SC MEDIA

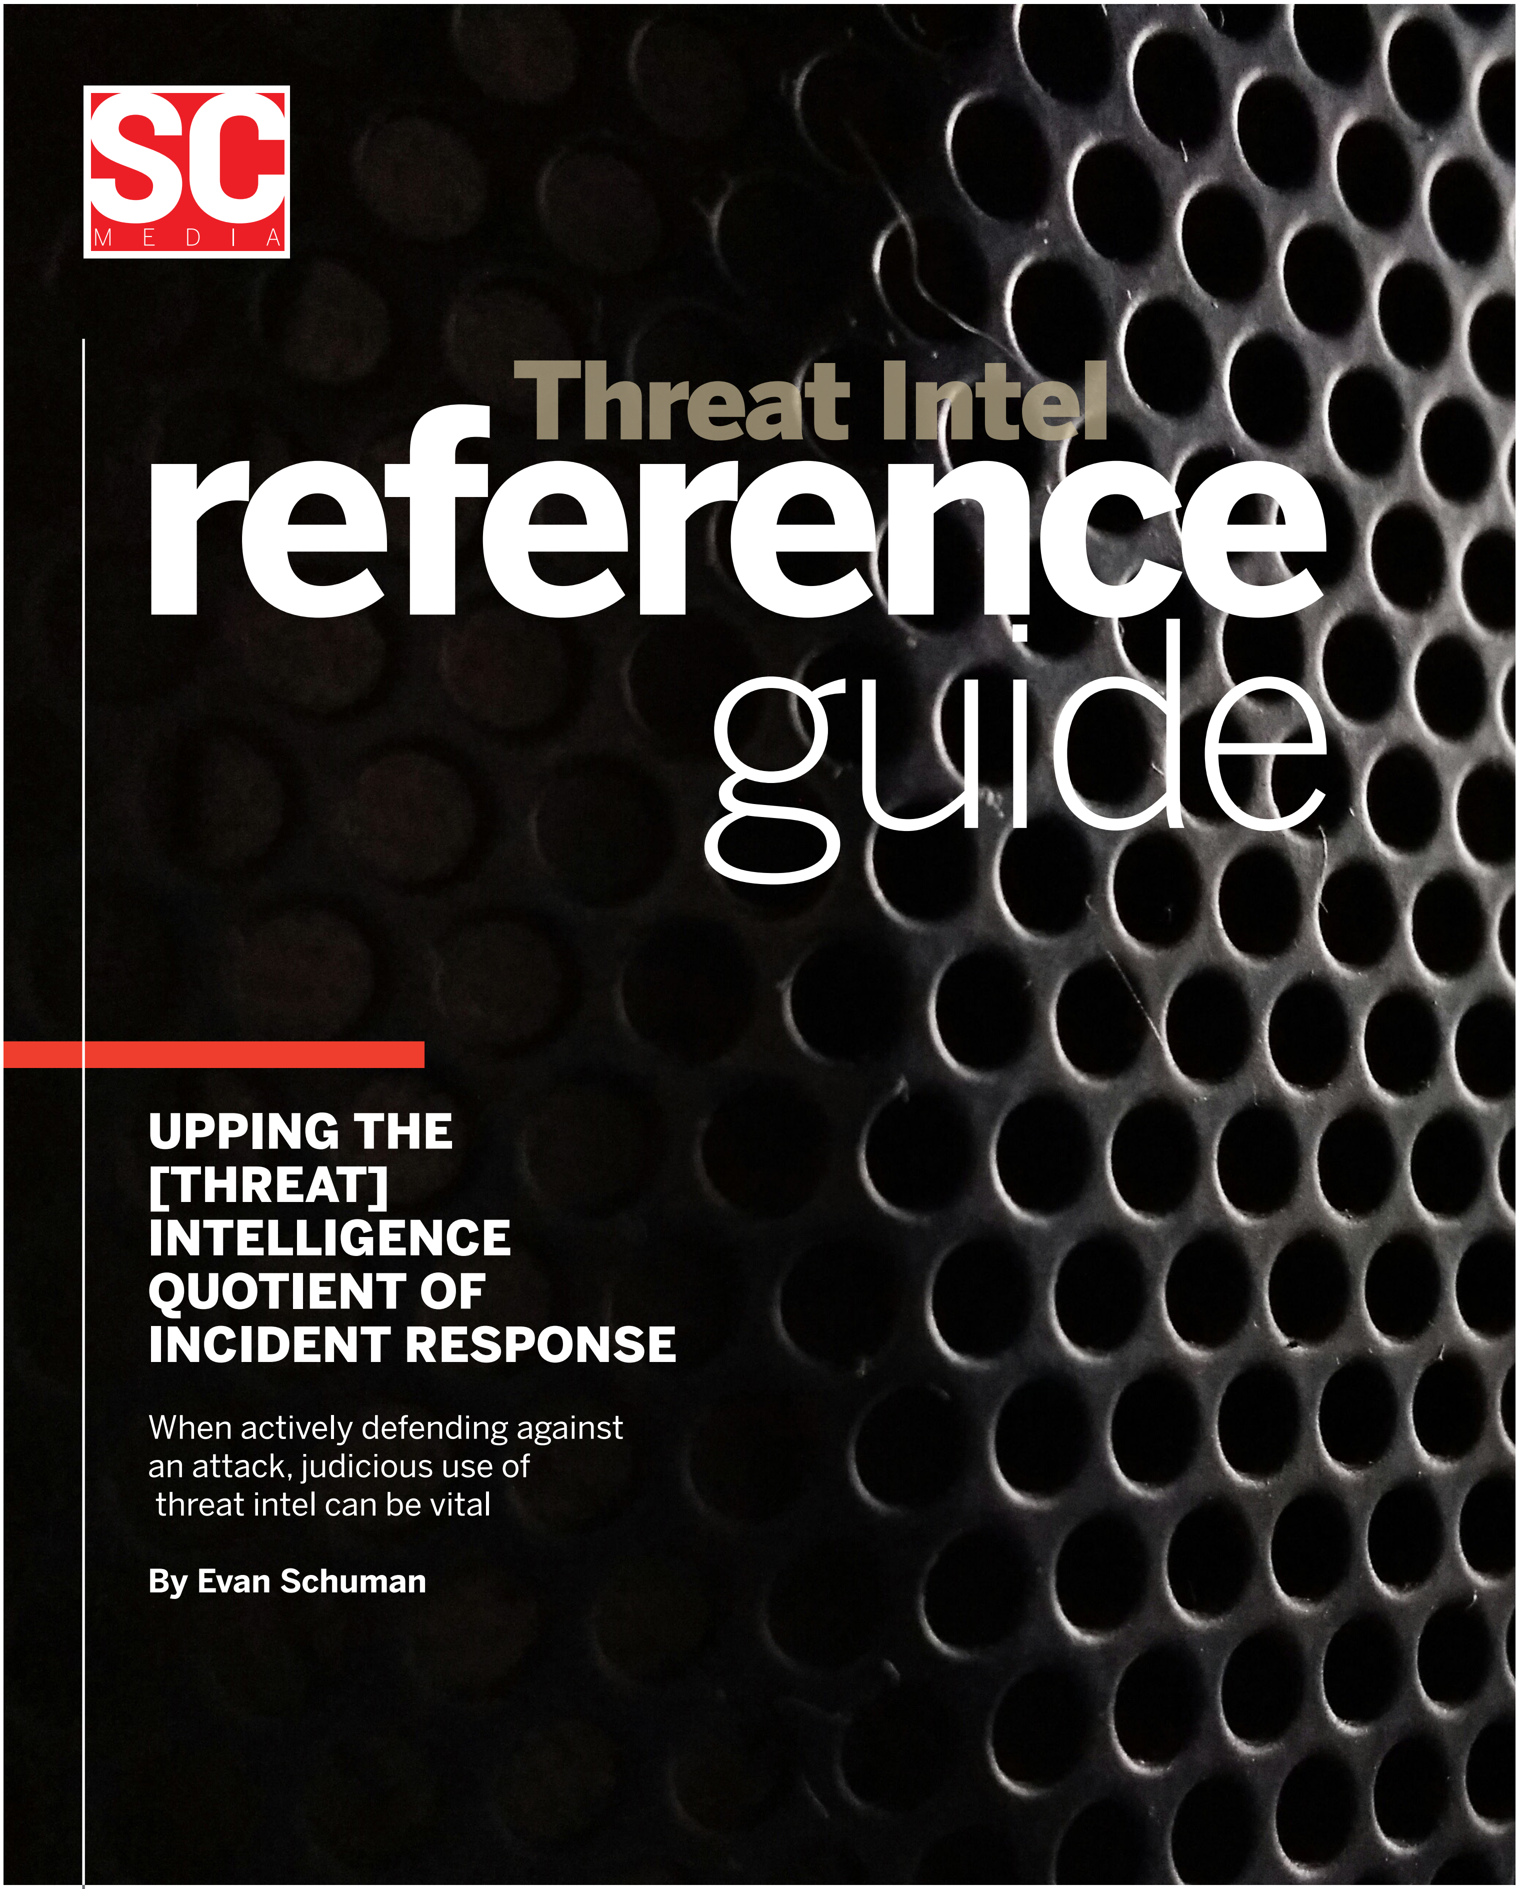

Threat Intel reference guide

UPPING THE [THREAT] INTELLIGENCE QUOTIENT OF INCIDENT RESPONSE

When actively defending against an attack, judicious use of threat intel can be vital

By Evan Schuman

BROUGHT TO YOU BY

ThreatConnect™

AI-ENHANCED THREAT INTEL POWERS UP INCIDENT RESPONSE

One of the most daunting security tasks is dealing with an active attack and trying to leverage an enterprise's collection of threat intelligence — including commercial feeds, open source feeds, internal logs, Information Sharing and Analysis Centers (ISAC) data, government alerts and media reports of other attack attempts — to fight off the attack.

The question CISOs ask themselves is: When should the attack be halted, especially by the most extreme measure of shutting down the network, and when should it be watched to learn more about the attacker? The most critical factor for determining the answer is often time.

Security analysts might have seconds to decide which actions to take when seeing evidence of an attack, but halting it a split-second before sensitive data is exfiltrated. Experienced attackers know this well and they will often game the system by acting innocuously for as long as possible before shifting into attack mode, whether for stealing data, planting malware to steal data later, or engaging in sabotage.

Making matters worse, the time urgency of fending off an active attack typically makes going back to review threat feeds pragmatically impossible. For some, that means artificial intelligence (AI) and machine learning (ML) must be used, to both analyze the attack and to then review an ocean of threat feeds to try and find some kind of pattern match — all in a couple of nanoseconds. But not all security specialists believe ML is ready for such a mission-critical task.

> “If somebody wants to commit fraud or some bad act, it has to manifest itself at some point in the lifecycle.”
>
> Johan Gerber, executive vice president, MasterCard

Some argue that continuous authentication is the answer. That process continually watches users after they have been granted access, not only to determine what actions they take and attempt, but the exact methodology they use. Those activity logs could prove crucial when trying to recognize an attacker's return, especially when the attacker tries to vary methods explicitly to avoid such early detection.

Intel vs incident response

“Once the attack is ongoing, you're heads down on the attack,” says Johan Gerber, Purchase, N.Y-based executive vice president at MasterCard who is in charge of the financial behemoth's security and cyber innovations efforts. Threat intelligence, at that moment, “is secondary in importance. When we're in the middle of an attack, TI (threat intelligence) goes out the window.”

Gerber says machine learning has strong potential, but ML “has its own challenges.” Cyberattackers have their own ML systems and that creates a frustrating situation where one ML system is trying to avoid creating the patterns that another ML system could detect, he notes.

“If I see a specific threat, we'll play with that threshold in an automated fashion. Even though I am looking for something similar, it doesn't have to be exactly the same IP address, a fake IP address, this kind of message,” Gerber says. “It is looking for a pattern rather than an exact match. [The attackers] may change just enough that it's hard for the ML to pick up.”

Johan Gerber, executive vice president, MasterCard

With a marriage of ML and continuous authentication, the odds can move back in the favor of the enterprise. Gerber argues that it is not necessarily a matter of logging and storing “every minutiae of every user's behavior.” Instead, he says, let the ML system determine what to log and what to ignore. “That's what ML does really well: It figures out what is important,” Gerber says.

MasterCard, ranked 112 on the Fortune 500, sees “more than 800 billion of these events a year,” Gerber says. Using ML to sort through those incidents in real time is critical because it is premised on a simple assumption: Even if the attacker engages in minutes of innocuous behavior to throw security analysts and their ML systems off track, “if somebody wants to commit fraud or some bad act, it has to manifest itself at some point in the lifecycle,” Gerber says.

That said, Gerber argues that “ML is not the silver bullet for everything” and that security must allow ML — presumably in unsupervised learning mode — to do its thing by learning your system and what normal interactions look like, such as payroll lookups, address lookups. “I don't know what bad looks like,” Gerber says, adding that it is the machine learning system's job to figure that out.

Gerber stressed that companies need to attach special security attention to areas that bad actors like to use, such

OUR EXPERTS

Johan Gerber, executive vice president, MasterCard

Gary Hayslip, director of information security, Softbank Investment Advisers

Paul Hill, senior consultant, SystemExperts

Joshua Motta, CEO, Coalition

Kim Resch, owner and principal consultant, Creative Commerce Group

Doug Saylors, director, ISG

Mike Sechrist, chief technologist, Booz Allen Hamilton

Salvatore Stolfo, computer security professor, Columbia University as the ability to change a customer’s physical address.

Kim Resch, owner and principal consultant at Creative Commerce Group in Scottsdale, Ariz., is also a strong proponent of using ML to get around the active threat assault dilemmas.

“The volume of incoming threats is taxing teams whose resources are precious. Advancing AI techniques utilizing a machine that can learn threat patterns and monitor trends quicker and more efficiently than a human is quickly becoming a mainstay in day-to-day security operation centers,” Resch says. “It might seem like the magic bullet to keep up with growing threats and low resource pools [but] it comes with a price.”

That price, Resch says, is going to force CISOs to tailor their teams to accommodate how ML behaves, rather than the other way around. “Criminals have become very savvy on obtaining data from companies and using it against them, injecting in the wrong data and actually taking over systems. Companies must spend the time and money on [getting] their staff to think through these scenarios.”

Kim Resch, owner and principal consultant, Creative Commerce Group

Although ML “can identify threats quicker and with less errors than humans, it is not a set it and forget it scenario. Resources much be ready for the forensics to pursue the items identified. ML works off of set parameters and you have to have the humans and the intelligence to load this. Companies are scrambling for these skills and they are quickly training their staffs. If you are not proactive in setting this up and a system up to maintain it, you have to trust third parties and machines.”

> **“The volume of incoming threats is taxing teams whose resources are precious. Advancing AI techniques utilizing a machine that can learn threat patterns and monitor trends quicker and more efficiently than a human is quickly becoming a mainstay in day-to-day security operation centers. It might seem like the magic bullet to keep up with growing threats and low resource pools [but] it comes with a price.”**
>
> **Kim Resch, owner and principal consultant, Creative Commerce Group**

But, Resch argues, CISOs cannot trust ML too much, given its propensity for false positives. “Team members can get reliant on the blind trust of the systems doing the heavy lifting,” Resch says.

Although external feeds have a massive value, many security operations put insufficient priority on leveraging internal resources. That is often because security staffs are overworked and under-resourced and it’s easier for them to rely on prepackaged external threat feeds than to create their own internal feeds. But, done properly, those internal feeds are already precisely tuned to that enterprise’s security environments, the kinds of attackers interested in them and the cloud and software packages the enterprise is using. In short, internal feeds are, by their very nature, precisely tuned for that one enterprise.

Salvatore Stolfo, a computer security professor at Columbia University, encourages enterprises to leverage not only their internal logs, but to tweak the environment specifically to capture as much data about attackers as possible. And Stolfo is not above encouraging a little trickery to turn the tables on the attackers.

Fighting fire with fire

“Attackers hide behind VPNs and stepping stones, etc., to hide their tracks and complicate tracking them. One valuable addition to the defense of systems is to learn how attackers present themselves when executing their attacks, whatever it may be,” he says. “One can continuously model various logs that capture data generated by the attackers when executing their attacks and use those models to identify similar behaviors later or elsewhere. This is essentially tracking the attacker’s behavior over time. Exchanging these models across collaborating sites would be especially useful,” he continues.

“The key, however, is to gather ground truth that the data acquired is actually data generated by the attacker. There are ways to identify ground truth data. By planting decoys, for example, taken up by the attacker and misused by the attacker, monitoring for that planted data assures the defender they have ground truth data about the attacker. Plant decoy credentials on pastebin and monitor for their misuse. These are sophisticated methods wholly dependent upon ML systems to efficiently and accurately model attacker behavior. It is obvious to me this cannot be done by a rule-based or manual process," he says.

Stolfo has a few favorite tactics for tracking a cyberthief. "Some solutions to this problem include traditional threat intel sources such as searching DNS registry, new domain name registration databases, and other information in globally accessible databases to determine if a new hosting site seems suspicious. 'MyBank.com' would be deemed suspicious if it is a new domain name not authorized or owned by Bank.com.

"This approach," he continues, "requires sophisticated analysis such as applying machine learning and natural language processing techniques to identify suspicious new domain names. The process can detect spoofed site names before the sites are exposed to victims, but may have a high false positive rate or a false negative rate (missing a suspicious domain name used as a spoofed domain)."

Stolfo adds: "In addition to domain name processing to identify likely new spoofed sites, [CISOs can] also embed beacons that can track whenever the website is browsed and rendered on an endpoint device. The logging of the monitored beacon signal can be used to identify whether the webpage is spoofed, a copy of a legitimate site being hosted on an illegitimate, unauthorized site. This is accomplished by comparing the legitimate webpage hosted IP address of the server with the server IP address captured in the beacon signal. An unauthorized server IP address immediately implies an unauthorized spoofed website has been detected."

The next step of this tactic is the victim company's response. "When the spoofed site is detected, mitigation and response are the next considerations. Ideally, most organizations attempt to quickly take down the illegitimate, spoofed site. This depends upon the willingness of ISPs and the owners of the server domains to take down the offending site," he says.

"Some ISPs may not comply," Stolfo explains. "Recently attackers have cleverly exploited legitimate but vulnerable server sites to embed their spoofed sites. In the latter case, taking down the server site may not be easily achievable. The time between detection of the spoofed site and the successful take down operation is an open window of opportunity for attackers to gather victim user credentials."

And fighting trickery with trickery

The next path is to fight trickery with trickery, he says. "Here we stuff the spoofed site with decoy credentials, poisoning what the attacker might have already stolen. Decoy credentials are fake but believable login credentials consisting of a variety of fields of information required to gain access such as first/last names, addresses, phone numbers, account numbers, and other tokens deemed necessary for authentication at user login.

"Having a mixture of real and decoy credentials is a conundrum for the attacker, he says. "They have little choice but to test their quarry to identify the fakes fed to them. We also monitor for the misuse of decoy credentials to gather additional information about the attacker. The decoy credentials serve as a means of gathering non-obvious data about the attacker while thwarting their ongoing attack against a victim enterprise and their customers."

Stolfo is hardly alone in being a fan of trying to trick the bad guys. Joshua Motta is CEO of Coalition, a cyber insurance provider. Motta argues that internal analysis is critical, with tricky being just one tactic, but it's a good one.

"Many of the creative techniques we see in use by Fortune 1000 CISOs and CSOs come down to link analysis," Motta says. "How you can take one piece of threat intel and pivot to

> **"Those (enterprises) collecting [ISAC data] often don't share it back. It is far more risky to share than not to share... It's easier to consume it and not to give back. 'My job as CISO is to protect my company, not to protect yours.'"**
>
> **Joshua Motta, CEO, Coalition**

others? A threat intelligence report on an advanced persistent threat (APT) may contain information on the command infrastructure used by the threat actor (such as IP addresses associated with the infrastructure), which can then be used to monitor communications between that command infrastructure and the corporate network.

"I'm aware of one such CISO that used these techniques to determine that an APT always presented a rare user agent string referencing an out-of-date version of Chrome when hitting the company's perimeter," Motta continues, "which they were then able to monitor for and detect, allowing the company to discover new IP addresses in use by the attacker that were not previously known by other threat intelligence sources. They were then able to configure their defenses to drop packets associated with the fingerprint they discovered."

But Motta says some CISOs opt to go further. "Perhaps the most creative, recent technique I've heard used by a Fortune 1000 CSO, which is certainly in a legal grey area, involved purchasing data directly from botnet herders (criminals who operate botnets). The botnet herder, believing the CSO was a fellow criminal, provided a list of all hosts compromised by the botnet, allowing the company to not only address infections in their own network, but also that of their dependent business partners," Motta says.

Protecting the supply chain

Those business partners also represent another line of defense. Gerber points out that when MasterCard detects an attack leveraging a security hole, one of the first things the security team does — after neutralizing the attack, of course — is to run threat scans on its vast network of banks and processors and other partners, to try and identify the same security hole.

"If we can find it, criminals can find it, too," he says. This allows them to flag the problem to those partners, making MasterCard a more valuable partner and ideally increasing partner loyalty.

Mike Sechrist, chief technologist, Booz Allen Hamilton

"The best cyberthreat intel you are going to get is intel on the things you are seeing directly. CTI (cyberthreat intelligence) is all about contextualizing information."

Mike Sechrist, chief technologist, Booz Allen Hamilton

That idea could easily extend to any Fortune 1000 company and its partners. Beyond potentially building loyalty, it is an interesting way to enforce contracts requiring all partners to adhere to the enterprise's security requirements.

From a business management perspective, enforcing security opens up some interesting possibilities. Imagine Boeing or Walmart, for example, alerting a supplier to a discovered security hole. This is also a potential threat/warning that the company writing the supplier checks is now aware of a serious security hole. Why didn't that supplier discover that security — and fix it — on their own?

This might be worth remembering when it is time to renew contracts. Ultimately, the impact on the supply chain could be that the suppliers' own security and business teams might make identifying and mitigating vulnerabilities a higher priority.

This partner tactic also depends on the nature of the attack. For example, if the enterprise is a major hotel chain and the attacker is just looking to grab payment card credentials, then searching for that flaw among partners is certainly useful, but special attention should be paid to those partners who have a large number of payment cards on file. A supplier who does not might not be germane as far as that specific attacker is concerned. Still, the hole needs to be fixed.

Another threat intelligence concern that Motta has involves ISACs, especially those that are focused on specific verticals. Compliance rules — more precisely, how CISOs are interpreting those rules, which might or might not be accurate — are giving enterprises even more reasons to accept data from an ISAC but not share its data back.

“Those (enterprises) collecting [ISAC data] often don’t share it back. It is far more risky to share than not to share,” Motta says. “Privacy considerations are paramount everywhere and you don’t know what a third-party will do with your data. It becomes a risk versus reward situation.”

> **“The fact is that every incident is intelligence. The biggest mistake is that they take the threat intel from the security devices they purchased and they don’t go to the next step. Just one threat feed isn’t enough.”**
>
> Gary Hayslip, director of information security, Softbank Investment Advisers

Much of this confusion stems from the European Union’s GDPR privacy rules, along with the California Consumer Privacy Act. GDPR, for example, considers IP addresses to be PII (personally identifiable information) that has to be secured and shared only with a consumer’s opt-in permission. Some have interpreted GDPR to mean that once data is shared, the originating company is responsible for how it is used by any partners.

“The richest data is the most useful because it’s the most actionable,” Motta says. “Consider what the URL was and what was requested. There is a whole bunch of information in a query stream that can inadvertently be captured and shared. (Some CISOs) simply just say that ‘It’s not my job and I don’t have time to do.’ It’s easier to consume it and not to give back. ‘My job as CISO is to protect my company, not to protect yours.’”

Negotiating internal politics

Mike Sechrist, chief technologist at Booz Allen Hamilton, is another proponent of prioritizing internal threat data over external sources. “The best cyberthreat intel you are going to get is intel on the things you are seeing directly. CTI (cyberthreat intelligence) is all about contextualizing information,” Sechrist says.

Today this burden is going to likely fall on CTI analysts. But CTI teams often find themselves siloed in the corporate structure, unable to access much of the data and resources that they need, Sechrist says. “I haven’t seen one (CTI team) that is seamless without any internal barriers. There’s also a political component to it of data hoarding.”

Paul Hill, senior consultant, SystemExperts

Sechrist sees various reasons for these corporate barriers, not the least of which is the relatively young age of almost all CTI groups.

Cutting across these corporate data-access restrictions “takes political muscle internally,” he continues says. “They are going to need access to the SIEM, different security tools such as for endpoint data and web applications. Cloud, too. That access is typically not immediately granted unless, of course, if you have a high-level edict that is laid out by a major C-suite [executive].”

Paul Hill, a senior consultant with SystemExperts, an IT security consulting firm located in Sudbury, Mass., gave an example of these internal roadblocks that he ran into with one enterprise client: “Some audit logs included the public IP address of a device, but other audit log sources gave the internal NAT (network address translation) IP address, and the SOC (security operations center) staff were not provided with any tools to map public IP addresses to NAT IP addresses. Instead, the SOC team had to ask the network infrastructure team about individual IP addresses via email or IM (instant messenger).

“The SOC team did not have the ability to reverse query an IP address in order to determine the DNS name of the system,” he continues.

“Instead the SOC team had to ask the network infrastructure team about individual IP addresses. While the organization had multiple automated asset inventory systems, the SOC staff was not granted the ability to query or generate reports from the asset inventory systems. A request had to be made to either the server operation team or the desktop support team for individual IP addresses or system names."

Hill adds: "The data centers were outsourced to a large colocation facility that also functioned as a managed security service provider (MSSP), including the management of firewall rules, IDS (intrusion detection system) alert rules, and IPS (intrusion prevention system) rules. The SOC staff was not provided with a point of contact at the colocation facility. Instead, all inquiries and configuration requests had to be submitted to the infrastructure team which would in turn submit them to the colocation provider."

Doug Saylors, director, ISG

Security teams accepting external threat data and not localizing the data was an oft-discussed concern. Gary Hayslip, director of information security at Softbank Investment Advisers, says his experience at multiple companies shows many security teams "tend to look at threat intel as something already included, already running in the network. The fact is that every incident is intelligence. The biggest mistake is that they take the threat intel from the security devices they purchased and they don't go to the next step. Just one threat feed isn't enough." Even though "they may be comfortable getting the threat feeds from the vendor that they have already purchased [other systems] from. Some think its usefulness [is less so] because it's baked into so many products."

Hayslip argues that CISOs need to have created a detailed "repository of prior attacks, noting 'this is what worked the last time.'"

Doug Saylors, a Dallas-area director at Information Services Group, a technology research and advisory firm based in Stamford, Conn., believes most enterprises today are not well prepared for an attack and that is mostly because their security defenses are predicated on security practices from five to 10 years ago. "Perimeter defenses, intrusion prevention — they are still focused on preventing those on the outside getting in. A lot of the threats we are dealing with today are from the inside."

Saylors bemoans the fact that many enterprises today do not have current and comprehensive asset inventory maps, despite this being one of the many GDPR requirements. "If you don't know what should be there and what shouldn't be there," it's a lot harder to quickly detect what should not be in your network, Saylors says. "I think that was an optimistic hope that GDPR efforts would solve that."

Saylors adds his name to those who think that machine learning will be a key security defense tool, but thinks that it needs time to mature. He added that a zero trust model is also the way to go, but says that true zero trust implementation is going to be a three- to five-year journey for most Fortune 1000 enterprises.

Another concern with handling an active attacker is simply: What is the CISO's objective. Is the goal to block the attacker from ever getting in again or is to have the attacker arrested and prosecuted? Those very different objectives would dictate very different approaches.

"It's a revenue versus risk calculation," Saylors says. "Are you going to pursue and prosecute? How much harm do they do in that time" that it takes to gather the evidence needed for prosecutors?

Saylors takes a contrarian view to continuous authentication for fighting an active attack.

"The overhead is pretty significant because you have to track people for a while. There are millions of access points a day to gather, analyze and store. There's a tremendous cost to that." ■

Masthead

EDITORIAL
VP, EDITORIAL Illena Armstrong
illena.armstrong@haymarketmedia.com
SPECIAL PROJECTS EDITORIAL DIRECTOR
Stephen Lawton
stephen.lawton@haymarketmedia.com
SPECIAL PROJECTS COORDINATOR
Victor Thomas
victor.thomas@haymarketmedia.com

DESIGN AND PRODUCTION
ART DIRECTOR Michael Strong
michael.strong@haymarketmedia.com

SALES
VP, PUBLISHER David Steifman
(646) 638-6008 *david.steifman@haymarketmedia.com*
VP, SALES Matthew Allington
(707) 651-9367 *matthew.allington@haymarketmedia.com*