# AI-Human Collaboration in Modern SOCs

SPONSORED BY TANDEMTRACE

## A SANS First Look

Written by **Mathias Fuchs** | March 2026

### Introduction

Enterprises face upwards of 3,000 security alerts daily, and according to the SANS 2025 SOC Survey, two-thirds of security operations center (SOC) teams cannot keep pace. In many organizations, nearly half of all alerts go completely uninvestigated. This paper examines whether SOCs that refuse to adopt AI-assisted capabilities can survive what's coming.

The math is stark: A mid-sized organization receiving 4,500 alerts daily needs nearly 94 full-time analysts working without breaks—just to give each alert 10 minutes of attention. Meanwhile, the global cybersecurity workforce gap has ballooned to 4.8 million people. Traditional SIEM and SOAR tools have hit a ceiling, with analysts spending 80 to 90% of their time chasing false positives.

Attackers aren't waiting for our consensus on AI adoption. Around 80% of phishing emails now incorporate AI-generated content. CrowdStrike's threat report puts average breakout time at 48 minutes, while traditional SOC workflows measure response in hours or days. Research on human-AI teaming suggests hybrid teams outperform both pure-human and pure-AI approaches by roughly 25%.

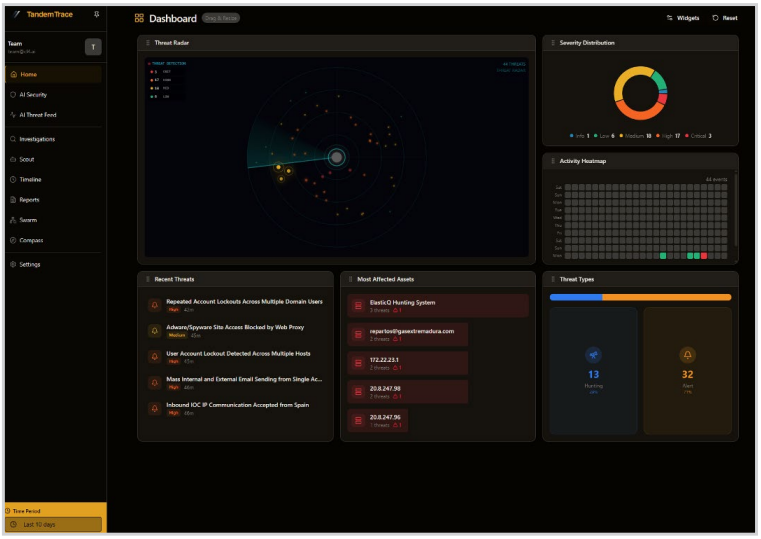


*Figure 1. Tandem Trace Dashboard*

Tandem Trace was built to close this gap by introducing a collaborative AI-human model that pairs experienced analysts with AI reasoning agents to investigate, correlate, and prioritize threats at machine speed, while keeping humans firmly in control of critical decisions. Instead of overwhelming teams with false positives and fragmented data, it transforms alert overload into structured, explainable investigations (see Figure 1).

This marks the next evolution of the SOC: not automation that replaces people, but

intelligent collaboration that amplifies them. By embedding AI across the entire defense chain (telemetry, analysis, decision-making, and response), Tandem Trace enables teams to shift from reactive firefighting to proactive threat prevention, demonstrating that the future of security operations isn't human or machine, but both working in tandem.

## AI in Defense Chains

Security operations function as a chain: Sensors generate telemetry, which feeds into analysis, which informs decisions, which drive response. Each link can benefit from intelligent augmentation. ML-based anomaly detection reduced time-to-discovery for zero-day threats by 43% compared to signature-based tools. Organizations using AI-powered security tools identify and contain breaches roughly 108 days faster, translating to approximately $1.76 million in reduced breach costs. However, data quality remains foundational, model drift requires constant attention, and human-in-the-loop validation is essential for any automated action with business impact.

Most SOCs operate in permanent firefighting mode, with the urgent perpetually crowding out the important. When adversaries can compromise an entire environment in under an hour, response workflows measured in days are functionally useless. The shift from reactive to proactive defense changes how we measure success, from "incidents closed" to "attacks prevented." Organizations making this transition report 30% reductions in mean time to detect and respond.

## Tandem Trace: Where AI and Proactive Defense Converge

Tandem Trace represents one approach to operationalizing the hybrid human-AI model. The core idea—pairing human analysts with AI reasoning agents in a collaborative framework—provides a concrete example of how these concepts translate into practice. What distinguishes frameworks like this from simpler automation is the emphasis on collaborative decision-making. The AI doesn't just process data and spit out verdicts; it reasons through investigations, builds evidence chains, and presents its work for human validation. Think of it less like a replacement analyst and more like a very fast, very thorough research assistant who never sleeps and never forgets to check the obvious things.

The multistage reasoning approach is worth highlighting. Rather than simple pattern matching, modern AI investigation frameworks can pursue investigative threads across multiple data sources, correlating indicators that might appear unrelated to a human analyst working under time pressure. They can cover the entire defense chain from intelligence gathering through detection and response, maintaining context that would be difficult for human analysts to track manually across complex incidents.

Dynamic confidence scoring addresses one of my biggest concerns about AI in security: the black-box problem. When the system tells you not just what it found but how confident it is and why, analysts can calibrate their trust appropriately. High-confidence findings on well-understood attack patterns deserve different treatment than low-confidence anomalies in novel contexts. Making that distinction transparent enables better human judgment, not worse.

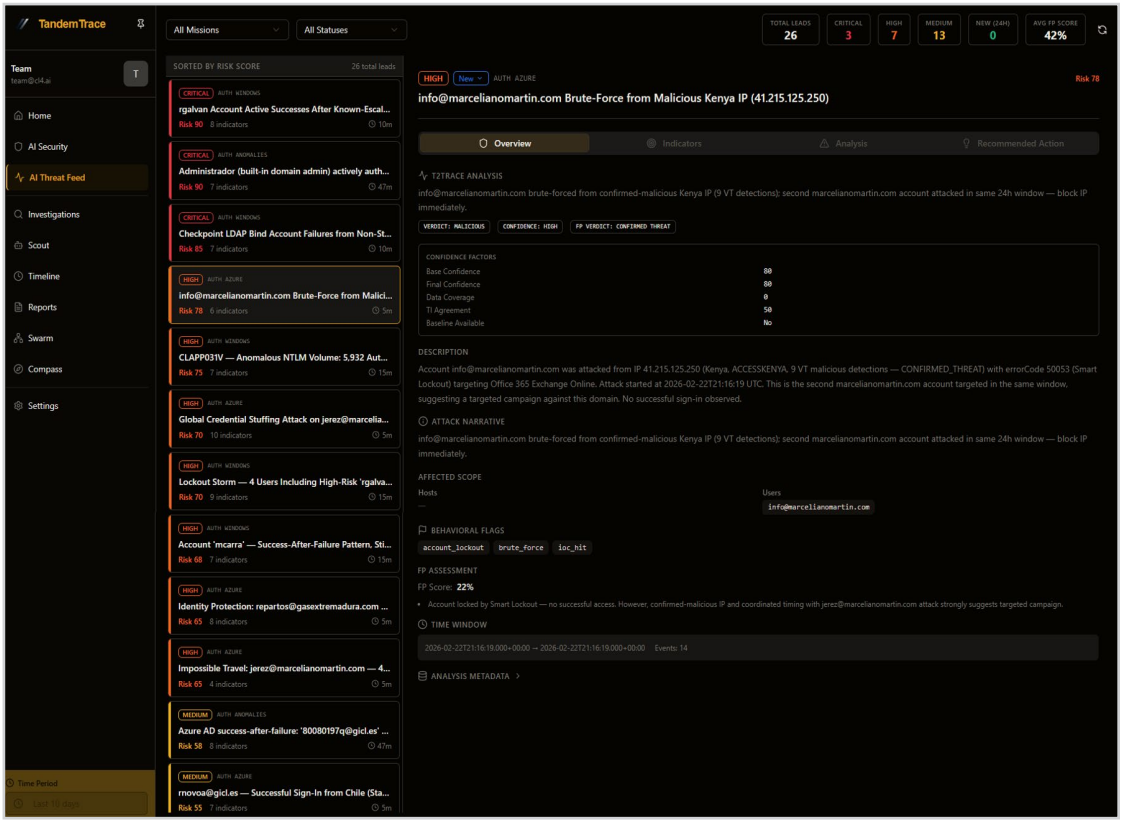


*Figure 2. AI Threat Hunting Feed*

Manual threat hunting requires elite talent that most organizations can't hire and can't afford. Tandem Trace democratizes that capability by pairing AI-powered detection with human expertise. The platform continuously hunts across your environment—correlating behaviors, chasing anomalies, and surfacing threats that signature-based tools miss entirely. When it finds something, it doesn't just raise an alert; it delivers a complete investigation brief: what happened, what's affected, what it means, and what to do next (see Figure 2).

## Humans Remain Essential

Frameworks like Tandem Trace don't make human analysts obsolete. If anything, they make human judgment more important, not less. The analyst's role shifts from mechanical alert processing to something more like conducting an orchestra, setting direction, validating conclusions, and providing the contextual and ethical judgment that AI systems fundamentally lack.

In a tandem cockpit, both pilots have full instrumentation and full controls—but the roles are distinct. The AI flies the data: scanning telemetry, correlating indicators, and flagging threats at machine speed. The analyst makes the decision: interpreting context, weighing consequences, and acting with authority. Neither is a passenger. Neither is redundant. That's tandem defense.

## Where This Is Heading

Industry analysts—Omdia, Gartner, and others—increasingly suggest that hybrid AI-human SOC models will become standard within the next few years. Purely AI-less SOCs will increasingly occupy a niche position. There may still be contexts where they make sense—highly regulated environments with strict constraints on automated decision-making or very small organizations with minimal attack surface. But for most enterprises facing sophisticated adversaries and high alert volumes, the hybrid model isn't just advantageous; it's becoming necessary for basic competitive defense.

## Conclusion

Can AI-less SOCs compete in the future? After examining the mathematics of alert volume, the acceleration of AI-enabled attacks, and the fundamental limits of human-only response times, I believe the answer is no. The path forward requires AI integration across the defense chain, a shift from reactive firefighting to proactive threat anticipation and hybrid collaboration models that preserve human judgment while leveraging machine speed and scale.

The future SOC isn't about choosing between humans and machines. It's about building something neither could achieve alone. The organizations that figure out how to combine these capabilities, thoughtfully, with appropriate safeguards and human oversight, will be the ones that successfully defend against tomorrow's threats. The time for debate is ending. The time for action is here.