the ability to scale, support complex environments and deliver consistently over time.

Smadari put it plainly: "Enterprises want stability. They want to know you can scale from five thousand endpoints to two hundred thousand—and still operate." That expectation shapes buying decisions as the AI SOC conversation accelerates.

It also helps explain why recent funding activity in the space has drawn attention without fundamentally changing the narrative. Torq just [announced a \\$140 million Series D](#) round at a \$1.2 billion valuation, led by Merlin Ventures with participation from existing investors. The capital signals continued confidence in approaches that are already operating at scale, including growing interest from federal and public-sector organizations with their own regulatory and oversight demands.

What Changes for Human Defenders

The most important shift may not be technical at all. It's how work changes for people.

When AI takes over repetitive investigation and first-pass triage, analysts don't disappear. Their roles evolve. Smadari reflected on his own path from analyst to executive and the toll of early SOC work. "If people don't have to do the boring stuff every day, they can go deeper—into places AI won't be able to."

That shift—from alert chaser to investigator, from responder to strategist—is already happening in organizations that deploy AI at scale.

The Real Story Behind the AI SOC

What's being framed as a new category is better understood as delayed recognition. AI didn't arrive in the SOC because it was trendy. It arrived because the old model stopped working.

As more startups enter the space, the market will sort novelty from execution. Security operations doesn't reward ambition alone. It rewards platforms that

operate under pressure, day after day.

That may be why, of all the places enterprise AI has tried to prove itself, the SOC is where it's starting to stick—not as a theoretical vision of the future, but as a response to a problem that already existed.

[Editorial Standards](#)

[Reprints & Permissions](#)



Find Tony Bradley on [LinkedIn](#). Visit Tony's [website](#).

Join The Conversation

Comments 0

One Community. Many Voices. Create a free account to share your thoughts. Read our community guidelines [here](#).

[See All Comments \(0\)](#)

Forbes

© 2026 Forbes Media LLC. All Rights Reserved.

[AdChoices](#) [Privacy Statement](#) [Your Privacy Choices](#) [Cookie Preferences](#) [Digital Terms of Sale](#) [Terms of Service](#) [Contact Us](#) [Send Us Feedback](#)

[Report a Security Issue](#) [Jobs At Forbes](#) [Reprints & Permissions](#) [Forbes Press Room](#) [Advertise](#)

