DARKREADING

MARCH 2024

Making Sense of Your Security Data: The 6 Hardest Problems

Maximizing security operations data takes not only an understanding of data sources and what they mean for risk, but also careful data management and cost awareness.

Brought to you by informa tech

INSIDE:

Next

FEATURE

Making Sense of Your Security Data: The 6 Hardest Problems

Maximizing security operations data takes not only an understanding of data sources and what they mean for risk, but also careful data management and cost awareness.

By Ericka Chickowski, Contributing Writer, Dark Reading

We *should* be in a golden era of cybersecurity visibility. Security monitoring capabilities are more prevalent than ever, and sources of security telemetry are plentiful. Unfortunately, most security operations centers (SOCs) are suffering from too much of a good thing when it comes to security data. With so many data sources to choose from — and with the velocity and volume of data generated from each of them scaling exponentially — operations teams are swamped. Many SOC teams are unable to effectively select the data sources that matter, not to mention cost effectively ingest and retain data within the technology stack in a manner that fuels detection and insights to enable swift threat response.

Here are some of the biggest problems that security operations teams face in managing and making sense of the data they have, and what security experts say it will take to overcome those challenges.

1. Security Data's Deadly S's: Sprawl and Silos

Many problems related to security analytics and security operations data are tied to the deadly s's of security data: sprawl and silos. There's a dark side to cybersecurity's penchant for leaning on "best-in-class" products to fill visibility gaps, solve niche detection problems, and chase new threats. Every new stand-alone product that should have been a feature, and every bell and whistle from these added solutions, contribute to the growing problem of tool sprawl.

The downstream data management problem caused by tool sprawl is far-reaching. The data streams that each of these products pump out are often locked in data silos that can be difficult to integrate into the existing security operations working stack. Most SOC analysts today must jump from tool to tool to get all the information and context they need from these various data streams

Previous Next

to get their work done.

"You have security tools and your applications that are all throwing off data that you need to consume to drive security outcomes, and they're in silos by default. Your logs are in one place, maybe you centralize some of it in a SIEM, but that's a really expensive and hard-to-manage proposition, and so you end up with corners you can't see around," explains Greg Notch, CISO of Expel, a managed detection and response (MDR) firm and a longtime security veteran who previously served as CISO for the National Hockey League.

Not only is security data typically siloed, but it's also extremely diverse and varying in quality, Notch says. This means security operations personnel have to perform a great deal of deep analysis and have a thorough understanding of the data's context so they can use it effectively — either manually or by building detection content — to drive detection, remediation, and automation of tasks.

"Bringing it together and providing context has always been the security operations challenge, but it's not getting easier with tool proliferation, let's put it that way," Notch says.

The combination of data integration and analysis issues that crop up from tool sprawl and data silos demands a lot of strategic and tactical planning from security leaders. The first step is getting more disciplined and selective about

what tools the SOC puts on its road map, laser-focusing on data compatibility and integration with the existing stack.

"For SOCs evaluating or deploying data-focused tools, the most important best practice is ensuring a tool's scalability and compatibility with existing systems and verifying that it provides actionable insights rather than just data collection," explains John Pirc, vice president at Netenrich, a security and operations analytics company.

Many CISOs pair discipline in their selection process with a drive to consolidate tools or migrate to platforms to minimize sprawl and data silos.

"Choose platforms over tools, offering more holistic capabilities and combining functionalities into a single platform where possible. And prioritize seamless integrations," recommends Balazs Greksza, threat response lead at Ontinue, an MDR provider.

2. SIEM Data Management and Compute Limitations

So, wait a minute: Isn't consolidating SOC analytical capabilities and moving to platforms the whole point and promise of security information and event management (SIEM) platforms? Why is fragmentation of data such a

Previous Next

problem after decades of SIEM adoption?

Unfortunately, SIEM evolution has been a cyclical ouroboros of the security world: The limitations of older SIEMs cause more sprawl and silos, while the new generation of SIEMs tries to fix the problem with slightly better integration and consolidated features. This SIEM progression has been going on for decades, but the serpent just keeps eating its own tail.

"The promise of the traditional SIEM is like, well, all your data's going to go here and give you a central place to do threat detection and incident response, but, largely because [of] the cost of data ingestion and the cost of retention, that never materialized," says John Bland, cybersecurity data cloud principal at Snowflake, a data cloud company. "And I think it's safe to say the market never delivered that 'single pane of glass.'"

As Bland alluded to, some of the biggest limitations of SIEMs in place today revolve around data ingestion and retention issues. On the ingestion side, many security operations teams struggle to smoothly add new data sources to their SIEM due to platforms' inability to parse inconsistent data source formats, including syntax and field scheme models. According to a recent survey by Gurucul, some 42% of security organizations report that it takes weeks or longer to add new sources to the SIEM.

Indeed, ingestion issues are not only driving up the cost of adding new data but also introducing operational friction

and the risk of a higher error rate in detection functionality.

"Building and maintaining customized parsers or waiting for the vendor to build a new one can take weeks or months, and the security team is left with an incomplete picture in the meantime," says Sanjay Raja, Gurucul's vice president of product. "Poor data ingestion means poor detection, even with analytical capabilities or analytics supported by machine learning and AI. Solutions that require a lot of custom engineering or development to support new data sources are a major detriment to the efficacy of the SOC."

Meanwhile, once a data source has finally been added to the SIEM, the data retention conundrum — namely, how long to keep it — begins. The cost dimension of data retention in the SIEM is huge — particularly as the variety, volume, and velocity of data streaming from security tools and the assets they monitor keep snowballing.

"You have to decide how long to keep it, how quickly you need it, how quickly you want to correlate it, and then you can't actually predict in some cases how much of it you're actually going to generate," Expel's Notch explains. "So, you've got this variable cost problem in your budget."

Budgets are finite, and data retention costs can get very expensive, very quickly in many SIEM environments. Part of the reason for this, Snowflake's Bland explains, is that many SIEM pricing models don't separate storage from compute capacity, which drives up retention costs. Similarly, SIEM contracts rarely build in flexibility or elasticity in compute capacity, which can make things very difficult when problems hit.

"Everyone vying for a static amount of compute resources causes problems when maybe there's a large investigation and it's all hands on deck and they can't elastically scale up to meet that need," he explains. "And, certainly, no one wants to renegotiate a three-year SIEM contract for a two-day investigation or because a new user came on board and enabled a few poorly written detections and now the performance is suffering across the whole organization with the SIEM."

3. Tough Data Architecture Choices

These SIEM limitations are what's driving a growing trend to

Previous Next

consolidate at the data layer with data lakes. A security data lake pools security data into a centralized, unstructured repository that can be directly queried or that other security analytics tools can be built upon.

"Security data lakes have helped push data retention beyond compliance use cases with more emphasis on the identification of key threats and trends — particularly when security practitioners can search across a year's worth of data at speed versus 30 to 90 days, if they were lucky, in the past," says Ken Westin, field CISO of cybersecurity firm Panther Labs. "Having access to large amounts of valuable security data at your fingerprints, with response times in milliseconds versus hours or days, security practitioners can develop hypotheses for new detections much more quickly than they have in the past."

Overall, the approach breaks down data silos that impede detection and response, Bland explains, and opens up use cases beyond basic search and investigation, such as advanced behavioral analytics and threat hunting.

But security data lakes alone typically aren't a direct replacement for SIEM, notes Oliver Rochford, a longtime security industry analyst and security futurist. As he explains, security data lakes don't give security operations and broader cybersecurity programs "the full shebang" of what they'd get from SIEM, most particularly when it comes to compliance reporting, ticket management, and security content that's built on top of the data.

"A lot of the companies I've spoken to in the past couple of years [are] not throwing out their SIEM and replacing it with a data lake; they're running both. They're doing certain things in a SIEM, they're trying to reduce the cost, but what they want in the data lake is this long-term historical threat hunting," Rochford explains. "They want to be able to build their own models. They're starting to move to a far more sophisticated detection engineering model, as well."

With that said, running SIEM systems and data lakes side by side — not to mention adding products such as extended detection and response (XDR) into the mix — is not a viable solution in the long term. Right now, the industry is in a temporary phase in which security teams are juggling a lot of variables. Notes Rochford, underneath most modern XDRs and SIEM is a data lake, which at its core is just an innovative way of looking at data storage, collection, and acquisition.

"But a SIEM is much more than that," he adds. "A SIEM is all of the security content that you have to build on the top."

XDR providers, on the other hand, seek to go extremely deep in quality data, but that depth comes at the cost of breadth: "As you move to the data lake, you have a huge amount of breadth, but what you don't have is the content," Rochford says.

Ultimately, the market will converge to the point where it won't be an either-or choice, and most of these vendors will be playing in the same market segment because they're all vying for the same security operations budget.

In the meantime, security operations leaders must determine the extent to which the data lake should be decoupled from the detection stack or the incident response stack. For flexibility's sake, an organization could have its data lake and data pipelines running separately, but then they lose the unifying power of integration.

"It's like having your headset separate from your body — it makes no sense at all," Rochford says. "So, basically, balancing these two extremes, maintaining that flexibility but also having strong integration and interoperability so that you can derive synergies to make the sum greater than the parts, that's the challenge for users."

Unfortunately, there's no easy answer, and the choice will depend on the company, the kinds of resources it has, and the security objectives that matter most to its leaders.

4. Data Quality Issues

Regardless of what direction an organization goes with security data architecture, data quality issues stand as a fundamental obstacle that must be overcome to foster good security outcomes.

"Security operations depends on quality information," says Shane Shook, venture partner at Forgepoint Capital. "Automation for data ingestion for analysis is widely available; however, automation for data quality is currently scarce. Simply excluding 'unfit' records in analytics pipelines leads to false negatives that can critically impede sequential analyses — such as the MITRE ATT&CK framework that is favored by many security organizations."

Shook says that before the industry puts the cart before the horse by fixating on correlation engines and AI/ML-powered analytics, it needs to start to get more serious about finding ways to ensure data completeness and quality across security data sources.

At the security operations practitioner level, this could mean focusing more on log parsing and defining default field schemes of log data, says Or Saya, cybersecurity architect at CardinalOps, a detection posture management company. Log parsing is the process of extracting information from logs and putting them into fields for faster searching and visualization of data.

"Security operations teams should implement robust log-parsing mechanisms to ensure that the data is structured correctly for analysis," Saya says. "This involves defining clear parsing rules for various log formats and ensuring consistency in log interpretation."

Security operations teams should also be establishing standardized default field schemes for log data within the organization, defining fields that every log entry should have, such as timestamp, source IP, destination IP, user,

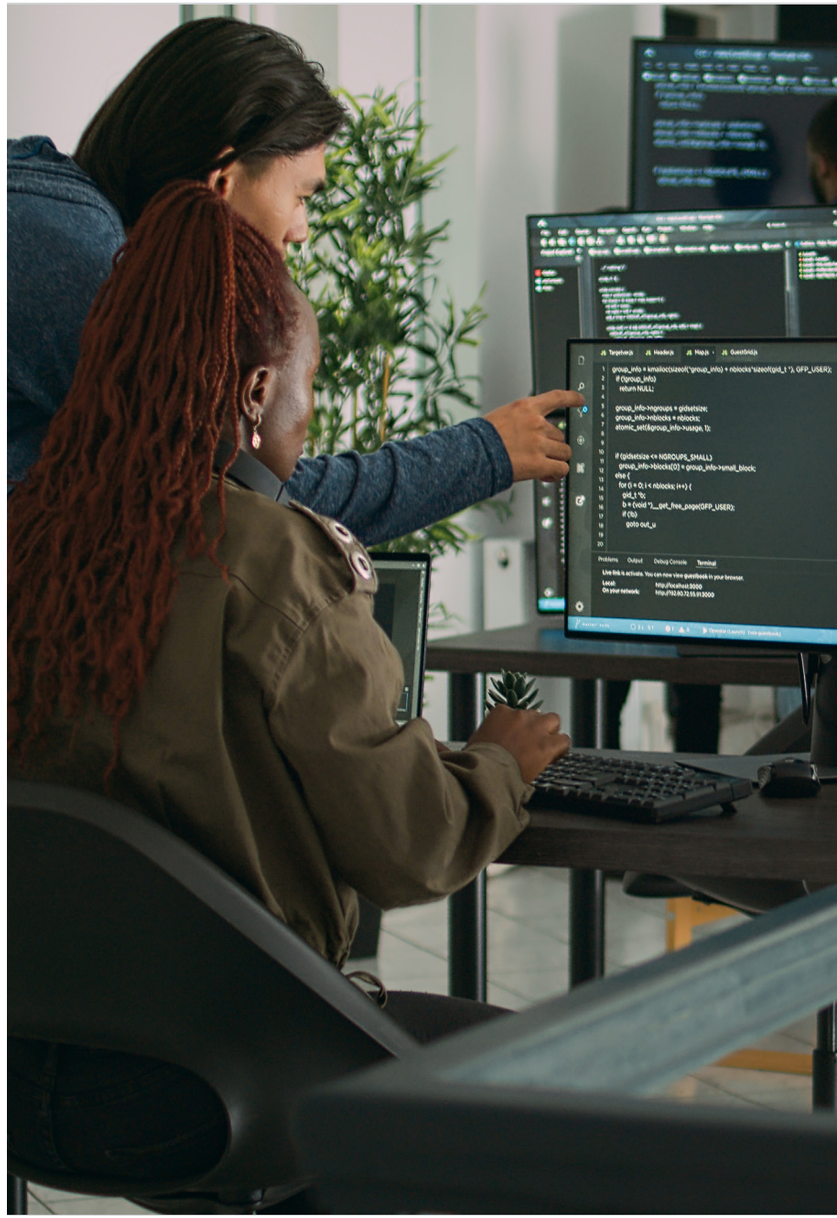

and action taken. This helps ensure consistency across different log sources and makes it easier to do correlation work and analysis.

However, analysts must understand that this scheme could cause unintended consequences.

"On one hand, a scheme can help analysts make sense of unfamiliar log sources by mapping them to an understandable model," Saya explains. "On the other hand, working with such a scheme requires continuous validation to ensure that the data is correctly normalized based on the scheme. Failure to do so can lead to blind spots that are very hard to identify."

5. Data Selection Dilemmas

The value of security data is a function not only of data quality and integrity but also of the contextual value it brings to an analyst in understanding the nature of a threat or an exposure in the environment. One of the biggest challenges security operations teams face is in selecting the right data sources for the analysis or detection job they're dealing with at any given time. As Notch explains, the vast majority of data in the SIEM today is "useless" from a detection and response perspective.

"Most of the data in there is not helpful for either security or business outcomes, but you have to keep it for either compliance reasons or for investigations," Notch says. "There is some valuable data in there, but the value is in the

correlation and the enrichment of it."

The operations team that supports his firm's detection and response services does a great deal of work in selecting the best data for the analysis or correlation content that they're building out.

"We really try to suppress garbage data from getting even near our environment," he says, explaining that a lot of network detections — unless they come from highly restricted environments — tend to fall into that category. Similarly, "flotsam and jetsam" from Windows environments beyond important authentication events tend to fall in that garbage category. "We've got very smart folks who are thinking about that data ingestion, what to take, what to leave behind, what things matter, how they fit together — so, how an alert from your EDR would fit together with an alert from your network connectivity, and only taking the pieces … that matter to make that correlation."

For most organizations, data choices will also be driven by costs. Organizations would do well to conduct cost analysis for the care and feeding and effective use of particular data sets.

"You have to model the cost of what you're doing in a way that you might not have 15 years ago," Notch says. "If you're a CISO, you need to be thinking, 'How do I get cost consistency and good unit economics out of security detections?'"

Saya says it's important to evaluate costs associated with ingestion into a SIEM versus the value of being able to run certain kinds of queries.

"Consider which cost model aligns better with the organization's budget and operational preferences. Some organizations may find it more cost-effective to optimize data ingestion processes, while others may prioritize efficient query execution," he says. "Striking the right balance ensures cost efficiency without compromising on the effectiveness of security operations."

6. Shortcomings in SOC Data Science Skills

Clearly, data selection requires a deep understanding of what the telemetry data actually means about security posture and

threat activity. But equally important to the team skillset is having the data science chops to not only build the detection rules and analytics around that selected data but also to do the cost optimization of data pipelines.

With most SOCs today already facing a security skills shortage, security leaders who want to get the most out of their security operations data will need to refocus SOC hiring and training on data-centered skill sets.

"SOC recruiting and professional development needs to evolve by incorporating data science training and principles into the skill sets, fostering a workforce adept in both security operations and data analysis techniques," says Netenrich's Pirc.

Rarely is a SOC going to snag a unicorn in this area — someone who is a grizzled security veteran and a data science whiz all in the same body. This means that security operations leaders will have to do some intentional team building to create the right mix of skills and then focus on cross-functional training. The key is picking up "lifelong learners" on the security expert side who are willing to integrate data science principles into their work and their professional development, Ontinue's Greksza says.

"Enabling them through job rotations and cross trainings is most effective," Greksza says. "Build cross-functional skills: Data scientists with security exposure and security experts with interest in data visualizations and statistics excel [when] delivering value together."

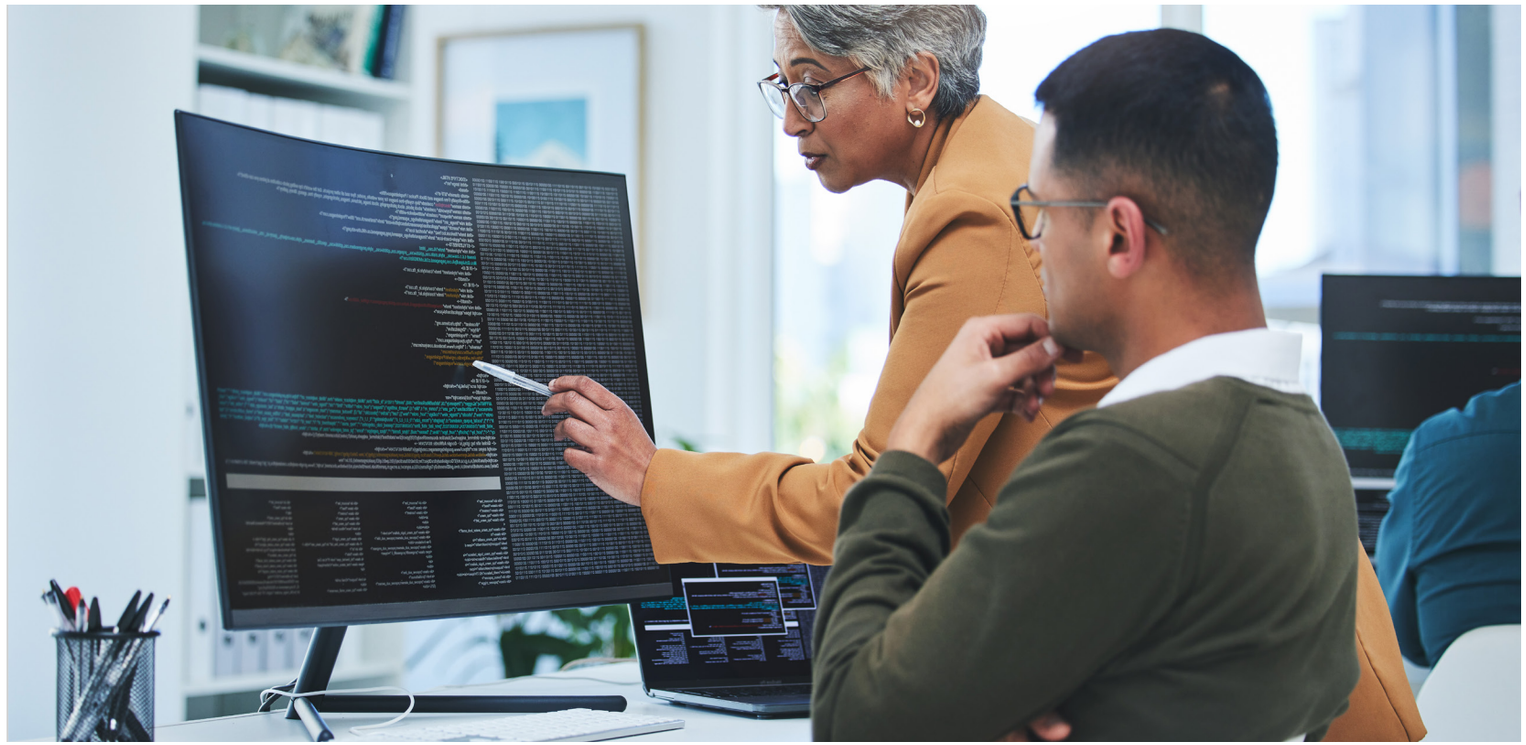

Data Strategy Matters

Ultimately, security operations requires a strong data management and data handling strategy to effectively balance objectives around response time and cost.

"Security is not a big data problem; it is a 'right information and intelligence at the right time' problem," Greksza says. "To come to the right conclusions, different tooling will bring different benefits to the table. A security data lake has a very different use case in comparison to XDR, SIEM, or even a compliance monitoring solution."

For security leadership to drive good outcomes (and to validate those up the corporate food chain), the key is to define clear objectives and requirements, and then tie that back to how data architecture strategy and data selection work is approached.

About the Author: Ericka Chickowski specializes in coverage of information technology and business innovation. She has focused on information security for the better part of a decade and regularly writes about the security industry as a contributor to Dark Reading.

Most Enterprise SIEMs Blind to MITRE ATT&CK Tactics

Organizations are largely deluded about their own security postures, according to an analysis, with the average SIEM failing to detect a whopping 76% of attacker TTPs.

By Elizabeth Montalbano, Contributing Writer, Dark Reading

Despite enterprises' best efforts to shore up their security information and event management (SIEM) postures, most platform implementations have massive gaps in coverage, including missing more than three-quarters of the common techniques that threat actors use to use to deploy ransomware, steal sensitive data, and execute other cyberattacks.

Researchers from CardinalOps analyzed data from production SIEM platforms from companies such as Splunk, Microsoft Sentinel, IBM QRadar, and Sumo Logic and found that they have detections for just 24% of all MITRE ATT&CK techniques. That means that adversaries can execute about 150 different techniques that can bypass SIEM detection, while only about 50 techniques are spotted, the researchers said.

This is despite the fact that current SIEM systems actually do take in sufficient data to cover potentially 94% of all these techniques, CardinalOps revealed as part of the company's "Third Annual Report on the State of SIEM Detection Risk."

Moreover, organizations are largely deluded about their own security postures and "are often unaware of the gap between the theoretical security they assume they have and the actual security they have in practice," the researchers wrote in the report. This creates "a false impression of their detection posture."

MITRE ATT&CK is a global knowledge base of adversary tactics and techniques based

Previous Next

on real-world observations that's aimed at helping organizations detect and mitigate cyberattacks. The report's data is the result of analysis of more than 4,000 detection rules, nearly 1 million log sources, and hundreds of unique log source types used in SIEM across a range of diverse industry verticals — including banking and financial services, insurance, manufacturing, energy, and media and telecommunications.

A Lack of SIEM Fine-Tuning to Blame for Detection Fails

The key issue contributing to the current state of SIEM efficacy (or lack thereof) seems to be that even though resources exist for organizations to use knowledge, automation, and other processes to detect adversaries and potential attacks on their environments, they still largely rely on manual and other "error-prone" processes for developing new detections, the researchers noted. This makes it difficult to reduce their backlogs and act quickly to fill gaps in detection.

Indeed, SIEMs themselves "are not magic" and rely on the organizations deploying them to do so correctly and efficiently, notes Mike Parkin, senior technical engineer at Vulcan Cyber, a security-as-a-service provider of enterprise cyber-risk remediation.

"Like most tools, they require fine-tuning to deliver the best results for the environment they're deployed in," he says. "These report results imply that many organizations have gotten the basics working but haven't done the fine-tuning necessary to take their detection, response, and risk management strategies to the next level."

In addition to the need to scale detection-engineering processes to develop more detections faster, one key issue that seems to be tripping up detection in enterprise SIEM deployments is that on average, they have 12% of rules that are broken, which means they will never file an alert when something is amiss, according to the report.

"This commonly occurs due to ongoing changes in the IT infrastructure, vendor log format changes, and logical or accidental errors in writing a rule," the researchers noted in the report. "Adversaries can exploit gaps created by broken detections to successfully breach organizations."

Why MITRE ATT&CK Matters

MITRE ATT&CK, created in 2013, has now "become the standard framework for understanding adversary playbooks and behavior," the researchers noted. And as threat intelligence has advanced, so has the wealth of knowledge the framework provides, currently describing more than 500 techniques and sub-techniques used by threat groups such as APT28, the Lazarus Group, FIN7, and Lapsus$.

"The biggest innovation introduced by MITRE ATT&CK is that it extends the traditional intrusion kill chain model to go beyond static indicators of compromise (like IP

Previous Next

addresses, which attackers can change constantly) to catalog all known adversary playbooks and behaviors (TTPs)," the CardinalOps researchers wrote.

Organizations clearly see the value in using MITRE ATT&CK to help them in their security efforts, with 89% currently using the knowledge base to reduce risk for security-operations use cases — such as determining priorities for detection engineering, applying threat intelligence to alert triage, and gaining a better understanding of adversary TTPs, the researchers noted, citing Enterprise Strategy Group (ESG) research.

However, using the framework to support SIEM efforts and using it well appear to be two very different scenarios, the report found.

Closing the SIEM Gap

There are steps that organizations can take to help close the gap between what a SIEM is capable of in terms of cyberattack detection and how they currently are using it, researchers and security experts said.

One key strategy would be to scale SIEM detection-engineering processes to develop more detections faster using automation, something that companies already use widely to great effect in "multiple areas of the SOC, such as anomaly detection and incident response," but not so much in detection, they noted in the report.

"The detection-engineering function remains stubbornly manual and typically dependent on 'ninjas' with specialized expertise," the researchers wrote.

Indeed, having a focus on automation is critical to achieving goals with limited human and financial resources, agrees one security expert.

"This includes expanding automated detection to include Internet of things (IoT) and operational technology (OT) attack vectors, as well as having plans already in place for automated threat remediation," says John Gallagher, vice president of Viakoo Labs at Viakoo.

One key challenge that organizations continue to face is that the current attack surface — which now includes large numbers of vulnerable network-connected devices as well as the typical enterprise network — has grown well past what the IT organization is currently capable of supporting or managing, Gallagher says.

"To defend and maintain the integrity of those assets requires IT working closely with other parts of the organization to ensure those assets are visible, operational, and secure," he says.

Indeed, Parkin observes, until organizations can get a clear picture of their threat surfaces, manage their risk, and prioritize events to focus on what matters most, there will be problems.

"We have the tools to make it happen," he says. "But it can be a challenge to get them deployed and configured for best effect."

About the Author: Elizabeth Montalbano is a freelance writer, journalist, and therapeutic writing mentor with more than 25 years of professional experience. Her areas of expertise include technology, business and culture.

Tel Aviv Stock Exchange CISO: Making Better Use of Your SIEM

If rule writing for SIEMs isn't managed properly, it can lead to false positives and misconfigurations, which create extra work for the SOC team.

By Dan Raywood, Contributing Writer, Dark Reading

For Gil Shua, getting the most out of the security information event management (SIEM) system for the Tel Aviv Stock Exchange comes down to getting the signal-to-noise ratio right. That, and writing the right rules.

Signal-to-noise ratio, as every radiofrequency engineer knows, boils down to the amounts of actual content (signal) to static and other sonic disruption (noise). For Shua, the goal is to minimize the amount of noise getting sent to the SIEM in favor of actionable content. He's looking for something that makes him get up from his desk with the realization, "We have a problem; we have something that we want to address now and fix it."

Shua has worked in various security positions at the Tel Aviv Stock Exchange (TASE) for more than a decade and was appointed CISO in 2022. During that time, he says, it's been a "constant chase for data resources" to ensure that the signal-to-noise ratio skews in favor of signal data to maximize the capabilities and benefits of the exchange's SIEM.

Filtering the Noise

Shua and his team have their work cut out for them since, with most SIEMs, "you see a lot of noise, and not a lot of signal." This leads to false positives and misconfigurations,

Previous Next

which, in turn, create extra work for the SOC team, reduce productivity, and are an impediment to trying to getting a SIEM working.

To minimize this, Shua says, the SOC team can write rules for how the SIEM handles incoming data, but creation of those rules takes up valuable SOC team time, as well.

Writing SIEM correlation rules is relatively easy if the SIEM solution already has predefined log parsing and rules for the reporting application, says Shua. But before rules can be written, the SOC team must:

- Figure out the data structure and identify relevant fields needed for the rule.
- Understand the logic of the reporting systems, as they may have their own log standards.
- Create an exact rule correlation and analyze exceptions.
- Perform quality assurance and testing.

These action items can take a few hours each, but if they're more complex, they can take days to complete, Shua adds.

"When you establish a SIEM, you have two concerns. One is, 'Do I have the rules that protect me against relevant attacks … am I covered with effective rules?' The second thing is, 'Do I get the information from our log sources and other telemetry that will trigger these rules?'"

The recent addition of CardinalOps' platform has improved Splunk Enterprise at TASE. Shua says the process of writing rules has been massively reduced, with 85 rules produced since the technology has been in use. "The team is more focused on implementing rules and testing them and not writing them, which was the most time-consuming process in the link," he adds.

So, are SIEMs worth the time and money spent on correlation and rules writing? Shua admits that maintaining a SIEM is a demanding task, as there's a need for constant updates and modifications. Despite all the effort, some attacks may go unnoticed due to lack of visibility or matching rules.

"I expect future solutions would adopt automation capabilities for autonomous rule creation and response, out of the box," Shua says.

And because SIEMs draw data from many sources, they must become more efficient with processing, analyzing, and storing data that's in different formats. "You have to make adjustments to maintain," Shua says, adding that improper change management means an organization is likely to miss some security events.

About the Author: With more than 20 years' experience in B2B journalism, including 12 years covering cybersecurity, Dan Raywood brings a wealth of experience and information security knowledge to the table. He has covered everything from the rise of APTs, nation-state hackers, and hacktivists to data breaches and the increase in government regulation to better protect citizens and hold businesses to account.

Omdia: Stand-Alone Security Products Outsell Cybersecurity Platforms

Cybersecurity platform vendors say enterprises want to buy fewer solutions from fewer vendors. Omdia research tells a different, more nuanced story.

By Eric Parizo, Principal Analyst, Omdia

In its many briefings with cybersecurity vendors, one of the most consistent themes Omdia hears is why enterprises need cybersecurity platforms.

Across nearly all segments of cybersecurity, the opening statement from large vendors always goes something like this:

"Enterprises have too many stand-alone security products. They're expensive to purchase, deploy, and manage; point solutions function in a silo because they aren't designed to work together; trained, experienced cybersecurity professionals are hard to find and harder to keep — hence, fewer products mean more efficient training and staffing for CISOs. Not to mention, they aren't working because look what happened with that latest big scary data breach!"

Instead, vendors claim, enterprises could get better outcomes if they give up their multitude of stand-alone products and instead purchase a cybersecurity platform solution, which rolls up the capabilities of many discreet products into an all-in-one offering from a single vendor.

CrowdStrike, Fortinet, Palo Alto Networks, Trend Micro, and many others have positioned themselves as cybersecurity platform vendors, employing go-to-market messaging that emphasizes the integration, single user interface, improved security efficacy, and better return on investment that their cybersecurity platforms provide.

On its face, this seems sensible. All the above-mentioned challenges that come with point solutions are very real. Omdia asserts enterprises need their cybersecurity products to work together, specifically by exchanging data and performing orchestrated functions. However, building and running an integrated ecosystem of best-of-breed security solutions provided by many vendors is a never-ending challenge — one that keeps security architects awake at night.

Previous Next

A Growing Number of Deployed Products

Today's enterprises really do have a lot of security products. Omdia research shows that a majority of enterprises have 21 or more stand-alone security products, and a third of organizations have 31 or more.

It's not hard to buy into the sales pitch from platform vendors that the fastest way for enterprises to improve their security is by buying fewer point products and shifting spending to cybersecurity platforms.

However, according to Omdia research, it's simply not happening.

According to data from the 2023 Omdia Cybersecurity Decision Maker survey, organizations indicated an increase in the number of stand-alone security products they are using, not a decrease.

Omdia research shows that from June 2022 to May 2023, more than 80% of survey respondents saw an increase in the number of stand-alone security products in their organizations (161 respondents). Furthermore, for 44% of respondents, it wasn't just a minor increase; in those enterprises, the number of stand-alone products increased 11% or more. Conversely, just 7% of respondents noted a decrease.

The numbers highlight a stark contrast between the perception that is being advanced by cybersecurity platform vendors and the reality being observed in enterprises. Despite the vendor-touted benefits of the platform approach, data indicates enterprises still live in a best-of-breed environment.

Omdia theorizes there are several possible explanations:

- **Messaging:** Platform vendors simply aren't effectively communicating the benefits of a platform-based approach. This is possible, but unlikely. Anyone who has viewed vendor websites or attended an industry trade show like Black Hat in recent years has received a heavy dose of platform-centric marketing.
- **Entrenchment:** Charge is hard, particularly in cybersecurity, and security teams are usually loath to abandon tools into which they have invested time and effort to achieve their desired outcomes. It is much easier for an incumbent vendor to win a renewal than it is for a rival to win a displacement.
- **Lock-in:** Furthermore, in many cases, the vision of a migration to a platform approach requires a commitment to a broad, expansive, and potentially disruptive change. Implementing a platform approach means making a long-term commitment to one vendor and forgoing choices down the line. Many enterprises may be reluctant to cede that level of control.
- **Efficacy:** Enterprises simply don't see cybersecurity platforms delivering the desired outcomes. Indeed, Omdia has observed that many if not most cybersecurity platforms are created through a series of point product acquisitions, each of which is then re-engineered to be a component of a platform offering.

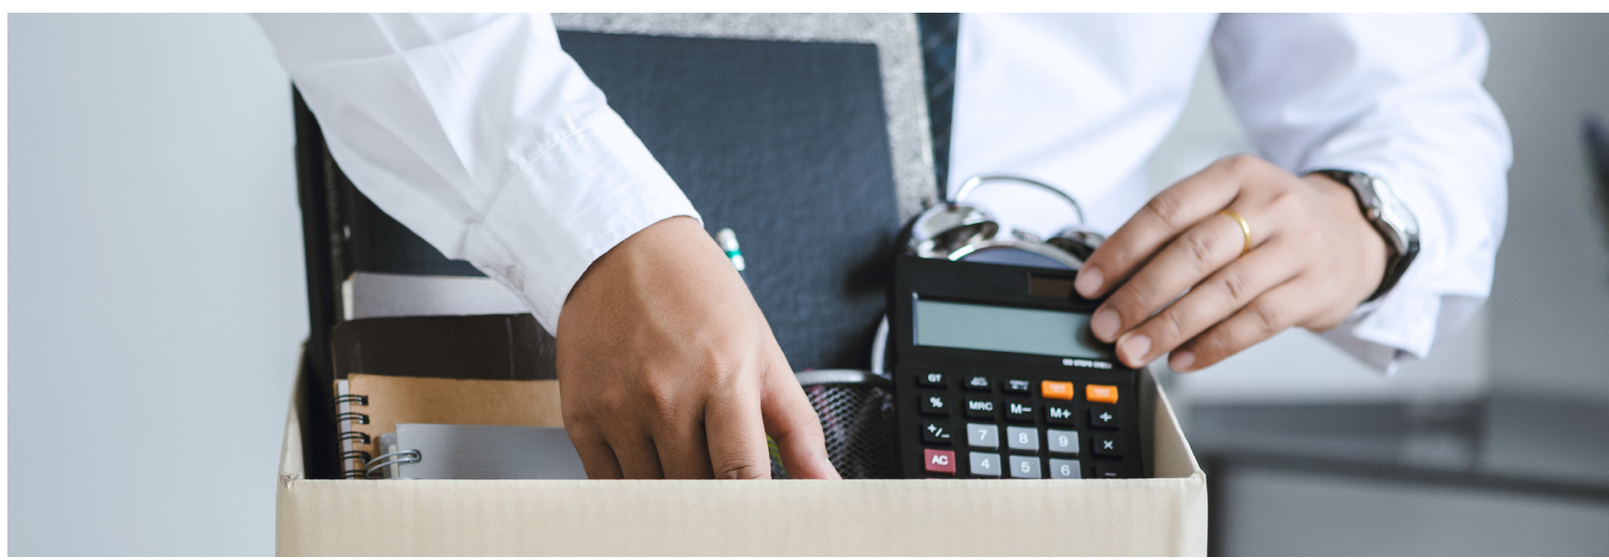

Previous Next

In practice, this means a single platform may include tools written in multiple programming languages, using different data formats, and requiring incongruent user interfaces. This creates a series of underlying technical challenges that negatively impact platform outcomes.

- **Specificity:** Enterprises purchase point products because they tend to be very good at solving a very specific problem, and that earns loyalty among customers. Niche vendors that successfully ease a key cybersecurity pain point can stand the test of time amid a tumultuous industry. Case in point, vendors AlgoSec and Tufin have been addressing multivendor firewall management going on two decades and counting, while cybersecurity titans like McAfee and Symantec have risen and fallen.

These are just a few of the possible reasons enterprises are using an increasing number of stand-alone security products, and Omdia intends to conduct further research in this area. But, for now, there are several notable takeaways.

Accept Enterprise Reality

For vendors, fostering a cybersecurity platform may be a sensible business strategy, but it is equally important to accept the reality that few enterprises are buying into single-vendor platforms. Meeting the needs of the market also requires catering to organizations living a best-of-breed approach. This means vendors should

not only support ease of integration through technology partnerships and open standards, but also minimize the finger-pointing when customers need help making rival vendors' offerings work together.

For enterprises, the best-of-breed approach may be familiar, but platform vendors are working hard to address many of the shortcomings that have hindered all-in-one cybersecurity platforms to date. When solution refresh cycles come up, it's worth stepping back and taking a broader look at whether the evolving platform landscape may offer a better long-term approach. For those focused on point solutions, be sure the requirements include actual examples of successful best-of-breed integrations and testimonials from customers that have achieved the desired outcomes from their integrated security architectures.

For service providers and channel partners, platform vendors' struggles are your opportunities. On one hand, working through (or entirely removing) the challenges of implementing and running a best-of-breed cybersecurity solution architecture isn't easy, but vendors and enterprises alike desperately need this assistance. On the other hand, cybersecurity platforms make for a compelling service offering, and evangelizing the benefits of these platforms is an area where vendors' reliance on the channel is only growing.

For more information, read "Omdia's Cybersecurity Decision Maker 2023: Overall Findings & Enterprise Cybersecurity Operations (SecOps)."

About the Author: As Principal Analyst, Eric Parizo leads Omdia's Enterprise Cybersecurity Operations (SecOps) Intelligence Service, its research practice focusing on threat detection, investigation, and response, as well as security operations center (SOC) issues. Eric also monitors global cybersecurity trends and top-tier cybersecurity vendors in North America.

Previous Next

Security Operations Data Analytics: What to Collect?

The data that enterprises decide to collect depends on security monitoring needs and use cases.

By Anton Chuvakin, Security Adviser at the Office of the CISO, Google Cloud

For years, organizations deploying security information and event management (SIEM) or similar tools have struggled with deciding what data to collect. So, the dreaded question lives on: What data sources do I integrate into my SIEM first?

The best answer to this question is "output-driven SIEM," where SIEM collection depends on your security monitoring needs and use cases, as well as on how you prioritize them based on risk. In contrast, a list of top log sources aggregated from many organizations will end up being useless for an organization with different security needs and challenges.

While the concepts behind the output-driven SIEM approach have been known for more than 10 years, many organizations are still looking for best practices in data collection before they decide how they will use the data. In fact, large organizations often make the decision to integrate a log source into their SIEM or SecOps platform based on factors other than pure security necessity.

Such factors often include:

- Necessity for detection
- Necessity for alert triage and incident response
- Necessity as context data for utilizing another log source
- Compliance requirements to collect and retain specific log types

- Compliance requirements to monitor a data source and/or system
- Ease of integration of the log source
- Collector and parser availability from the vendor
- Ability to transfer log data to a SIEM
- Other planned log sources that compete for attention
- Data volume of the log source

If a SIEM product charges per volume of data collected, the cost of introducing a new data source into the platform may be one of the deciding factors. For example, will you include a data source that will consume 10% of your overall SIEM license if you plan to use it only as context — valuable though it may be — for another data source? (In other words, you don't plan to write any detection rules or apply other detection logic based on this telemetry.) A popular example here would be DHCP logs: How many detections rely solely on DHCP logs? None or very few, at most.

In fact, experiences with SIEM deployments (going back to 2002) have shown us that few organizations include DNS or DHCP logs during their initial phases of SIEM rollout. In fact, some never include them in their SIEM at all. When asked why, those people usually explain that while they are convinced of the general utility of DNS logs, they do not see much value in each individual message that costs money to collect. These logs are essentially

"sparse value logs," where the value is in getting the bulk rather than in getting some particularly valuable messages — for example, Windows Security Event ID 1102. As a result, SIEM operators have doubts about paying for the inclusion of this data into their SIEM.

This mindset has resulted in an architecture model where one product is used for high-value logs while another product augments it by storing more voluminous logs. This kind of setup works if each product has good APIs — to perform functions such as querying one telemetry repository from another — but it is useful to remember that the model does not offer advantages other than cost.

At the same time, top log sources change over time, and the firewall and server logs flooding SIEM tools in the early 2000s have been supplemented with critical sources such as:

- Cloud logs (including AWS CloudTrail, Google Cloud Audit, and Amazon VPC flow logs)

Previous Next

- Sysmon and EDR telemetry
- Identity provider logs (such as Okta, Ping, and Microsoft Entra ID)
- Microsoft 365 and Workspace logs, and other key SaaS application logs
- API access logs from various applications and platforms
- Development environment logs, CI/CD pipeline logs, HashiCorp Terraform logs
- Container system Kubernetes logs (such as Kubernetes audit logs)
- Enterprise browser logs

At the same time, some of the classic sources remain very popular and useful:

- IT and security tool management console access logs (from VPN, UTM, EDR, SIEM, and SOAR to other management tools)
- VPN and zero-trust system logs
- Web proxy logs

Also, some log sources qualify as "newly popular," even though organizations have been collecting them for years, if not decades. These include:

- Business application logs
- DLP and other data-aware security technologies (such as emerging data detection and response)

- Email logs (likely overlapping with popular SaaS application logs)

Finally, if you integrate a new log source type, make sure that you monitor for the log telemetry actually arriving into your SIEM.

To optimize your SIEM:

- Practice "output-driven" SIEM, as this approach increases the chance of collected log data being useful for your detection and response efforts.
- Include logs that are of key investigative value and those useful as context (such as DNS and DHCP logs).
- Review your current collection posture, and align it with your detection use cases.
- Evolve the collection based on changes related to needs, risk, and IT. (For example, add cloud logs when cloud use for the business increases.)

About the Author: Anton Chuvakin works for the Office of the CISO of Google Cloud, where he arrived via the Chronicle Security (an Alphabet company) acquisition in July 2019. Before that, he was a research vice president and distinguished analyst for the Gartner for Technical Professionals (GTP) Security and Risk Management Strategies team.