KPMG | Google Cloud

Challenges to AI-driven security

Overview

The modern security operations center (SOC) is less of a well-oiled machine and more of a frantic triage unit in a hospital that's perpetually on fire. The sheer volume and sophistication of threats are overwhelming teams, and this pressure is forcing a hard look at artificial intelligence (AI), not as a luxury, but as a potential lifeline. Organizations are turning to AI to drag their SecOps out of the digital Dark Ages and address some painfully familiar challenges.

The SOC is facing increasing pressure to manage a growing volume and complexity of security alerts, while also addressing sophisticated cyber threats. This pressure, combined with existing limitations, is driving the need for AI adoption. Organizations are looking to AI to enhance their SecOps capabilities and address several key challenges.

Introduction

In the rapidly evolving landscape of cybersecurity, SOCs find themselves caught in a paradoxical time warp. Despite two decades of technological advancement, today's SOCs grapple with many of the same fundamental challenges that plagued their predecessors in the early 2000s. The persistent issues of manual workflows, overwhelming data volumes, false positive fatigue, and analyst burnout continue to hamper security operations, creating an increasingly unsustainable environment. As cyber threats grow more sophisticated and numerous, these long-standing challenges have evolved from operational inconveniences into critical vulnerabilities that demand immediate attention.

This historical persistence of SOC challenges, coupled with the exponential growth in both threat sophistication and data volume, has created an inflection point in the industry. Organizations can no longer afford to maintain the status quo, yet many struggle to break free from these deeply entrenched operational patterns. As we examine these challenges in detail, it becomes clear that transformative solutions, particularly through AI integration, are not only advantageous but also necessary for the future of security operations.

Market gaps and opportunities

Many of today's SOCs are affected by largely the same problems as their peers from 10 or even 20 years ago. Manual efforts, data overload, false positives, and analyst burnout plagued the SOCs of the 2000s and 2010s and they continue to affect today's organizations.

Slow manual efforts

Traditional SOC operations often involve manual tasks such as analyzing logs, investigating alerts, and responding to incidents. These manual efforts are time-consuming and can lead to delays in detection and response, leaving organizations vulnerable to attacks for extended periods.

Despite the buzz around automation, a shocking number of SecOps teams are still stuck in a reactive, manual slog. This reliance on manual effort for critical tasks like alert triage and data enrichment leads to glacially slow response times. In fact, nearly a third of teams admit it takes them hours to respond to threats, a lifetime in the world of cyberattacks where adversaries can break out in under a minute. This isn't just inefficient; it's a gaping vulnerability.

High cost of SOC operations

Operating a SOC requires significant resources, including personnel, technology, and infrastructure. The high cost of maintaining a 24/7 operation with skilled analysts is a major concern for many organizations. AI can help automate tasks and improve efficiency, potentially reducing operational costs.

This perpetual struggle to justify security spending often leaves SOCs under-resourced and unable to implement the very improvements that could demonstrate their value.

Lack of focus on high-priority tasks

With a flood of alerts and data, SOC analysts may struggle to prioritize and focus on the most critical threats. AI can help by automating triage, identifying high-priority incidents, and enabling analysts to focus on complex and strategic tasks.

SOC analysts are buried under a mountain of alerts and data, much of which is just noise. More than half of security teams report that false positives are a massive problem, and many are simply overwhelmed by the sheer volume of data. This "data deluge" makes trying to find a real threat akin to finding a specific needle in a continent-sized haystack of other needles. The result? Fatigued analysts spend their days chasing ghosts instead of hunting actual adversaries, and critical threats get missed.

Why AI/AI-agentic SOC is the solution – Benefits or what are we solving for?

AI/AI-agentic SOC solutions offer the potential to improve threat detection, accelerate incident response, and optimize resource utilization. By automating routine tasks and providing intelligent insights, AI can enhance the effectiveness and efficiency of SOC operations.

AI/AI-agentic SOC solutions also offer the potential to automate the mundane, prioritize alerts with something resembling intelligence, and accelerate incident response. By sifting through massive data sets in real time to spot anomalies a human might miss, AI can, in theory, boost the efficiency and effectiveness of a beleaguered SOC.

Identified business challenges

Resistance to operational change

Implementing AI in the SOC often requires changes to existing workflows and processes. Resistance to change from analysts and other stakeholders can hinder successful adoption.

Implementing AI means changing how things are done, and people are creatures of habit. Many SOC teams resist new models, clinging to familiar processes, even if inefficient. This resistance can stem from a fear of job displacement, a lack of understanding, or the simple fact that teams are "too busy" fighting daily fires to even think about transformation. Some organizations even believe every security incident is a unique snowflake that cannot be handled by a playbook, making automation a nonstarter in their view.

Resistance to change (specifically within a SOC) manifests in multiple layers. People feel comfortable with consistency, so they tend to gravitate towards established practices, providing a sense of control and reliability. The introduction of AI threatens this comfort zone, and the consistency people have become accustomed to. This threat of change intensifies in the SOC environment where even the smallest mistakes can yield severe consequences. Even when a change like agentic AI is presented, basic human psychology in the chaotic world of SOC prevents people from adopting the new cutting-edge solutions in favor of a known, outdated solution.

Aside from the psychology behind change, other factors (like case management) are preventing teams from pivoting their SOC approach. Case management issues materialize in the form of case overload and case uniqueness. Case overload is a paradox, a vicious cycle of inefficiency and overwork. SOC teams are drowning in the amount of case-related work they need to manage on a day-to-day basis, making them too busy to implement AI solutions. The very tools/solutions that could provide relief are rejected due to the immediate pressure of current operations. Case uniqueness on the other hand is the belief that security incidents are too unique and complex for automation. While true for some, many security incidents follow predictable patterns that AI can effectively handle. Case uniqueness concern ties to a deeper seeded lack of trust with AI solutions, and the reluctance to change.

Implementing AI in SOC can transform the way security is handled in your organization, but it also brings a number of risks along. When resistance to change goes unaddressed, an organization will see poor adoption rates of AI tools, leading to a reduced return on investment (ROI) on AI investments and minimal improvement on operational efficiency. On top of these consequences, team morale can decrease, causing increased friction within the team.

Lack of a standardized framework for AI adoption

Currently, there is a lack of standardized frameworks for AI adoption in SecOps. This can make it difficult for organizations to integrate AI solutions into their existing environment and ensure compatibility.

The typical security environment is a chaotic mess of "tool sprawl," with numerous disparate, nonintegrated tools. This fragmentation is often a symptom of organizational silos, where different teams buy their own tech without a cohesive strategy. Trying to layer an AI solution on top of this disjointed architecture is a recipe for failure, hindering visibility and driving up complexity and costs.

Tool sprawl creates a complex web of challenges that extends beyond the issue of integration. Organizational silo and building out individual tool stacks for each business unit presents a significant barrier for AI implementation. From an architectural lens, tool sprawl complicates AI adoption through inconsistent data formats, incompatible application programming interfaces and integration points, and disconnected security workflows.

A disjointed organizational approach creates cascading fractions across several areas of the organization. From an architectural standpoint, this lack of alignment leads to redundant technology investments, conflicting priorities, varying levels of AI-readiness among existing tools, poor resource allocation, and breakdowns in communication. These challenges then create their own fractions, such as elevated integration costs, increased maintenance overhead, reduced ROI on AI investments, and unexpected compatibility issues. Tool sprawl presents a cascading effect across the organization, challenging the ability of organizations to implement and adopt AI tooling.

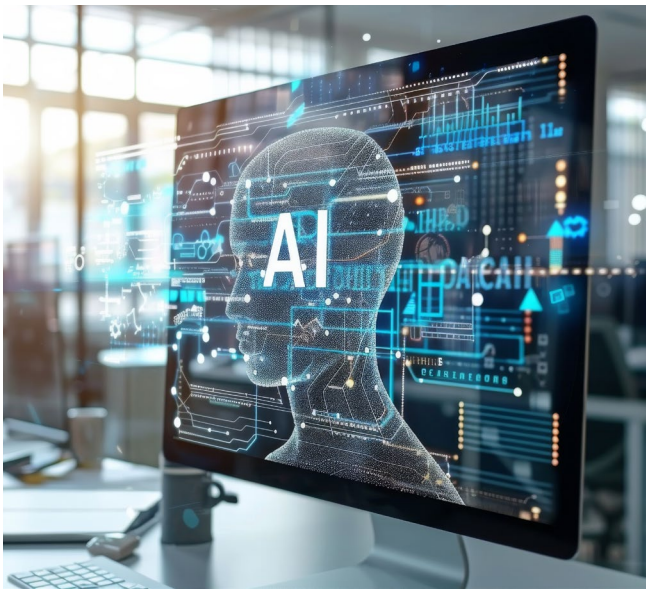

Talent management and workflow changes needed for AI integration

Integrating AI into the SOC requires skilled personnel with expertise in AI and cybersecurity. Organizations may need to invest in training and development to ensure that their staff can effectively use and manage AI solutions. Workflow changes may also be necessary to optimize the use of AI.

Integrating AI successfully requires people with the right skills—not only in cybersecurity, but also in AI, cloud architecture, and data science. There's an acute shortage of this specialized talent. Furthermore, there's a serious risk that over-reliance on AI could lead to the "erosion in foundational security analysis skills" among human analysts, leaving them unable to handle novel threats when the machines fail. Workflows must be redesigned to augment human expertise with AI efficiency, not just replacing it wholesale.

The integration of AI into SOC is a multifaceted talent management challenge that extends beyond simple hiring and training practices. The market is seeing a shortage in professionals with this specialized skillset, forcing organizations to hire from a limited pool of candidates, leading to inflated compensation demands and extended vacancy periods. AI-driven security operations require a deep talent pool. As AI systems take on more routine tasks, there's a growing need to ensure that security analyst skills are continuously developed and valued. Without intentional upskilling and role design, organizations risk creating gaps in human expertise that are critical for interpreting, validating, and guiding AI outputs.

Talent based challenges are not the only risk AI-integration presents. The introduction of AI demands a complete reimagining of the SOC workflow. Traditional processes must evolve into dynamic workflows with clear delineation between human and AI responsibilities. Quality control, oversight mechanisms, and decision-making frameworks all need to change to incorporate AI insights effectively. The AI-driven SOC revamp demands other activities be met, such as upskilling

Fragmented security landscape and its impact on business operations and costs

The fragmented nature of the security landscape, with various tools and technologies, can make it challenging to implement and integrate AI solutions. This fragmentation can also impact business operations and costs, as organizations may need to invest in multiple AI solutions to address different security needs.

Similar to tool sprawl, modern security environments have evolved into point solutions that address a specific security need but operate in isolation. These siloed solutions create a complex web of security solutions, each with their own interface, data formats, and operational requirements. Technological sprawl presents the challenge that each AI implementation must navigate an increasingly intricate web of security tools. Investment in multiple AI solutions or integrations complicates operations on the technical side while increasing spend on the financial side. Integration issues can result in multiple AI solutions and solution management, redundant capabilities/ redundant resource allocation, additional integration expenses, and overhead management of multiple vendor relationships.

The technological and financial challenges result in further operational challenges. SOC teams now need to navigate multiple interfaces and workflows, leading to an increase in required knowledge base and reduced efficiency. Disjointed integration between tools can also cause operational blind spots, where critical security information may be missed or delayed due to the complexity of correlating data across different platforms. This complexity challenges the daily operations and the organization's ability to maintain a consistent compliance program. Policies and procedures across different solutions will vary, much like the evidence collection for each, resulting in more overhead and potential compliance violations.

Key security concerns related to AI adoption include data privacy, model security, and the potential for adversarial attacks against AI systems. Organizations must carefully consider these concerns and implement appropriate safeguards to ensure the security and integrity of their AI-driven SOC operations.

Furthermore, the challenges of AI adoption in SecOps are viewed differently by various stakeholders within an organization.

SOC operations

SOC analysts are concerned about the impact of AI on their roles and responsibilities, as well as the accuracy and reliability of AI-driven insights.

Adopting AI requires a fundamental shift in mindset. Organizations must prioritize addressing underlying people and process issues rather than simply acquiring new tools. The goal is to build a cohesive security architecture that guides technology acquisition, not the other way around. This means focusing on integrating tools wisely and maturing processes before expecting AI to solve everything.

The primary concern for the frontline SOC analysts revolves around how AI will reshape their daily responsibilities and the potential for critical aspects of their work to be automated. SOC personnel have spent years honing their skills and now begin to question if their expertise will remain relevant in an AI-driven environment.

Job displacement isn't the only concern as the accuracy of AI-driven insights is a worry for these teams. Security analysts understand that threats are constantly evolving and question whether AI can adapt as quickly as human analysts to new attack vectors. Even if it can, trust issues can arise from AI-driven decisions due to the nature of the "black box" AI solutions come with. The uncertainty of how decisions are being made is a cause for skepticism that some teams may not be willing to look past.

Risk and regulatory compliance

Risk and compliance teams are focused on ensuring that AI solutions comply with regulations and address security risks effectively. AI implementation in security operations presents a complex web of regulatory considerations and risk management challenges. These teams must grapple with an evolving regulatory landscape where guidance on AI usage in security operations often lags behind technological advancement.

Risk and regulation teams encounter a unique challenge: ensuring AI solutions meet existing regulatory requirements while anticipating future regulations. This challenge evolves in complexity when teams need to consider the multiple dimensions of regulations—from data privacy concerns to industry-specific regulations. Data governance is emerging as a priority for compliance teams due to the nature of AI and the vast amount of data they require. The organization needs to establish appropriate data collection and storage policies while maintaining increasingly complex audit trails.

Aside from the ability to control, risk and compliance teams encounter the increased complexity of AI's role in decision making. New questions around liability, accountability, and transparency are emerging. Their responsibilities are evolving to include regular validation of AI model accuracy and reliability. As AI systems introduce new operational dynamics, organizations will need to adopt updated documentation and assurance practices to maintain compliance and readiness.

Executive level

Executives are interested in the ROI of AI solutions and their impact on overall business strategy and risk posture. The adoption of AI in security operations represents a significant strategic decision that extends far beyond technological implementation. C-suite leaders must weigh the substantial investment required against potential returns while considering how AI adoption aligns with broader business objectives and risk management strategies.

Executives have to balance multiple priorities when deciding to utilize AI within their organization. AI solutions can be financially demanding and require a considerable up-front investment. Aside from the technology acquisition, an organization needs to upgrade infrastructure, acquire/train talent, and prepare for business disruptions. It will be a challenge for them to demonstrate the true value of these integrations with traditional ROI calculations struggling to quantify the benefits of preventing incidents or improving efficiency.

Leaders must also consider strategic alignment as they build their AI integration strategy. AI security solutions should enhance rather than hinder business agility and innovation, which presents the challenge of developing robust security measures without creating bottlenecks for business operations. On top of the strategic alignment of processes, leaders need to balance the strategic alignment of people, navigating how AI adoption could affect relationships with customers, partners, and stakeholders who have concerns around AI-driven security.

Cliffhanger – Transforming challenges into opportunities

While the challenges facing AI-driven security operations are substantial, they are not insurmountable. The landscape of AI security solutions is rapidly evolving, offering groundbreaking capabilities that directly address these pressing challenges. Our next white paper, "AI-driven capabilities," will explore leading tools and technologies that are revolutionizing how organizations approach security operations.

From advanced security orchestration platforms that seamlessly integrate disparate tools to intelligent automation solutions that enhance analyst capabilities rather than replace them, the next generation of AI-driven security tools promises to transform the SOC landscape. We'll examine how these solutions are breaking down silos, streamlining workflows, and enabling more efficient threat detection and response.

As we transition from understanding the challenges to exploring solutions, we'll demonstrate how organizations can leverage these emerging technologies to build more resilient, efficient, and effective security operations. The journey to AI-driven security may be complex, but with the right tools and approach, organizations can navigate this transformation successfully. Join us in our next white paper as we unveil the technological innovations that are reshaping the future of security operations.

Contact us

Steve Barlock
Principal, Advisory
E: sbarlock@kpmg.com

Anton Chuvakin
Security Advisor at Office of the CISO, Google Cloud
E: chuvakin@google.com

Niranjan Girme
Director, Advisory
E: ngirme@kpmg.com

Ash Elahi
Manager, Advisory
E: ashelahi@kpmg.com

Justin Horbacz
Sr Associate, Advisory
E: justinhorbacz@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**