Picture credit: AdobeStock

THE SOC FIGHTBACK

Merlin Gillespie *reveals how AI and machine learning can enhance SecOps*

In the ongoing battle against threat actors, security operations (SecOps) are looking increasingly strained. Drowning in alerts, many are buckling under the pressure. Every day at the coal face could be the day a serious threat sneaks through. And if tooling can't effectively discern the false from true positives, the likelihood of a worst-case scenario continues to grow.

But it doesn't need to be this way. Consider how AI agents might give rise to a new model: the autonomous security operations centre (SOC). According to one analyst house, it could become standard practice for CISOs within one to two years. The key to their success will be how well they integrate these machines with the most important asset in any SOC: its people. Because we forget the 'human in the loop' at our peril.

According to IBM, the mean time it took organisations to identify and contain breaches last year was 241 days. That's over seven months per breach, during which time threat actors were able to escalate privileges, move laterally, map the target network and exfiltrate/encrypt data. If SOCs were doing their job, the dwell time of an attacker would ideally be measured in hours, not days. Yet many security operations (SecOps) teams are under tremendous pressure. Without the tooling and workflows to prioritise alerts accurately, teams can quickly become submerged. It's why many organisations outsource their entire SOC to an expert third party, which has the required resources and expertise.

Nearly half of security and risk professionals report that they or their staff are experiencing burnout

Those that choose not to must contend with a tough jobs market. SecOps skills remain in high demand, and as pressure increases on teams, a growing number of analysts are succumbing to burnout. According to a report from last year, nearly half (47 percent) of security and risk professionals report that they or their staff are experiencing some level of burnout. And over 10 percent describe their condition as acute – "very burned out" or on the verge of leaving the profession altogether. As SOC experts leave without anyone to replace them, the pressure increases on the team members that are left.

In an attempt to mitigate these stresses, and restore a little work-life balance to the SOC, most (78 percent) organisations reduce their analyst numbers by 50 percent or more during holidays and weekends. Nearly one in 10 don't have anyone at all staffing the SOC outside of the regular working week. This matters, because over half of organisations struck by ransomware over the past year were hit during the weekend or a public holiday.

At the same time, adversaries are making the most of AI tools and 'as-a-service' offerings to improve success rates. The National Cyber Security Centre (NCSC) warns that technology advances will "almost certainly" make offensive operations more efficient and effective, "leading to an increase in frequency and intensity of cyber threats". Tactics, techniques and procedures (TTPs) as varied as victim reconnaissance, vulnerability research and exploit development (VRED), social engineering, basic malware generation and exfiltrated data processing could receive a boost.

This is where machine learning (ML) can be a useful ally for time-poor SOCs overwhelmed with alerts they can't contextualise. It supercharges Security Orchestration, Automation and Response (SOAR) to assist with everything from qualifying an alert through to telling the SOC analyst how to deal with it. In this way, it's possible to automate up to 90 percent of investigations.

Start with enrichment. ML-powered SOAR automates the time-consuming process of enriching alerts with contextual threat intelligence. This speeds things up – as each alert has two to three enrichment opportunities, each of which may take an analyst a couple of minutes to work through. But it also helps to improve the quality of the investigation.

At the same time, the SOAR will automatically aggregate alerts relating to the same entity (eg users, devices, IP addresses). This is vital in unearthing and piecing together the 'low-and-slow' attacks which try to stay under the radar. Individually, these alerts may not raise any red flags. But together, they start to illuminate a pattern of coordinated behaviour. ML also automates repetitive steps such as SIEM searches, data gathering and sandboxing – once again accelerating time-to-detection and response while improving consistency.

ML optimises SOC operations by empowering teams to build pre-defined playbooks featuring selective logic. These can be designed to help guide Level 1 analysts through investigation steps and even automatically contain, remediate and close cases based on defined criteria. That's a great way to lower the cost of SOC operations and the skill threshold for entry-level analysts. A SOC run in this way might process millions of automation actions per day before analysts even get involved in cases. This frees them to focus on what's most important and prevents a threat response bottleneck.

AUTONOMOUS SOC STILL REQUIRES ANALYSTS AS AI CAN GET DETECTION AND RESPONSE WRONG

In contrast to ML, generative AI (GenAI) has yet to come of age in the SOC, despite the proclamations of many vendors. It still largely lacks the precision needed to replace ML. While the latter provides a consistent outcome every time, a GenAI bot might respond differently each time to the same question. That's less of a problem during investigations, when outcomes can be more subjective. But it's no good in a detection context where there is zero room for error in diagnosing attacks.

Yet at the same time, innovation is happening at pace and GenAI is already proving its worth in certain areas. Note taking and summarisation is perhaps the most obvious example. If, for instance, the SOC deals with around 150,000 alerts every month, every alert requires a written summary when it's investigated. The vast majority of the time, all those comments will be identical, because it's the same thing that's happening. Writing two or three paragraphs on the alert takes a surprising amount of time, but AI can reduce that duplication.

GenAI can also lower the technical skills required for a number of critical SOC tasks, including natural language explanations of alerts and investigations, and analysis of large security data sets. GenAI can be set to work writing reports through LLM prompts, and can even support free-text threat hunting – by translating questions into syntactical queries of SIEM and EDR data. It can also support DevSecOps engineering by enabling the development of playbooks and detections quickly and intuitively for users.

Agents could take things a step further, by working autonomously through entire workflows – reasoning, planning and investigating, without the need for human involvement. As an assistive technology it has great potential not only for detecting threats, but also suggesting remedial actions. The latter are often the same across multiple detection types, so auto-remediation like this could save analysts a significant amount of time.

The kind of continuous learning, adaptive decision making and contextual reasoning that agents could introduce bodes well for the future of the SOC. AI can be set to work processing, understanding and interpreting log data. ML will help to support

graphs and detection engines. And automation can be unleashed everywhere to reduce costs, increase productivity and improve outcomes.

This combination of AI, ML and automation will soon help to deliver on the vision of the autonomous SOC. It will dynamically build and run playbooks on the fly for investigations, determine the verdict of those investigations, and even remediate and contain threats automatically. In so doing, it may imperil the role of the Level 1 analyst. But that's not to say the SOC analyst will become redundant. Human oversight will be even more important as machines start to play a bigger role in SecOps.

GENAI CAN LOWER THE TECHNICAL SKILLS REQUIRED FOR A NUMBER OF CRITICAL SOC TASKS

Why will the autonomous SOC still require analysts? One very good reason is that AI can still get detection and response wrong. During recent trials, a model not only misinterpreted the threat but then continued to go on to produce a fictitious kill chain and mitigation advice. In short, it's still a work of progress and agents should initially be restricted to autonomy only in low-risk workflows like creating tickets. There are other reasons for caution. Agentic systems require continuous tuning and adjustment to work effectively and mitigate any adversarial attempts at manipulation – something that will need to be done by experts. Costs are also high. We estimate that AI would double the cost of MDR, which includes SOAR, a 24x7 SOC and engineering capabilities. That means organisations need to be selective about where they deploy it and find the use cases which align with business outcomes and perform a cost-benefit analysis before embarking on projects.

FIND THE RIGHT BALANCE

The ultimate goal then is to balance the strengths of AI, ML and human talent to optimise SecOps. What you'll then get is a kind of 'security cyborg' capable of working smarter, more autonomously and more effectively to intuitively detect and response to threats at speed and scale.

In this new era, the role of the SOC analyst will start to evolve. On balance, this should be a welcome development. With agents doing more of the Tier 1 grunt work, analysts can focus on the kind of tasks that first attracted them to the role – like threat hunting and complex analysis. They'll have more effective tools to assist them in this, without being overwhelmed with alerts. That should minimise burnout and improve the productivity of the SOC.

Organisations must begin to consider how they will embark upon their journey to the autonomous SOC because AI-driven attacks are already happening. The cyber arms race never went away. It just evolved into an AI-cyber arms race ●

Merlin Gillespie, is CTO at Cybanetix.

Human oversight will be even more important as machines start to play a bigger role in SecOps