

Autonomous SOC's Need Humans

What 138,475 Alerts Reveal About Agentic AI in SecOps

Introduction

Security leaders are right to be skeptical of “Autonomous SOC” claims. The question is not whether agentic AI can produce answers — it is whether those answers are trustworthy enough to change how analysts work, and whether the time savings are real enough to show up in SOC capacity. Stellar Cyber took an unusual step: instead of publishing benchmark lab results, the company aggregated data from customer end-of-trial reports across 124 days of production use — 138,475 alerts processed, every verdict reviewed by a human analyst, every outcome documented. This paper examines what that data shows, why “analyst replacement” is the wrong frame, and what it means for SOC teams, MSSPs, and security leaders trying to cut through the agentic AI noise.

Why “Autonomous SOC” Messaging Fails

Every major security vendor now promises AI that investigates, triages, and closes alerts autonomously. That framing triggers the wrong response from every audience that matters: analysts hear a job-replacement story and disengage; CISOs hear unverifiable marketing math and discount it; boards hear hype and ask harder questions.

The more credible frame is not autonomy. It is verified trust. An AI that makes triage decisions is interesting. An AI whose decisions human analysts reviewed at scale and agreed with 99.7% of the time is something categorically more valuable — because it is auditable, not just claimed. Stellar Cyber’s data makes exactly this case: the finding is not that the AI replaced analysts, but that analysts checked the AI’s work and almost never disagreed.

138,475	Total alerts processed across 124 days of customer trial environments
99.7%	Rate of AI-and-human agreement when analysts reviewed AI triage verdicts
19 min	Analyst time returned per hour of active alert work
8,047	False positives confidently closed by AI triage (64% of total alert volume)
1,875	True positives escalated to analysts for investigation and response

What the Data Shows

Three proof points form the spine of Stellar Cyber’s argument, and each does different work.

Trust: 99.7% AI-Human Agreement

The strongest claim is not that the AI made decisions — it is that humans reviewed those decisions and agreed almost every time. That transforms a capability claim into a trust claim. Security leaders evaluating autonomous triage are not primarily asking “can the AI do this?” They are asking “can I trust what the AI did?” A 99.7% agreement rate tied to documented customer trial reports answers the trust question, not just the capability question.

Time: 19 Minutes Returned Per Analyst Hour

Nineteen minutes returned per analyst hour compounds to roughly one full day per week per analyst — and approximately 1.5 full-time analyst equivalents recovered across a modestly staffed SOC without adding headcount. Security teams are not understaffed because hiring is impossible. They are understaffed because qualified analysts spend most of their time on low-value triage that produces no meaningful security outcome. Returning 19 minutes per hour is a structural shift, not a convenience improvement.

Signal: Separating Noise from Risk

Of the 138,475 alerts processed, 64% (8,047) were false positives confidently closed by AI triage, 15% (1,875) were true positives escalated for analyst investigation, and 21% were classified as informational. The AI is not summarizing alerts or suggesting next steps — it is making a determination with sufficient confidence that analysts reviewing those calls agreed 99.7% of the time.

The question is not: Can AI act alone? The question is: Can AI reliably return time to humans while improving their focus on what actually matters?

What This Means by Audience

MSSPs: AI-assisted triage that returns 1.5 analyst equivalents without adding staff is a margin protection story. The 99.7% agreement rate also addresses the quality-control concern that has historically made MSSPs cautious about AI autonomy: if analysts validate verdicts at that rate, the service quality risk is demonstrably low.

Enterprise SOCs: Analyst burnout is driven by the volume of low-value alert work that feels disconnected from actual security outcomes. Removing false-positive triage from the analyst's day is a retention argument as much as a productivity one — analysts redirected toward real escalations have a materially different job experience.

CISOs and Boards: Regulators and auditors are increasingly asking how AI-driven security decisions are being reviewed and documented. Stellar Cyber's model — AI verdict, human review, logged outcome — is inherently auditable. That gives CISOs a defensible board answer: not "our vendor says it's accurate," but "our analysts reviewed 138,475 decisions and disagreed 0.3% of the time."

The most important word in "Autonomous SOC" may not be "autonomous." It may be "auditable."

Conclusion / EMA Perspective

Before accepting any Autonomous SOC claim — including Stellar Cyber's — security leaders should ask the following questions. If a vendor cannot answer all of them, the data should be treated as incomplete.

1. How many alerts were evaluated, and over what time period?

2. Were results drawn from real customer environments or controlled lab conditions?
3. Were human analysts able to review and disagree with AI verdicts?
4. What percentage of alerts were closed, escalated, or marked informational?
5. Can every published number be traced back to a customer-facing report or audit log?
6. What was the false-negative rate — how often did the AI miss something a human would have caught?

Stellar Cyber's data answers most of these questions directly. The sixth — the false-negative rate — is the one that deserves continued scrutiny as the dataset matures. No AI triage system is without miss rate, and the long-term credibility of the Autonomous SOC argument depends on vendors being as transparent about what the AI gets wrong as they are about the 99.7% it gets right.

The early market debate around AI in security operations has been framed as humans versus machines — a contest in which one side wins and the other is made redundant. Stellar Cyber's 124-day dataset suggests a more accurate and more useful model: humans plus machines, with each doing the work the other cannot do well at scale.

AI handles repetitive triage at machine speed and with consistency that no human team can sustain across 138,475 alerts. Analysts apply judgment, context, and accountability where it matters most — on the 1,875 true positives that actually represent risk, on the exceptions the AI flags for human review, and on the governance layer that makes AI-assisted decisions defensible to regulators and boards.

If the numbers are valid — and the methodology behind them is more rigorous than most comparable claims in the market — the Autonomous SOC is not the end of the analyst. It is the beginning of a better analyst experience: less noise, more focus, faster decisions, and more time for the security work that actually requires a human being.

The real question security leaders should be asking is not whether to trust autonomous AI, but whether their current vendor can show them 138,475 reasons why.



About EMA

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Learn more about EMA research, analysis, and consulting services at www.enterprisemanagement.com or follow EMA on [X](#) or [LinkedIn](#).

4765.073026