ACTIVANT RESEARCH

The Long Road to Automating the SOC

Why it's time to shift the focus to acceleration

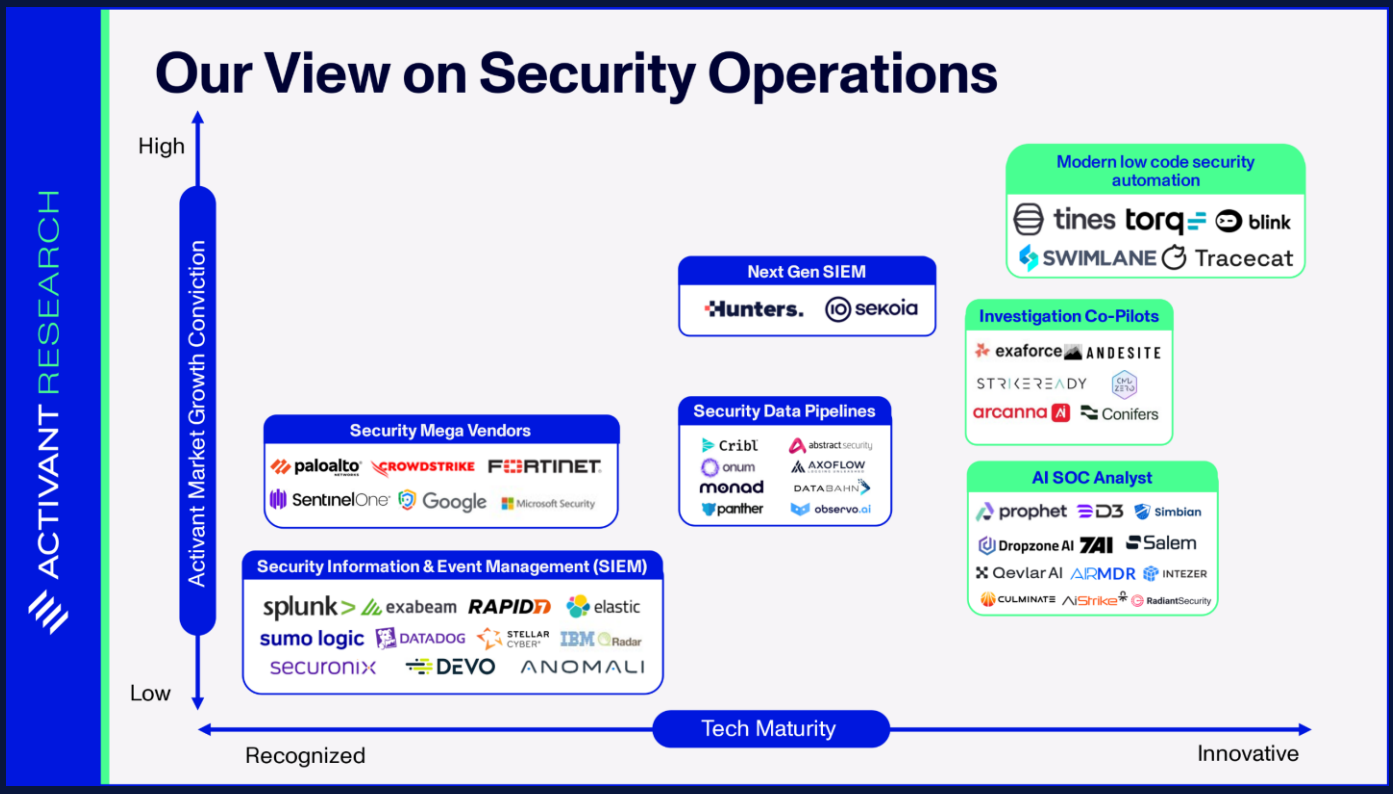

Q2 2025

Jono Vickery

It's past midnight. The Security Operations Center (SOC) analyst on duty checks his screen: "Suspicious login attempt... phishing email flagged... malware detected on an endpoint." His organization has invested in a highly sophisticated security platform built from best-of-breed tools, all flagging issues for the SOC to remediate. Bleary-eyed, scanning a sea of red indicators, the SOC analyst is drowning in alerts. Unfortunately, this is not the start of a work of fiction – **the average SOC team receives four and a half thousand alerts, every day.**[1]

What's needed is a system that can intelligently filter, prioritize, and remediate alerts. However, from SIEMs (Security Information and Event Management) to SOARs (Security Orchestration, Automation, and Response) and beyond, threat volumes have consistently risen faster than these systems can handle. The result is that despite all these tools, SOC analysts are bogged down in repetitive, manual processes. Automation helps with only 17% of alerts, leaving SOC analysts with too many low-quality alerts to possibly sift through.[2] **Consequently, 67% of these alerts are ignored completely**, making it no surprise that we read about new breaches on an almost daily basis.[3]

In this article, we ask why SOC tools often fall short and how pragmatic AI combined with targeted workflow automation can shift us from 'autonomous' hype to action. The goal: an accelerated SOC with less grunt work and fewer breaches. Let's dive in.

What's a SOC?

Modern IT sprawls across many surfaces: cloud and on-prem, SaaS and internally developed applications. Vulnerabilities can arise in networks, endpoint devices, internal code, and even the configurations of cloud resources. To address this, security tools have similarly proliferated to cover these expanding attack surface. As we can see below, the ecosystem of security tools has become highly fragmented, with the average team now using as many as 50 different tools.[4] **The key issue is that each of these systems generates its own security alerts.** Managing this fragmented landscape in a systematic way is what gave rise to the SOC.

[1] Vectra, State of Threat Detection, 2023

[2] Palo Alto Networks, Investor Presentation, 2021

[3] Vectra, State of Threat Detection, 2023

[4] IDC, How Many Security Tools Do Organizations Have, and What Are Their Consolidation Plans?, 2024

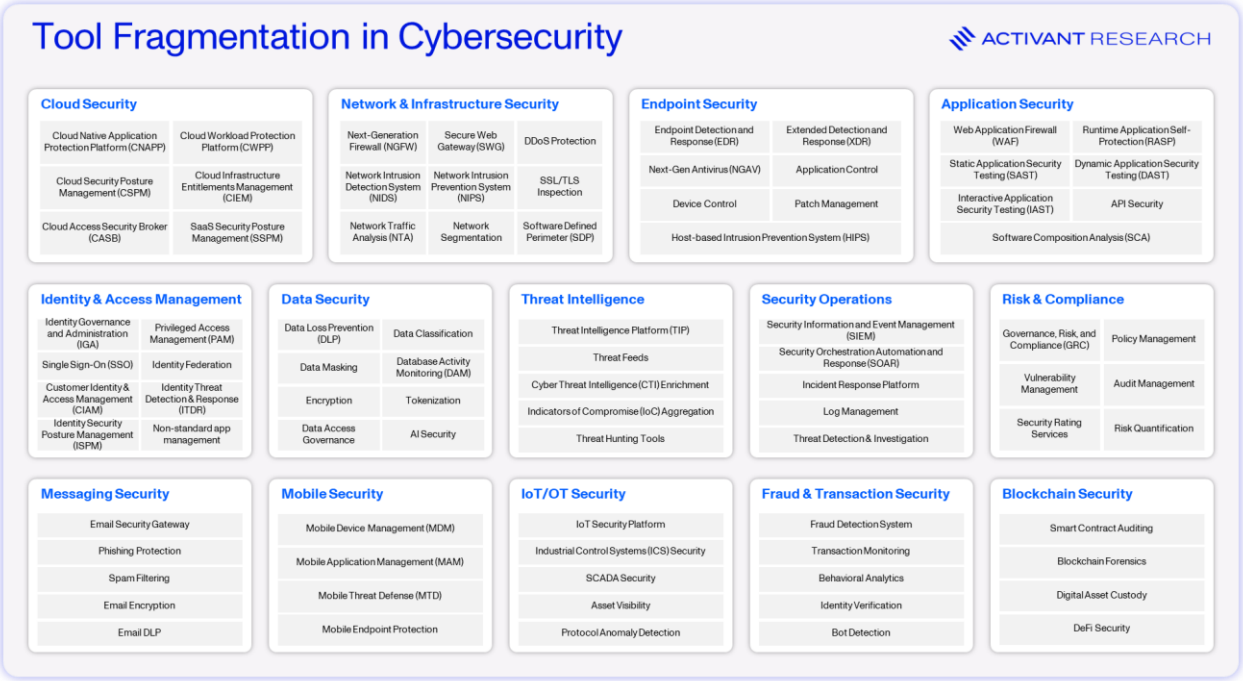

The SOC centralizes the security management function to bring order to this chaos. Beyond real-time monitoring, the SOC is tasked with threat analysis, incident response, prevention, and management of the full security tooling stack. Performance is often measured by how quickly it can detect and remediate threats, as well as how well it minimizes system downtime and business disruption. While large enterprises often run in-house SOCs, many smaller firms outsource the function to Managed Security Service Providers (MSSPs) or run a combined security/IT team.

Today the SOC is all about humans – manpower-heavy and organized in tiers:

- **Tier 1 Analysts**, or *Triage Specialists*, act as the first line of defense – monitoring alerts, performing initial triage, and escalating serious incidents.
- **Tier 2 Analysts**, or *Incident Responders*, take over escalated cases to conduct deeper investigations, correlate data, hunt for threats, and develop new detection use cases.
- **Tier 3 Analysts**, or *Threat Hunters*, handle the most complex threats, lead proactive threat hunting, oversee vulnerability assessments, and fine-tune security tools to enhance the SOC's overall effectiveness.

Below, we visualize the SOC process flow, and the key tools utilized to get the job done.[5]

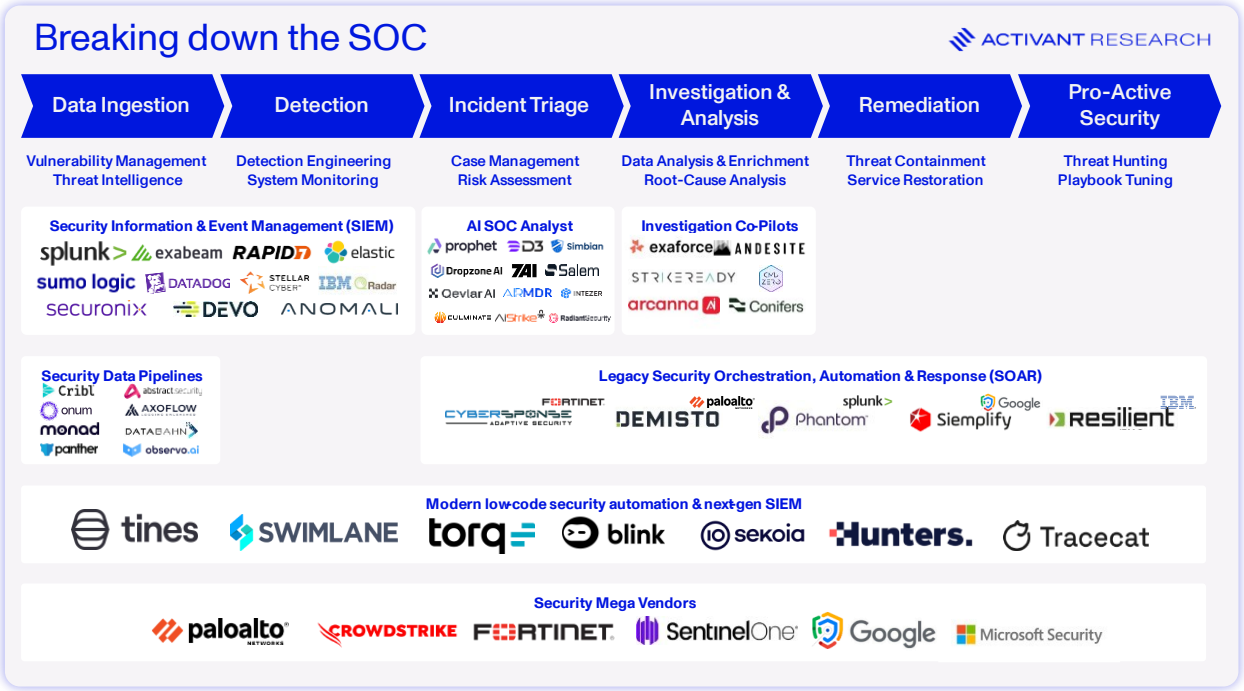

In theory, this seems like a sensible and effective approach to managing security operations. In practice, it's a disaster.

SOCs Suck

Piping every security alert from across an entire organization into a small team (normally 2 - 10 people) produces some unfortunate but predictable problems, namely alert fatigue, tool overload, and staffing issues.[6]

1. **Alert Fatigue:** A typical SOC faces 4,500 alerts daily, and at Fortune 100s, the number runs into the millions.[7,8] Most alerts are repetitive, low fidelity "noise" but must be cleared manually. Analysts believe that at least half of their tasks could be automated.[9] Worse still, somewhere

[5] Note that legacy SOAR is no longer an independent category and is included for illustrative purposes only. All legacy SOAR vendors were acquired except for Swimlane, which has adopted modern automation tooling.

[6] SANS, SOC Survey: Facing Top Challenges in Security Operations, 2024

[7] Vectra, State of Threat Detection, 2023

[8] Palo Alto Networks, Investor Presentation, 2021

[9] Cymulate, 4 Ways to Save Your SOC Analyst from Burn Out, 2025

between 50% and 99% of these alerts are *false positives* – alerts which turn out to be benign when investigated.[10,11] Imagine starting your morning with the impossible task of sifting through thousands of repetitive events, knowing most will be a waste of time. **Analysts get desensitized, demotivated, and miss true threats hiding in the noise – it's no surprise then that up to 67% of alerts are completely ignored.[12]**

2. **Tool Overload:** For the average security team, resolving alerts requires gathering context and taking action across 50 different security tools.[13] Swivel chair syndrome, as it is unaffectionately known in the industry, makes it even harder to make critical decisions and stay productive. Further, it can be extremely costly to build and maintain the connections required to string these tools together, so they often remain siloed. This lack of visibility into enterprise IT ranks among the top 5 challenges SOCs face.[14]
3. **Staffing issues:** The total security industry employs an estimated 5.5 million people yet declares a workforce gap of ~4mn.[15] On average, every security employee is doing the work of ~1.7. Unsurprisingly, many believe that burnout and attrition is highest in security staff.[16]

The result: fatigued, short-staffed SOC teams without the right tools at their disposal. It's no wonder SOCs underperform in preventing, detecting and responding to actual breaches. **There were over 10,000 reported breaches in 2024; and more alarmingly, the average enterprise takes an average of 287 days to detect and respond to one of these.**[17,18] That's right – many enterprises are carrying unresolved breaches for as long as a human pregnancy. The SOC is crying out for automation.

---

[10] Critical Start, The Impact of Security Alert Overload, 2020

[11] Usenix, 99% False Positives: A Qualitative Study of SOC Analysts' Perspectives on Security Alarms, 2022

[12] Vectra, State of Threat Detection, 2023

[13] IDC, How Many Security Tools Do Organizations Have, and What Are Their Consolidation Plans?, 2024

[14] SANS, SOC Survey: Facing Top Challenges in Security Operations, 2024

[15] ISC2, Cybersecurity Workforce Study, 2023; *Workforce gap represents "difference between the number of cybersecurity professionals that organizations require to properly secure themselves and the number of cybersecurity professionals available for hire"*

[16] Enterprise Strategy Group, SOC Market Trends Report, 2023

[17] Verizon, Data Breach Investigation Report, 2024

[18] Palo Alto Networks, Investor Day Transcript, 2021

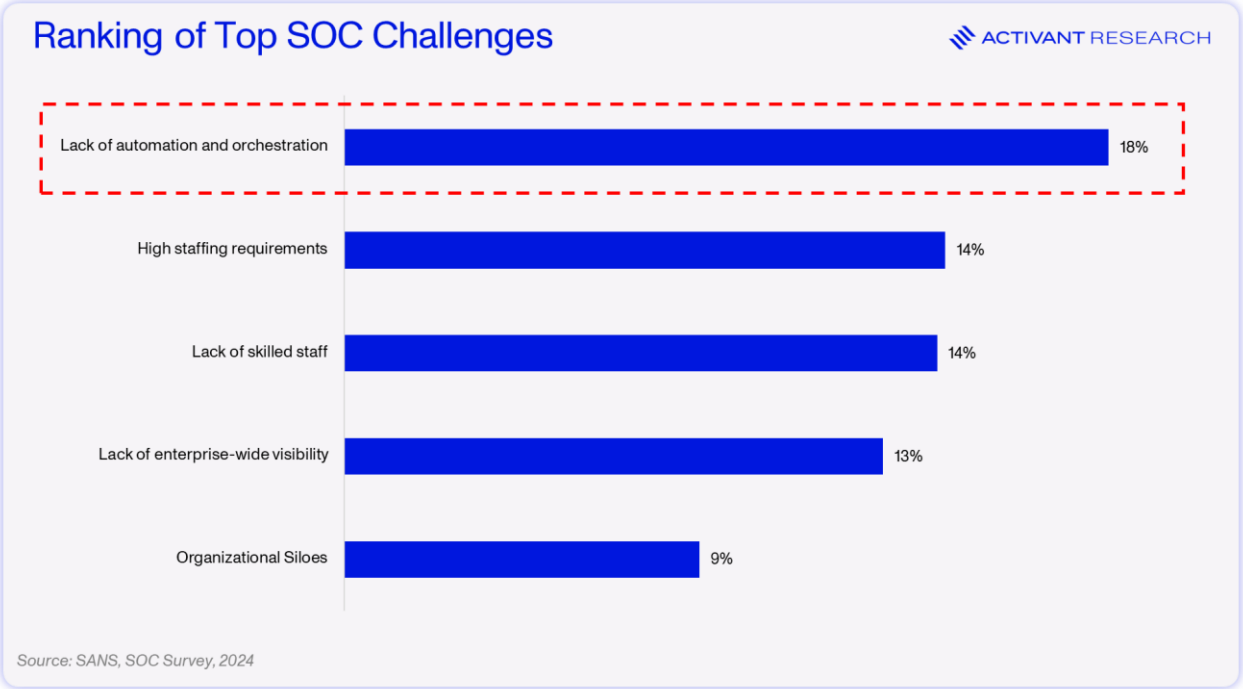

It's not that no one has ever taken on the challenge of automating the SOC, but considering the statistics above, it's evident that the status quo– the SIEM and SOAR combo – has failed.

SIEMs, SOARs, and not a lot of progress

As we entered the internet era in the late 90s and early 2000s, potential security vulnerabilities expanded dramatically. As demonstrated below, that same period saw an explosion in the number of security tools on the market, growing 11.2x in the 5 years between 1995 and 2000. This meant that security alerts and important context behind them were ever more distributed across systems.

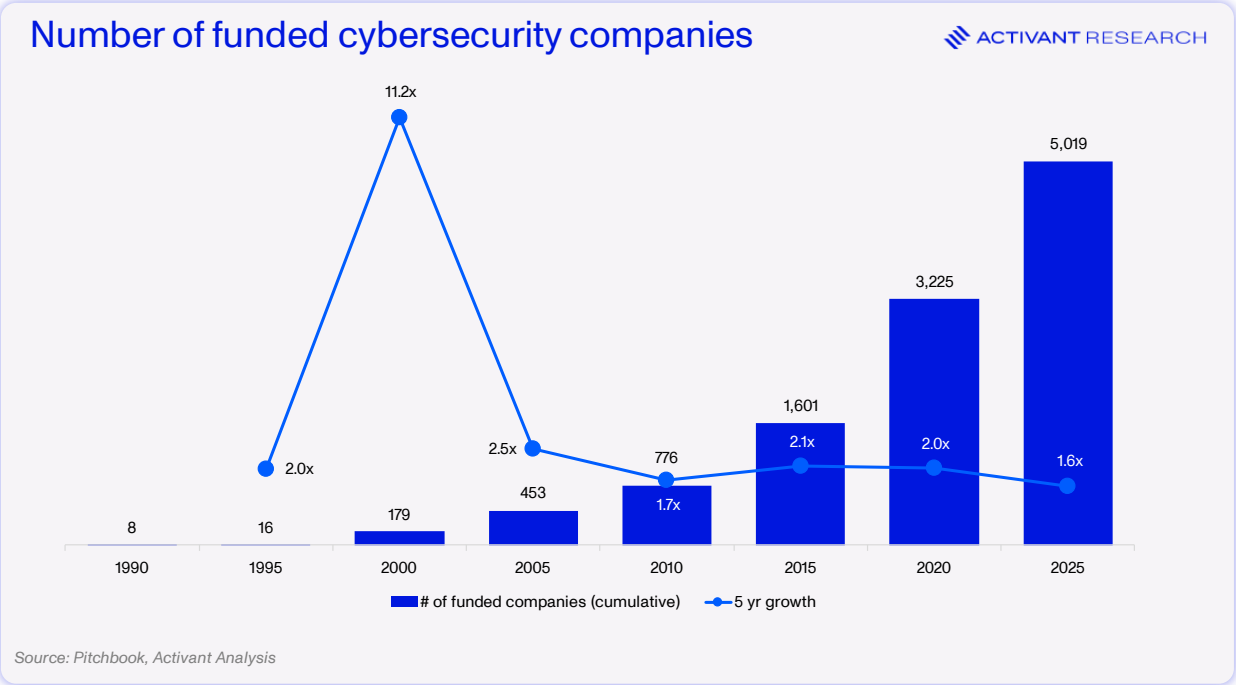

SIEMs like Splunk, ArcSight, and QRadar emerged in the early 2000s, promising to aggregate logs and alerts from across all those disparate tools into one place so that analysts could detect threats through correlation and have a "single pane of glass." In practice, SIEMs did centralize detection – but they also ended up centralizing the alert deluge and it was still up to humans to investigate and respond. The SIEM didn't fix the *response* bottleneck.

By the mid-2010s, the industry's answer to this was SOAR platforms. Startups like Phantom, Demisto, and Swimlane built systems to automate incident response via "playbooks" – essentially, codified response procedures spanning multiple tools. A SOAR playbook could, for example, automatically take an alert from your SIEM, enrich it with data from threat intel feeds, and – if certain conditions were met – execute actions like quarantining a device or creating a ticket. The vision was seductive: if SIEM is the eyes, SOAR would be the hands, automating the routine so analysts could handle the exceptions.

Reality, however, proved more complicated. **Early SOAR implementations were code-heavy and brittle**. Standing up a playbook meant custom scripting (often in Python or JavaScript) and painstaking integration work for each tool and use case. SOC teams, already overloaded with alert volumes, struggled to find the spare time to write complicated automation scripts, and lacked the technical skills required. Many SOCs bought SOAR tools but struggled to fully use them – one might automate a few common tasks, but the long tail of incidents still required human intervention. Over time, some playbooks fell out-of-date as systems changed, leading teams to

"babysit" their SOAR workflows. Compounding the issue, most SOAR vendors were acquired by SIEMs in a bid to strengthen their competitive positioning. The result was a monolithic, legacy product with exorbitant pricing: Splunk costs 3x – 5x more than modern logging and storage systems.[19]

In the end, SOAR ranks in the bottom ten security tools by user satisfaction and Gartner have declared the category obsolete.[20,21] The vision of the SIEM with a SOAR on top automating the SOC failed, not because the strategy was wrong, but because the SOAR was *inaccessible*. That's where workflows come in.

Less Code and More Automation

SOAR became a toxic acronym, and new vendors rebranded. From "next-gen SIEM" to "hyperautomation", the point was the same – automate the SOC. This era of security automation looked to drop the code and make it quick and easy for any level of security professional to start automating their jobs. Even junior team members could create meaningful automations in a few hours.[22]

The low-code era of SOC automation was not about fundamentally rethinking the SOAR, just accelerating it. Instead of writing Python scripts, teams could build drag-and-drop workflows through a visual interface. Tines, an Activant portfolio company, cut all enterprise automation down to just seven actions (and now with AI – eight), enabling companies like Elastic to deploy 49 workflows in 12 months and save 750 days of analyst time.[23] An action– for example, a webhook – could grab alerts like malware reports from the SIEM or other security system, correlate them with other data, apply various rules and, if passed, then trigger an action such as quarantining an infected endpoint automatically.

While Tines, Torq and Mindflow were many years behind earlier versions of workflow automation like Zapier, they were able to capture the security team by focusing on enterprise SLAs that security teams require, security-specific integrations out of the box, and a general seriousness that comes with dealing with security workflows.

---

[19] Hydrolix, Pricing Comparison Page, 2025; *Assumes 1TB/day processed with 12-month retention period*

[20] SANS, SOC Survey: Facing Top Challenges in Security Operations, 2024

[21] Gartner, Hype Cycle for Security Operations, 2024

[22] Tines, Snowflake Case Study, Accessed 2025

[23] Tines, Elastic Case Study, Accessed 2025

Security was how these players found their wedge, but for Tines and Mindflow, they weren't confined to this. Doing the hard thing first gives you an edge with teams as diverse as HR, finance and IT. And, once teams across these disparate divisions start to automate their own tasks with the same generalizable tool, that tool can start to tie all of them together. It's a formidable position to be in – a tool with infinite use cases, that is easy to implement and creates connectivity across the entire org. Players like Tines – rightly – formulated ambitious visions of becoming enterprise-wide automation ***and orchestration*** platforms. So why then is a lack of automation still the number one issue in the SOC today?

To cut low-code some slack, enterprise adoption simply takes time. What we call legacy is often just being rolled out at non-tech Fortune 100s. But in truth, low-code automation isn't a panacea. It removed much of the heavy lifting of coding, but the workflows created were still fundamentally deterministic logic. A low-code playbook could efficiently orchestrate **known** responses (reset a password, isolate a device, open a ticket, etc.), but it couldn't improvise or "think" outside the predefined path.

This also meant that if the number of potential paths was large, i.e. if the workflow was complex, then the chances of the workflow breaking increased quickly. While building workflows was easy, building one that could handle every possible edge case was a huge lift. In effect, low-code platforms solved the *accessibility* problem of SOAR, but they left a **cognitive gap**. Every flow depended on rules that a human had seen, understood, and taken the time to build. If an alert was novel, the automation cracked, and a human took the wheel.

Of course, that's where AI comes in.

Climbing the Ladder to Autonomy

AI Agents close that **cognitive gap**. Instead of executing a fixed script, Agents act dynamically. An AI Agent, with memory and access to key data and tools, can analyze context, correlate information from different sources, solve problems, and make complex decisions. Agents don't need extensive pre-written rules or logic; they just need clear objectives and constraints. They have the promise to be powerful yet require only a few minutes of prompting.

We're already seeing solutions that promise to automatically triage 90% of alerts with AI and improve SOC productivity by 10x. These are not empty promises. In a study by Splunk, organizations with advanced Generative AI programs dropped MTTD disruption-causing incidents from 34 days to 21– a 38% improvement – and MTTR from 5.7 days to 1.8 – a 68% improvement.[24] In Palo Alto Network's own SOC, AI has taken their analysts from doing "mostly" alert triage to spending 70% of their day threat hunting and running attack simulations.[25] It's no surprise that six in ten security leaders see AI as a game changer, with the ability to drive value across a wide variety of security use cases.[26]

These early success signals are exactly the breath of fresh air that the SOC has been crying out for. The catch: **hallucinations**. Security can't be trusted to AI *just yet*. Practitioners are rightly skeptical and, while we complained of how poorly SOAR ranked on user satisfaction, it's worth pointing out that Generative AI ranks dead last in that same survey.[27]

---

[24] Splunk, 2024 State of Security: The Race to Harness AI, 2024

[25] Palo Alto Networks, Security Operations in 2025 and Beyond, 2024

[26] KPMG, Survey: C-Suite Cyber Leaders Optimistic about Defenses, but Large Percentage Suffered Recent Cyber Attack, 2024

[27] SANS, SOC Survey: Facing Top Challenges in Security Operations, 2024

Shifting systems from deterministic to probabilistic is significant, but it's also a major issue for industries like cybersecurity where the cost of an error could be as much as $5 million.[28] Further still, systems that follow defined paths are great for auditing, compliance and adherence to best practices like those from NIST. In these high accuracy, standards-based, regulated industries, handing over the keys to AI might not be an option.

Any critique of these tools must acknowledge that costs are coming down relentlessly and capabilities are improving rapidly. Innovations like fine-tuning, RAG (retrieval-augmented generation), and prompt-engineering have proven that they can and will continue to boost performance. The response to early tools may have been negative, and security industry-specific characteristics are definite barriers, but we think that Generative AI has an important place in security. However, we need a framework for sensible implementation. This is where autonomous vehicles can help.

Levels of SOC Automation

ACTIVANT MARKET MAP

	Autonomous Vehicle Hierarchy	Autonomous SOC Hierarchy	Use Cases	Representative Vendors
Level 0	No Automation	SIEM Only	n/a	n/a
Level 1	Driver Assistance	AI Co-Pilot	Case Management Data Analysis & Enrichment	exaforce, ANDESITE, STRIKEREADY, Conifers, arcanna.ai, CMU ZERO
Level 2	Partial Automation	AI Triage & Investigation	Risk Assessment	7AI, prophet, RadiantSecurity, Dropzone AI, D3, AiStrike
Level 3	Conditional Automation	Low-Code Automation	Threat Containment Service Restoration	tines, torq, SWIMLANE, blink
Level 4	High Automation	Agentic AI w/ guardrails	Threat Hunting Vulnerability Management	n/a

Level 0: No Automation – Humans handle every task including log review, triage, enrichment, response. Playbooks live in wikis, not code.

[28] IBM, IBM X-Force 2025 Threat Intelligence Index, 2025; *Figure refers to the average cost of a data breach*

Level 1: Driver Assistance or AI chatbots and co-pilots – Humans still need to be engaged in the task at hand but AI speeds up their work. Exaforce allows users to query contextualized security data with natural language, and automatically provides predictive insights, charts and graphics. In Tines' Workbench, analysts can take actions from within the chat experience. Level 1 is a long way from a truly autonomous SOC, but that's not a statement of *value*. Accelerating investigations is high value and will provide meaningful leverage on the budget allocated to the SOC.

Level 2: Partial Automation – As with Tesla's Autopilot, the SOC analyst is still required to be in the driver's seat and ready to take control at any time, but for the most part they no longer need to think about mundane tasks like **alert triage**. DropzoneAI automatically turns every SIEM alert into an investigation, performs triage, and reports a summary to the analyst with escalation for critical issues. From a SOC analyst's perspective, they still need to take all required actions for remediation, but D3's Morpheus is starting the push towards action by recommending remediation steps for analysts.

This is where we see most tools branding as "AI SOC," performing the work of a Tier 1 analyst. That's not a knock on these tools – this is exactly the type of work that LLMs are great at: summarizing massive amounts of information from disparate sources and leveraging a well-structured base of documentation and playbooks to make sense of it. And, while some may say that accuracy can't be trusted, we say that 100% coverage at 80 - 90% accuracy is incredible when contextualized with the data that non-automated SOCs ignore up to 67% of alerts.[29]

Level 3: Conditional Automation – The system can handle most tasks when conditions allow. Analysts still need to step in, but in a Level 3 SOC they would mostly focus on threat hunting because alert triage and remediation are largely automated. Combining low-code workflow automation with AI enables this shift. For example, players like Torq have followed what's happening in the "AI SOC" with their own triage agent, Socrates. It similarly triages 90%+ of Tier 1 alerts which can be linked to the existing workflow engine to take action across the entire IT stack. Today, a human in the loop is required to trigger the workflows, so this is not **yet** Level 3 automation, but shows the potential: once remediation can be automatically triggered, the SOC really starts to drive on its own, *under specific conditions.* With these tools, analysts can do phishing response, malware containment, fraud account locks and remediate CSPM findings.

Level 4: High Automation – Waymo can operate without a driver in the seat, but is legally required to retain a steering wheel and can only do so in certain locations. In the SOC, this level of automation is not yet possible. Pure LLM-driven agents have well-known downsides: they hallucinate facts, misjudge novel situations, and lack true accountability. To safely automate high-

[29] Vectra, State of Threat Detection, 2023

stakes security work, AI needs *scaffolding* around it – fallback mechanisms, oversight layers, and human checkpoints. We believe that placing the power of AI Agents *inside* of deterministic workflows could be the innovation that takes us there. Rather than expecting an AI Agent to run automation end to end, it's preferable to tightly control it inside of the existing workflow canvas. This could be through using an AI Agent to run an atomic task inside of a workflow, while also using AI to create workflows on the fly. AI generated workflows, when trained to include proper failback mechanisms, combine the autonomy and generalizability of AI with the reliability of defined workflows. For players like Tines, Torq and BlinkOps, workflows were the glue to tie the enterprise together, and with AI, those connections become richer, easier to build, and more impactful.

Level 5: Full Automation – The truly autonomous SOC, running 24/7 and auto-remediating 100% of alerts, with zero humans. It's a SOC where AI remediates every alert, at 100% accuracy, all the time. Additionally, the AI can autonomously trigger and execute tasks like vulnerability management. To save from digressing into the AGI or human-level intelligence AI debate, we'll just say that this is not the future that we see at Activant.

So, what does this tell us about the future of SOC automation? **Well, it's not about automation, it's about acceleration.** AI triage agents free up analysts from alert fatigue, speed up the upskilling of Tier 1 analysts and focus the whole SOC more on Tier 2 and Tier 3 work. AI co-pilots significantly speed up investigations, complex decision making, threat hunting and more. Finally, AI Agents will make it easier to build automations, handle edge cases and remediate more issues in an end-to-end automation, but in a more limited and human-in-the-loop manner than what could be described as "autonomous." AI means going faster; it's about *accelerating* the SOC.

And, with attackers constantly increasing their own sophistication, AI adoption will become non-negotiable for security teams. The tools are here already, and enterprises should start experimenting with Level 1 at the very least.

The Activant View

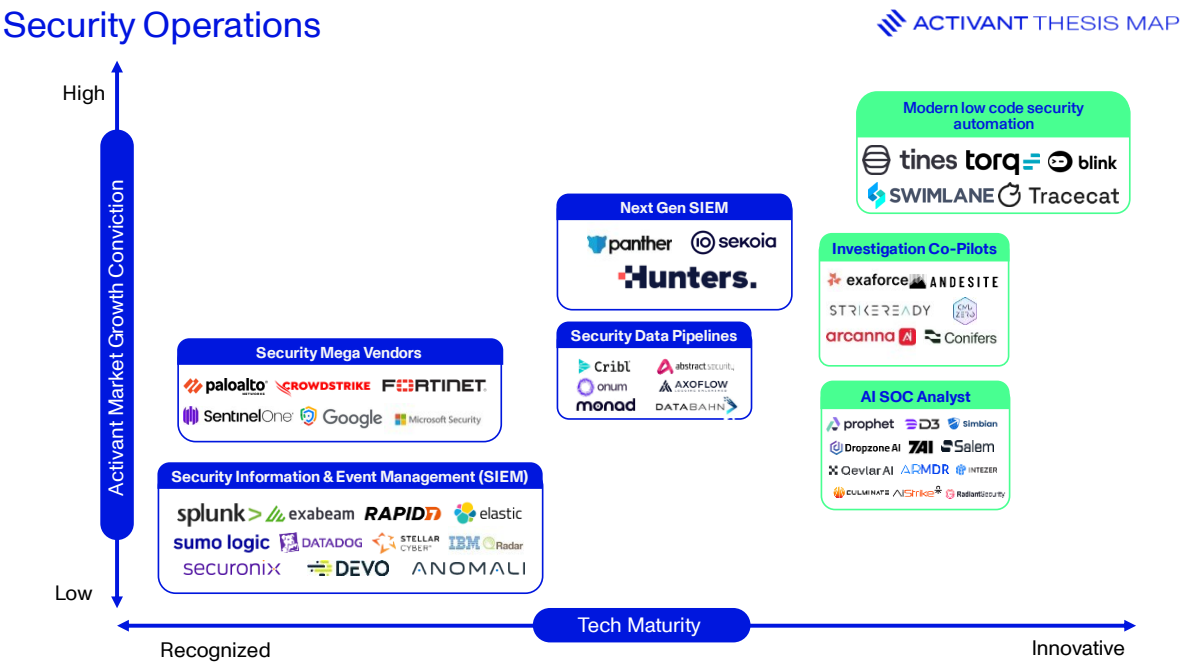

We see three competitive vectors in the bid to take the SOC into its next level of autonomy.

1. **AI Natives:** For **AI SOC Agents** like Dropzone and D3 Security, we expect strong early traction as they solve for the acute pain point of alert fatigue. The same goes for the value of **Investigation Co-Pilots** like Exaforce and Andesite. However, as outlined above, AI needs scaffolding. Companies who progress from this category will do so because they were able to make their AI trustworthy, with strong guardrails and other supports. If they can execute on this roadmap, AI-first automation could disrupt the need to build workflows and prove our thesis on level 4 automation wrong. However, the whole category needs to contend with the risk that customers' expectations progress more quickly than the capabilities of the AI models that underly them.
2. **Large platform vendors:** Players such as CrowdStrike and Palo Alto Networks are executing on a platformization strategy, where it might not matter that specific tools (like their workflow engines) are not as strong as Tines and Torq. Customers get the opportunity to reduce complexity by standardizing on their offerings, potentially reduce total cost of ownership, and benefit from a holistic solution that is more tightly integrated. 67% of CrowdStrike's customers

use over 5 different modules in the platform, a figure that is up 500bps in 2 years.[30] 1,150 of Palo Alto Networks' 80,000 customers have standardized on their offerings.[31] Even in just the SOC, these companies have a huge scale advantage: Palo Alto Network's reports over $1bn in ARR for their SecOps tool, Cortex and CrowdStrike's Next-Gen SIEM has already reached $330mn in ARR, growing more than115%.[32] Concerning AI, CrowdStrike sees "trillions" of security events every week, giving them an advantage in one of the most critical inputs to AI success – data.[33] Their real weakness, though, is one that the industry knows well: what gets advertised as a tightly knit, all-in-one platform often becomes unwieldly, difficult to implement, slow to change, and not best-of-breed in any of the components.

3. **Workflow automation tools implementing AI**: This includes companies like Tines, Torq and BlinkOps. As we've highlighted, we believe that these platforms have a real opportunity to take SOC automation to the next level. Path dependence suggests that their existing platforms might be the perfect place for AI Agents to sit in the SOC. They will need to execute in order to make sure that AI SOC players don't give their early AI efforts the required guardrails first. Critically, as AI improves workflows and makes it easier to string together best-of-breed security tools, the value prop of an integrated platform like those offered by the mega vendors declines. With an extremely intelligent and powerful orchestrator at the center, it makes sense to pick flexibility, ease of use and depth of integration over one monolithic platform. Additionally, Tines has one unique advantage: they are not just for the SOC. Tines wins security first to prove that they can handle important workflows, and then moves into HR, sales, finance, and more. Over time, their point of competition is not just on how tightly they can knit together the SOC, but the entire enterprise; this is where mega security vendors can't compete.

Conclusion

Google indexed the world's data, which made it an order of magnitude easier to perform research and write a book than in the pre-internet era. However, writing a book still required a high degree of skill and dedication. With AI, those limitations are gone – anyone can write a book in a matter of minutes. SIEMs provided a Google-like leap for SOC analysts doing information discovery but left a level of drudgery to resolve alerts. AI holds the promise to automatically resolve security alerts in minutes as easily as it can write a book. However, authors must still ask - is the book that AI wrote in a few minutes any good? The same applies for SOC analysts. The next era of the SOC is

[30] CrowdStrike, Investor Presentations, 2023 - 2025

[31] Palo Alto Networks, Q2 FY25 Investor Presentation, 2025

[32] Ibid

[33] CrowdStrike, Form 10K, 2025

not about automating everything; it's about freeing up human analysts to do the work that matters and making them more efficient in doing so.

If you're building in this space, please get in touch. We'd love to engage with you.