SPRING 2023

PALO ALTO NETWORKS

CYBER PERSPECTIVES

WILL ARTIFICIAL INTELLIGENCE REPLACE YOUR SOC?

COLORING PAGE
INCLUDED IN THIS ISSUE

LATEST RESEARCH
The State of Cloud-Native Security Report 2023

AI
The Value of AI/ML in Security Environments

Cyber Perspectives Magazine

A print magazine featuring critical research and content for cybersecurity executives from industry experts

Subscribe to receive each new issue every quarter for **FREE** at register.paloaltonetworks.com/cyberperspectivesmagazine

paloalto® NETWORKS

ARTIFICIAL INTELLIGENCE

SECURITY AS A BUSINESS ENABLER

ZERO TRUST

LATEST RESEARCH

The Value of AI/ML in Security Environments

By Matt Kraning

Artificial intelligence (AI) and machine learning (ML) are terms that are heard everywhere across the IT security landscape today, as organizations and attackers are both seeking to leverage these advancements in service of their goals. For the bad actors, it's about breaking down defenses and finding vulnerabilities faster. But what value can AI and ML offer when you're working to secure an organization?

It would be great to say that these technologies are an end to themselves for your cybersecurity and that merely adopting them means your organization is fully protected. But it's not that simple. Not all uses of AI and ML are created equal. And—spoiler alert—it's not all about using the latest algorithms.

However, in order to meet the challenges and speed of today's threat landscape, AI and ML are vital parts of a holistic security solution and should be focused on the ultimate outcome of preventing every type of attack you can and responding as fast as possible to the ones you can't.

AI Alone Is Not an Answer

Artificial intelligence itself is not a differentiator for security. In fact, there are many different AI frameworks and models in common usage today. Generally speaking, those frameworks come from academia and are open-source, public implementations available to everyone. So it's not the AI framework that makes a difference. What differentiates is how the AI is used and what data is available for AI to learn from.

What Makes AI Better and Smarter for Cybersecurity?

Regardless of the purpose, AI that learns how to act via machine learning needs high-quality data and as much data as possible to be effective. It's through that abundance of good data that AI comes to have an understanding of possible scenarios. The more real-world data it acquires, the smarter it becomes and the more experience it can leverage.

So think about this through the lens of cybersecurity. Learning from just one deployment or threat vector isn't enough. What's needed is a solution that learns from all deployments and a tool that leverages information from all its users—not just a single organization. The bigger the pool of environments and users, the smarter the AI. To that end, you also need a system that is able to handle both large volumes—and different kinds—of data.

AI is about more than just simply doing math with a computer. While data is a critical component for AI to be effective, the AI and ML itself also need to be baked into operational processes. AI and ML should not be thought of as stand-alone technologies but rather as enabling technologies that bring value to security processes and operations.

The most successful AI techniques are the ones that combine large-scale statistical pattern matching from ML to learn, along with other techniques integrating things like domain knowledge to provide a hybrid system. Statistical techniques derived solely from ML are generally unable to adapt to newly developed, previously unseen threats that by definition have little to no baseline statistics associated with them. Similarly, domain expertise can be leveraged to create logic (often partly derived from large-scale data analysis) that effectively prevents and detects specific attacker tactics and techniques.

However, aggregating these insights using expert systems results in unbalanced and skewed error rates across deployments. What's needed is an AI system that uses statistical insights from ML together with domain-driven insights from other parts of the system that can generalize to novel attacks while maintaining consistent and low-error rates for all.

The Value AI and ML Truly Provide for Cybersecurity

At a fundamental level, using AI and ML well in your organization's security enables security operations center (SOC) teams to do a lot more effectively, with fewer people. It's a multiplying »

factor that strengthens an organization's capacity and allows analysts' skills to be put towards the right work to leverage their experience.

A common use case for AI and ML in security is to help establish a baseline of normal operations and then alert a team to potential anomalies. AI and ML can also be used to improve operational effectiveness by identifying the more mundane tasks that people are doing all the time. The technology can create or suggest automation playbooks that will save time and resources.

AI and ML also help inform and power automation—which is the key to scalability in environments where staff and resources are always constrained. Every SOC today needs to address more threats that are more sophisticated, with fewer people. At the end of the day, the goal of AI and ML is to help provide a good security outcome in a way that specifically makes rapid use of very scarce resources.

How AI and ML Can Improve Security Outcomes

With security operations, there is never just one problem that needs to be solved, but rather a series of problems that are often coupled. With AI and ML helping to improve automation and remove manual processes across security operations, it can be possible to prevent more risks from becoming security incidents. If you prevent more risks, then the organization can respond more effectively, as it will be responding to fewer actual security incidents. AI and ML give you the benefit of focus and the power to scale with the threat landscape by leveraging the same tools as the attackers, strengthening your organization's overall security posture.

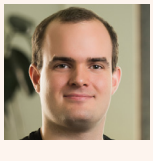

Matt Kraning is chief technology officer of Cortex at Palo Alto Networks

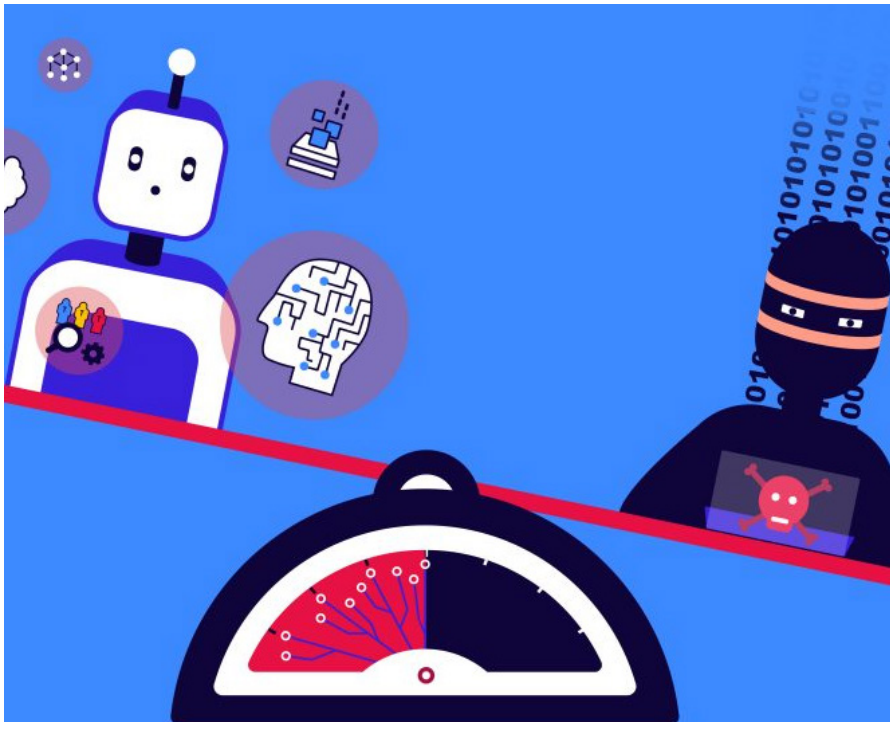

2 Innovations That Can Tip the Balance in Cybersecurity

By John Davis

What critical innovations can change the balance in cybersecurity, providing those of us responsible for defending our organizations with more capabilities against those who would do us harm?

This is not just a theoretical exercise. It is something all of us in cybersecurity need to understand—and a key national security priority.

I've given this question considerable thought in my role advising many of my former colleagues and other leaders in the U.S. government. In my view, there are two key interrelated developments that can shift the cybersecurity paradigm. They are:

1. Innovations in automation.
2. Software-based advanced analytics—including big data, machine learning, behavior analytics, deep learning and, eventually, artificial intelligence.

I'm not saying these innovations can reverse the historical advantage offense has had over defense. But improved use of automation—combined with software-based advanced analytics—can help level the playing field.

Cyber threats are increasingly automated using advanced technology. Unfortunately, defense has continued to employ a strategy based mostly on human decision-making and manual responses taken after threat activities have occurred.

This reactive strategy can't keep pace against highly automated threats that operate at speed and scale. The defense has been losing—and will continue to lose—until we in the cybersecurity community fight machines with machines, software with software. »

Prevention Is Key

Any good defensive strategy should be comprehensive with protection, detection, response, recovery and resilience. Prevention is key, especially in today's complex environment. That is where we have not invested enough—and where automation and advanced analytics can make an enormous difference.

First, let me define what I mean by prevention, starting with understanding the basic cyberattack process, sometimes referred to as the cyber threat lifecycle. This process consists of seven steps:

1. Probing;
2. Developing a delivery mechanism to get to a victim or target;
3. Exploiting a vulnerability in the network environment;
4. Installing malicious code;
5. Establishing a control channel;
6. Escalating privileged access;
7. Moving laterally within the network environment.

These steps usually occur in that order, but not always. The final step defines a successful attack, which could be encrypting data for ransom; exfiltrating sensitive data; exposing embarrassing information, or disrupting/destroying targeted systems, devices or data.

Modern cyber threat actors can work their way through the attack process more quickly than ever with advanced software and machines.

But the process still takes time—allowing defenders to see and stop a threat at any step in the process. To do so, however, defenders must have complete visibility across their network environment and be able to deliver protections everywhere automatically. Therefore, they need both sensors and enforcement points. Just seeing malicious activity without being able to stop it won't change the dynamic between offense and defense.

Tackling Speed and Scale

Automation lets security teams fight machines with machines and save their most precious resource (people) to do things that only people can do better and faster than machines. This includes hunting and deep, high-end analysis. Any other approach will never keep pace with the speed and scale of modern cyberthreats.

Software-based advanced analytics enable security teams to fight software with software. They make it possible to deploy sensors and enforcement points in all critical places in a network environment. More importantly, they enable the integration between the sensors and enforcement points.

With advanced analytics, any type of suspicious behavior in a network environment can be quickly matched to the attack process used by all known threat actors or organizations. Analytics can even identify a threat never seen before or a possible threat not directly matched to a known bad signature or activity.

Using machine learning algorithms, a decision can be rendered in near real-time—less than 10 minutes is state-of-the-art today—and a protection can be delivered automatically to stop the threat everywhere in the organization's enterprise environment without the need for any human intervention.

Defenders have access to an enormous amount of data from networks, endpoints and clouds. The right kind of data includes cyber threat indicators of compromise as well as contextual information. It does not include traditional policy and legal landmines such as personally identifiable information, protected health information, intellectual property or surveillance-related data.

Leveraging this data, it is possible to act at speed and scale with a very high degree of precision, achieving false positive rates of less than one percent. The key to this kind of effective defense is complete, continuous and consistent visibility and security controls across all elements of an organization's network environment—from the network to the cloud (public, private, hybrid, multi, SAAS) to endpoint and IoT devices.

Stopping Threats, Mitigating Risk

Cybersecurity protections that leverage automation and advanced analytics are available today and getting better as time goes by, with more of the right kinds of data to drive automated decisions and protections.

Best case, the use of these two innovations enable security teams to see and stop cyber threats before they are successful, providing an advantage for the defense. Worst case, they let security teams limit the damage of a successful attack to something determined to be an acceptable level of risk.

Why is this so important? Eliminating or reducing the advantage that cyber offense has over defense is critical to creating a more stable cyberspace. Traditionally, when offense has the advantage, it creates enormous instability. When defense has the advantage, it creates a more stable environment.

We're living in a world with an unacceptably high level of instability in the cyber domain. The risks of miscalculation, misinterpretation or even a plain mistake are just too high. Effective use of automation and software-based advanced analytics can help level the playing field between offense and defense, and create a much more effective cybersecurity posture for any organization.

John Davis is vice president and federal chief security officer at Palo Alto Networks

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

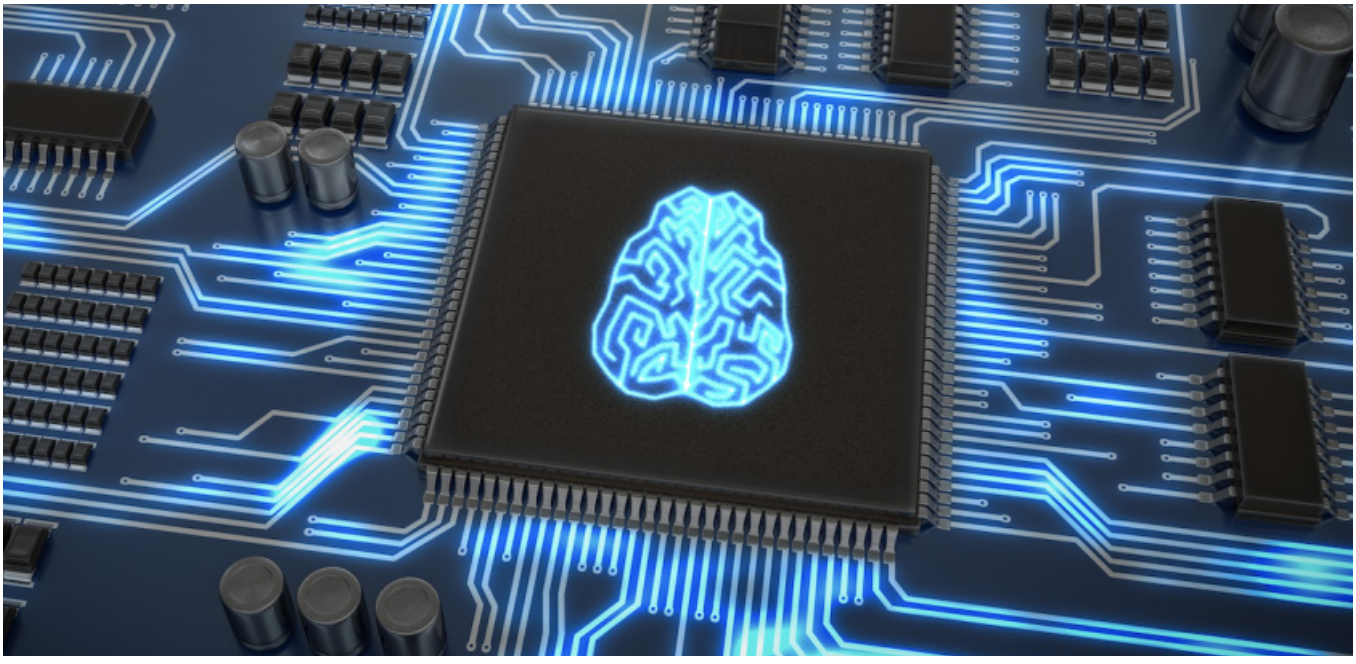

When Bad Guys Use AI and ML in Cyberattacks, What Do You Do?

By Mercedes Cardona

Artificial intelligence (AI) has become vital in defending organizations against crippling cybersecurity attacks. If your organization is not leveraging AI as a core component of your Security Operations Center (SOC), you risk falling behind those who are out to do you harm.

Companies are investing heavily in AI. The market for cybersecurity AI is expected to grow by about 23% annually, to $38.2 billion in sales by 2026, from $8.8 billion in 2019.[1] In 2023, Gartner predicts that global spending on security and risk management will grow by more than 11%.[2]

At the same time, adversaries are also using AI—as well as its subset, machine learning—to mount more automated, aggressive, and coordinated attacks. In today's environment, it's not just about using automation to fight machines with machines; it's about using the right tools and technologies to fight intelligence with intelligence.

By using AI and machine learning, hackers can be more efficient in developing a deeper understanding of how organizations are trying to prevent them from penetrating their environments.

For example, natural language processing, the same technology that makes many customer-service applications more effective, can also create better phishing emails. Recent attacks have used machine learning to mimic the usual digital activity of users, all the better to cover the tracks of perpetrators, according to.[3]

"You can buy a service that will create phishing emails and send them out for you," said Keri Pearlson, executive director of Cybersecurity at MIT Sloan (CAMS). "That service isn't necessarily designed to create bad emails. It's designed to create emails."

Hacking as a Service

We're all getting those emails, often unnoticed in our daily inbox flow. "You probably get six or seven thousand a day, as well as quite legitimate ones," Pearlson adds. "Those tools are not designed for nefarious purposes, but they can be turned into something for nefarious purposes."

Dark Web marketplaces also offer a number of AI and machine learning hacker tools—and even help-desk support to turn them into working attacks, noted Pearlson. CAMS conducts research on the operations and strategic issues that affect cybersecurity and recently published a report[4] noting the emergence of cyberattacks-as-a-service (CaaS) on the Dark Web.

The report said hackers can use AI to leverage personal information collected from social media to automatically generate phishing emails than can rack up open rates as high as 60%. That's higher than the typical spear-phishing campaigns where hackers do the job manually. »

Data Quantity and Quality Are Key

Data is the first line of defense to keep your own AI safe and to fight back against cybercriminals who may be using the technology. AI and ML tools and applications are only as good as the models and data they are built on, so keeping those models up to date with the latest threat intelligence is key.

The more data you can collect, the better. But for AI and ML to be most effective as defensive weapons, the data must have complete, relevant, and rich context collected from every potential source, whether that is at the endpoint, on the network, or in the cloud. You also have to focus on cleaning the data so you can define outcomes.

In this context, AI and machine learning are changing the paradigm in cybersecurity by enabling organizations to:

- Be more efficient in preventing malware and other attack modes from entering your environment.
- Detecting sophisticated threats, including new modes of attack that have never been seen before.
- Automating the response and using learning from each incident to prevent the same or similar threats from successfully penetrating your defenses again.

Information Sharing and Security-by-Design

Security-by-design[5] and information sharing[6] are additional tools in your arsenal, particularly if your cybersecurity teams are using AI-based tools and technologies for automated prevention, detection, and response.

With a security-by-design model, AI and machine learning can be implemented across the entire organization, including applications, devices, networks, and clouds.

"That's something that we're still figuring out—how to deal with some of the AI-enabled tools," said Josephine Wolff, assistant professor of cybersecurity policy at the Tufts Fletcher School of Law and Diplomacy.

Experts recommend security professionals learn about AI, encourage a security-by-design mindset when deploying AI and machine learning applications in their networks, and share threat information internally and with other organizations.

A number of industry groups are trying to enable information flow,[7] such as the Forum of Incident Response and Security Teams[8] and the Cyber Threat Alliance.[9] Many key industries—such as aviation, financial services and healthcare—have information sharing and analysis centers (ISACs) to help enterprises collaborate on cybersecurity and work together through the National Council of ISACs.[10]

Adding AI or any other technology can both solve problems and create new attack surfaces for hackers to exploit, said Wolff. That trade-off has to be managed. "Every time you implement some new security control, you've potentially created a new attack surface, even if you've closed off one, as well," said Wolff. "You have to adjust your sense of what the threats are."

New Tools, New Attack Surfaces

In a recent paper[11] for the Brookings Institution, Wolff noted that building better AI means using more data to fine-tune the systems; but with more data accessible, bad actors could figure out the AI models.

Security needs to navigate those risks. "It's good to have a deep understanding of what the application is. That's what lets you build the threat model of what could go wrong," said Wolff.

AI as a whole is often treated like a black box, said Pearlson. Most automation can be tested simply, by entering data and watching for the expected result, which means the model is working as expected. But AI and ML are designed to learn, which complicates the testing.

"They're designed to find the needle in the haystack. So how do you know that the output is actually the right needle and that the system hasn't been compromised?" she asked. "You're designing to find something you wouldn't have seen. Yet you don't know whether that's because it's the actual answer or because something was hacked along the way."

Adversaries share information on the Dark Web, Pearlson noted. "They're really good at sharing information. And we outside (the Dark Web) are not good at sharing information."

Organizations keep quiet to avoid damage to their reputations, their customers, their supply chains, and other consequences. "People are afraid they put a target on their back if they tell how somebody got in," she said.

Information sharing is a powerful tool, even if many organizations find it counterintuitive. Adversaries have no trouble sharing tips and best practices for taking down networks. As defenders, organizations need to become more open to sharing information, both to help others and to learn before the next attack.

The richer the data you have, the more you will need to rely on artificial intelligence to strengthen your defenses. But keep in mind your adversaries are doing the same, so it's smart to always stay at least one step ahead.

1. *Artificial Intelligence in Cybersecurity Market*, MarketsandMarkets, December 12, 2022.
2. "Cybersecurity Budgets Are Going Up. So Why Aren't Breaches Going Down," The Hacker News, February 2, 2023.
3. *The Next Paradigm Shift: AI-Driven Cyber-Attacks*, Darktrace, November 2018.
4. Huang et. al, "Casting the Dark Web in a New Light," MIT Sloan School of Management, June 2019.
5. "How to manage cyber risk with a Security by Design approach," Ernst & Young, February 7, 2020.
6. Ryan Olson, "Your Business Depends on It: Shared Threat Intelligence," SecurityRoundtable.org, accessed March 31, 2023.
7. Ibid.
8. https://www.first.org/
9. https://www.cyberthreatalliance.org/
10. https://www.nationalisacs.org/
11. Josephine Wolff, *How to improve cybersecurity for artificial intelligence*, The Brookings Institution, June 9, 2020.

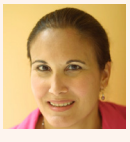

Mercedes Cardona is an editorial consultant and founder of Commerce and Reads

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

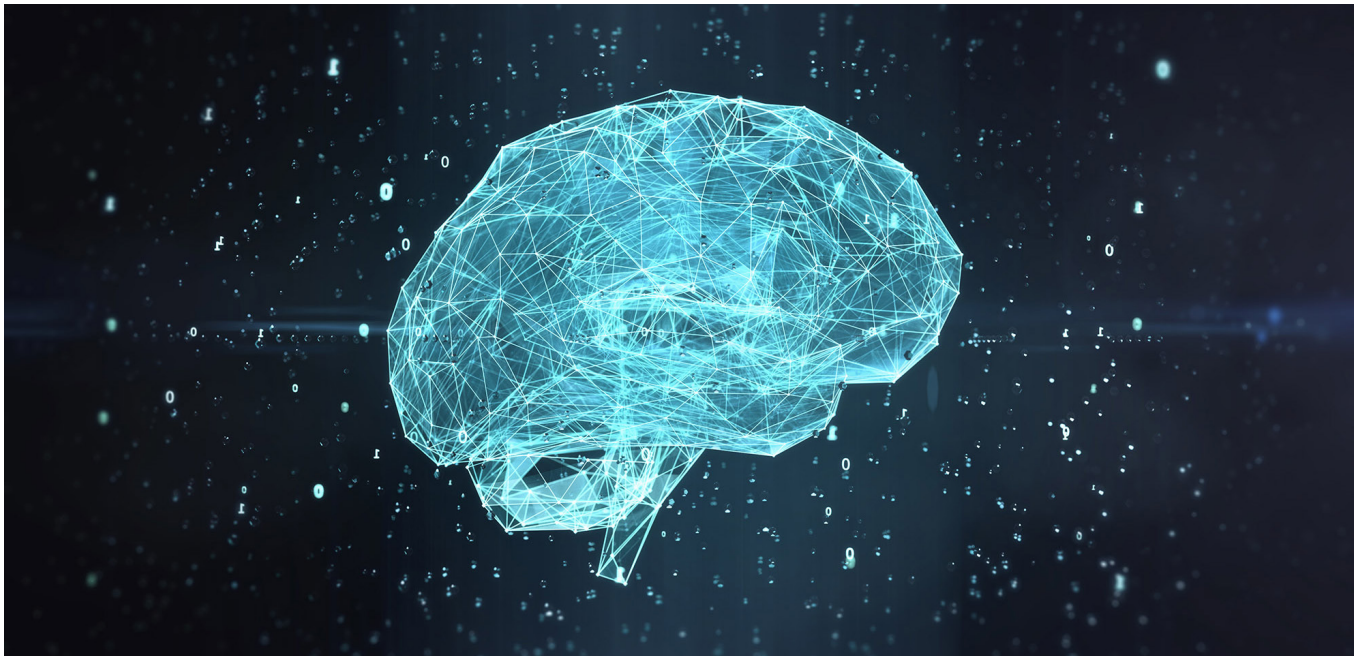

AI and Machine Learning: Do You Know the Difference?

By Al Perlman

Machine learning and artificial intelligence (AI) are transforming how organizations modernize their approach to cybersecurity. The two interrelated technologies simplify cybersecurity operations, increase efficiency and reduce risk by helping security teams detect known and unknown attacks.

Given these benefits, it is no wonder that the market for AI in cybersecurity is expected to grow at an annual rate of more than 31% through 2025, reaching more than $34.8 billion.[1] But before your organization can truly maximize the benefits of these technologies—even if you already have some knowledge of what is machine learning and what is artificial intelligence—it is important to understand a couple of things:

- Machine learning and artificial intelligence are often used as interchangeable terms, but they are not the same thing. They are related in that machine learning is a subset of AI, but each delivers different capabilities. For decision-makers in business, IT and cybersecurity, you can set proper expectations for what each can and can't accomplish.
- Both machine learning and AI can bring huge benefits to your organization—but only if they are fed the right data,[2] which means the data must have complete, relevant and rich context that is structured in a common language.

In evaluating how machine learning and artificial intelligence can simplify operations and make your organization safer from cybersecurity threats, it is important to understand the capabilities, similarities and differences of the two terms. Here's what you need to know.

What Is Machine Learning?

There are different definitions used to describe machine learning, but they fundamentally say the same thing. This is a simple definition courtesy of Stanford University:

Machine learning is the science of getting computers to act without being explicitly programmed.

Machine learning requires a large and rich data set and the use of algorithms to learn from the data. It allows a computer to make predictions about new data it has never seen before, based on patterns it has seen in the past. The viability of any machine learning algorithm is only as strong as the data modeling behind it, according to Giora Engel, former vice president of product management at Palo Alto Networks.

"The actual algorithm in use only plays a secondary role," Engel says. "If the selected data parameters do not contain parameters that can predict the result, you can use fancy algorithms, but the accuracy of the results will be very low." »

To ensure that the selected data parameters provide the data you need, you should be asking three fundamental questions of your teams:

1. **Can you see everything?** You need to be able to see data from everywhere—across cloud, network and endpoints.
2. **Can you analyze it quickly?** Once you've established that you can see something, you need to be able to quickly analyze it, which goes well beyond just storing data and applying manual actions or additional tools to analyze it. Analytics must be baked into your processes so they can be done in real time.
3. **Are you leveraging automated response capabilities?** Ask your teams how many people and tools it will take to respond and how long it will take to immediately stop an attack. If you don't have automated response with machine learning, you are facing unnecessary and incremental risk.

Bottom line: If it requires multiple people and tools to manually make sense of the data and act on it for response and future protection, you are not leveraging your data, you are not maximizing machine learning and you are not simplifying anything.

What Is Artificial Intelligence?

As with machine learning, there are multiple definitions for AI. Here are a few that are resonant, culled from a simple internet search:

- Artificial intelligence makes it possible for machines to learn from experience, adjust to new inputs and perform human-like tasks.
- Artificial intelligence is an area of computer science that emphasizes the creation of intelligent machines that work and react like humans.
- The core problems of artificial intelligence include programming computers for certain traits, such as: Reasoning, perception, and ability to manipulate and move objects.

Artificial intelligence requires the same care and attention to data collection and management as machine learning, so similar questions apply in order to maximize AI to simplify operations and reduce risk. But, as noted, artificial intelligence is a broader concept than machine learning and thus has the potential to deliver different benefits to your organization.

What Are the Differences?

To highlight the differences between machine learning and AI, we turned to Navneet Singh, product marketing vice president at Palo Alto Networks, who provided guidance as well as direction to several general respected industry reference points. Here is a summary:

- Artificial intelligence is the broader concept of machines being able to carry out tasks in a way that we would consider "smart." Machine learning is a current application of AI based around the idea that we should really just be able to give machines access to data and let them learn for themselves.[3]
- Machine learning is one of the ways we expect to achieve AI.[4] Machine learning relies on working with large data sets, by examining and comparing the data to find common patterns and explore nuances.
- You can think of deep learning, machine learning and artificial intelligence as a set of Russian dolls nested within each other, beginning with the smallest and working outwards. Deep learning is a subset of machine learning, and machine learning is a subset of AI,[5] which is an umbrella term for any computer program that does something smart.

Business leaders may think of machine learning and AI as future technologies that can be deployed sometime down the road. When it comes to cybersecurity, that is not the case, particularly with machine learning. If your cybersecurity teams are not taking advantage of machine learning today, you are exposing your organization to greater risk of being successfully attacked.

You are also probably spending more time and money on personnel resources doing jobs that can be better done with the help of machines. This doesn't mean reducing staff; it means making your teams more efficient by having them use modern, effective and smart tools to simplify operations, reduce the time needed to respond and leverage automation to fight machines with machines.

Cybersecurity is emerging as one of the critical use cases for machine learning and AI.[6] If your organization is not using these technologies, be forewarned—the bad actors who would do harm to your business are using them. When it comes to getting smart about cybersecurity, the future is now.

1. *AI in Cybersecurity Market*, MarketsandMarkets, September 12, 2022.
2. Al Perlman, "The Growing Role of Machine Learning in Cybersecurity," SecurityRoundtable.org, last accessed March 31, 2023.
3. Bernard Marr, "What Is The Difference Between Artificial Intelligence And Machine Learning?" Forbes, December 16, 2016.
4. Roberto Iriondo, "Machine Learning (ML) vs. Artificial Intelligence (AI) — Crucial Differences," Medium, October 15, 2018.
5. David Petersson, "AI vs. machine learning vs. deep learning: Key differences," TechTarget, January 24, 2023.
6. Gerd Leonhard, "It's Time to Get Real About Artificial Intelligence," SecurityRoundtable.org, last accessed March 31, 2023.

Al Perlman, co-founder of New Reality Media, is an award-winning technology journalist

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

Will Artificial Intelligence Replace Your SOC?

By Sergej Epp

Artificial intelligence no longer is the "next new thing." AI, machine learning, deep learning and other forms of algorithmic-based, automated processes are now mainstream and on their way to being deeply integrated into a wide range of front office, back office and in-the-field operations. And we certainly have seen a lot of great examples of AI being used to spot potential cybersecurity threats and preventing their infection on an organization.

As business leaders, you have given at least some consideration to the notion that artificial intelligence will completely replace your security operations center (SOC). After all, you've probably calculated the money it takes to run your SOC 24/7/365, and what it means when your CISO comes to an executive lunch or the board meeting and explains that we need more resources – i.e., people, technology and money – to fight new and more security threats. I can hear and feel the dollar signs spinning in your heads … because I've been there.

My advice to you is this: Don't rely solely on technology to protect your organization, but assess instead how AI can help to *complement* your SOC.

To help you understand why you will not be able to replace your SOC with AI, let me give you a real-world lesson from the world of competitive chess.

Deep Blue's Unorthodox Move Against Kasparov

Most of you know that in 1997, chess grandmaster Garry Kasparov played—and lost to—IBM's famous AI machine, Deep Blue. What you may not realize is that Kasparov was winning a key game when Deep Blue made what was then considered an unusual move, confusing Kasparov to the point where he lost his rhythm and, ultimately, the game. Deep Blue's unorthodox move, however, was not a calculated step to trip up the chess master. Instead, it was later discovered that Deep Blue ran into a bug[1] and made a random, rather than meticulously thought-out, move.

While Deep Blue's victory was hailed as a milestone in the evolution of AI, the "bug" influencing the outcome of the key game should be a cautionary tale in not putting all our eggs in the AI basket. In fact, sometimes you have to think and to act outside of the box (like the error in Deep Blue) and not based on predefined rules to win the game. This is true, especially when it comes to cybersecurity.

In other words, let's not get caught up in the hype around AI and machine learning[2] and assume that it is ready to replace our SOCs and the dedicated, resourceful and critically essential security engineers and analysts. »

How Is Cybersecurity Benefiting From AI?

AI and ML have demonstrated the ability to automate many tasks previously done either by SOC personnel or earlier-generation tools. And AI is a great way to automate many decision-making processes about cybersecurity. But AI will always be limited in its ability to replace human intelligence in an area that is changing as rapidly and dramatically as cybersecurity threat identification and management.

Ask any CISO or SOC analyst across a wide range of industries and geographies, and you'll get widespread agreement on a key challenge in implementing artificial intelligence in the SOC, let alone having it replace the SOC: We often don't know what the threat is and what its impact can be until it is actually spotted.

Consequently, often it is not possible to train a machine in advance to recognize completely unknown patterns. Machines, like humans, have found it extremely difficult to sort out the signal from the noise, the real threats from the false positives. Why do you think that we still have so many unfounded intrusions even in an era profoundly influenced by automated and algorithmic tools?

As we continue to compete as cybersecurity grandmasters, we look for ways to get ahead of the threats by tapping into the massive and still-growing public data set coming from threat intelligence services and other surveillance methods.

Analyzing recent incidents, participating in cybersecurity discussion groups, setting up honeypots or crafting red-team exercises all help and can become the training set basis for an AI-driven defense. But training our machines using this data is very difficult, and far from gap-proof.

Teaching the Machines

What machines are great at doing, of course, is recognizing patterns based on input and learning from human sources. I can teach a machine how to recognize a chair by showing it billions of pictures of different sizes, shapes and formats. But what happens to our machine learning when someone develops a completely new form of chair, like those ergonomic chairs in the form of a large rubber ball or some product of whimsy like a chair shaped like a farm animal or a piece of sporting apparatus like a baseball glove?

In those cases, the human brain is going to make the connection between this never-before-seen format and the functionality of a chair, while any machine will immediately fail to understand that you can sit on it unless it looks like a chair.

We still need our clever SOC analysts to teach the algorithms how to recognize it is a chair—just as they would teach the AI system to recognize a new piece of malware for the threat it is.

So, while AI and ML are not going to replace your SOC, those technologies are going to play an increasingly important role in automating decision processes at light-speed in such areas as:

- Network traffic analytics
- File or mail classification
- Endpoint protection
- User behavior analytics
- Source code analysis
- Application or database request analysis
- Process behavior (think of credit card fraud or other forms of identity theft)

Is Your Organization AI Ready?

Before being able to consume artificial intelligence, organizations often forget to transform both cybersecurity technologies and the SOC itself. The success of AI is defined by the automation and integration level of your security controls. Technologies and tools designed to block bad network traffic, quarantine a machine, remediate a problem or roll out a patch must be available and implemented beforehand—as an automated application programming interface across an entire enterprise. The advantage of rapid decisions by AI is otherwise useless if you can't act in an automated way.

AI is going to have an important impact on SOC analysts—but not the job-killing impact that news reports and pundits would have you believe. AI will actually enrich the role SOC analysts play by freeing them up to become data scientists and security architects. In those roles, they will focus on re-architecting core operational processes, ensuring that the right data is being collected and is of the highest quality and coming up with innovative "hunting" techniques and creative new ways to spot problems unique to individual industries, organizations or job functions. And the SOC analyst will sooner or later evolve into those roles.

So when business executives and boards start thinking about the role that artificial intelligence plays in supplementing and extending—not replacing—the SOC, it's important to understand that AI is going to reduce your risk, but also transform your SOC personnel.

Consequently, executives need to focus on AI's ability to automate the decision processes when machines are working under the direction of SOC personnel to ensure full threat visibility, access to the full range of relevant data and the instrumentation of controls.

Finally, remember that there are new bad actors popping up all the time, and they don't play by the rules the machines have learned and mastered. So you'd better have your own cybersecurity grandmasters at hand to ensure you can thwart the attackers as they invent new rules.

1. Klint Finley, "Did a Computer Bug Help Deep Blue Beat Kasparov?" Wired, September 28, 2012.
2. Mike Perkowski, "Don't Get Caught in the Hype: Cut Out the BS in AI and ML," SecurityRoundtable.org, last accessed April 3, 2023.

Sergej Epp is chief security officer, Central European region, for Palo Alto Networks

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

It's Time to Get Real About Artificial Intelligence

By Gerd Leonhard

Artificial intelligence has become a global buzz-phrase that gets massive media attention yet remains broadly misunderstood, at least in terms of what is real about so-called AI today and its potential impact on the future.

As a business leader justly concerned about the potential role of AI in cybersecurity, you can't afford to fall into the hype trap, regardless of whether it's utopian or dystopian.

What is AI? First, let's state that artificial intelligence is neither artificial nor intelligent (in the human sense of that word). Every machine, algorithm or technology is "artificial," i.e., not organic or biological. So this is hardly unique to AI. And true, broad intelligence in the human sense—intellectual, emotional, kinesthetic—is still a faraway goal for even the most advanced machines.

IA, not AI as in 'Ex Machina'

Therefore, I would propose to simply use "smart machines" or "intelligence assistance" (IA) as more suitable terms.

If you buy into the much-hyped fears around AI – that machines are somehow on the verge of taking over the world – you may miss a big opportunity to use the technology to make your organization more secure and innovative.

On the other hand, if you over-inflate what AI[1] or IA can actually do for you today, you may make your organization more vulnerable to cybersecurity threats by minimizing the importance of the human element. Dehumanization is never going to be a benefit!

To find the right balance, let's start with three fundamental guidelines to follow in approaching AI today:

1. **Don't let fear run your decisions—but keep asking the right questions**. We are many years away from the Hollywood worlds of robots using their superior intelligence to do away with humans. Not that such a potentiality isn't possible or dangerous (it is—which is why I am all for regulating Artificial General Intelligence/AGI); it's just that the technology is probably 30 to 50 years down the road. In the meantime, what I currently characterize as IA, can be an immensely powerful tool in driving business decisions and improving cybersecurity protections.
2. When we deploy smart machines, we humans can't get too lazy. There is a tendency to assume that if intelligent machines are on the job, there is no need for human intervention or decision-making. This is abdication and will be quite dangerous. Humans have many qualities (I call them the androrithms)[2] that machines don't, particularly in examining complex decisions that require context, ambiguity, nuance, intuition, judgment, empathy, imagination and other human characteristics. An intelligent machine might do a better job than a doctor at analyzing 5 billion images of skin cancer, but should the doctor simply accept what the machine tells her and prescribe medication without talking to the patient and examining other factors, such as the quality of the patient's life? »

2. We should not allow the use of AI to result in increasing dehumanization. We must always remember that it is our humanity that will make all the difference in a world dominated by algorithms and smart machines. As I discuss in my book "Technology vs. Humanity"[3] our future is to become more human, not less, and today the biggest danger is not that machines will eliminate us, but that we will become too much like them. If we as business leaders choose to fire people or subjugate people to the reductionism of binary algorithms, we are not using these technologies to their full potential. Machines are there to help people, to assist them, not to substitute them. Technology is not what we seek but how we seek!

AI in Cybersecurity Today

I stated earlier that a more accurate definition for what we call "artificial intelligence" would be "intelligent assistance." Computers are very good at binary information—limited and targeted, but infinitely scalable. They are invaluable in memorizing the facts in oncology reports or analyzing cybersecurity attack patterns. But it is difficult if not impossible for computers to do more human things, like reading body language or understanding sarcasm—simply because they lack human context and because they don't exist. As some philosophers would say, they have quanta but not qualia. Machines are binary, humans are multinary!

The type of intelligence and machine learning that is available in today's technology make it particularly well suited to provide intelligence assistance to the typical cybersecurity challenges that organizations face, as long we keep awesome humans in the loop.

For example, software can analyze patterns and learn from experience to recognize new and similar patterns. AI machines have the capacity to scale and examine infinitely more patterns than humans. They can tell us they've seen a certain pattern before and recommend methods to respond to both existing and new attack vectors.

They bring a computational value that humans can't match. But that doesn't mean they should replace humans. One of the paradoxes of computers is that when you have too much information that is unstructured, ambiguous and constantly changing, the machine will not be able to make a decision (the so-called intractable problem). To use them effectively, particularly in cybersecurity, we have to be careful of assuming that they are always right, and we should employ a healthy skepticism of what they are showing us.

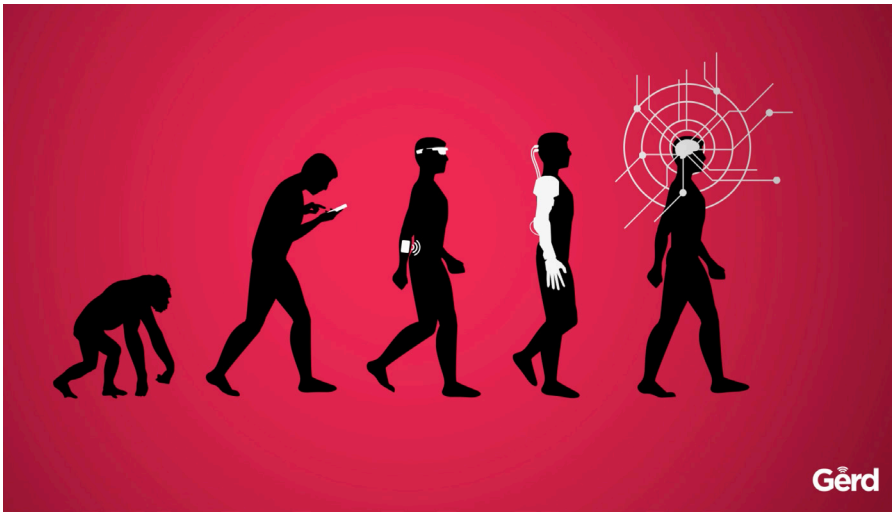

AI's Vast Potential

Having offered those provisos in terms of where AI is today, it is hard to overstate the dramatic impact that AI technology will have on the world in the coming years, including the world of cybersecurity. In my brief film called "We Need to Talk About AI,"[4] I note that AI will be the greatest wealth generator of our time, enabling us to cure diseases, enable smart cities, redefine poverty and tackle our foremost environment challenges.

As business leaders and as citizens, we all have a rare opportunity to begin building an ethical framework for using AI technology to help humanity achieve these formidable achievements. Cybersecurity should be an important area of innovation and investment. It is also an area where digital ethics is now moving centerstage. As we look ahead and begin the path towards a world in which machines increasingly act like humans, we must avoid the trap of surrendering to the algorithm and having humans behave more like machines.

1. Mike Perkowski, "Don't Get Caught in the Hype: Cut Out the BS in AI and ML," SecurityRoundtable.org, last accessed April 3, 2023.
2. http://androrithms.com/
3. https://www.techvshuman.com/
4. Gerd Leonhard, "We need to talk about AI - a film by Futurist Gerd Leonhard: thoughts on artificial intelligence," YouTube video, May 21, 2018.

Gerd Leonhard is a leading global futurist keynote speaker and author of five books

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

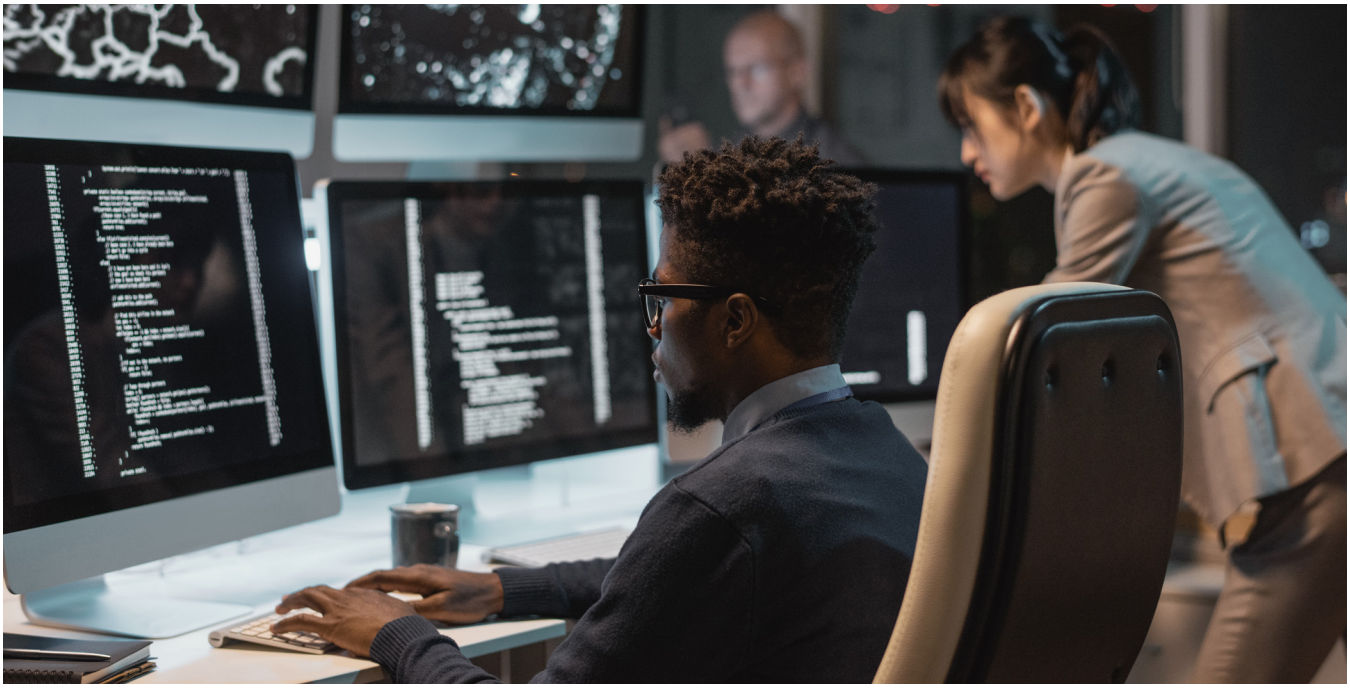

Cybersecurity Automation: Levelling the Playing Field

By Leonard Kleinman

Many things challenge how we practise cybersecurity these days. Digital transformation has brought significant adoption of new technology and business models, including cloud solutions, e-commerce platforms, smart devices, and a significantly more distributed workforce. These, in turn, have brought with them an increase in new threats, risks, and cybercrime.

As organisations emerge post-pandemic, many of the risks and uncertainties manifested during that period will persist, including the hybrid workforce, supply chain risk, and other cybersecurity challenges.

Let's look at some of these cybersecurity challenges and how automation can level the playing field.

Problem: Not Enough Cybersecurity Talent

A major contributor to the growing spate of cyberattacks is the lack of skilled cybersecurity personnel. The overall global numbers of experienced cybersecurity practitioners are low compared to the need for such practitioners to handle the cyberthreats that manifest across all industry sectors. While demand for practitioners continues to escalate, the growth in actual numbers is low, leading to the increasing deficit between demand and supply.

This contrasts significantly with the global cybersecurity market, which is expected to expand at a compound rate with more demand for solutions and products. The increasing number of cyberattacks, digital transformation changes, and talent shortages are contributing to this growth, and organisations are expected to acquire/deploy more advanced security solutions to detect, mitigate, and reduce the risk of cyberattacks.

Automation, AI, and Vocation

Automation systems are everywhere—from the simple thermostats in our homes to hospital ventilators—and while automation and AI are not the same things, much has been integrated from AI and machine learning (ML) into security systems, enabling them to learn, sense, and stop cybersecurity threats automatically. So instead of just alerting us to a threat, an automated system would be able to act towards neutralising it.

At its core, automation has a single purpose: to let machines perform repetitive, time-consuming, monotonous tasks. This, in turn, frees up our scarce human talent to focus on more important things or simply things that require the human touch. The result is a more efficient, cost-effective, and productive cyber workforce. »

Even threat actors are themselves using automation to facilitate their attacks.[1] The MyDoom worm,[2] one of the fastest-spreading pieces of malware on the internet, uses automation to propagate and is estimated to have caused around $38 billion in damage. It is still spreading, but the surprising part is MyDoom is not new. Released in 2004, it can still be seen trolling the internet.

A persistent fear in cybersecurity is that automation is here to replace humans. While somewhat justified, the reality is that automation is here to augment humans in executing security operations and, in some cases, help organisations supplement and address the growing talent gap. As advanced as it may be perceived, automation will always be reliant on humans, completely configurable and under the supervision of the security team. If anything, automation and AI are bringing forth new cybersecurity roles such as Algorithm Bias Auditor or Machine Risk Officer.

The Benefits of Automation

Automation can do many things, from detecting potential threats to containing and resolving threats. These actions take seconds and are largely independent of human intervention. Provided via security orchestration, automation, and response (SOAR), automation gives SOCs a significant boost in execution, significantly improving productivity and response. The Cost of a Data Breach 2022 Report highlights the role of automation in halving the cost of a data breach and reducing the time to identify and contain by 77 days.[3]

Orchestration provides the ability to activate the many tools in your operational environment, seamlessly connecting them via playbooks to undertake specific actions. This allows for a consistent, repeatable response process together with all the necessary information for your cyber practitioner, all in one place.

Additional efficiencies are derived from the AI/ML engine within SOAR, which can learn attributes from alerts and use that knowledge to prevent future attacks. Every alert and event handled are learned from for future purposes. Automation plays a significant role in terms of enabling an agile, proactive cybersecurity capability.

Most importantly, automation provides a better quality of life to your cybersecurity team, reducing alert fatigue and frustration and giving them back precious time. In the age of the great resignation, retention has become a significant issue.[4] Retaining staff allows you to increase your ROI on people—acknowledging the significant investment organisations make through recruitment, ongoing training, and tacit knowledge learned on the job.

Automation helps organisations address the talent challenge. It also enables a greater ROI on your current tools and technology, bringing them into play as part of the orchestration process.

Where to Start?

A prerequisite for automation begins with gathering and correlating data. Any good automation system requires good data to work efficiently and effectively. The more data sources, the better the quality of operations.

Aim to gather data from all aspects of your business environment, such as endpoint, network, and cloud. The AI/ML system within the automation platform makes analysing and correlating all this data easier. These two components are what make cybersecurity automation possible.

Next, analyse your current standard operating procedures (SOPs), looking for regularly recurring activities/processes—ones that reduce workload and the risk of an overlooked alert. Look for tasks that do not deviate or vary in an unpredictable manner. These are prime candidates for automation.

Now, identify the tools that need to be orchestrated within those processes, along with the required APIs (or create them) to enable the integrations.

Finally, create your playbook. This gives you control over the process, providing you with the ability to consistently replicate and improve the process over time. Include any specific actions you require, the tool/s to perform, and any other additional tasks, e.g., block, notify, contain, etc.

Don't Drop the Ball on Automation

Cybersecurity is essential for any business in a digitally transformed world, protecting company data, its people, and its customers. However, just the implementation of cybersecurity will not be enough as our adversaries continue to innovate and get craftier in their approach.

As organisations continue to pursue digital transformation initiatives coupled with technology advances, the automation of cybersecurity is not just recommended—it is mandatory in levelling the playing field.

1. Danny Palmer, "Cybersecurity warning: 10 ways hackers are using automation to boost their attack," ZDNET, March 25, 2020.
2. Brad Duncan, "MyDoom Still Active in 2019," Unit 42, July 26, 2019.
3. *Cost of a Data Breach 2022 Report*, IBM Security, July 2022.
4. Paula Morgan, "Top Five Tips For Retaining Employees During The Great Resignation," Forbes, August 4, 2022.

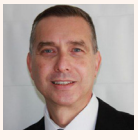

Leonard Kleinman is the field chief technology officer and evangelist of Cortex at Palo Alto Networks

Enabling Business and Security Together: The Value of SASE

By Josh Dye

Security is often thought of as an insurance policy—something you get to manage risk. What's less talked about is the ways our security choices can actually enable business and expand what's possible within organizations. Security can be a tool for resiliency and innovation.

Learning to maximize the value security can deliver typically comes from a couple of key sources: The first of which is from looking inside your own organization and partnering with other leaders to understand both the business needs and workflows.

Interviewing leaders across business units to understand their needs is a powerful first step. The knowledge you gain enables more informed decisions and helps you implement security technologies that support, and even accelerate, success. The other source of learning comes from educating yourself on challenges outside of your organization, specifically focusing on entities that have been victims of attacks.

As security professionals, we must take the hard lessons others have faced and put them to use for ourselves. Looking at how others' incidents are handled gives you the opportunity to reflect—continually examining and re-examining your security strategies—informed by history rather than theory.

Beyond security "insurance" to business enablement

Business leaders who take the old "security-as-insurance-policy" approach mostly view their cybersecurity as an expense line item and a cost of doing business. The reality, though, is that security investments can also be an important cost avoidance method for organizations, protecting the value of the business.

It's hard to provide a true dollar accounting for the value of a specific cybersecurity investment in advance—if I buy X tool it will save me Y amount of money. However, leaders need to think about their security in ways that go **beyond the purchase cost**.

This involves not only looking at what an organization stands to lose if they are breached by being unsecured, but also considering what soft costs the organization saves—associated with everything from deployment to ongoing management or usage—if they make sound choices. Some modern cybersecurity solutions enable organizations in ways that go beyond business continuity. One of those is SASE.

SASE isn't just about security

Fundamentally, SASE provides two core things: connectivity enablement and security. On the connectivity enablement side, the deployment of SASE provides ways of working remotely in a highly resilient manner. The security piece is that SASE provides integrated security deeply embedded into the same technologies that provide the connectivity and enablement capabilities. So SASE isn't just about giving business users security tools, it's about creating new opportunities to work securely.

At Jefferies,[1] for example, our traders had been using a trading turret system in our offices, which for the longest time had been a very physical thing. In the past, a turret system was always a big, bulky device that would sit on a desk on a trader workstation. However, by using SASE in conjunction with technology from the turret vendor we were able to devise a software-based version that our traders could use remotely, outside of the office.

Today we have a touch tablet that sits in our trader's home offices next to their main laptop connected via a SASE architecture using Prisma Access[2] from Palo Alto Networks. This new approach has enabled our IT team to support traders working remotely, without requiring massive amounts of hardware to be installed and deployed to every single user. It's been a great example of how security—and SASE specifically—can enable the business to do something that was never possible before.

Beyond supporting traders working at home, SASE also helps to enable our branch offices. Previously, our company had a hub and spoke model for network connectivity where branch offices were backhauling traffic out to the closest regional data center. As a result, critical SaaS applications such as email, purely on Office 365, and everything else were going to the data center to access the internet.

If there was ever an issue with the data center or somewhere in between the branch and the data center, that could have meant real trouble. Not to »

mention the fact that application performance from the branch was often degraded. With a SASE model, our branches access applications in the cloud directly, while still being fully secured.

Not only that, but by moving our branches to SASE we have been able to improve internet resiliency and remove the dependencies that created a central choke point, which is our data centers. In doing so, we also decreased latency for SaaS-based applications and internet connectivity. It's been a great win for us.

How to pitch SASE to the business

Getting your organization to buy into SASE and invest isn't about pitching the term or detailing the technology behind SASE. On the contrary, that's a path that will lead to blank stares, more likely than not. It's also not about doom and gloom, telling leadership that something must be done to protect the organization. That approach may work to a certain extent, but it really plays into the old insurance policy mindset.

Instead, educate your leadership on the things that are going to benefit them: explain it in a way that doesn't even talk about security. Focus your conversations with leaders and end users on the things that they really care about, which is the **ability to get their jobs done**. Few people care about security the way that security engineers do—so we need to explain it to them in ways that highlight how it will help them be more productive.

It goes without question that SASE has helped our business to be more resilient, and more successful. In times of natural disasters like a hurricane or human health challenges like the pandemic, SASE technology has helped to enable our security, remote connectivity, and business continuity. And that's a lot of value to derive from making the right security decisions.

1. https://www.jefferies.com/
2. https://www.paloaltonetworks.com/sase/access

Josh Dye is senior vice president of Information Security at Jefferies Group

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

3 Fundamentals to Truly Secure Remote Workers

By Christian Aboujaoude

In the pre-pandemic days, security solutions could be more basic. Securing the perimeter could be likened to locking the door of your house. But with remote workers taking devices off premises and sometimes using their own, securing the workplace requires a new approach. Sophisticated threats come in from every angle, and preparing a complete defense is vital.

We are in an environment of constant change and unexpected events. Just when many people started thinking we might be in a post-pandemic world, cases started rising again, and the need to apply proper controls, governance, education, and tools for remote workers has once more become top of mind for many cybersecurity leaders.

For CISOs and their teams, the challenge is to build a culture that facilitates the ability to adapt to change on an ongoing, continuous basis. This requires a new mindset in securing all users—remote users in particular. It also means evolving your approach so that cybersecurity is no longer viewed by business management as a cost center, but rather as a means of **competitive differentiation and innovation** for the organization.

In my view, there are three critical aspects to changing the culture and mindset to adapt to current and future cybersecurity challenges, particularly as remote work becomes more deeply ingrained as a business requirement.

Education

Develop a deep understanding of every aspect of your organization and spend a lot of time and attention »

on education—for everyone, whether they are on your security teams, in your executive suite, front-line workers on-premises, remote workers, or anywhere else in your ecosystem.

Technology

Even in some larger organizations, basic technologies such as multi-factor authentication or secure VPN are not given the priority necessary to allow remote workers to operate in a more controlled environment. It is important to have the basics under control before adding innovations such as Zero Trust.

Procedures and practices

It is vital to maintain a philosophy of ongoing education along with continuous evaluation of the technology your organization is using or, in some cases, not using. From a procedural perspective, you have to understand everything in your environment. Once you understand it, you can assess and address its impact on your current risk and overall risk profile.

1. Leveraging education to secure remote workers

The reason education tops my list is that **over 80% of cybersecurity events** are people-related. Everyone needs to truly understand what cybersecurity is—and know that it's not just a password or two-factor authentication. Cybersecurity is an approach, a mechanism. It's how you go about conducting work. Achieving a strong cybersecurity posture takes cultural change, behavioral change, and constant learning.

When users were largely on-premises, most organizations could compensate for potentially dangerous behavior by having multiple controls to help protect them. However, when those same people go remote, there's a bit of a loss of control and governance. There are technologies to help cover user behavior, but it is better when the behavior doesn't exist in the first place.

This means that we have to educate folks on cyber hygiene, making sure they understand that the steps they take at work may not be the steps they take when they are working remotely or from home. This is especially critical in this very open-ended environment, when a user's device may be used by other people in the home.

2. Leveraging technology to secure remote workers

Strong foundations are also important from a technology perspective. You must make sure you have controls, processes, and governance for multi-factor authentication and secure VPN. It's those things that pave the way for Zero Trust.

My best advice is to approach everything from the bottom up, understanding not just your inventory but every single behavior that takes place from a public-facing standpoint. This is especially important for remote workers. A good place to start is by asking yourself and your team key questions:

- Do we know what our environment actually contains?
- Are we aware of all the devices and services running in our environment?
- Do we have an inventory of all of our IoT devices?
- Do we understand the needs and potential risks of all of our users?
- Do we know the needs of each application and user based on key criteria such as performance, availability, resilience, data usage, and, of course, security?

Fundamentally, you need technology tools that are able to exist on your network and can identify all connected devices. I'm talking about tools that are able to actually interrogate the network, understand packets, capture specific metadata for insight into the imprint of each device, and how it lives on the network.

3. Leveraging procedures and practices to secure remote workers

If you haven't figured it out by now, I'm a huge stickler for inventory. From a process standpoint, you have to understand what your inventory is, what it means, why it matters, and what its impact is on your business as well as your security posture.

So, from a procedure standpoint, you need something in place that is able to identify what you have in your environment. Then you have to relate and correlate that information to any situation, to the point where you can say about any device: "This device is connected to this application that lives here and does that."

From there you can build a configuration management database (CMDB) approach to really understand your environment and have processes in place to integrate with your ITSM tool so you can execute the specific actions you need to take.

Maintaining ongoing processes also relates back to my first point: education. CISOs need to ensure training and education are continuing when people work from home or remote locations, and they need to have tests, controls, processes, and governance to continuously identify and correct non-malicious but potentially dangerous behavior. Quick hit trainings without repetition rarely are effective.

My advice for CISOs and other cyber leaders

If I could leave CISOs and other cybersecurity leaders with a key takeaway from this article, it would be this: **Every CISO should figure out how to balance the business operations of their organization with a security mindset that is not destructive to the business** but is, in fact, built into the fabric of the business. In order to do that, I urge all security professionals »

to take the time to understand as much as they can about the business in which they work.

Note the emphasis on *the business*, not cybersecurity. Most security professionals know security exceptionally well. But if they don't have an equally exceptional understanding of their business or organizational needs, they are potentially setting themselves and their organizations up for failure.

Whether you are the CISO or anyone on the security team, you need to be able to go to the people in any department and have detailed conversations with them related to their protection and their business needs. It may start with something simple: "We saw that you have these devices. They are not in compliance with our security posture, and we need to take this action or we will be forced to put it offline."

Of course, the immediate reaction will be: "You can't do that!" And the response is: "Yes, we know. That's why we have to fix the problem." A solution-focused and service-focused mindset is key.

The opportunity ahead

Remote work is here to stay. To make it successful, you have to make it secure. Cybersecurity leaders and their teams have an opportunity to make huge contributions to their organizations over the next few years by developing cyber-aware cultures that are both agile and responsive to the changing needs of their organizations.

By focusing on the fundamentals, CISOs can prepare themselves, their teams, and their organizations to be ready for whatever comes next. As we've learned all too well over the past few years, **the only constant in cybersecurity is change**. Be ready.

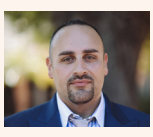

Christian Aboujaoude is chief technology officer at Keck Medicine of USC

Reprinted from Security Roundtable by Palo Alto Networks (SecurityRoundtable.org).

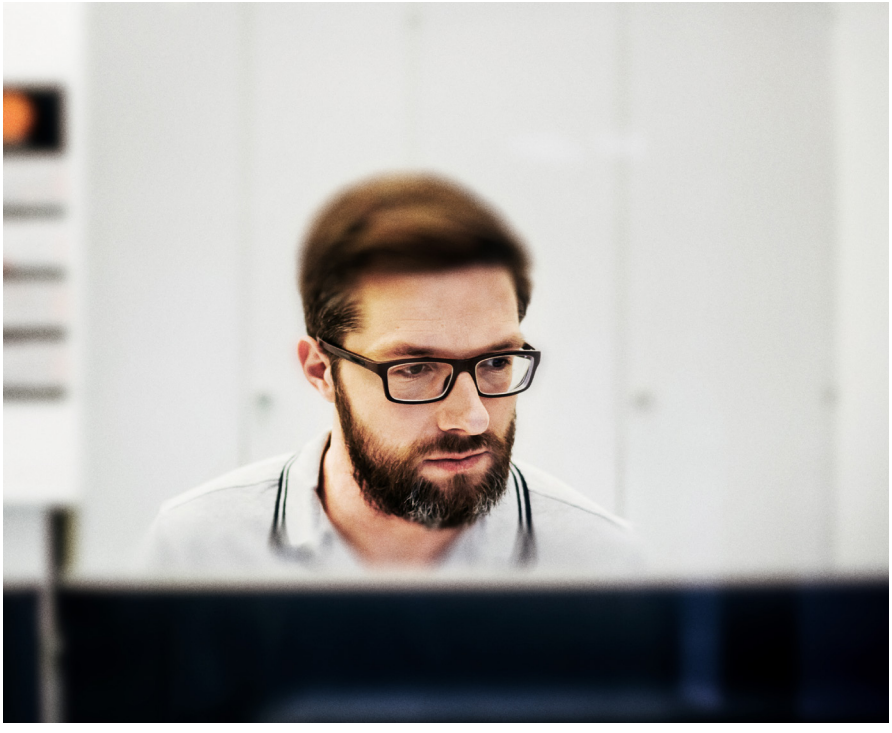

Consolidation: The Secret to Supercharging Your SOC

By Niall Browne

Most CISOs have a clearly defined goal: to ensure smooth, reliable, and gap-free security operations.

But for many security leaders today, that goal feels increasingly out of reach as they tack on more and more point products to address every security need. In fact, large organizations today manage 31.5 security solutions on average, taking up hundreds of their precious hours every month.[1]

So how can security teams pivot away from a low-value procurement and maintenance role and focus on more mission-critical tasks such as keeping threats at bay? That's where a cybersecurity consolidation comes in.[2]

What Is Cybersecurity Consolidation?

Securing today's large IT environments is becoming increasingly complex. The latest attacks target vulnerabilities in different networks, clouds, and endpoints, bypassing the traditional cyber defenses used by organizations.

In response, security teams deploy multiple point products to cover every issue. A solution for network firewall, another for IoT devices, another for SD-WAN—the list goes on. The problem is **the point product approach creates gaps and complexities in an organization's security posture**. Products from different vendors work in silos, making prevention, detection, and response difficult.

Watch Palo Alto Networks founder Nir Zuk on cyber consolidation:[3]

Cybersecurity consolidation protects all attack surfaces with a single security platform (and vendor) for an organization's entire IT environment. The architecture covers every defense vector— »

Streamlines Operations

Building security policy with fewer vendors is **3 or 4 times easier** than upgrading a security policy for each different one

— Transportation and Logistics

Automates Response

Without a platform, we'd probably have to **increase the SOC team by potentially 50-75%**

— Oil and Gas

Simplifies Procurement

The entire procurement process actually has to go through a bunch of different steps…we saved about **40-50% of the procurement time**

— Insurance

from threat prevention to intelligence and response.

Let's say your organization uses a separate tool for intrusion detection, data loss prevention, DNS security, and remote user access. A single, unified platform that provides all of these services simplifies your security architecture, reduces complexity, and improves your overall risk posture.

How Does Consolidation Enhance Security Operations?

Security tools in a point product ecosystem can do a great job in their dedicated functions—and together, they might form a comprehensive cyber defense. But when those tools use different datasets, analyses, and UIs, organizations can't see the bigger picture.

A consolidated security platform covers the entire gamut of security requirements with the added benefit of shared intelligence. Information gets shared seamlessly across the platform and gives organizations a united front against threats.

Key benefits of a consolidated platform:

- Reduced time to respond to incidents
- Reduced time to update security policies
- Reduced time to complete threat analysis
- Decreased number of security alerts
- Reduced overall risk exposure

For example, let's say your security team wants to analyze your organization's risk exposure in-depth. Security products that report metrics differently mean more time manually piecing together disparate information—and a chance of employee error.

A consolidated security platform shares that intelligence across all critical areas. Teams can analyze reports on a single dashboard, standardize security policies, supercharge risk posture, and, most importantly, rapidly detect and block attacks across the entire platform.

What Palo Alto Networks Customers Think

We chatted with several organizations that use our consolidated platform solutions. Their responses are in the figure above.

Security leaders also reported several other strategic benefits:

- Standardize and unify security policies.
- Facilitate and speed up reporting to management.
- Decrease the likelihood of human errors.
- Alleviate the impact of staffing deficits.

According to a customer, "Everyone is on the same page. There are no longer different skill sets, different platforms, and different versions of those platforms. People have the same capabilities and are able to back each other up." … "We can generate reports with a summation of all the changes that have happened, either to management or to security operations, and they know what is happening."

A Platform Approach: Better Security with Less Complexity

The solution sprawl only complicates your operations and exposes them to threats. A platform approach keeps your entire portfolio under one roof and supercharges your security posture.

1. "What's Next in Cyber: A Global Executive Pulse Check," Palo Alto Networks, December 2022.
2. https://www.paloaltonetworks.com/paloaltonetworks-portfolio.
3. Palo Alto Networks, "Security Consolidation Q&A with Nir Zuk," YouTube video, February 8, 2023.

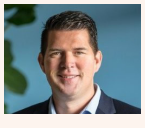

Niall Browne is a chief information security officer at Palo Alto Networks

What Executives Should Know About Shift-Left Security

By Zachary Malone

What Does Shift-Left Security (Really) Mean?

The term "shift left" is a reference to the Software Development Lifecycle[1] (SDLC) that describes the phases of the process developers follow to create an application. Often, this lifecycle is depicted as a horizontal timeline with the conceptual and coding phases "starting" the cycle on the left side, so to move any process earlier in the cycle is to shift it left. "Shift-left security" is the concept that security measures, focus areas, and implications should occur further to the left—or earlier—in the lifecycle than the typical phases that used to be entry points for security testing and protections.

How Did the Term Shift-Left Security Originate?

Shift-left security spawned from a broader area of focus known as shift-left testing. The term was first coined by Larry Smith in 2001.[2] Since then, the concept of shift-left security has continued to gain traction as organizations' reliance on the cloud has grown and as higher-profile cyberattacks increasingly target development tools and pipelines for apps that are cloud-delivered and/or SaaS.

Why Is Shift-Left Security Important in Cybersecurity?

Simply stated, while the advancements of cloud services for developer and product teams provide incredible speed and breadth in delivering applications, it has also led to some extreme challenges in maintaining regulation and control. Security needs to keep up with the fast-paced growth and agility of development cycles and be flexible enough to support a broad array of cloud-delivered solutions.[3] The only common denominator in these new development workflows is the code that underlies everything from application to infrastructure is open and manipulatable to the development teams. As such, bringing security all the way "left" to the coding phase wraps security around the source of what malicious actors attempt to attack, leading to the greatest reduction in risk of exploits possible.

What Is the Spin Around This Shift-Left Security Buzzword?

Like many cybersecurity buzzwords, many vendors treat shift-left security as "the only thing you need to be secure," as if it were a panacea to security issues when in reality, this breaks the idea of Zero Trust as you would be implicitly trusting the developer/s and their coding abilities. Also, there is a distinct lack of consistent understanding and standard practice for how application development should work in a modern DevOps department—such as code supply chain (open source packages and drift) or integration tools (Git, CI/CD, etc.). This creates risks. For example, if an organization believes, "Our data storage is freely open to everyone on the internet, but that's not an issue because all the data is stored in an encrypted format," this belief allows attackers to simply make a copy of the data and then work to either brute force the decryption, or look for the keys in whatever storage place they happen to be.

What Executives Should Consider When Adopting Shift-Left Security

Shifting security left in your SDLC program is a priority that executives should be giving their focus to. The pervasive reach given to development teams to not only create business-critical applications via code but also to handle every step, from coding the application to its compilation, testing, and infrastructure needs with additional code, is an extraordinary amount of control and influence for a department that is singularly focused. Extending security into all the workflows »

that development teams are moving into is the core ideology of shift-left security. However, it would be exceptionally risky to abandon or discredit the security programs that remain in the later or "right-side" stages of the lifecycle. Security needs to be wrapped around the entire lifecycle, from building the code to staging the surrounding deployment to, ultimately, the application and environment handling it.

Here are some questions to ask your team for a successful shift-left security adoption:

- How can we envelop all of the phases of our SDLC into our security program without creating a massive overhead of new tools to learn for each step covered?
- How do we enable our development team to correct simple security mistakes without delaying or blocking their ability to release critical applications and updates?
- We must integrate into the tools and workflows that our development uses to code, aggregate, test, and deploy. How do we accomplish this while still meeting the needs listed above?
- Suppose something does happen to be deployed insecurely. How do we send the request for a fix back into the workflow that our developers utilize with actual coding changes included automatically?
- Are there any platforms that can handle our need to shift left, protect our runtime environment, and feed into our security operations, governance, and compliance; infrastructure architects' workflows to provide visibility, protection, and auditing layers for our entire application landscape?

1. https://en.wikipedia.org/wiki/Systems_development_life_cycle
2. Larry Smith, "Shift-Left Testing," Dr. Dobb's, September 1, 2001.
3. Matt Chiodi, "Cloud Supply Chain Security," Palo Alto Networks, last accessed April 4, 2023.

Zachary Malone is a systems engineering manager, SE Academy at Palo Alto Networks

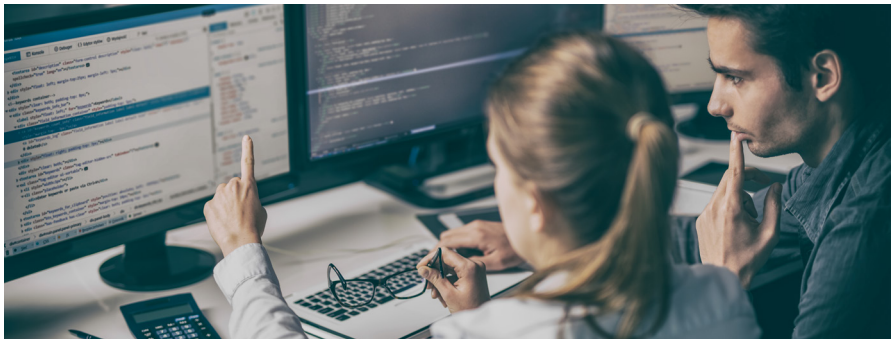

Improving Cybersecurity Outcomes

By Sean Duca

In recent years, organizations of all sizes have been collecting increasing volumes of traffic and application telemetry data from different devices, logs and services. Much of it is leveraged to inform operational and strategic decisions. However, this same data also has the potential to significantly strengthen an organization's security posture—but only if it's processed and used effectively.

To strengthen cybersecurity, there is plenty of data that organizations can and do collect to understand what's happening inside their environments. It comes from log files, system events, network traffic, applications, threat detection systems, intelligence feeds and myriad other sources. However, the sheer volume of this data can pose a significant challenge as organizations look to extract value from what they're gathering to inform security policy, threat detection and risk mitigation.

If your systems can't process the data you collect, they won't be able to make sense out of it and correlate what's going on. In that case, you're really just sitting on some dead logs. Adding to this challenge is the fact that collected data is often siloed in ways that can keep a security professional from connecting the dots to identify potential issues. Analysts should not have to look at 25 different screens trying to make manual connections, which takes additional time and effort that distracts from the primary goal of actually identifying threats.

As an industry, cybersecurity created this world where there are so many different point solutions out there that organizations have been effectively forced into becoming plumbers, connecting all these different solutions together. I think it's time that we start to think about how we find a way that's more automated and integrated because a lot of the tools that people are using were never designed to interoperate and work together.

Extracting Greater Value from Data with Automation and Playbooks

Collecting the right data and extracting the highest value from it is not a single task or operation. Rather it's a journey that involves multiple components.

Technology. From a technology standpoint, have a look at what you've actually got. For starters, are the tools capable of actually identifying the modern threats? If they are not, then you've got a challenge there because you're likely not going to be collecting any logs and telemetry to actually make an informed decision.

Automation also plays a critical role in extracting more value from data. With the volume of data that is being collected, even if it's all the right data, individual humans simply cannot keep up. Automating the identification of higher value incidents from data that correlates and enriches simple log data and provides insight is a critical component.

People. Automation ties in directly with the people's perspective on getting »

the most value out of data. Many organizations have security operations centers (SOCs) staffed with IT professionals working eight-hour rolling shifts, clicking on refresh all the time and simply chasing the logs. That's not really going to help them find anything.

Adding further insult to injury, the first line of defense and analysis for data is typically a level-one analyst, who often will burn out within a year after the monotony of sifting through endless logs and deciding what needed to be escalated. Think about the logic: The least experienced and lowest paid person, is actually making a call to escalate an incident to a more senior person. It doesn't make sense, and it's time to change the model.

When automation is leveraged to handle the deluge of data, becoming the first line of the decision on what needs to be escalated, human talent can focus on the more intricate challenges like threat hunting. The easier a threat hunter's life—where we can start to link all the disparate data sources to help chase potential risks, rather than just having to sift through alerts and large logs—the better.

Process. Finally, process is the key to continuous improvement and always optimizing the value from data. We need to go back to the drawing board all the time and keep on refining the data and technology that's already in place. Organizations need to keep on creating playbooks to help aid automation. Anything that's a repeatable task, organizations should be automating as much as possible.

With all the sources of security data available to the modern enterprise, it can be overwhelming to figure out what to do. By first understanding what security data sources the organization has, streamlining processes with automation and playbooks, and tying things together with technology to create a unified view, it's possible to dramatically improve security outcomes.

Sean Duca is vice president and regional chief security officer for Palo Alto Networks

What Executives Should Know about Zero Trust

By Zachary Malone

Invented in 2010 by Forrester Research,[1] Zero Trust is a cybersecurity model enterprises can leverage to remove risky, implicitly trusted interactions between users, machines and data. The Zero Trust model provides a process for organizations to protect themselves from threats no matter what vector the threat originates from—whether from across the world or from Sandy down the hall. The three main principles to follow to realize the benefits of this model were:

- Ensure that all resources are accessed securely, regardless of location.
- Adopt a least-privileged strategy and strictly enforce access control.
- Inspect and log all traffic.

After 11 years, these ideas and principles have matured in the face of growing digital transformation, remote work, and bring-your-own-device proliferation. New principles have developed in light of the U.S. Federal Government mandating Zero Trust, codified in the NIST 800-207[2] with further details in the NCCoE's Zero Trust Architecture.[3] Those principles are:

- Shift from network segmentation to protecting resources such as assets, services, workflows, and network accounts.
- Make authentication and authorization (both subject/user and device) discrete functions performed on every session, using strong authentication.
- Ensure continuous monitoring.

Why Is This Important in Cybersecurity?

The move toward Zero Trust has been one of the more significant shifts »

in how business approaches security. Before adopting a Zero Trust mindset, most companies tried to manage security as a gated function. Once a transaction was validated in the gated area, it was innately trusted.

This approach presents a problem because threat vectors do not always originate outside that area. Also, the world at large continues to adopt digital transformation and hybrid workforces, nullifying the concept of resources only existing behind a gate. Zero Trust methods require validating each element of every interaction continually—no matter where they occur—including all users, machines, applications, and data. There is no area of implicit trust.

What Is the Spin Around This Buzzword?

Many vendors today productize Zero Trust, naming their products as "Zero Trust solutions" in and of themselves, rather than acknowledging that Zero Trust is a model and strategic framework, not a product solution. When looking at the cybersecurity market, you'll see vendors try to claim a supposed title is "THE Zero Trust player."

On closer inspection, however, those vendors typically only address a single principle of Zero Trust. For example, creating tunneling services between users and applications. This aligns with the second original principle: adopt a least-privileged strategy and strictly enforce access control. However, that same vendor might fail on the first principle: ensure that all resources are accessed securely, regardless of location. When they implicitly trust that the user is not a threat vector, they do not scan for malware or exploits inside the tunnel.

Others may cover only some of the aspects of the first original principle, like trying to claim identity and authorization checks are what make Zero Trust. Vendors may also suggest that only web-based traffic needs to be scanned. However, when only partial coverage of the model is implemented, companies risk creating an implicit trust that opens them up to vulnerabilities that would be otherwise covered in the remaining principles.

Our Advice: What Should Executives Consider When Adopting Zero Trust?

The first step is to reframe your thinking on how enterprises should be secured, moving from a gated approach to one that continuously validates all interactions. To help make that shift:

- Define the resources your company needs to protect, where they exist, and what interactions should be flowing around, into, and through them
- Remember users, applications, and infrastructure/devices must all be covered for every interaction they create.
- Understand that interactions consist of identity, access, device/workload, and transactions

Next, enact changes with a plan, beginning with your enterprise's most critical users, assets, and interactions. Those will be your crown jewels and things that may be related to finance or intellectual property. Then, over time, expand your purview to include all interactions. The plan should cover how the users, applications, and infrastructure go through each of the four parts of an interaction when requesting a resource.

The final step in this transformation is really a recurring event: maintaining and monitoring.

- Leverage continuous monitoring to account for everything happening versus intermittent checks.
- Look for ways to improve the current model as standards continue to evolve while covering more and more interactions.

Questions to Ask Your Team to Successfully Adopt Zero Trust

- What are our system-critical datasets, applications, and functionalities?
- How can we secure each of the four parts of every interaction to these resources, no matter who or what is requesting them?
- What is our plan to continuously monitor important events like logs to facilitate baselines and detect anomalous behavior?
- What is our strategy for selecting vendors that will assist us with our Zero Trust goals, and what more will we need to do that products cannot cover?
- What is the strategy for going from covering one resource to fully covering all resources, and what sort of scalability of products and people will we need to do this?

1. *No More Chewy Centers: Introducing The Zero Trust Model Of Information Security*, Forrester Research, September 14, 2010.
2. Rose et. al, "Zero Trust Architecture," National Institute of Standards and Technology, August 2020.
3. https://www.nccoe.nist.gov/projects/implementing-zero-trust-architecture

Zachary Malone is a systems engineering manager, SE Academy at Palo Alto Networks

PRISMA CLOUD | paloalto NETWORKS

The State of Cloud-Native Security 2023 Report

Executive Summary

Everything about cloud-native application development is ephemeral, from IP addresses to containers, which is why legacy security vendors are struggling. As cloud infrastructure evolves, each year delivers us to the brink of a new horizon—this year is no different. In this report, we explore the findings of a survey spanning seven countries and five sectors of industry to inquire about practices affecting cloud-native development and security outcomes. Of our 2,500+ survey respondents, greater than 50% represent enterprise-sized organizations. Approximately 50% of our population hold executive roles and 50% practitioner roles. From our findings, two themes emerge: the application lifecycle moves faster than we would have imagined 12 months ago, and complexity across a spectrum of challenges is the ultimate threat to securing our environments. Our goal in this report remains the same, that is, to provide you with data-based insights to inform your cloud-native security decisions in 2023 and beyond.

Cloud Expansion and Strategy

- Over half (53%) of cloud workloads are hosted on public clouds, an increase of 8% in the past year. Platform as a service (PaaS) and serverless were the dominant application execution environments. Both are expected to experience further growth, as 70+% of respondents reported an expected increase in usage of PaaS and serverless in the next 24 months.
- On average, organizations spent 24% of their total cost of ownership on application migration costs. The breakdown in methodology for application deployment was cloud native (37%), lift and shift (36%), and refactor or rebuild (27%). This is the first time cloud native is at the forefront of application development.
- More than 70% of organizations moved at least 30% of workloads into the cloud, while over a quarter of enterprises (26.6%) moved at least 50% of workloads. Overall, survey respondents reported a 13% increase in workload moved to the cloud compared to the previous year.

Application Velocity

- More than 75% of respondents from this year's survey are deploying new or updated code to production weekly. As fast as that is, 38% are committing new code daily, and 17% of respondents are deploying code multiple times a day. Add to that the ratio of 10 developers for every security professional, and the potential for challenges in scale and complexity is high.
- A third of enterprises reported operating with internal service-level objectives of less than a day of lead time for changes, and 38% expected service restorations within a day. Along with 68% of respondents reporting increased deployment frequency, 64% also reported increased lead time for changes.

Obstacles to Cloud Expansion

- Technical complexity ranked as the No. 1 challenge when moving to the cloud. What's more, the next four challenges all point to technical complexity. Complex environments require higher levels of talent and adaptiveness to changing technology, are more difficult to fully secure, more difficult to gain visibility across, and result in greater compliance challenges.
- While 81% percent of respondents say they understand their responsibility to deliver security across the development lifecycle, organizations have distributed responsibility for designing and implementing cloud security policies across teams, and 47% of respondents say that the majority of their workforce doesn't understand their security responsibilities.
- Upwards of 75% of survey respondents reported that the number of cloud security tools they use creates blind spots that affect their ability to prioritize risk and prevent threats.
- At 90%, the vast majority of respondents say their organization cannot detect, contain, and resolve threats in less than one hour.

Download the full report.

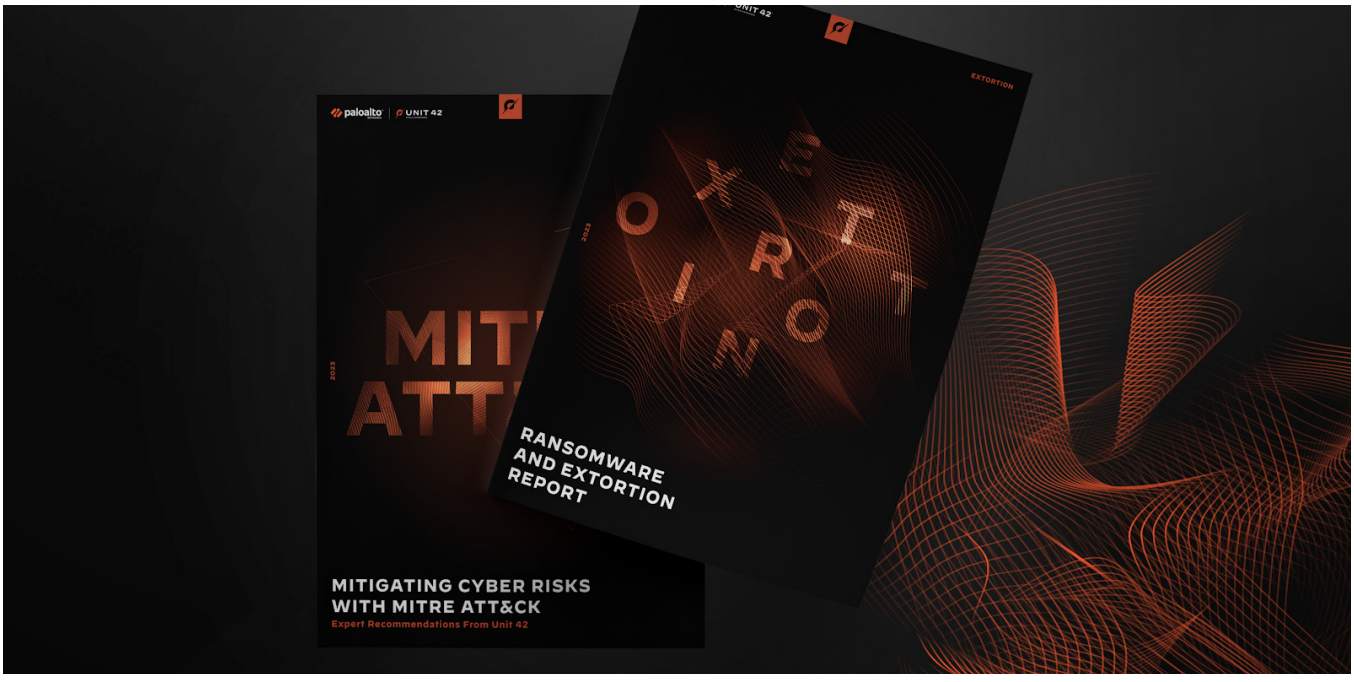

Key Takeaways from Unit 42: Ransomware and Extortion Report

Introduction

While much attention has been paid to ransomware in recent years, modern threat actors increasingly use additional extortion techniques to coerce targets into paying—or dispense with ransomware altogether and practice extortion on its own.

Organizations, in turn, need to evolve defenses to address the various methods threat actors use to apply pressure. Incident response plans today need to involve not only technical considerations but also safeguards for an organization's reputation and considerations for how to protect employees or customers who may become targets for some of extortionists' more aggressive tactics.

Our 2023 Unit 42 Ransomware Threat Report explores recent incident response cases, as well as our threat intelligence analysts' assessment of the larger threat landscape. It also offers predictions for how we believe threat actors will use ransomware and extortion tactics going forward.

Key Findings From the 2023 Unit 42 Ransomware and Extortion Threat Report

Multi-Extortion Tactics Continue to Rise

In Unit 42 ransomware cases, as of late 2022, threat actors engaged in data theft in about 70% of cases on average. Compare this to mid-2021, and we saw data theft in only about 40% of cases on average. Threat actors often threaten to leak stolen data on dark web leak sites, which are increasingly a key component of their efforts to extort organizations.

Harassment is another extortion tactic we see being used in more ransomware cases. Ransomware threat actor groups will target specific individuals in the organization, often in the C-suite, with threats and unwanted communications. By late 2022, harassment was a factor in about 20% of ransomware cases. Compare this to mid-2021, when harassment was a factor in less than 1% of Unit 42 ransomware cases. »

Extortion Gangs Are Opportunistic–But There Are Some Patterns in the Organizations They Attack

Based on our analysis of dark web leak sites, manufacturing was the most targeted industry in 2022, with 447 compromised organizations publicly exposed on leak sites. Unit 42 believes this is due to the prevalence of systems used by this industry running on out-of-date software that isn't regularly or easily updated or patched—not to mention the industry's low tolerance for downtime.

Organizations based in the United States were most severely affected, according to leak site data, accounting for 42% of the observed leaks in 2022.

Large, Multinational Organizations Can Be Lucrative Targets for Threat Actors

Attacks on the world's largest organizations represent a small but notable percentage of public extortion incidents. In 2022, 30 organizations on the Forbes Global 2000[1] list were publicly impacted by extortion attempts. Since 2019, at least 96 of these organizations have had confidential files publicly exposed to some degree as part of attempted extortion.

Predictions for What to Expect From Extortion in the Coming Year

Unit 42 experts have put together predictions for what we expect to see from extortion groups in the coming year. Our predictions for 2023 include:

- A large cloud ransomware compromise
- A rise in extortion related to insider threats
- A rise in politically motivated extortion attempts
- The use of ransomware and extortion to distract from attacks aimed to infect the supply chain or source code

Recommendations to Improve Ransomware and Extortion Preparedness

Prepare a Playbook for Multi-Extortion

During an active extortion incident, rapid support from your incident response partner and outside legal counsel is critical. From a mitigation perspective, having a comprehensive incident response plan with corresponding crisis communication protocols will greatly reduce uncertainty. It's important to know which stakeholders should be involved, and the process to make decisions promptly (e.g., whether or not to pay, or who is authorized to approve payments). »

1. "The Global 2000," Forbes, May 12, 2022.

The crisis communication plan should also cover what to do (or avoid doing) in the event that employees or clients are being harassed. Ransomware harassment awareness training should be delivered to an organization's staff to equip them with tools and processes to follow during an active harassment incident.

Organizations should conduct a post-mortem compromise assessment to validate that any remnant backdoors or other indicators of compromise (IoCs) (e.g., scheduled tasks or jobs) have been removed. This ensures that the threat actor cannot easily conduct a follow-up attack after an initial breach.

How Unit 42 Can Help

Read the full 2023 Unit 42 Ransomware and Extortion Threat Report for more ransomware and extortion insights, trends and recommendations for best practices.

Under Attack?

If you think you may be subject to an active ransomware or extortion attack or have an urgent matter, get in touch with the Unit 42 Incident Response team by calling Toll-Free: 866.486.4842 (866.4.UNIT42)

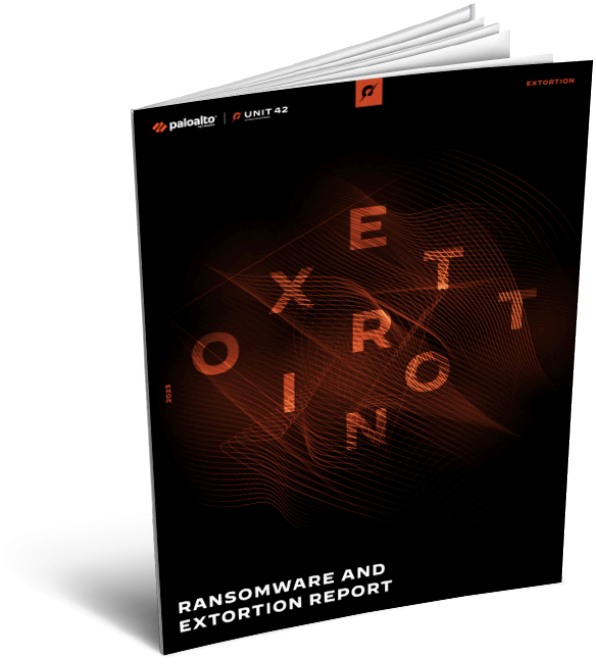

Scan the QR code to share with a colleague.

Don't Panic.

Did you know? Coloring can help adults sleep better at night, reduce stress levels, calm their nerves, decrease anxiety, and even help pull them out of depression.

Protecting your organization can be stressful.

We've got you covered.

UNIT 42
BY PALO ALTO NETWORKS

SECURITY ROUNDTABLE

BY PALO ALTO NETWORKS

A business magazine for executives on cybersecurity management, trends, and best practices.

Read more at **securityroundtable.org**

paloalto NETWORKS | UNIT 42 BY PALO ALTO NETWORKS

Harassment as an Extortion Tactic Is up 20X. Are You Prepared?

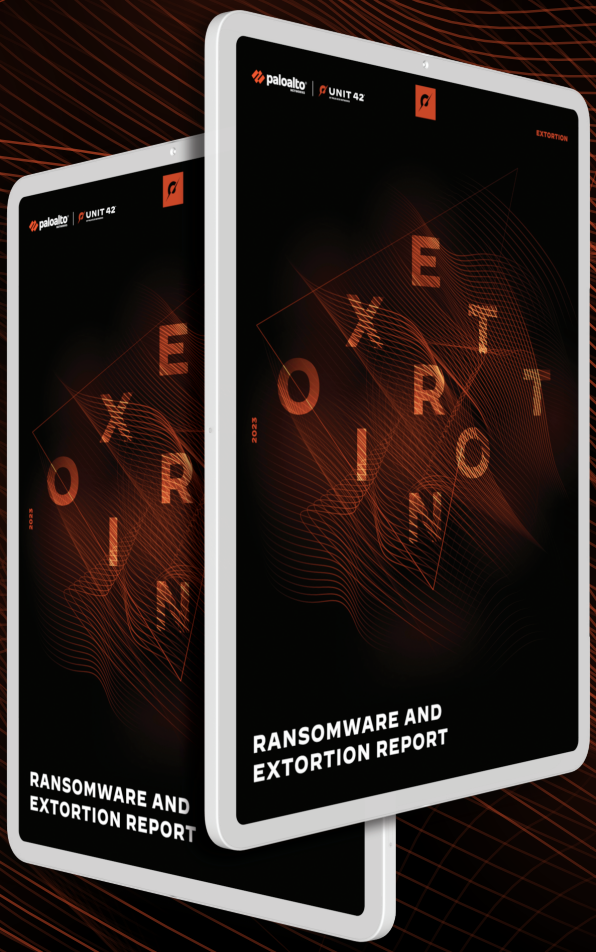

Read the report

Find out in the

2023 Ransomware and Extortion Report