



FIELD SECURITY REPORTS

SentinelOne Launches Agentic AI Investigations

Purple AI's new Agentic Investigation feature aims to close cybersecurity's real bottleneck: not detection, but investigation capacity.

FEATURING



SentinelOne

AUTHOR Vatsal Sharma

PEER REVIEWED BY Pinki Sharma, Devesh Taneja

PUBLISHED July 20, 2026

CONTACT vatsal@fieldciso.com · fieldciso.com

SentinelOne Launches Agentic AI Investigations

Purple AI's new Agentic Investigation feature aims to close cybersecurity's real bottleneck: not detection, but investigation capacity.

The Job Nobody Sees

It's 2 a.m. and an alert just fired on a domain controller. This is what happens next, and why it doesn't scale:

- **It's manual.** A SOC analyst pulls process trees by hand, checks identity logs in one tab and cloud activity in another, piecing together, alert by alert, whether this is real or just noise.
- **It's slow.** Industry research puts the average investigation at around 30 minutes. At that pace, one analyst can meaningfully work through roughly 15 alerts a shift, in a world generating thousands.⁶
- **It's inconsistent.** Two analysts can look at the same alert and land at completely different depths of investigation, shaped by how much they've seen before, how tired they are, and what they happened to check first. There's no standard for "how deep is deep enough," just tribal knowledge passed analyst to analyst.

Getting good at this work is also its own barrier:

- **The talent pipeline can't keep up.** The global cybersecurity workforce is short by an estimated 4.8 million people, and demand for information security analysts in the US is projected to grow 29 percent through 2034, far outpacing the pipeline of people who can actually do the job.^{8,9}
- **Even entry-level hiring is slow.** More than a third of organizations say hiring takes three to six months, and for the first time, the industry now says the skills gap, not headcount, is the bigger problem: 60 percent of security leaders name it the greater challenge, against 40 percent citing pure staffing shortages.^{7,10}
- **The scarcity shows up in pay.** The US median wage for information security analysts sits at \$124,910, climbing past \$186,420 at the 90th percentile.⁹

None of this is a secret to the people living it. ISACA found that 66 percent of cybersecurity professionals say their role has grown more stressful over the past five years. Separately, 9 in 10 security professionals say false positives are actively hurting their teams.^{7,6} Manual. Inconsistent. Hard to break into. Burning people out on the way in. That's the job, today, before any AI touches it.

Now, SentinelOne Says AI Can Fix This

SentinelOne's answer is Purple AI Agentic Investigation, opened to all customers on June 17, 2026 as a complimentary trial, the next version of the company's AI-based SOC capability.¹ Here's what it actually does:

- **It's part of a bigger shift.** Security vendors are moving from AI assistants that answer questions to AI agents that independently investigate, reason, and act. CrowdStrike, Cisco, and Palo Alto Networks all shipped agentic SOC tools at RSAC 2026; Microsoft widened autonomous triage into identity and cloud alerts; Google added threat hunting agents at Cloud Next.^{2,3}
- **It investigates on its own.** When an alert crosses a set threshold, Purple AI opens an investigation without an analyst triggering it, correlating telemetry across endpoint, identity, cloud, and third-party data already inside the Singularity Platform, and builds an attack timeline.¹

- **It renders a verdict, with evidence.** Real threat or false positive, complete with an evidence chain an analyst can actually review. Depending on how much autonomy an organization allows, a confirmed threat can trigger an automated response or get surfaced as a recommendation for a human to approve.¹
- **It runs on what SentinelOne already has.** Nothing new to deploy. It reasons using a mix of Anthropic's Claude, OpenAI's GPT, and SentinelOne's own "Ultraviolet" models, and every action is admin-controlled, role-based, and reversible.¹
- **It's free to try.** New and existing customers get a free trial allotment of the company's new Singularity Credits, no payment method required, before paying for further use through partners, direct billing, or SentinelOne's store.¹

"Investigation capacity has become the binding constraint of the modern SOC: detections climb, alerts queue, and verdicts wait on analyst availability. Purple AI's Agentic Investigation capability is designed to remove that constraint by making investigations automatic, continuous, and immediate."

— Chris Corde, Chief Product Officer, SentinelOne¹

Democratizing the SOC, or Just Making Experts Faster?

This is the real question, and it's genuinely unresolved. The case for democratizing: a junior analyst who can't yet reconstruct an attack timeline from raw logs can now read one the AI already built, watch how it reasoned from evidence to verdict, and learn the craft by example instead of by years of trial and error. Zero-config deployment means a five-person SOC gets the same investigative depth as a five-hundred-person one. That could meaningfully lower the entry barrier this report just spent three paragraphs describing.

The case against: someone still has to know when the AI is wrong. An evidence chain is only useful to a reader who already knows what to look for in it, an inexperienced analyst can rubber-stamp a bad verdict just as easily as a good one. The judgment calls that actually require expertise, business context, escalation thresholds, adversarial thinking about what an attacker does next, aren't automated by this. It's plausible that Agentic Investigation mostly hands time back to the people who were already senior enough to spend it well, and leaves the skills gap exactly where it was for everyone else. Both things are probably true, unevenly, across different SOC's. Our honest read: it will make good analysts significantly faster before it meaningfully democratizes the floor.

What We Liked. What We Didn't.

What we liked: the zero-configuration pitch is credible, since it genuinely rides on telemetry SentinelOne already collects. The evidence chain behind every verdict is the right design choice for a tool making autonomous calls, and the generous, no-payment-method trial lowers the risk of finding out whether it actually works for a given environment.

What we didn't like, or simply don't know yet: SentinelOne hasn't published a false-positive or false-negative rate for Agentic Investigation, so its accuracy at scale isn't yet a matter of public record. The credit-based pricing model also doesn't come with published per-investigation unit costs, so what this looks like on a large customer's bill isn't clear from the announcement alone. In our own use, a handful of investigations came back with an inconclusive or unknown verdict rather than a clear real-threat/false-positive call, though overall we didn't run into any major issues.

Alert fatigue is why this matters at all. One Ponemon Institute study found surveyed organizations generated an average of 9,854 false-positive alerts a week.⁴ Expel says world-class SOC's keep false-positive rates below 10 percent, while many organizations tolerate rates above 90.⁵ If Purple AI closes that gap without a hidden cost in accuracy, it solves a problem that's been quietly burning out this industry for years. If it doesn't, it's a faster way to be confidently wrong. The real measure of success will be its accuracy, explainability, and performance in complex enterprise environments, where false positives, incomplete telemetry, and hybrid infrastructures are the norm.

References

1. **SentinelOne, Inc.** "SentinelOne Opens Purple AI Agentic Investigation to All Customers, Bringing Frontier AI Directly Into the SOC." Press release, June 17, 2026.
<https://www.sentinelone.com/press/sentinelone-opens-purple-ai-agentic-investigation-to-all-customers-bringing-frontier-ai-directly-into-the-soc/>
2. **VentureBeat.** "CrowdStrike, Cisco and Palo Alto Networks all shipped agentic SOC tools at RSAC 2026 — the agent behavioral baseline gap survived all three." April 1, 2026.
<https://venturebeat.com/security/rsac-2026-agentic-soc-agent-telemetry-security-gap>
3. **Prophet Security.** "Top 5 AI SOC Analyst Platforms of 2026." June 11, 2026.
<https://www.prophetsecurity.ai/blog/top-5-ai-soc-analyst-platforms>
4. **N-able.** "8 Steps to Reduce False Positives with AI Threat Detection." March 4, 2026 (citing Ponemon Institute 2024 data).
<https://www.n-able.com/blog/reduce-false-positives-ai-threat-detection>
5. **Expel.** "How do you reduce false positives in SOC operations?" CyberSpeak Glossary. Last updated June 19, 2026.
<https://expel.com/cyberspeak/reduce-false-positives-in-soc-operations/>
6. **Panther.** "AI False Positives in the SOC: Why They Happen and How to Reduce Them." May 11, 2026.
<https://panther.com/blog/ai-false-positives-soc>
7. **ISACA.** "Nearly Two-Thirds of Cybersecurity Pros Say Job Stress Is Growing, According to New ISACA Research." Press release, October 1, 2024.
<https://www.isaca.org/about-us/newsroom/press-releases/2024/nearly-two-thirds-of-cybersecurity-pros-say-job-stress-is-growing-according-to-new-isaca-research>
8. **ISC2.** "2024 Cybersecurity Workforce Study." October 2024 (global workforce gap of 4.8 million).
<https://www.isc2.org/Insights/2024/10/ISC2-2024-Cybersecurity-Workforce-Study>
9. **U.S. Bureau of Labor Statistics.** "Information Security Analysts." Occupational Outlook Handbook, May 2024 wage data.
<https://www.bls.gov/ooh/computer-and-information-technology/information-security-analysts.htm>
10. **Industrial Cyber.** "SANS 2026 report flags cybersecurity skills crisis, putting critical infrastructure and OT sectors at measurable breach risk." April 6, 2026 (covering the SANS/GIAC 2026 Workforce Report).
<https://industrialcyber.co/reports/sans-2026-report-flags-cybersecurity-skills-crisis-putting-critical-infrastructure-and-ot-sectors-at-measurable-breach-risk/>

Author: Vatsal Sharma | Peer Reviewed by: Pinki Sharma, Devesh Taneja | FieldCISO Advisory | vatsal@fieldciso.com | fieldciso.com

This report summarizes and analyzes a third-party vendor announcement for informational purposes, alongside independently sourced workforce and industry data. Claims attributed to SentinelOne reflect the company's own press release. Observations in the "What We Liked, What We Didn't" section reflect FieldCISO Advisory's own limited hands-on use of Agentic Investigation and should not be read as a comprehensive or statistically rigorous evaluation; opinions there and in the democratization section are FieldCISO Advisory's own editorial assessment. Product capabilities, pricing, and availability are subject to change; readers should confirm current details directly with SentinelOne before making procurement decisions. FieldCISO Advisory is not affiliated with SentinelOne, Anthropic, or OpenAI. All product names and trademarks belong to their respective owners.