PALO ALTO NETWORKS

CYBERPERSPECTIVES

ISSUE 3 | 2024

PREPARING FOR A BRAND-NEW FIGHT

HOW TO BE A CYBERSECURITY WARRIOR

5 THINGS THAT KEEP CIOs UP AT NIGHT
Hint: It's Not Always the Obvious Stuff

FIGHTING BACK AGAINST CYBERATTACKERS
How to Counter AI with AI

GEN AI IN CYBERSECURITY
Threats and Defenses

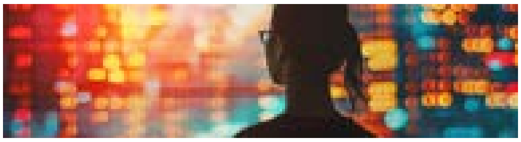

AI-DRIVEN CYBERDEFENSE TO FIGHT ADVERSARIAL AI

LEADERSHIP CHALLENGES AND STRATEGIC INSIGHTS

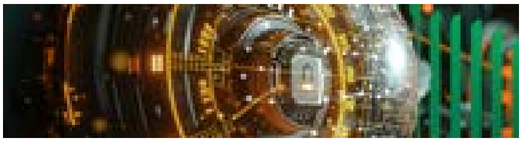

STAYING AHEAD OF EVOLVING THREATS

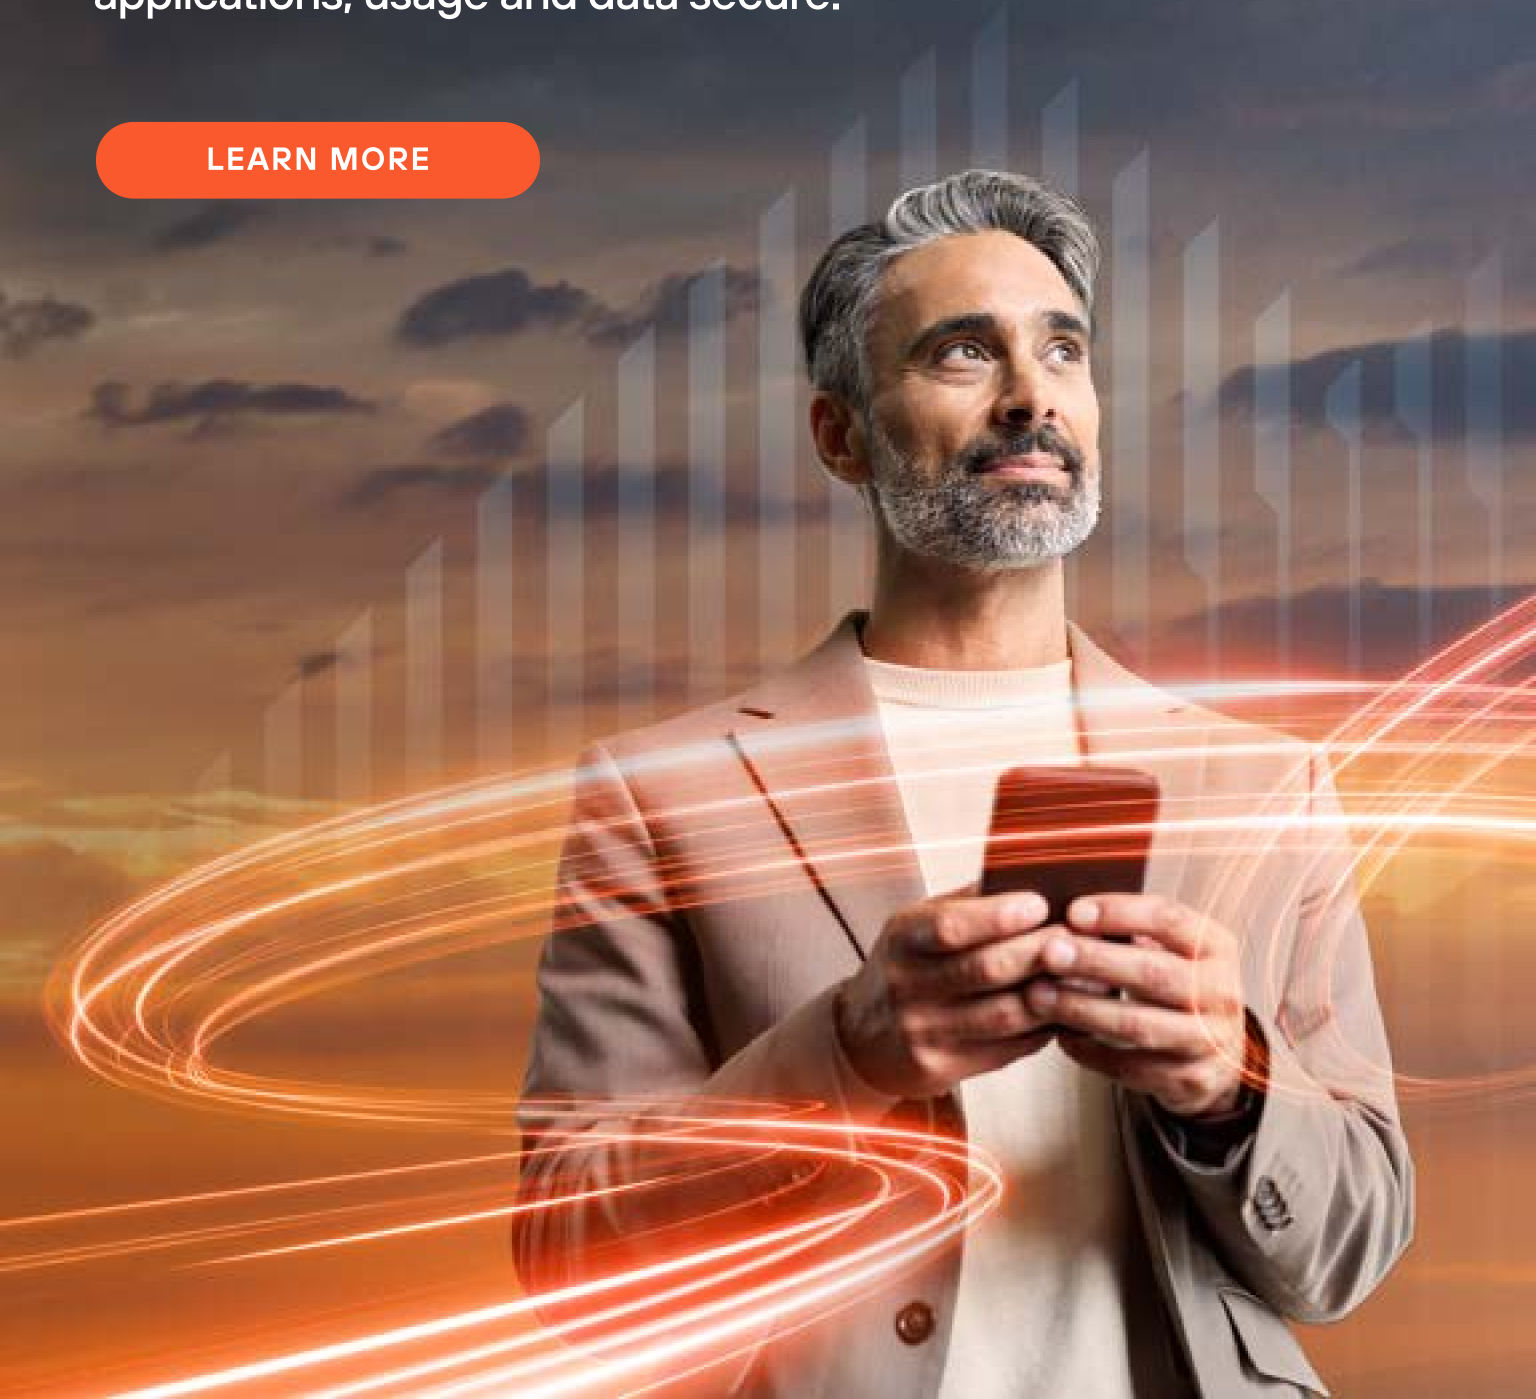

Fighting Back Against Cyberattackers: How to Counter AI with AI

Cyberattackers are increasingly leveraging artificial intelligence (AI) and machine learning (ML) to execute more advanced and sophisticated threats, amplifying the scale and impact of their attacks. As a result, the perception among many organizations is that they see the scales tipping in favor of the attackers. In fact, a recent Enterprise Strategy Group survey[1] noted a sobering statistic: 76% of organizations believe adversaries benefit the most from generative AI (GenAI), while only 24% think defenders have the upper hand.

To change this perception, organizations must adopt a proactive AI-driven defense strategy. This includes developing an AI-integrated cybersecurity plan and implementing business-aligned tactics to counter the growing AI-powered threats.

Strategies for Leveraging AI in Cybersecurity: Choosing the Right AI Models

Understanding and navigating the complexities of AI can be daunting even for seasoned IT professionals, security experts, and data scientists. The constant state of change in AI methodologies, technologies, risks, and requirements means that deeply specialized knowledge is essential to leverage AI's full potential in cybersecurity.

GenAI stands out among the various forms of AI as the most-talked-about option as well as the most widely adopted in cybersecurity. GenAI's ability to simulate and train on cyberattacks has captured significant attention, making it a crucial tool for enhancing cybersecurity measures. Predictive AI is another emerging approach, helping organizations pinpoint when and where attacks are most likely to happen due to its pattern recognition capabilities. Also, causal AI is gaining momentum because of its ability to map relational patterns between cyberattacks and responses; this allows security teams to anticipate and counter threats with unprecedented speed.

But perhaps the most exciting strategic AI commitment may be Precision AI™. This framework helps create trustful AI outcomes empowering organizations to make mission-critical decisions with confidence. Precision AI uses rich data, honed from years of data capture and analysis by Palo Alto Networks tools and systems to create a security-specific model. This proprietary model is the key to automatically and intelligently detect, prevent, and remediate potential threats.

An important part of Precision AI is its ability to handle these requirements using contextually relevant data. This contextual relevance makes Precision AI a purpose-built AI model for cybersecurity. By combining generative AI, deep learning, and machine learning, Precision AI identifies and utilizes the right data for exactly the right use cases, including threat detection, anomalous behavior analysis, and Zero Trust implementations.

In addition to identifying and implementing the right AI model, an organization's AI-powered cybersecurity strategy should include:

- **Continuous monitoring and threat detection:** Implement AI-driven tools that offer real-time monitoring and detection of emerging threats.
- **AI-specific governance:** Establish clear governance policies to manage AI applications, ensuring compliance and reducing risks.
- **Data integrity and protection:** Secure sensitive data used in AI training and operations against leaks, poisoning, and unauthorized access.
- **Model auditing and validation:** Regularly audit and validate AI models to ensure accuracy, fairness, and robustness against adversarial attacks.
- **Human-AI collaboration:** Foster a security culture that integrates human expertise with AI capabilities for more effective threat management.

Developing and implementing these strategic steps can't be left solely to the chief information security officer (CISO) and their team. Cybersecurity is a collective effort, requiring vigilance and input across the organization, including even nontechnical stakeholders. Effective AI strategies for cybersecurity must have the unwavering support and active involvement of the C-suite and board members. Creating a collaborative approach ensures that decision-making is well rounded and not unduly influenced by any sole perspective; this is critical for ensuring a comprehensive cybersecurity approach.

Tactics for Using AI Against the Other Side: Use Cases That Make a Difference

Even when all sides come together, there still are many tactical questions that need to be answered. For instance:

- Should an organization build its own model using its own data, or is it more expeditious to use a third-party, off-the-shelf model?
- Which software tools, frameworks, and methodologies are best?
- Is the right AI infrastructure in place to support compute-intensive applications?
- Are budgets sufficient in size, scale, and flexibility (remember, new AI advances are appearing daily)?
- Does the cybersecurity team possess the appropriate experience and expertise to understand AI-powered threats and leverage AI for more efficient and effective cybersecurity?
- Is there a full understanding of where AI is already being used inside the organization, including "rogue AI" efforts that are surreptitiously launched without official knowledge, backing, and support?

Addressing each of these questions from a tactical perspective is essential in using AI for good in cybersecurity. But perhaps the key tactical decisions to get the most from deploying AI for cybersecurity come down to selecting the most appropriate "high gain" use cases and applications. According to industry researcher Enterprise Strategy Group[2], AI already is invaluable in use cases that "improve security team productivity, accelerate threat detection, automate remediation actions, and guide incident response."

One of the key benefits in using AI for a wide range of use cases is its ability to limit and even overcome the negative effects of both the cybersecurity skills gap and the AI skills gap. Each on its own has been a major drain on organizations' efforts and a bottleneck » in getting the job done right. Filling in those two gaps has created a challenge of Grand Canyon–esque proportions, requiring executive commitment to allocating the proper resources.

This doesn't mean that organizations should jettison their hiring plans for both AI experts and cybersecurity engineers simply because AI adoption provides tangible benefits. Plenty of both will still be needed, but leveraging the key AI use cases for cybersecurity will rely heavily on the technology's innate automation and contextual awareness.

Here are a few specific use cases where AI will make a big difference in cybersecurity effectiveness (getting the job done in any way possible) and, especially, efficiency (doing so as quickly, frictionlessly, and cost-efficiently as possible) that should be in the consideration set for your tactical plan:

- **Advanced malware detection:** Cybercriminals are getting more creative in their use of AI to create and launch malware attacks. Cyber defenders, on the other hand, can use signature-based detection to extend the capabilities of traditional antivirus software, using signature data that leverages data on emerging threats.
- **Threat intelligence:** Even though most organizations subscribe to one or more threat intelligence services, the impact of AI on hackers' ability to introduce new threats faster than ever means threat intelligence tactics must similarly move ahead. AI provides more accurate and precise data analysis based on huge data volumes, as well as offering predictive analytics to spot problems before they emerge and to have the right response and remediation plans in place.
- **Real-time threat monitoring:** Continuous monitoring of system logs, network traffic behavior, user activity, and security infrastructure health is essential and AI makes that an integral part of overarching cybersecurity frameworks.
- **Anomaly detection:** AI algorithms—especially those with contextual awareness, such as Precision AI—are great at rooting out and surfacing abnormal, unexpected data or user behavior that could signal a vulnerability, threat, or active attack.

Next Steps Toward Successful Use of AI in Cyberdefense

While many organizations are already taking steps to use AI for cybersecurity-related use cases, the strategies and tactics are fluid, dynamic, and always changing. But here are a few tips to help you get started—or to improve your chances of success:

1. Cybersecurity is a strategic initiative, so AI-powered cybersecurity absolutely must be a critical aspect to an overarching cybersecurity framework.
2. Don't wait to get started. If you haven't put a plan in place, you're already way behind the curve, and your risk profile is expanding by the minute. Conveying a sense of urgency is critical throughout the organization, including at the C-suite level and with the board of directors.
3. Make sure you have the right people on the strategy team. They should represent the full spectrum of the organization, not just the technical side. And strategy development must include representatives from business units, such as sales, marketing, legal/compliance, finance, and operations.
4. Your strategic plan for AI in cybersecurity should be a living document, evaluated and updated regularly and frequently to reflect the breakneck pace of technological improvements and the frightening speed with which new AI-powered attacks are launched.
5. Don't try to boil the ocean when it comes to use cases. Especially in your early stages of introducing AI for cybersecurity, pick a few use cases that will be relatively easy to implement and learn from, balanced with a handful of more challenging but big-impact use cases that really move the needle toward cybersecurity resilience.

Learn more about how to fight AI with AI at paloaltonetworks.com/precision-ai-security.

[1] Gruber, Dave, et al. "Generative AI for Cybersecurity: An Optimistic but Uncertain Future." Enterprise Strategy Group by Tech Target, 15 April 2024.

[2] Ibid.

paloalto® NETWORKS | CYBERSECURITY PARTNER OF CHOICE

Unify your data empowering AI to block more attacks.

That’s the power of **platformization.**

EVOLVE YOUR SECURITY

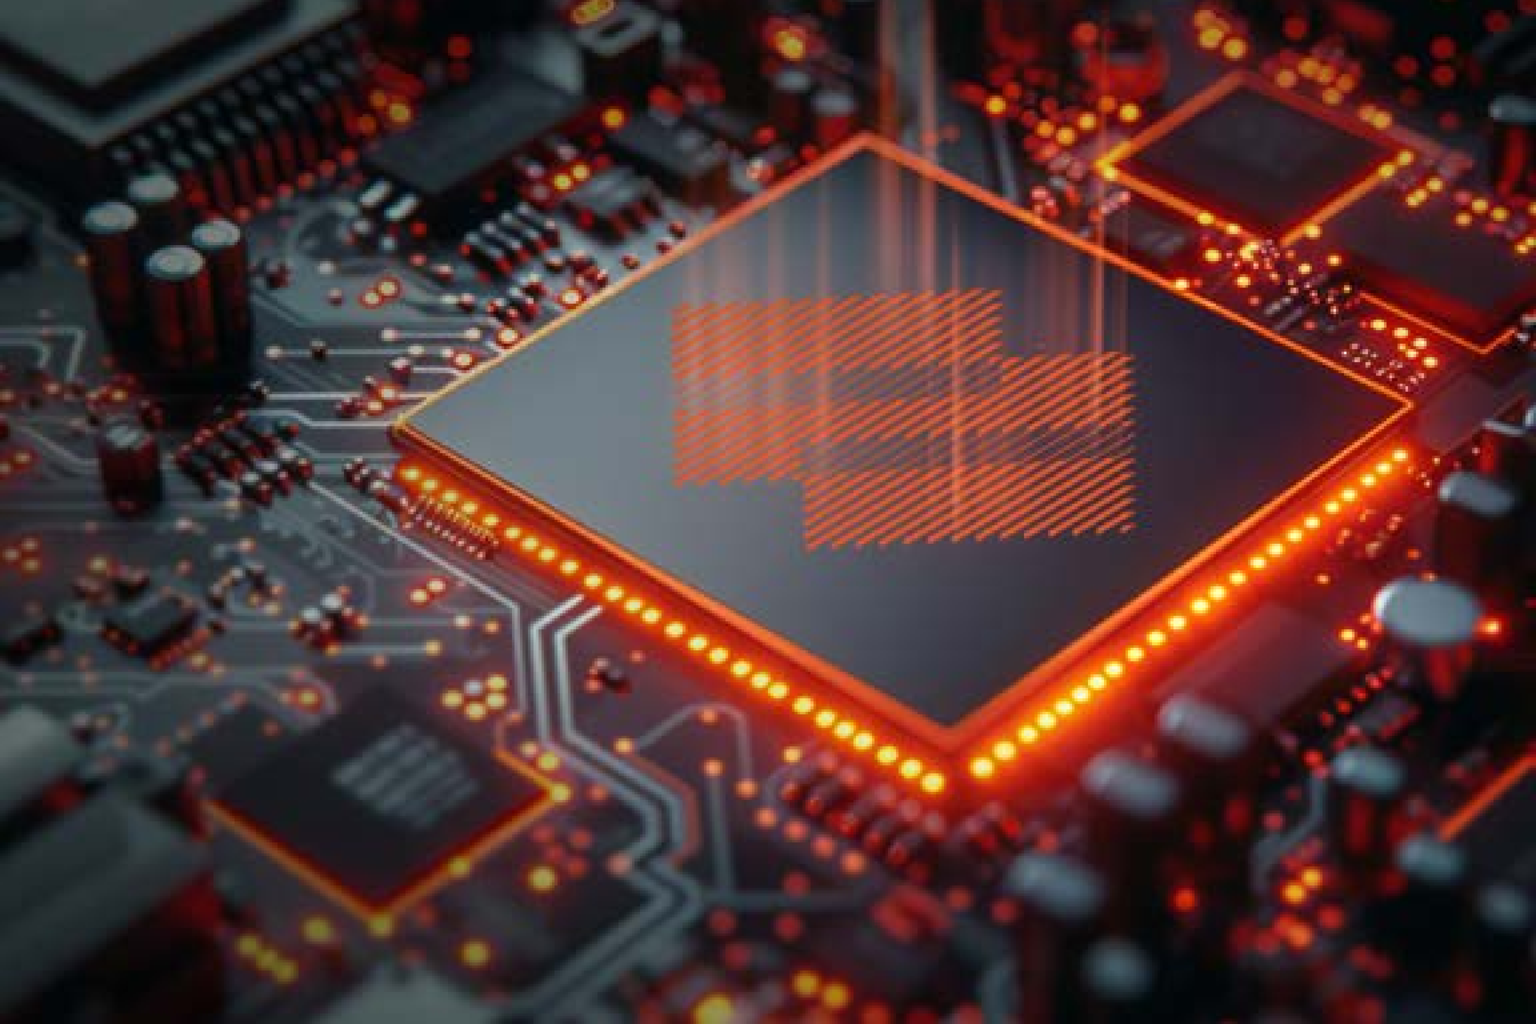

Leveraging AI as a Force Multiplier for Attack Surface Management

By Andrew Scott

Director, Product Management for Cortex, Palo Alto Networks

Attack surface management has grown more complex as attack surfaces increase in size and threats multiply in number, intensify in speed, and expand in diversity, giving the attackers a decided edge. But organizations can use artificial intelligence in new and innovative ways to gain a sustainable edge in controlling attack surfaces.

Attackers recently developed new tools to identify our rapidly growing and complex attack surface. To make things even more problematic, they're using a technology that is as readily available to them as it is to us as defenders: artificial intelligence (AI).

This has changed the rules of the game when it comes to attack surface management (ASM), an increasingly complex intersection of technologies, risks, methodologies, and potentially devastating outcomes. And here's the reality for chief information security officers (CISOs): AI has made the future of ASM extremely complex to understand and address.

Over the years, there have been three big challenges associated with using AI for cybersecurity, and until very recently, no cybersecurity solution has overcome all three. They are:

- Customers typically do not know large and substantial fractions of the IT assets they own. They often know at most about two-thirds.
- There has been limited data integration between different types of data and different cybersecurity tools, impeding the creation of efficient and unified views of organization-wide security risks.
- Attackers are innovative in their tactics and very good at changing tactics over time, which makes it hard to pinpoint their attack vectors and exploit choices.

Any attempt to make ASM more efficient and effective had to consider those three challenges.

The Role of AI in ASM

Building a truly world-class and world-leading solution for attack surface management is challenging because you have to build several components simultaneously. We believe you need to build and integrate as many as five components to do justice to ASM. But what happens when you do a great job on four but miss the mark on the fifth?

Unfortunately, you don't get 80% of the benefit because you've flawlessly done four of the five components. You get substantially less of the potential benefit—maybe as little as 10%, which means your ASM defense doesn't work because of the

It’s tricky to identify and overcome challenges associated with around-the-clock monitoring of essential IT and operational technology functions. One reason centers on the notion of ownership; specifically, we have long struggled with the idea of who owns what on the internet.

need for all this to be tightly integrated into a single solution.

Fortunately, ASM can be enhanced by using AI. ASM is enhanced by solutions like the Palo Alto Networks Cortex Xpanse platform in this way. The technology offers numerous benefits for ASM, including automation, scalability, performance (at scale), accuracy, and much more. It is a powerful, efficient, and reliable force multiplier of our cybersecurity teams, technologies, and ASM-directed processes.

Challenges of Continuous Monitoring of Digital Assets

In particular, AI aids in the efficient operation of ASM by providing high-quality, reliable, and consistent data collection. That’s important because while it’s fairly simple to gather attack-related data on the internet once in a “point in time” manner, it’s far more difficult to do that regularly and consistently, time after time.

It's tricky to identify and overcome challenges associated with around-the-clock monitoring of essential IT and operational technology functions. One reason centers on the notion of ownership; specifically, we have long struggled with the idea of who owns what on the internet.

For example, assume we identify a company, and that company has a subsidiary. We may know that the subsidiary bought another company six years ago, and they had a web server configured in a specific way. Today, that web server is running on Azure, exposing a database server to the internet insecurely. That’s complex, but it’s also relatively simple. But when this scenario is extrapolated across our customer base, it becomes substantially more complex.

Another issue to consider is the knowledge graphs built to understand customers and their assets over time. That knowledge graph must be perfectly accurate for ASM to work properly. This can only be done efficiently and comprehensively with AI.

Finally, keep in mind that attackers have figured out that they can use AI to expand their attack surface in the same way as we seek to defend it: automated, scalable, with high performance and precise accuracy. They want to use AI to maintain the competitive edge they built up over the years against cyber defenders.

In order to prevent that from happening, we have to take AI to the next level to fulfill ASM's potential.

An AI-Charged ASM Solution That Puts Defenders in the Lead

A key element of ASM is its ability to provide real-time, comprehensive visibility. However, historically, ASM hasn’t been as good as providing the ability to do anything meaningful and actionable with the results we tell our customers. Suppose we leverage a huge amount of technology built natively within Xpanse and our company’s broader portfolio of security solutions. In that case, we can tell our customers that they can know not just about all their digital assets but especially where attacks are taking place—or are likely to occur—and how to find and fix the problem.

»

This integrates data, AI, and workflows that give our Xpanse platform the power and performance CISOs need. It also enables the integrated technologies to come together in powerful ways to make ASM more intuitive, actionable, and successful at spotting and blocking threats. And, as important as it is to remediate the impact of threats, we need to focus on preventing exploits. Attackers have gotten so good at using AI and other tools to infiltrate and exploit our systems that we don't have nearly enough time to respond to threats before an incident occurs.

In some ways, this is a new, important personification of the "shift left" mentality that has taken the software development lifecycle by storm, and for good reason. However, there is an important distinction between software development and AI. For instance, software is a fairly static thing that gets updated periodically with new versions and patches; it exists close to its original form, doing similar things over time.

Conversely, AI is part of a dynamic process—especially when it's done properly and appropriately. We call this a "data flywheel," where in order to have really great AI that can help ASM do more and stay ahead of the attackers, it is part of a continual development process. Data comes in from customers, is understood and processed, and this is done in a way that is normalized with other data. If you've chosen your data correctly, even basic regressions of the right data can have a profound impact.

Another important aspect of our work with the Cortex Xpanse platform is the use of what we call Precision AI. This unique combination of technology and processes from Palo Alto Networks enables Xpanse to detect and prevent attacks in real time. By combining machine learning, deep learning, and generative AI—all trained on the largest security data lake in the world among pure-play cybersecurity leaders—cybersecurity strategies like ASM are enhanced immeasurably. This happens by:

- Stopping threat variants in real time, without signatures.
- Understanding unstructured data and preventing leakage of sensitive information.
- Continuously improving detection rates—and reducing false positives—by creating new attacks using GenAI.

These capabilities have been integrated into Cortex Xpanse to create a stronger, more intelligent, and more resilient ASM defensive framework. This is done by leveraging the most diverse and highest-quality data, a platform-based approach that makes data shareable and accessible, and the right AI and cybersecurity expertise to create the most accurate and actionable models.

Looking Ahead

It is important for CISOs and all cybersecurity practitioners to anticipate what attackers may try next. Fortunately, we have the benefit of looking back and learning that the attackers are truly committed to their mission and have the skills and tools to make it happen. It's our job to ensure that we have the right tools, strategies, and processes to prevent that.

We must look beyond the innovative technology available to us and take a more strategic approach to anticipating and thwarting attacks. CISOs, working with their C-suite partners and boards of directors, must adopt a much more proactive mindset. We talked earlier about the "shift left" movement, but we need to think and act proactively. Remember that by the time an asset is breached, it may be too late to do any meaningful remediation.

This is especially true if you're talking about unmanaged assets, which unfortunately tend to be about 30% of all assets for Fortune 500 organizations. If you have an unmanaged asset that's on the internet, you have no time to respond. You're not likely to know when it has been exploited, and therefore, all the lateral movement that has happened since the time of the exploit now happens more quickly without being monitored.

Of course, making the right financial and human resources investments is a given for proper ASM and efficient cybersecurity solutions, especially when it comes to spotting and dealing with a new breed of AI-based attacks. Given the mission-critical nature of the assets we must protect, this is a time for fresh thinking and creative solutions.

We have the opportunity to get out ahead of the attackers and stay ahead of them. Let's not let that opportunity pass us by.

Crush It, Don't Get Crushed — Combat SOC Analyst Burnout with AI

Dena De Angelo

Content Marketing Manager, Product Marketing, Cortex, Palo Alto Networks

Anyone who works in cybersecurity knows that it's full of rewards and challenges, with threat actors keeping folks on their proverbial toes. And with artificial intelligence (AI) permeating cybersecurity at seemingly warp speed, it's critical for practitioners to stay up to date on the latest developments and ensure they are integrating AI responsibly into their security protocols.

It's also crucial for current and aspiring security analysts who work in the trenches to understand its impact and prepare for their futures. Let's explore how AI is reshaping SOC analyst roles, address the critical issue of burnout, and discuss practical advice for thriving in this new era.

The Current State of SOCs — Challenges and Burnout

Today's SOC analysts face a myriad of challenges that contribute to high stress levels and burnout. The sheer volume of data they must process is overwhelming, often described as finding needles in ever-growing haystacks. The complexity of managing multiple, disparate security tools further exacerbates these issues, leading to implementation challenges and inefficiencies.

The psychological toll of these challenges cannot be overstated. The repetitive nature of investigating false positives can be soul-crushing, leading to burnout[1] and high turnover rates. In fact, research shows[2] that 65% of IT security operations personnel admitted that the stress levels within the SOC environment had led them to contemplate switching careers or leaving their current jobs. This turnover not only affects security teams but ripples through entire organizations, impacting overall cybersecurity effectiveness.

Organizational conflicts, such as decentralized operations and tensions between IT and infosec teams, further complicate the SOC analyst's role. These challenges collectively create an environment where burnout is not just a risk but an increasingly common reality.

The AI-Powered SOC — a New Paradigm

Artificial intelligence significantly enhances the cybersecurity toolkit, offering powerful solutions that

»

can mitigate many of the challenges that contribute to professional burnout. AI is revolutionizing SOCs by expediting threat detection, automating triage processes and enabling intelligent incident response.

AI's ability to process vast amounts of data at unprecedented speeds allows for the identification of patterns and anomalies that human analysts might miss. The enormous data quantities that machine learning can analyze is beyond human capacity, creating exponential scale for the SOC. This capability facilitates near real-time threat detection, significantly reducing the time between initial compromise and discovery. Moreover, AI systems can automatically categorize and prioritize alerts, drastically reducing the flood of false positives that often overwhelm Tier 1 analysts.

In incident response, AI-powered systems can suggest or even automate response actions based on historical data and learned patterns, accelerating resolution times. Additionally, AI excels at data enrichment, providing deeper context and understanding of security events, which can help analysts quickly grasp the full picture of an incident.

The Evolving Role of SOC Analysts

As AI takes on more routine tasks, the roles of SOC analysts at all levels are evolving:

- **Tier 1 analysts** are seeing their focus shift from alert triage to deeper investigation of potentially malicious activities. They're developing skills in AI tool operation and interpretation of AI-generated insights, while cultivating critical thinking to validate and contextualize AI findings.
- **Tier 2 and 3 analysts** are becoming experts in AI systems, understanding their inner workings and limitations. They're developing advanced programming skills to customize and optimize AI tools, leading AI-driven initiatives, and focusing on complex investigations and threat hunting that AI can't fully automate.
- **SOC managers** are gaining a deep understanding of AI capabilities to make informed strategic decisions. They're developing skills in translating AI-derived insights into business-relevant actions and fostering a culture of continuous learning and adaptation to AI advancements. Importantly, they're championing AI integration and communicating its value to executive teams.

Preparing for an AI-Driven Cybersecurity Career

For those starting or advancing their careers in cybersecurity, preparing for an AI-integrated future is crucial. Embracing continuous learning is key, with a commitment to ongoing education in both traditional security concepts and emerging AI technologies. Developing a strong foundation in networking, operating systems and security principles remains essential, as AI will augment these skill areas rather than replace them.

Cultivating AI literacy is also increasingly important. While you don't need to become a data scientist, understanding how AI works in cybersecurity contexts is valuable. Enhancing data analysis skills is vital, as the ability to interpret and act on AI-generated insights becomes more central to the role.

As AI handles more routine tasks, focusing on problem-solving and critical thinking becomes even more important. These skills are needed for tackling the complex security challenges that AI can't solve alone. Building soft skills like communication, teamwork and strategic thinking is equally pressing, as these human-centric abilities become more valuable in an AI-augmented workplace.

Seeking hands-on experience with AI-powered security tools, either through internships, projects or even home labs, can provide practical knowledge and a competitive edge. Staying informed by following cybersecurity news, attending conferences and participating in professional networks helps professionals stay current with AI advancements in the field.

The Future — Toward a Self-Healing Utopia

While the future capabilities of AI are unknown, one possible scenario might be the integration of AI in SOCs moving toward greater automation and even "self-healing" systems. This future state could include automated remediation of more incidents without human intervention, and more comprehensive AI-driven orchestration across IT, security and compliance functions.

While this level of automation will take time to develop and earn trust, it has the potential to significantly reduce analyst burnout by handling routine tasks and allowing human experts to focus on more strategic, fulfilling work. The vision is of a system that can predict, prevent, detect and respond to threats with minimal human intervention, thus freeing analysts to focus on higher-level strategic work.

> **AI excels at processing data and identifying patterns, but it lacks the intuition, contextual understanding and creative problem-solving abilities that human analysts bring to the table.**

The Human Element — More Important Than Ever

Despite the advancing capabilities of AI, the human element in cybersecurity remains front and center. AI excels at processing data and identifying patterns, but it lacks the intuition, contextual understanding and creative problem-solving abilities that human analysts bring to the table. As AI systems become more prevalent, cybersecurity professionals who can effectively work with AI, interpret its outputs, and apply human insights will be in high demand.

The future of cybersecurity lies not in replacing humans with AI, but in creating powerful synergies between human expertise and AI capabilities. While AI tackles the vast majority of threats in an automated process, skilled analysts can focus on the most advanced threats, creating a more fulfilling role and career path.

Embracing the AI-Driven Future

The integration of AI into cybersecurity operations presents both challenges and opportunities. By embracing this change, continuously updating skills, and focusing on areas where human insight is irreplaceable, professionals can position themselves for successful and rewarding careers in the evolving world of cybersecurity.

Remember, AI is a powerful tool, but it's the human professionals who will drive innovation, make critical decisions, and ultimately secure our digital future. As you navigate your cybersecurity career, embrace AI as a partner in your mission to protect and defend against ever-evolving threats.

Learn More

Combat burnout and elevate others to new heights of effectiveness and job satisfaction.

Download our new SOC Analyst Career Guide and listen to our podcast, Tackling SOC Analyst Burnout.

[1] Taylor, Chalsley. "Cybersecurity Leader Burnout — Causes and Solutions." Gartner, 10 January 2023,

[2] "Ponemon Institute and Devo Technology Study Reveals 65% of Cybersecurity Analysts Consider Quitting Due to Burnout, Lack of Visibility." Devo, 29 July 2019

5 Things That Keep CIOs Up at Night (Hint: It's Not Always the Obvious Stuff)

The recent global digital outage[1] did more than disrupt commerce and business operations in multiple industries. It sent a harsh and frightening reminder to chief information officers (CIOs) that new threats will emerge around any corner at any moment—and they must be prepared.

Many CIOs lose sleep[2] anticipating a middle-of-the-night cybersecurity alert. Their concerns run the gamut from the newest ransomware attack to the expanding threat vectors of edge computing and the internet of things. Add in big-impact issues such as artificial intelligence as a double-edged cybersecurity sword[3], the rising costs of data breaches, and a deepening and widening cybersecurity skills gap, and there's no wonder why sleeplessness plagues these C-suiters.

But while the threats and their implications weigh heavily on the minds of technology executives, they are actually symptomatic of bigger, more strategic issues. Here are the five that stand out as relentless causes of insomnia and anxiety.

1. Reducing Complexity Before It Strangles the Organization

It may sound like belaboring the obvious, but complexity is a substantial and still-growing threat to solid cybersecurity. In fact, complexity is the No. 1 factor cited by countless studies and research reports[4]. Fortunately, organizations are investing heavily in digital infrastructure to make their operations more efficient and transform how they compete. However, much of the purchasing, deployment, and management of new technologies and tools is done on an ad hoc, transactional basis. As a result, there are huge amounts of incompatibilities and inefficiencies in everything from antimalware to ransomware detection.

These disparate cybersecurity silos inhibit communications about threats, bad actors, and potential remedies, making the day-to-day life of a cybersecurity leader more challenging every day. As a result, many organizations are embracing new approaches, such as integrated cybersecurity platforms. This move toward platformization is a major step forward for CIOs who are looking for any way possible to reduce complexity.

2. Building a More Beneficial and Honest Relationship with the C-Suite and Board

As cybersecurity grows in complexity and criticality, boards of directors are becoming more interested and involved. The old practices of making periodic presentations at board meetings have been replaced with more frequent, strategic, and at times "urgent" communications. The prevalence of communications tools like Zoom and Microsoft Teams makes it far easier to discuss cybersecurity threats and challenges in real time, rather than waiting for a quarterly board presentation.

While these conversations can be great ways to demonstrate their business acumen to the board in framing cybersecurity in an operational, marketing, or financial light, the reality is that these meetings are often fraught with uncertainty and complexity. That's why it's more important than ever to position the CIO as a critical, integral, and strategic part of the leadership team in the eyes of the board. That's not to say it's easy.

The term "boardroom politics" is widely used and understood by savvy executives, creating further stressful scenarios. CIOs need to think about their language for the board, with a focus on 360-degree problem-and-solution discussions in an honest, candid manner. Although many of those conversations focus on risk, it's also important for CIOs to talk about how cybersecurity strategies and tactics actually increase business opportunity.

3. Becoming an Expert in Resource Management

This is actually an amalgam of challenges, requiring deft skill in assessing and prioritizing challenges even as new ones emerge. Consider these facts:

- The cybersecurity skills gap is huge, and despite a surge in hiring around the world, that gap is still growing. The World Economic Forum[5] noted that the global cybersecurity workforce grew by more than 12% in 2023—but there still is a global shortage of more than 4 million cybersecurity-related workers.
- Organizations continue to allocate an increasing portion of their technology budgets to cybersecurity readiness and defense—yet budgets still often aren't sufficient to keep up as threats expand in volume, velocity, creativity, and effectiveness.
- Automation and intelligence-based tools are helping immeasurably, but new threats are emerging faster than ever, blunting the impact of these productivity improvements. Research from AAG[6] indicates that a cybercrime is committed every 37 seconds, putting CIOs on a continuous treadmill of threat and response.

The most successful CIOs are the ones who understand how to accumulate, assess, and marshal the right resources at the right time to make the most out of what they already have. Questions like "can we leverage AI copilots to glean operational efficiencies and reduce the outsourcing budget or at least mitigate the need to expand it?" are complex to answer but critical to success. They must also fine-tune their abilities to go to the gatekeepers, lobby effectively, and use political guile to point out the costs of not putting together the proper and sufficient cybersecurity resource portfolio.

4. Prioritizing Data Storage and Security for AI Usage

A 2022 study[7] reported that 63% of the respondents used the cloud "heavily." The use of cloud services in 2021 was 59%, and 53% in 2020. A 2023 Cloud Security Alliance report, surprisingly, revealed that 98% of organizations worldwide use cloud services.[8] Not only do CIOs need to be more than cognizant of data storage, it needs to be a top priority.

Not long ago, data storage and security were already complex issues, but AI has added a new layer of urgency. The vast amounts of data required to train and run AI models weren't on most CIOs' radars 18 months ago, yet today, these models demand large-scale, secure data environments. CIOs now face the challenge of balancing data security, privacy regulations, and managing the growing storage needs that AI technologies require—all while ensuring that sensitive data isn't exposed or mishandled.

This added burden requires CIOs to rethink their storage and security strategies, introducing new protocols and tools for safeguarding data across these AI-driven environments. What was once an afterthought is now a critical need for staying ahead of threats and maintaining operational integrity.

5. Taking Care of Professional Needs and Personal Well-Being

The responsibilities and challenges facing a CIO are significant and complex. Sure, those are prestigious, generally well-paying positions. But the pressure they are under has always been intense and often unacknowledged. Now, that pressure is building in direct relationship to the stakes of a failure. It's important to keep in mind that CIOs may »

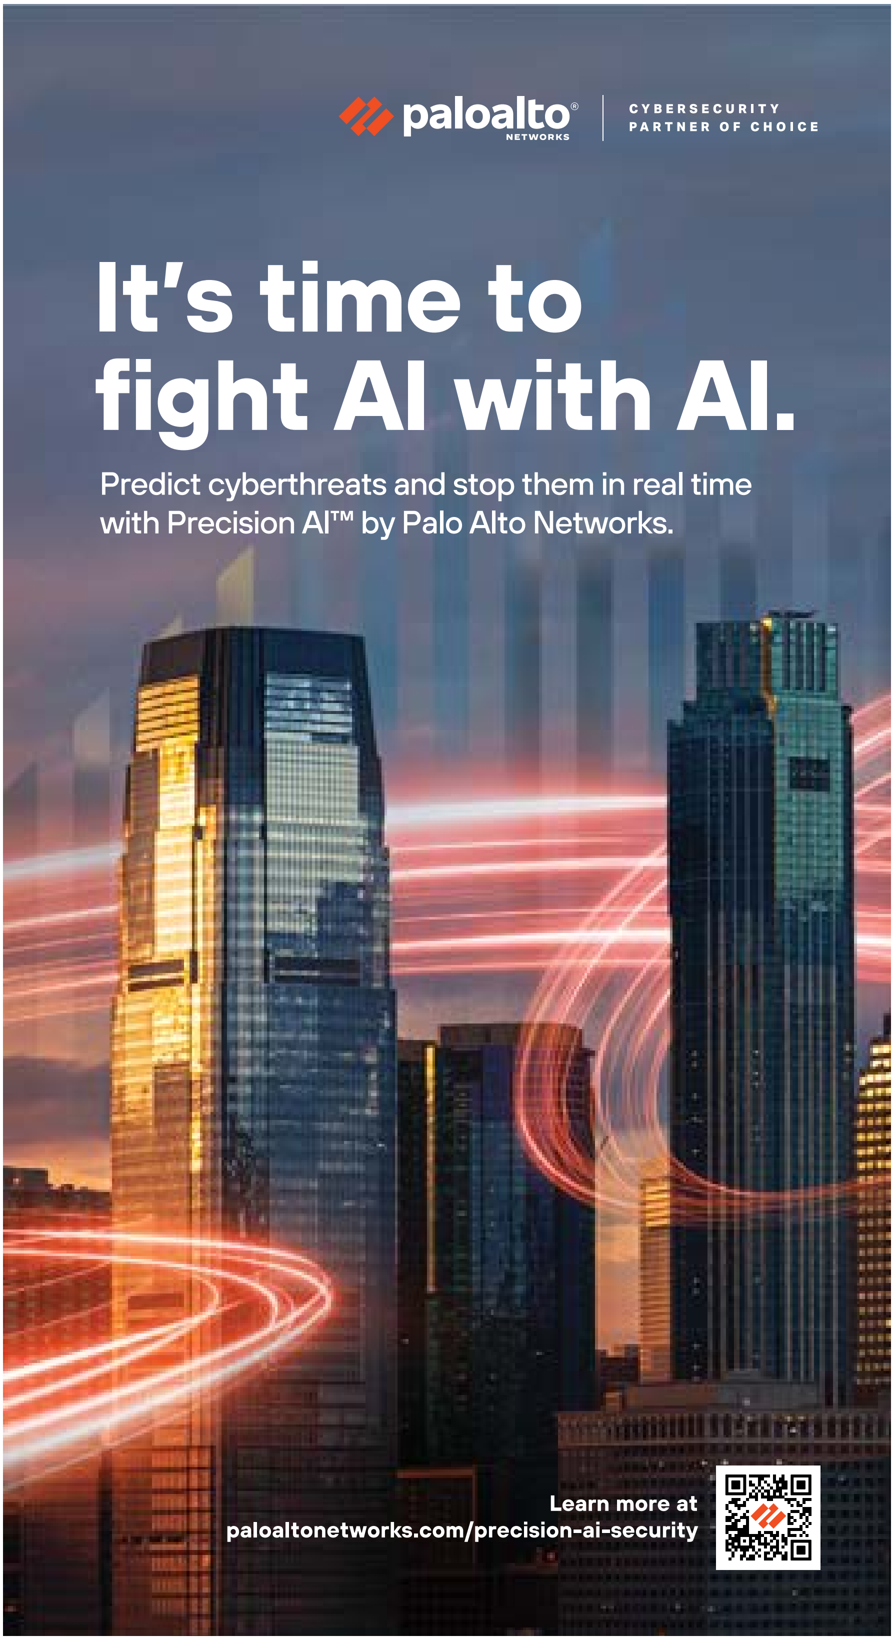
paloalto NETWORKS
CYBERSECURITY PARTNER OF CHOICE
It's time to fight AI with AI.
Predict cyberthreats and stop them in real time with Precision AI™ by Palo Alto Networks.
Learn more at paloaltonetworks.com/precision-ai-security

have trouble sleeping at night because they are, unlike their AI counterparts, human.

The life of a CIO can be tremendously satisfying and a source of great professional pride. It also can be terrifyingly confusing and fraught with risk. Those executives—and the organizations that employ and rely on them—need to understand and focus on big picture issues and not get caught up in the technical minutiae. While CIOs certainly need to focus on how to turn data into business opportunity and competitive advantage, they can't do that if this critical asset is in perpetual jeopardy.

A multitude of CIOs may find more nights sleepless than others, but embracing proactive techniques and acknowledging the risks, rather than ignoring them, could be the first steps toward getting a good night's sleep.

For more information on how you can address the skills shortage gap, check out our Cybersecurity Perspectives article.

[1] Weatherspoon, Saba, and Zhenwei Gao. "The Great IT Outage of 2024 is a wake-up call about digital public infrastructure." Atlantic Council, 6 August 2024.

[2] Unit 42. "5 Concerns for CISOs and How to Address Them." Palo Alto Networks, 2024.

[3] De Angelo, Dena. "AI's Offensive & Defensive Impacts." Palo Alto Networks, 1 May 2024.

[4] Oltsik, Jon. "The Life and Times of Cybersecurity Professionals." TechTarget, July 2023.

[5] Meineke, Michelle. "The cybersecurity industry has an urgent talent shortage. Here's how to plug the gap." The World Economic Forum, 28 April 2024.

[6] Griffiths, Charles. "The Latest Cyber Crime Statistics (updated July 2024)." AAG IT Support, July 2024.

[7] "How Many Companies Use Cloud Computing in 2024? [10 Statistics and Insights]." Edge Delta, 17 May 2024.

[8] "State of Financial Services in Cloud | CSA." Cloud Security Alliance, 5 June 2023.

Why Resilience — Not Prevention — Should Be Your Cybersecurity Goal

By Jason Koler

Deputy Chief Information Security Officer, Eaton Corp

At some point soon, our organizations will be hacked. We won't know exactly when, where, or how, but it will happen.

For all the hard work, dedication, long hours, and a substantial amount of money we all put into preventing cybersecurity threats from damaging our organizations, it simply doesn't change the fact—yes, the fact—that successful cybersecurity attacks are inevitable. As important as strong preventative measures are, we can't pretend that a strategy built solely or primarily around preventing attacks is enough.

What's the alternative? Fortunately, we have a good Plan B, which I think should become almost every organization's new Plan A. It's cyber resilience, and it starts with the assumption that all organizations will be hacked at some point. A cyber resilience strategy ensures that when our defenses are penetrated, and our data is exfiltrated, we can recover quickly and completely, thus limiting damage. Most importantly, cyber resilience ensures that we continue to operate on a nearly continuous basis with little or no downtime and with minimal negative impact.

Let me anticipate your next question: No, I am not advocating a policy of digital appeasement, where we let the bad actors do their thing and focus entirely on damage control. Far, far from it, in fact. Prevention must remain an essential part of our cyber defenses; we can't let the attackers run amok inside our systems. We must make life difficult for hackers, thieves, and digital miscreants. No one wants to be an easy target.

But a prevention-centric, or prevention-first, approach to cybersecurity almost guarantees that someone or something will get inside our walls and wreak havoc, impacting our ability to keep mission-critical systems up and running. That outcome carries serious financial, operational, reputational, legal, and regulatory impacts. In some industries, downtime or data exfiltration can carry life-or-death implications.

Why Cyber Resilience Is a Better Approach

I like cyber resilience as the cornerstone of a cybersecurity framework for several reasons. First, although much of the responsibility falls on the IT and SecOps teams, emphasizing sustained operations rather than threat prevention puts the onus across the entire organization. It becomes—in the case of manufacturing—a site-specific strategy for a managed recovery. It's not an IT or security problem but an organizational imperative. Finance, operations, R&D, supply chain, customer service—every core part of our business is involved because the emphasis is on sustaining critical business operations.

Second, it shifts the need for continuous visibility from cyberthreats to how those threats may impact the business. C-suite » executives, board members, and even our partners and customers should know that we are maniacally focused on minimizing risk to revenue, profits, and brand reputation rather than on blocking phishing attempts.

Third, resilience doesn't mean you put aside the basic blocking and tackling of cybersecurity. It reinforces the need for intelligent cyber hygiene, regular testing, a competent backup and recovery strategy, and a focus on real-time communication inside and outside the organization. Resilience doesn't replace traditional detection-and-response requirements; it augments and cements it.

Fourth, cyber resilience puts aside the often misplaced prioritization of compliance as a proxy for good cybersecurity. As crucial as regulatory compliance is as a byproduct of cybersecurity best practices, it is not an end. Cyber resilience makes compliance easier and more efficient to demonstrate. Going beyond the bare minimum of compliance, cyber resilience streamlines the demonstration of adherence by fostering repeatable and documented processes alongside a well-maintained, up-to-date, and tested incident response plan.

Momentum for Cyber Resilience Is Building

There's good news for those reading this article and wondering how you can sell a mind-shift change like this to your management and colleagues. The concept of cyber resilience has rapidly gained ground in recent years. This idea took root in the past decade and has steadily gained acceptance and adoption.

Market research organizations—another important influencer group in promoting emerging industry trends and educating executives on language and processes—also have picked up the banner of cyber resilience. Enterprise Strategy Group (ESG), a leading market-watcher and publisher of important cybersecurity market data, noted that "better resiliency" was the number-two benefit achieved by organizations in building cloud-native applications.[1] In another report, ESG said "improving cybersecurity and resiliency against cyberattacks" was the number-one consideration in justifying IT investments in the next 12 months—language matters.[2]

Manufacturing companies must pay obsessive attention to resilience because the cost of downtime or data loss can be massive to us—and not just the financial costs. It wasn't quick or easy for us, and I won't try to tell you it will be for your organization. What I will tell you is that the effort was well worth it for us, and it will be for you, too.

Next Steps Toward a Resilience-Based Cyber Strategy

Let me close with a few suggestions that may help you make this journey toward cyber resilience successful for you and your organization:

- Practice, practice, practice your incident response plan after you've built in resilience. Don't assume everything will run according to plan.
- Take time to do a thorough analysis of how cybersecurity has impacted—and could further impact—your business operations. Measure everything, and don't assume any part of your business operations is fully protected all the time.
- Educate the entire organization, from the C-suite and the board to the newest hire. Cyber resilience best practices should be communicated at every opportunity and followed with religious zeal.
- Keep asking yourselves: "How can we recover faster and more completely?" Your answers are likely to change over time.
- Help the line-of-business leaders (and their entire teams) understand how this benefits them. You need their buy-in, but then much more. You need their vision, creativity, innovation, and honesty. You won't succeed without it.

It's quite a journey to become a cyber-resilient organization. But every step is well worth it. As I said earlier, you can't stop every incident. But cyber resilience will ensure that you will be better positioned to avoid potential disasters and get back into the game immediately.

[1] Paul Nashawaty and Christian Perry, Research Report: Distributed Cloud Series: The Mainstreaming of Cloud-native Apps and Methodologies, ESG, July 21, 2023.

[2] Christian Perry and Bill Lundell, "2024 Technology Spending Intentions Survey," Enterprise Strategy Group, February 13, 2024.

Why CISOs Need to Fully Embrace AI

By Ali Khan

Chief Information Security Officer, Better

I'm not the first to say it, but it's a statement worth repeating: Artificial Intelligence (AI) is a game-changer for cybersecurity. If we trust AI and use it well, it will be a vital addition as a technologically advanced partner to the humans on our cybersecurity teams. If we fail to fully embrace AI, bad actors will use it as a weapon to overwhelm us.

AI is a double-edged sword in cybersecurity. Right now, the edge seems to be on the side of the cybercriminals. They use AI to accelerate and scale attacks, making many tactics, techniques, and procedures (TTPs) far more effective.

For example, by using off-the-shelf tools like ChatGPT, scammers can make their business email compromise (BEC) and phishing attacks more targeted and realistic. Would-be criminals with lower levels of expertise are using AI to launch far more sophisticated attacks than they would be if left to their own devices. What we're seeing now is just the tip of the iceberg.

Although AI has been around for a long time, it is still in its relative infancy in cybersecurity. Remember, ChatGPT has only been available since November 2022. Much of AI learning will be around the application stack as it evolves. It's only a matter of time before hackers use AI extensively to exploit software code.

What AI Can Do to You

Cybercriminals are always searching for innovation and new attack methods. It's a simple business model: Whatever works. They already use AI effectively and are eager to use it more extensively.

Given these challenges, many CISOs are already wary of what AI may do to them instead of what AI may do for them. A London-based cybersecurity consulting company conducted a recent survey of 250 CISOs. Eighty percent of respondents said they believe AI is the most significant cyber threat to their businesses—and a whopping 81% believe the risks of AI outweigh the advantages.[1]

This thinking, while understandable, is also short-sighted. We can't control the tools and technologies used by cybercriminals. And they will use AI exponentially—especially if we don't build strong defenses against AI-generated and AI-based attacks. If we rely solely on humans to defend against AI-based attacks, we are fighting with one hand tied behind our backs.

How do we build those defenses? In my opinion, the only way is to fight AI with AI.

»

What AI Can Do for You

The reality is that AI can do specific jobs faster and more accurately than humans can. It can iterate more quickly than any human; it can craft code in less time than a human would. A single AI can do the job of 10 people if used for the right functions. AI is becoming more capable with each passing day.

In today's environment, we use AI to get alerts, monitor, and react to threats. Within the next year, we will use AI to monitor alerts and take independent actions.

Why? Because humans are not fast enough to keep pace with the threat landscape. We have to allow AI to do what it does with appropriate limitations. To allow AI to take independent actions, we must define its rights and responsibilities. Perhaps most importantly, we must trust it to do its job, just as we place trust in the humans on our cybersecurity teams.

Another Member of the Team

If you don't use AI right, it's a useless tool that could make your organization less secure. As CISOs, I think one of our most important responsibilities is to create and nurture the culture of our organizations. If we embrace AI as a necessary tool, our teams are much more likely to embrace it.

I suggest an interesting way to think about AI—as another employee, a co-worker on the cybersecurity team with a lot of autonomy and the ability to execute important tasks comprehensively, accurately, and with enormous speed.

CISOs and our teams must be willing to trust AI to do many of the rote tasks that can bog down human workers. That means carefully evaluating what AI can do effectively, now and into the future, as the technology continues to evolve. Where does it supplement our humans, where can it take over the work of humans, and where can humans use it to be more effective? What tasks may be susceptible to human error that can be performed faster, better, and with less risk by AI?

We also must ensure that we are using AI responsibly. That means carefully crafting audits to expose errant actions. For companies in highly regulated markets, as we are at Better, we must take extra precautions when using AI. The same goes for other companies in finance or fields like healthcare. There are nuances to using AI. In our field, for example, we are concerned with potential discrimination. In healthcare, a medical misdiagnosis can have devastating consequences.

Time to Join the Bandwagon

The ideas and even the cautions I've raised in this article are all factors to consider in determining how to use AI most effectively as a proactive, preventive, and, in many cases, defensive tool in the arsenal of our cybersecurity teams. The key point is that we should be thinking about how to use AI and not whether to use AI.

Fortunately, as CISOs, we are not alone. We have a community of support and shared knowledge. We have solid and innovative partners in the vendor community who have the resources—and the motivation—to invest in AI-powered solutions that can help us all.

We already see AI being used positively in real-time threat intelligence and response, automated incident response, behavioral analytics, security information and event management (SIEM), fraud detection, and more.

Looking ahead, I see tremendous opportunities to use AI in code patching and code generation, identifying bad code and fixing vulnerabilities automatically, completing those tasks faster and more accurately than the humans on our team. Humans free from those tasks have more time to build features that can generate revenue.

There will be a learning curve in using AI to strengthen our defenses, just as there is a learning curve for cybercriminals exploiting our vulnerabilities. We can't afford to let the bad guys get too far ahead of us on the curve. If they do, they will use AI to overwhelm us.

Now is the time for CIOs and our teams to embrace AI, recognize that AI is a game-changer in cybersecurity, and build a culture around securely adopting and trusting it.

Better is a pioneer in using digital technology to make home ownership simpler, faster, and more accessible.

[1] Chris Middleton, "Why 80% of CISOs see AI as the biggest threat to their business," Diginomica, October 11, 2023.

The Hidden AI Risk Lurking In Your Business

Anand Oswal

SVP, Products, NetSec, Palo Alto Networks

Today, there are thousands of Generative AI (GenAI) tools available on the market with dozens of new AI applications being launched every month. The truth is, more than half[1] of your employees are likely already using it to increase productivity at work, and that adoption is expected to grow as more AI apps become available for more use cases.

The problem is that most of these third-party GenAI apps have not been vetted or approved for use at work, which exposes companies to serious risks. There's a reason IT and InfoSec teams vet and approve third-party applications being used within their company's ecosystem of technologies they need to understand what apps are being used, whether they are safe, and what sensitive company data, if any, is making its way into these applications. They also consider (among many other things) how the app developer handles issues, like vulnerabilities, and what controls are in place to limit or control access to only what is needed for employees to do their jobs.

The adoption of unsanctioned GenAI applications can lead to a broad range of cybersecurity issues, from data leakage to malware. That's because your company doesn't know who is using what apps, what sensitive information is going into them, and what's happening to that information once it's there. And because not all applications are built to suitable enterprise standards for security, they can also serve malicious links and act as entryways for attackers to infiltrate a company's network, giving them access to your systems and data. All of these issues can lead to regulatory compliance violations, sensitive data exposure, IP theft, operational disruption and financial losses. While these apps provide enormous productivity potential, there are serious risks and potential consequences associated with their adoption if not done securely.

Take for example:

- Marketing teams using an unsanctioned application that uses AI to generate amazing image and video content. What happens if the team loads sensitive information into the app and the details of your confidential product launch leak? *Not the kind of "viral" you were looking for.*
- Project managers using AI-powered note-taking apps to transcribe meetings and provide useful summaries. But what happens when the notes captured include a confidential discussion about this quarter's financial results ahead of the earnings announcement?

»

- Developers using copilots and code optimization services to build products faster. But what if optimized code returned from a compromised application includes malicious scripts?

These are just a few of the ways that well-intentioned use of GenAI results in an unintentional increase in risk. But blocking these technologies may limit your organization's ability to gain a competitive edge, so that isn't the answer either. Companies can, and should, take the time to consider how they can empower their employees to use these applications securely. The following are a few considerations.

Visibility

You can't protect what you don't know about. One of the biggest challenges IT teams face with unsanctioned apps is that it's difficult to respond to security incidents promptly, increasing the potential for security breaches. Every enterprise must monitor the use of third-party GenAI apps and understand the specific risks associated with each tool. Building on the understanding of which tools are being used, IT teams need visibility into what data is flowing in and out of corporate systems. This visibility will also help detect a security breach so it can be identified and rectified quickly.

Control

IT teams need the ability to make an informed decision on whether to block, allow or limit access to third-party GenAI apps, on either a per-application basis or leveraging risk-based or categorical controls. For example, you might want to block all access to code optimization tools for all employees but allow developers to access the third-party optimization tool that your information security team has assessed and sanctioned for internal use.

> **The problem is that most of these third-party GenAI apps have not been vetted or approved for use at work, which exposes companies to serious risks.**

Data Security

Are your teams sharing sensitive data with the apps? IT teams need to block sensitive data from leaking to protect your data against misuse and theft. This is especially important if your company is regulated or subject to data sovereignty laws. In practice, this means monitoring the data being sent to GenAI apps, and then leveraging technical controls to ensure that sensitive or protected data, such as personally identifiable information or intellectual property, isn't sent to these applications.

Threat prevention

The potential for exploits and vulnerabilities can be lurking underneath the surface of the GenAI tools being used by your teams. Given the incredibly fast rate at which many of these tools have been developed and brought to market, you often don't know whether the model being used was built with corrupt models, trained on incorrect or malicious data, or is subject to a broad range of AI-specific vulnerabilities. It is a recommended best practice to monitor and control data flowing from the applications to your organization for malicious or suspicious activity.

While AI tools bring the incredible potential to maximize employee productivity and enable your organization to grow its top line while at the same time improving the bottom line, these tools also harbor new and more complex risks than we've ever seen before. It's on business leaders and their IT teams to empower their workforce to confidently use AI tools while ensuring they are protected with awareness, visibility, controls, data protection and threat prevention. Once your security teams know what's being used and how, they can prevent sensitive data leaks and protect against the threats lurking inside insecure or compromised AI platforms.

This article originally appeared on Forbes.

[1] "More than Half of Generative AI Adopters Use Unapproved Tools at Work." Salesforce, 15 November 2023,

GenAI in Cybersecurity — Threats and Defenses

Michael Graven

Director, Global Consulting Operations, Palo Alto Networks

In the Unit 42 Threat Frontier: Prepare for Emerging AI Risks report, we aim to strengthen your grasp of how generative AI (GenAI) is reshaping the cybersecurity landscape. We explore how attackers are leveraging GenAI to support their efforts, and how you can formulate appropriate guardrails and protections for your organization.

With this knowledge, you'll be better equipped to fully leverage this powerful technology without creating unnecessary risk. As GenAI adoption outpaces previous enterprise technologies, understanding these developments is crucial for protecting your assets and maintaining your competitive edge. This overview will provide you with key insights to lead your company safely through the AI revolution in cybersecurity, ensuring you're not just keeping pace, but staying ahead of emerging threats.

The Evolving Threat Landscape

GenAI is rapidly reshaping the cybersecurity landscape. Defenders and attackers alike are harnessing this technology to boost their capabilities. This report will help you grasp how attackers use GenAI and how to defend against these evolving threats.

Attackers have already started using GenAI to speed up and enhance their operations. We've witnessed threat actors extracting massive data volumes in record time. The Muddled Libra group has even deployed AI-generated deepfakes in their intrusions. While significant, these changes represent an evolution rather than a revolution in attack techniques.

> AI serves as a capable co-pilot for less skilled attackers and can regenerate or impersonate certain existing types of malware

Offensive Security with GenAI

Our offensive security team now incorporates GenAI into red team engagements. We use it to bypass defenses, automate reconnaissance, generate authentic-looking content and create convincing deepfakes. These techniques showcase the potential capabilities of AI-equipped attackers. »

Defensive Strategies in the AI Era

Defending in the AI era demands both conventional and new approaches. Zero Trust architecture, rapid patching and other foundational security practices remain crucial. However, you must also adopt AI-specific defenses to outpace attackers.

Secure AI by design from the start. Monitor external AI usage, secure the AI application development lifecycle, and map the data pathways in your AI systems. Adopt AI safely by tracking AI application usage, scanning for sensitive data and implementing granular access control.

Addressing Shadow AI

Prepare for Shadow AI. Your organization likely uses AI tools already, whether you know it or not. Establish governance and rules of engagement for AI tool usage, tailored to your existing data security requirements.

Leveraging AI for Defense

Use AI to empower your defense team. Deploy AI and machine learning to uncover patterns in your logs, detections and other records. This will help your SOC scale up to match the increasing speed and volume of attacks.

GenAI and Malware Creation

Our research into GenAI and malware creation shows that while AI can't yet generate novel malware from scratch, it can accelerate attackers' activities. AI serves as a capable co-pilot for less skilled attackers and can regenerate or impersonate certain existing types of malware.

Action Steps for Executives

To stay ahead of these threats, take the following steps:

- Follow our CISO's AI Journey Checklist to guide your organization's AI adoption.
- Engage Unit 42 for an AI Security Assessment to secure employee use of GenAI and harden AI-enabled application development.
- Implement our AI-driven security products, such as Cortex XSIAM, AI Runtime Security, and AI Access Security.
- Study our Unit 42 Incident Response Report for insights from hundreds of engagements.
- Apply our recommendations to mitigate ransomware and extortion risks.
- Set up a Unit 42 Retainer for proactive and reactive consulting engagements.

Remember, GenAI adoption outpaces any previous enterprise technology. By understanding these threats and implementing appropriate defenses now, you can harness AI's power while minimizing risks to your organization.

Learn more about GenAI and security, access the Unit 42 Threat Frontier: Prepare for Emerging AI Risks report.

Unit 42 Incident Response Retainers Enhance Organizational Resilience

Wendi Whitmore
SVP Unit 42, Palo Alto Networks

Cyberattacks have increased in speed, scale and sophistication in the past year, as is highlighted in our 2024 Unit 42 Incident Response Report[1]. We have continued to see the threat landscape evolve faster than most organizations can keep pace:

- In about 45% of our cases in 2023, attackers exfiltrated data in less than 24 hours after compromise. This means that organizations must respond within hours to stop them.
- Exploitation of internet-facing vulnerabilities increased to 39% and became the top initial access vector in our incident response cases. This jump is related to several large, automated intrusion campaigns that swept across the internet in 2023.
- Attackers are more organized, with specialized teams for different parts of the attack. They're more knowledgeable and able to use IT, cloud and security tools as weapons of offense. And they're more efficient, using processes and playbooks to quickly achieve their goals.

To illustrate how these dynamics play out in real-world scenarios, let's examine two Unit 42 incident response cases that provide valuable insights into how today's adversaries operate and the strategies that are needed to defend against them effectively.

Speed & Scale

In just 13 hours, a telecom provider was devastated by a fast-moving ransomware[2] attack that encrypted files across tens of thousands of systems, exfiltrated sensitive data, and brought half of their business operations to a standstill. The client urgently engaged Unit 42 to contain the attack, prevent further data exfiltration, and help restore their operations. Within 2 hours of being called, Unit 42 began assessing the situation, quickly uncovering that the Black Basta[3] ransomware had been deployed via a phishing email, leading to widespread unauthorized access.

Given the speed of the attack, rapid deployment of Cortex XDR across the impacted environment within 96 hours was critical to containing the threat, allowing Unit 42's Managed Detection and Response[4] team to begin 24/7 monitoring and threat hunting. As part of their response, Unit 42 negotiated an 80% reduction from the initial ransom demand and successfully implemented the decryption keys to recover encrypted data. Further investigation revealed gaps in network segmentation, credential control, endpoint security and security visibility. To mitigate future risks, Unit 42 deployed additional firewalls and access control technologies, reinforcing the client's defenses against the speed and agility of evolving threat actors.

Sophistication

During a recent engagement, Unit 42 responded to a sophisticated cyberattack[5] orchestrated by the threat actor Muddled Libra[6]. Over one week, the client endured five targeted attacks that showcased the adversary's ability to adapt and exploit new pathways, even leveraging the client's own security tools for lateral movement and further compromise.

Unit 42 was swiftly brought in to investigate and respond, focusing on a holistic security approach that included containment and remediation. Drawing on deep knowledge » of Muddled Libra's tactics, Unit 42 conducted a comprehensive assessment to identify unauthorized access and determine the full scope and impact of the attacks. The team advised the client on immediate actions, including securing compromised accounts, isolating affected systems, reconstructing Active Directory, changing passwords and hardening firewalls.

With the priority of restoring systems to a secure state, Unit 42 applied patches and reinforced network defenses. This collaboration not only mitigated the immediate threat but also helped the client enhance their long-term security posture through improved practices, awareness training and regular security assessments.

What It Means to Have Unit 42 on Retainer

In today's rapidly evolving threat landscape, organizations need more than just a reactive response strategy. They need a partner who can proactively identify vulnerabilities and provide a quick, strategic response when incidents occur. This is where Unit 42 comes in. By having Unit 42 on retainer, organizations gain access to a wealth of expertise and resources that go beyond simply returning to normal operations; they gain a partner dedicated to transforming their security posture for the long term.

Unmatched Visibility and Expertise

Unit 42 delivers unparalleled visibility into the latest attack trends and tactics, combined with deep expertise in countering them. Backed by extensive telemetry data from more than 80,000 Palo Alto Networks enterprise customers worldwide and one of the industry's largest threat intelligence databases, our team has access to broader telemetry than any other cybersecurity company.

Industry-Leading Incident Response

Our incident response team is recognized as one of the best in the industry, handling more than 1,000 cybersecurity engagements annually. Named a leader in The Forrester Wave for Cybersecurity Incident Response[7], Unit 42 is known for its speed, precision and effectiveness in containing and mitigating incidents. But we don't just stop there. Our approach also focuses on helping organizations build resilience by transforming their security strategies and operations post incident.

The Power of Palo Alto Networks and Precision AI

Leveraging the advanced capabilities of Palo Alto Networks product platforms, powered by Precision AI, we bring a level of automation and insight that keeps us, and our clients, steps ahead of threat actors every time. This combination of human expertise and AI-driven technology ensures a comprehensive, proactive approach to cybersecurity.

Exclusive Offer for Palo Alto Networks Customers

Recognizing the growing need for rapid, expert intervention in today's threat environment, Unit 42 is pleased to offer our no-cost Unit 42 Rapid Incident Response Retainer program, exclusively to qualified Palo Alto Networks customers. This retainer ensures that when every second counts, you have a trusted partner ready to jump into action, minimizing impact and helping you recover with confidence.

Having Unit 42 on retainer means more than just access to top-tier incident response; it means having a partner committed to your organization's security success. Don't just react to threats, stay ahead of them with Unit 42.

The No-Cost Unit 42 Rapid IR Retainer

For qualified Palo Alto Networks customers, the Unit 42 Rapid Incident Response Retainer offers a suite of benefits:

- The initial 250 hours of Unit 42 Incident Response services
- A 2-hour response time SLA for incident response
- 24/7/365 access to the Unit 42 Incident Response team
- Expertise in threat intelligence from Unit 42

Contact your Palo Alto Networks account manager to put Unit 42 on speed dial. If you believe you are under attack, contact Unit 42 directly.

1. Unit 42. "Incident Response 2024 Report." Palo Alto Networks.
2. Unit 42. "Telecom Provider Contains Ransomware Attack & Restores Operations." Palo Alto Networks.
3. Elsad, Amer. "Threat Assessment: Black Basta Ransomware." Unit 42, 25 August 2022.
4. "Unit 42 Managed Detection and Response." Palo Alto Networks.
5. "Global business defends against multiphased Muddled Libra cyberattack." Palo Alto Networks.
6. Russo, Kristopher, et al. "Threat Group Assessment: Muddled Libra (Updated)." Unit 42, 8 March 2024.
7. Whitmore, Wendi. "Unit 42 — A Leader in The Forrester Wave for Cybersecurity Incident Response." Palo Alto Networks, 10 June 2024.

Using Time in Your Favor During a Ransomware Attack

Michael Graven

Director, Global Consulting Operations, Palo Alto Networks

Slow-Playing the Attackers

When you face extortion, there are battle-tested strategies to put the attackers on their back foot and give your team time to get ahead. The best way to mitigate a ransomware attack is by preventing it outright. But, you can't stop everything, so having robust strategies to handle an active incident minimizes the damage and recovery time.

Ransomware attacks are urgent for both sides. Organizations must quickly respond, recover and mitigate the damage while attackers need a swift payoff to move on to their next target.

But, it's not all about speeding up. Slowing attackers down is just as important.

After helping hundreds of organizations overcome ransomware attacks, we've learned that buying time can change the balance of power and set you up for a more successful resolution.

Stall During Negotiations

Ransomware attackers run businesses. They have English-speaking customer service representatives ready to facilitate payments. Their strategy is to pressure you to pay quickly to regain access to your data.

They also have an incentive to uphold their reputation. If they don't deliver on their promise to restore access after payment, word will spread, and potential victims will be less likely to pay them in the future.

Our Incident Response (IR) Report[1] showed that attackers tend to keep their promises when paid 67% of the time. While, 20.6% of attackers don't keep their promises at all, 7.8% of cases were unknown, and 3.9% partially fulfilled their promises.

It's worth engaging in communication with attackers, but not on the off-chance they will keep their promises. Instead, use the opportunity to buy your team crucial time to respond to the attack.

Negotiations are high-stakes. If you're in a high-pressure situation, call us.

During negotiations, several strategies can be employed to stall and gain valuable time:

1. Indicate willingness to pay but need more time — One effective stalling tactic is to indicate that you are willing to pay but need time to collect the necessary »

resources and gain executive approval. This not only buys you time but also keeps the attackers engaged and less suspicious.

2. Negotiate a lesser ransom — Tell the attackers that you cannot afford the amount they are asking. This can lead to extended discussions that delay the process and grant more time to focus on recovery.
3. Ask questions about the compromise — Use the negotiation period to ask questions that might reveal details about the compromise. Understanding the scope and specifics of the attack can be critical for your investigative and recovery efforts. Questions also stall the attackers as they take time to respond.
4. Play dumb — Act confused and be confusing in your responses, which force the attacker to engage with you continuously to clarify their intent and demands. By creating a back-and-forth exchange, you can significantly prolong the negotiation process.

Employing these stalling techniques can buy you the essential time needed to respond effectively. It's best to focus on recovery efforts, determine whether sensitive information has been stolen, and glean whatever information you can from the attackers.

Remember, the goal is to use every available advantage to mitigate the impact of the attack and improve your chances of a successful recovery.

Focus on Recovery While Attackers Focus on Payment

Recovering while being pressured into paying a ransom involves a multifaceted approach to help your organization bounce back as quickly and safely as possible.

Isolating the Breach

First, isolate the compromise. Identify which systems have been compromised and how the attackers gained their foothold.

Tools like Cortex XDR can be invaluable in helping you to investigate the breach and answer critical questions about the attack vectors and extent of the compromise. Additionally, this tool provides a comprehensive overview of your environment, enabling you to investigate the source of the attack and take appropriate action to contain it.

Remediation and Restoration

Next, focus on remediating vulnerabilities. Patch systems and revoke any account privileges that could potentially have enabled the compromise.

Addressing these vulnerabilities prevents further exploitation. Make sure your patch management processes are comprehensive and that all systems are up to date with the latest security fixes.

Rely on your disaster recovery strategy and use offline backups for the recovery process. Restore your systems carefully, verifying that the backup data is untampered and functional.

The recovery phase requires continuous monitoring and assessment of your systems. Keep an eye out for any signs of residual compromise or further malicious activity.

Post-Attack Preparation

Slow playing the attackers gives teams (others than your emergency responders) time:

- To understand if data has been stolen, not just locked down.
- Prepare to report the incident to the SEC, if required.
- In case of harassment, they can prepare reactive statements or support for employees or customers.

By concentrating on recovery while engaging in stalling tactics with the attackers, you'll manage the scope and depth of damage caused.

Make Time Your Ally

Even as attacks become more rapid and sophisticated, using the negotiation stage to your advantage can put the ball back in your court.

If a ransomware attack happens, commit to communicating with the attackers as a means to take control of the situation. Slow-playing them lets you execute your incident response plan effectively, gather valuable information, and focus on recovery.

If you're interested in preparing your defenses in advance, equip your organization with a trusted incident response partner like Unit 42, and be ready for any scenario.

[1] Unit 42. "Incident Response 2024 Report." Palo Alto Networks, 2024.

AI in OT Security — Balancing Industrial Innovation and Cyber Risk

Dena De Angelo

Content Marketing Manager, Palo Alto Networks

Whether defensive or offensive, cybersecurity is in constant flux. And in today's industrial landscape, the convergence of operational technology (OT), industrial control systems (ICS) and information technology (IT) is reshaping manufacturing and critical infrastructure. This convergence, while bringing unprecedented efficiency and innovation, also exposes traditionally isolated systems to new security risks, creating a complex ecosystem where AI is emerging as a powerful ally in securing these environments.

We recently interviewed Del Rodillas, distinguished product manager at Palo Alto Networks, who focuses on OT and ICS security, developing solution roadmaps and working closely with the product teams. His expertise extends to collaborating with sales teams, enabling them to better serve clients, and educating customers. Del's long-standing experience and insights make him a valuable asset in navigating the multifaceted landscape of the OT-IT convergence and emerging cybersecurity challenges in the manufacturing and industrial sectors.

The Changing Face of OT Security

Today, the manufacturing sector is embracing digital transformation at an unmatched rate. By 2026, industrial organizations are expected to employ over 15 billion new and legacy assets connected to 5G[1], the internet and cloud. As one might expect, this rapid adoption of new technologies is not without risk. The attack surface of a typical manufacturing organization becomes exponentially broader as more devices are deployed.

This expanded attack surface, coupled with the inherent vulnerabilities of legacy OT systems, creates a perfect storm for cybercriminals who are now setting their sights on these systems, »

leveraging advanced, AI-enhanced techniques to launch attacks. As Rodillas points out:

"OT-IT convergence plays a massive role in the cyberthreat landscape because it enables attackers with a more sophisticated playbook or set of tools that makes their capabilities more advanced, but it also increases the velocity and volume of their attacks."

This increased digital tangle of connectivity has made OT systems prime targets for cybercriminals. In 2021, 35% of reported OT cyberattacks had physical consequences, with an estimated damage of $140 million per incident[2]. These alarming statistics underscore the critical need for robust OT security measures that can keep pace with evolving threats.

AI — a Game-Changer in OT Security

As in other areas in cybersecurity, AI is proving to be a formidable ally in the fight against cyberthreats in OT environments. Rodillas emphasizes the importance of AI in addressing the unique challenges of OT security:

"AI plays a massive role in the cyberthreat landscape... I think AI is changing the mindset that it's not relevant to OT. It's very much relevant because of an integrated IT-OT attack lifecycle. From a sophistication standpoint, I think particularly on the social engineering phase, so people have to remember that attacks to OT primarily are ones that originate from IT and then pivot to OT.

Generative AI particularly can be used to automate the research and email generation to have a more targeted and more convincing spear phishing campaign. And the adaptability, I think, is another thing for the threat landscape, where the malware can constantly evolve, making it harder to detect and neutralize. I don't think it would be a stretch to say that AI will be applied to have more efficient and stealthy lateral movement in OT, thus shortening the time to compromise a critical asset."

Given these evolving threats, AI is not just a tool but a necessity in modern OT security.

Key Areas Where AI Is Making an Impact Protecting Industrial Environments

Enhanced Threat Detection and Response

AI-powered tools are revolutionizing how organizations detect and respond to threats in manufacturing settings. Rodillas highlights the importance of User and Entity Behavior Analytics (UEBA), stating,

"In manufacturing, the device aspect of UEBA becomes very interesting because now you're talking about OT devices, industrial IoT devices, IoT devices, IT devices, a lot of devices on the shop floor."

By leveraging machine learning algorithms, these tools can establish baselines for normal behavior and quickly identify anomalies that may indicate a security threat. This capability is particularly crucial in OT environments, where traditional IT security tools may not understand specialized industrial protocols.

Bridging the IT-OT Security Gap

One of the most significant challenges in OT security has been the disconnect between IT and OT teams. AI is helping to bridge this gap by providing a common language and unified view of the security landscape. Rodillas explains:

"Organizations are better off because there's that increased connectivity between the two environments. OT is becoming more like IT from a technology standpoint... AI can be one of these types of technologies, kind of a unifying capability."

By applying AI analytics across both IT and OT environments, organizations can detect threats earlier and map attacks to frameworks like MITRE ATT&CK, enabling better identification of threat actors and more effective response strategies.

Addressing the Skills Gap

The cybersecurity skills shortage is particularly acute in the OT sector. AI is helping to alleviate this limitation by automating routine tasks and enabling less experienced staff to handle more complex security operations. As Rodillas notes, "You need AI to take this burden off of humans and AI can do it 24/7 automatically, and it can only involve your personnel when there's a critical and high fidelity signal that is better handled by a human."

This automation not only helps to address the skills gap but also allows security teams to focus on strategic initiatives rather than getting bogged down in day-to-day alert management.

Challenges and Considerations

While AI offers tremendous potential in OT security, it's not without challenges. One of the primary concerns is the risk of false positives leading to unnecessary operational disruptions. Rodillas cautions, "If you act on a false positive and shut something down, and it causes a downtime and or some safety concern, that's like, ‘ the cure is worse than the problem’ kind of scenario."

To mitigate this risk, Rodillas suggests implementing decision assistance mechanisms that provide context and recommended actions to human operators, rather than relying on fully automated containment.

Looking Ahead — the Future of AI in OT Security

As we peer into the future, several advancements in AI are poised to have a significant impact on OT and ICS security:

- Improved accuracy in threat detection, reduced false positives
- Enhanced operational risk assessment capabilities
- Integration of AI with digital twin technologies for more effective security simulations

These digital twins, which are virtual replicas of physical systems, allow organizations to simulate and analyze potential security scenarios without risking their actual infrastructure. By applying AI to these simulations, companies can predict vulnerabilities, test response strategies, and optimize their security posture in a safe, controlled environment. This approach is particularly valuable in OT settings, where testing on live systems could lead to costly disruptions or safety risks.

Rodillas also sees potential in the application of large language models (LLMs) in OT security, particularly in querying and analyzing complex, interconnected datasets across OT and IT systems.

The convergence of OT and IT, coupled with the rise of AI, is ushering in a new era of industrial cybersecurity. While challenges remain, the potential benefits of AI in securing critical infrastructure and manufacturing environments are immense. By leveraging AI-powered tools and strategies, organizations can enhance their threat detection capabilities, bridge the IT-OT security gap, and address the persistent skills shortage in the cybersecurity field.

As we move forward, it's clear that AI will play an increasingly central role in safeguarding our industrial systems. Organizations that embrace these technologies and integrate them thoughtfully into their security strategies will be best positioned to thrive in the evolving threat landscape of the OT world.

Learn More

Download our State of OT Security Report-2024 to learn more about securing industrial environments.

1. "5G IoT Market Size & Trends, Growth Analysis, Industry Forecast [2030]." MarketsandMarkets, March 2023.
2. Alissa, Ayman, et al. "Enhancing Operational Technology (OT) cybersecurity | McKinsey." McKinsey & Company, 23 March 2023.