SIEM Kung Fu

Version 1.4
Released: April 21, 2016

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051
info@securosis.com www.securosis.com

Author’s Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on the Securosis blog, but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

This report is licensed by Intel Security.

www.intelsecurity.com

McAfee is now part of Intel Security. With its Security Connected strategy, innovative approach to hardware-enhanced security, and unique McAfee Global Threat Intelligence, Intel Security is intensively focused on developing proactive, proven security solutions and services that protect systems, networks, and mobile devices for business and personal use around the world. Intel Security is combining the experience and expertise of McAfee with the innovation and proven performance of Intel to make security an essential ingredient in every architecture and on every computing platform. The mission of Intel Security is to give everyone the confidence to live and work safely and securely in the digital world.

Copyright

SIEM Kung Fu

Table of Contents

SIEM Fundamentals

Another SIEM research paper? Really? Why are we still talking about SIEM? Isn't it old technology? Hasn't it been subsumed by new and shiny security analytics products and services? Be honest — those thoughts crossed your mind, especially because we have published a lot of SIEM research over the past few years. We previously worked through the basics of the technology and how to choose the right SIEM for your needs. A bit over a year ago we looked into how to monitor hybrid cloud environments.

Security monitoring needs to be a core, fundamental, aspect of every security program. SIEM — in various flavors, using different technologies and deployment architectures — is how you do security monitoring.

Security monitoring needs to be a core, fundamental, aspect of every security program. SIEM — in various flavors, using different technologies and deployment architectures — is how you do security monitoring. So it's not about getting rid of the technology — it's a question of how to get the most out of existing investments, and ensure you can handle modern advanced threats.

We understand how SIEM got its bad name. Early versions of the technology were hard to use and required significant integration just to get up and running. You needed to know what attacks you were looking for, and unfortunately most adversaries don't share their attack playbooks before they come knocking on your door. Operating an early SIEM required a ninja DBA, and even then queries could take hours to complete — even days for full reports. Adding a new use case with additional searches and correlations required an act of Congress and a truckload of consultants. It's no surprise organizations lost patience with SIEM. So the technology was relegated to generating compliance reports and some very simple alerts, while other tools were used to do 'real' security monitoring.

But as with most other areas of security technology, SIEM has evolved. Security monitoring platforms now support a bunch of additional data types, including integration with threat intelligence, reputation services, and network packet capture. The architectures have evolved to scale more efficiently and have both built-in fancy new 'Big Data' analytics engines as well as integrating with

3rd party analytics to improve detection accuracy, even for attacks you haven't seen before. Threat intelligence is integrated into the SIEM directly, so you can look for attacks affecting other organizations, so you are ready if/when they hit you.

This *SIEM Kung Fu* paper will provide what you need to know to get the most out of your SIEM, and solve the problems you face today by increasing your capabilities (the promised Kung Fu). But first let's revisit SIEM's key use cases and what is typically available out of the box with SIEM tools.

Alerting

The original use case for SIEM was security alert reduction. IDS and firewall devices were pumping out too many alerts, and you needed a way to figure out which required attention. That worked for a little while, but then adversaries got much better, and learned to evade many of the simple correlations available with first-generation SIEM. So the key objective of using the SIEM needs to evolve to getting actionable alerts.

Many different techniques are available to detect attacks. You can hunt for anomalies that kinda-sorta look like they could be an attack, or you can perform very sophisticated analytics on a wide variety of data sources to detect known attack patterns. What you cannot do any more is depend on simple signature-based detection because modern attacks are too complicated. You need to analyze inbound network traffic (to find reconnaissance), device activity (for signs of compromise), and outbound network traffic (for command and control / botnet communications) as well. And that's a simplified view of how a multi-faceted attack works. Sophisticated attacks require sophisticated analysis to detect and verify.

> You need to analyze inbound network traffic (to find reconnaissance), device activity (for signs of compromise), and outbound network traffic (for command and control / botnet communications) as well.

Out of the box a SIEM offer a number of different patterns to detect attacks. These run the gamut from simple privilege escalation to more sophisticated botnet activity and lateral movement. Of course these built-in detections are generic and need to be tuned to your specific environment, but they can give you a head start finding malicious activity in your environment. This provides the quick win which has historically eluded many SIEM projects, and builds momentum for continued investment to add more advanced use cases.

SIEM technology has advanced to where it can find many attacks and alert you to areas of interest without a lot of integration and customization, including brute force login attempts, suspicious egress traffic, privilege escalation, critical system file changes, log source unavailability and/or volume spikes, web application misuse, and hundreds of others. But to detect advanced and targeted attacks by sophisticated adversaries, a tool can only get you so far. You need to evolve how you use security monitoring tools. You cannot just put a shiny new tool in place and expect to find advanced adversaries.

Forensics

Once you have determined an attack is under way — more accurately once you have detected one of the many attacks currently taking place in your environment — you need to investigate and determine the extent of the damage. We have documented the incident response process, especially within the context of integrating threat intelligence, and SIEM is a critical tool to aggregate data and provide a platform for search and investigation.

Out of the box a SIEM will enable responders to search through aggregated security data. You still need a talented responder to really dig into an attack and figure out what's happening. Although some tools offer visualizations to help users see anomalous activity and figure out where certain events occurred on the timeline, which certainly makes the responder more efficient. No tool can take incident response from cradle to grave. A SIEM will not be the only tool your incident responders use. But in terms of efficiently figuring out what's been compromised, the extent of the damage, and an initial damage assessment, the SIEM should be a keystone of your process. Especially given the ability of a SIEM to collect, correlate and analyze on many types of critical threat intelligence including threat feeds and user and network behavior, providing more granularity and enabling you to build a timeline of what really happened.

Compliance

Finally, SIEM remains instrumental for generating compliance reports, which are still a necessary evil to substantiate the controls you have in place. This distinctly unsexy requirement seems old hat, but you don't want to go back to preparing for your assessments by wading through reams of log printouts and assembling data in Excel, do you? SIEM tools ship with dozens of reports to show the controls in place and map them to compliance requirements so you don't need to do it manually.

> This distinctly unsexy requirement seems old hat, but you don't want to go back to preparing for your assessments by wading through reams of log printouts and assembling data in Excel, do you?

Another reason the compliance use case is still important is the skills gap every security team struggles with. If you have valuable and scarce security talent generating reports to make an auditor go away, they aren't verifying and triaging alerts, tuning detections to find new attacks, or investigating incidents. Automating as much compliance as possible remains an important SIEM use case.

As we mentioned in earlier SIEM research, a lot of these basic use cases can (and should) be implemented during a PoC process. That way you can have the vendor's sales engineers help kickstart your efforts and get you up and running with their out-of-box capabilities. But a sophisticated attacker targeting your organization will not be detected by basic SIEM correlation. Through the rest of this paper we will dig into more complicated use cases, which require pushing the boundaries of what SIEM does and how you use it.

Advanced Use Cases

Given the advance of SIEM technology, the use cases described above are very achievable. But between the availability of more packaged attack kits leveraged by better organized (and funded) adversaries, and the insider threat, you need to go well beyond what comes out of the SIEM box, or can be deployed during a one-week PoC, to detect real advanced attacks.

As we dig into more advanced use cases we will tackle how to optimize your SIEM to both detect advanced attacks on your employees, and also monitor application technology stacks to prevent attackers from getting in through holes in your application infrastructure. In the past we have grouped use cases by which adversaries they are focused on, but that never worked out very well because there are great similarities between detecting an external actor and a malicious insider. They all want your stuff. Sure, you have different inspection points depending on whether you are dealing with internal or external actors, yet in almost every successful attack, the adversary gains presence on the network and therefore technically becomes an insider.

Instead we will break up advanced use cases by *target*. The most common path nowadays is to compromise devices (typically through an employee), escalate privileges, and move laterally to achieve the mission. Alternatively an attacker might target the application stack directly from the outside, to establish a path to the data center which does not require any lateral movement. Fortunately a properly utilized SIEM can detect both.

> There are great similarities between detecting an external actor and a malicious insider. They all want your stuff.

Attacking Employees

The most prominent attack vector we see in practice today is the so-called advanced attack, which involves a multi-stage process of attacker doing reconnaissance, building specific exploits to gain presence on a network, communicating with a central location to get further instructions, moving through the target's environment to achieve the mission, and then exfiltrating the stolen data. These attackers are well funded and patient. They figure out what is going to work to achieve their mission and they do it.

These advanced adversaries don't use typical attacks that you've seen before. They evade your IPS devices and make mincemeat of the traditional endpoint protection on your devices. They compromise devices and move laterally within your organization, compromising more devices, burrowing deeper until they eventually access the information they've been tasked to steal.

Detecting this kind of attack requires a different approach, involving looking for anomalous behavior at a variety of levels within the environment. Fortunately employees (and their devices) should be reasonably predictable in what they do, which resources they access, and their daily traffic patterns.

In a typical device-centric attack an adversary follows a predictable lifecycle: perform reconnaissance, send an exploit to the device, and escalate privileges; then use that device as a base for more reconnaissance, more exploits, and to burrow further into the environment. We have spent a lot of time on how threat detection needs to evolve and how to catch these attacks using network-based telemetry.

Leveraging your SIEM to find these attacks is similar; it involves understanding the trail the adversary leaves, the resulting data you can analyze, and patterns to look for. All the clues are changes from the normal state. During any attack the adversary changes something on the device under attack. Whether it is device configuration, creating new user accounts, increasing account privileges, or just unusual traffic flows, the SIEM has access to all this data to detect attacks. The key is to set a baseline of *normal* activity when implementing the SIEM and keeping that baseline up to date over time to pinpoint the anomalies.

Initial usage of SIEM technology was entirely dependent on infrastructure logs, such as from network and security devices. That made sense because SIEM was initially deployed to stem the flow of alerts streaming in from firewalls, IDS, and other network security devices. But that very limited view of activity eventually become easy for adversaries to evade. So over the past decade many additional data sources have been integrated into SIEM to provide a much broader view of your environment.

- **Endpoint Telemetry:** Endpoint detection has become the new shiny thing. There is a ton of interest in performing forensics on endpoints, and if you are trying to figure out how the proverbial horse left the barn, endpoint telemetry is great. Another view is that devices are targeted in virtually every attack, so highly detailed data about exactly what's happening on an endpoint is critical for both incident response and detection. And this data (or the associated metadata) can be instrumental when watching for the kind of change that might indicate an active threat actor that requires immediate triage and action.

- **Threat Intelligence:** Finally, you can leverage external threat data and IP reputation to pinpoint egress network traffic headed places you can recognize are bad. Exfiltration now typically includes proprietary encryption, so you aren't likely to catch the act through content analysis — instead you need to track where data is headed. You can also use threat intelligence indicators to watch for specific new attacks in your environment, as we have discussed *ad nauseam* in our threat intelligence and security monitoring research.

- **Identity Information:** Once an adversary has a presence in your environment, they will almost inevitably go after your identity infrastructure, because that is usually the path of least resistance for access to valuable data. You need access to identity stores so you can watch for new account creation and new privilege entitlements, which are both likely to identify attacks in process.

- **Network Flows:** The next step in the attack is to move laterally within the environment, and move data around. This trail can be detected in network flows. Full packet capture provides the same information with finer granularity, in exchange for more demanding data collection and analytics.

> Keeping your baseline current and minimizing false positives are critical to using SIEM for this use case. You need ongoing effort and tuning.

As mentioned above, the key to using this data to find advanced attacks is to establish a profile of what's normal within your environment, and then look for variation. Anomaly detection remains one of the top ways to figure out when attackers are having their way in your environment. Keeping your baseline current and minimizing false positives are critical to using SIEM for this use case. You need ongoing effort and tuning. Of course no security monitoring tool *just works* — so go in with your eyes open to the amount of work required.

Multiple Data Points

Speaking of minimizing false positives, how can you? More SIEM projects fail due to alert exhaustion than for any other reason, so don't rely on any single data point to determine that an alert is legitimate and demands investigation. Reduction of false positives is even more critical because of the skills gap which continues to flummox security professionals. Using a SIEM you can link together seemingly disconnected data sources to validate alerts and make sure the alarm is sounded only when it should be.

But what does that look like in practice? You need to make sure a variety of conditions are matched before an alert fires. And increase the urgency of alerts which trigger more conditions. A simplified example illustrates what you can do with the SIEM you likely already have.

1. *Look for device changes:* If a device suddenly registers a bunch of new system files installed, and you aren't in the middle of a patch cycle, there may be something going on. Is that enough to pull the alarm? Probably not, so you'll want to look for more attack indications.

2. *Track identity:* Next you see a bunch of new accounts appear on the device, and then the domain controller comes under attack as a means to dig deeper within your environment. Once the domain controller falls, it is pretty much game over, because the adversary can then set up new accounts and change entitlements; thus getting alerts on domain controller attacks is essential.

3. *Look for internal reconnaissance:* Finally you'll see the compromised device scanning everything else on the network, both so the attacker can gain his/her bearings, and also for additional devices to compromise. Traffic on internal network segments should be pretty predictable, so variations from typical traffic flows usually indicate something funky.

But can any of these individual data points conclusively indicate an attack? There's no smoking gun. But if you see a cluster of multiple indicators that's not great and definitely warrants investigation.

The current generation of SIEMs come with a variety of rules or policies to look for common attack patterns out of the box. These platforms also have more advanced analytics that can identify attack patterns happening in the environment by taking a baseline and then looking for anomalous activity.

These capabilities are helpful for getting started. Although to really take advantage of the SIEM and find uncommon attacks, you'll need to customize the tool and the analytics for your environment. When necessary, advanced adversaries will use malware and social engineering tactics you haven't seen before and therefore you won't have a detection built into the SIEM. In this case, you'll use the behavioral and data analytics within the SIEM to identify these unknown attacks, refining your thresholds for alerts; and increasing accuracy and reducing false alarms.

Application Stack Attacks

We alluded to this above, but to us an "application stack attack" is not just a cute rhyme, but the way a sophisticated adversary takes advantage of weaknesses within an application or another part of an application stack to gain a foothold in your environment to access data of interest. There are a number of application stack data sources you can pump into a SIEM to look for application attacks. These include:

- **Machine Data:** The first step in monitoring applications is to instrument it to generate "machine data". This could be information on different transaction types or login failures, search activity, or almost anything that can be compromised by an attacker. Determining how and where to instrument an application involves threat modeling to make sure the necessary hooks are built into the app. The good news is that as more and more applications move to SaaS environments, a lot of this instrumentation is there from the start. But with SaaS you get what you get, and that doesn't include much influence on which information is available.

- **APIs:** Applications are increasingly composed of a variety of components, residing in a variety of different places (both inside and outside your environment), so watching API traffic has become essential. We have researched API security, so refer back to that paper for specifics about authentication and authorizing specific API calls. You will want to track API usage and activity to profile *normal* activity for the application, and then start looking for anomalies.

- **Database Tier:** This last part of the application stack is where the valuable stuff lives. Once an attacker has presence in the database tier, it is usually trivial to access other database tables and reach the stuff they are looking for. So ingest any database activity logs or monitors available, and watch for triggers.

Each application is unique (like a snowflake!), so you won't be able to get prebuilt rules and policies from your SIEM provider. You need to look at each application to monitor and profile it, building rules and tuning thresholds for each application. This is why most organizations don't monitor their applications to any significant degree… And also why they miss attacks which don't involve traditional malware or obvious attack patterns.

Each application is unique (like a snowflake!), so you won't be able to get prebuilt rules and policies from your SIEM provider. You need to look at each application to monitor and profile it, building rules and tuning thresholds for each application.

Developer Resistance

Collecting sufficient *machine data* from applications isn't something most developers get excited about. Applications have historically not been built with instrumentation in mind, and retrofitting instrumentation into an app is more delicate plumbing than it is designing cool new features. We all know how much developers love to update plumbing. You may need to call for senior management air cover, in the form of a mandate, to get the instrumentation you need into the application. You can only request air support a limited number of times, so make sure the application is sufficiently important first.

More good news: as new applications are deployed using modern development techniques (including DevOps and Continuous Deployment), security is increasingly built into the stack at a fundamental level. Once the right instrumentation is in the stack, you can stop fighting to retrofit it.

Additional Leverage

We like to focus on attack detection relative to a SIEM because that's the most pressing need given the rise of sophisticated adversaries. But that's not the only area where a SIEM is going to provide value for your security program. Keep in mind the value of the SIEM is streamlining audit preparation, which is certainly a lot less fun than chasing attackers but still a pretty important part of your job description. After over a decade of helping customers prepare for audits, the SIEM has all the compliance reports you'll likely need already built-in. These reports are comprehensive enough that you don't need to do major customization and polishing of the reports. So the days of taking the SIEM report and having to do heavy massaging in Excel is over. Though we know you'll miss it.

Another key value add is the ability for a SIEM also to provide forensics and investigation capabilities to help analysts triage attacks and understand potential damage during an incident response. The SIEM will give the responder pre-defined aggregated views of user, event, or server activity, pulling this information together automagically. In the early days of SIEM, you are manually looking through log files or enriching existing records to fill out the full activity but that's no longer the case.

The SIEM allows you not only to gather the data and link it together, but also to make sense of the information in a structured fashion to accelerate identification of the root cause of any attack. Built-in usage profiles for activity baselining are essential, the ability to use threat intelligence to provide clues as to where to look for attacker activity, as well as advanced query facilities for quick and easy ad hoc analysis for response.

Getting Started and Sustaining Value

Many failed SIEM projects over the past 10 years have not been *technology* failures. More often organizations stumble over a failure to understand the time and resources needed to get value from the SIEM in early deployments, and the amount of ongoing effort required to keep it current and tuned. So a large part of SIEM Kung Fu is just making sure you have the people and process in place to leverage the technology effectively and sustainably.

Getting Started

As a matter of practice you should be focused on getting quick value out of any new technology investment, and SIEM is no exception. Even if you have had the technology in place for years, it's useful to take a fresh look at the implementation to see if you missed any low-hanging fruit that's there for the taking. Let's assume you already have the system up and running, are aggregating log and event sources (including things like vulnerability data and network flows), and have already implemented some out-of-the-box policies. You already have the system in place — you are just underutilizing it.

Adversaries

For a fresh look at SIEM we recommend you start with adversaries. We described adversary analysis in detail in the CISO's Guide to Advanced Attackers (PDF). Start by determining who is most likely to attempt to compromise your environment, and the likely objective (mission). Then profile potential adversaries to determine the groups most likely to attack you. At that point you can get a feel for the most likely Tactics, Techniques, and Procedures (TTPs) for adversaries to use. This information typically comes from a commercial threat intelligence service, although some information sharing groups (typically associated with industries, like financial services and retail) also offer technical indicators to focus on. From an integration standpoint, this threat intel can be easily integrated using standards such as STIX/TAXII.

Armed with these indicators you engage your SIEM to search for them. This is a form of *hunting*, which we will detail later in this paper, and you may well find evidence of active threat actors in your environment. This isn't a great outcome for your organization, but it nicely proves the value of security monitoring.

At that point you can triage the alerts you have received from SIEM searches to figure out whether you are dealing with false positives or a full-blown incident. We suggest you start by searching for the attacks of your most likely adversaries. Odds are you'll find plenty of things if you search for anything and everything. An initial focus on adversaries restricts your search to attack patterns most likely to be used against you.

Two Tracks

Once you have picked the low-hanging fruit from adversary analysis your focus shifts toward putting advanced use cases into a systematic process that is consistent and repeatable. Let's break up the world into two main categories of SIEM operations to describe the different usage models: reactive and proactive.

Reactive

Reactive usage of SIEM should be familiar because that's how most security teams function. It's the alert/triage/respond cycle. The SIEM fires an alert, your tier 1 analyst figures out whether it's legitimate, and then you decide how to respond — typically via escalation to tier 2 if it's not just noise. You can do a lot to refine this process so even your reactions are more efficient. Here are a few tips:

1. **Leverage Threat Intel:** As we described earlier, you can benefit from the misfortune of others by integrating threat intelligence into your SIEM searches. If you see evidence of a recent attack pattern (provided by threat intel) within your environment, you can get ahead of it. We described this in our Leveraging Threat Intel in Security Monitoring paper. Use this method of security monitoring — it works.

2. **User Behavioral Analytics (UBA):** You can gain insight into the severity of an attack by tracking it back to user activity. This involves monitoring activity (and establishing the baselines/profiles described previously) by device, but also aggregating data and profiling activity for *individuals*. For example, instead of just monitoring the CEO's computer, tablet, and smartphone independently, you can look at all three devices to establish a broader activity profile. Then if you see any of her devices acting outside that baseline, that would trigger an alert you can triage and investigate.

> We (still) don't believe you can get ahead of threats or detect zero-day attacks, or any other such security marketing nonsense. What you can do is shorten the window between when you are attacked and when you know about it.

3. **Insider Threat:** You can also optimize some of your SIEM rules around insiders. During many attacks an adversary eventually gains a foothold in your environment and becomes an insider. You can optimize your SIEM rules to look for activity specifically targeting things you know would be valuable to insiders, such as sensitive data (both structured and unstructured). UBA is also useful here because you are profiling an insider and can watch for strange reconnaissance, or possibly moving an uncharacteristically large amount of data.

4. **Refine Analytics:** As you use the SIEM, you should have a process in place to continually refine the analytics to provide more accurate and actionable alerts. This involves making the investment to actually dig into false positives and figure out how the rules and analytics need to change to minimize the likelihood that the alert would fire again under the same scenario. You can and should be tuning the thresholds and adding additional data sources (and use cases) over time to make sure you are getting maximum value from the analytics engines already built into the SIEM.

5. **Threat Modeling:** Yes, advanced SIEM users still work through the process of looking at specific, high-value technology assets and figuring out the best ways to compromise them. This isn't about understanding what attacks adversaries are using or other indicators available from threat intel feeds. This is about understanding how you (or a penetration tester) would attack the asset/technology stack/application and then building SIEM rules from those attack patterns to detect when that critical assets is being targeted.

Keep in mind that you need to consistently look at your SIEM ruleset, add new attack patterns/use cases, and prune rules that are no longer relevant. The size of your ruleset correlates to the performance and responsiveness of your SIEM, so you need to balance looking for everything (and crushing the system) against your chance of missing something.

This is a key part of the ongoing maintenance required to keep your SIEM relevant and valuable. Whether you get new rules from your SIEM or threat intelligence vendor, drinking buddies, or conferences, new rules require time to refine thresholds and determine relevance to your organization. So we reiterate that SIEM is not a "set it and forget it" technology — no security analytics tool is. Anyone telling you different is selling you a bill of goods.

Proactive

Before we dive into proactivity we need to spend a minute on our soapbox about the general idea of "getting ahead of the threat." We (still) don't believe you can get ahead of threats or detect zero-day attacks, or any other such security marketing nonsense. What you *can* do is shorten the window between when you are attacked and *when you know about it*. That is the main objective of SIEM Kung Fu: to shrink this window by however you can.

The reactive approach is to set SIEM rules to fire alerts based on certain conditions, and then react to them. The *proactive* approach is to task a human with finding **actual** attacks in process, which haven't triggered an alert. These folks like to be called **hunters** — probably because that sounds much better than "Senior SOC analyst".

Why wouldn't the SIEM alert have fired already? Maybe the attack is just beginning. Maybe the adversary is only performing recon, and mostly hiding to evade detection. Whatever the cause, the rules you set in the SIEM haven't triggered yet, and a skilled human may be able to find the attack before the monitoring tools you've deployed.

The hunter's tool set is more about threat intel, effective search and analytics, and a lot of instinct for what attackers will do, than a flexible SIEM rules engine. For example a hunter might see that a recent business partnership your company announced has irritated factions in Eastern Europe. So the hunter performs a little research and finds a new method of compromising a recent version of Windows by gaining kernel access and then replacing system files — which is currently being used by attackers in that region.

Then the hunter searches to see whether egress traffic is headed to C&C channels associated with them, and also searches your endpoint telemetry for instances where those system files were changed recently. Of course you could set a rule to look for this activity moving forward, but that's too late to get this initial attack. The hunter is able to mine existing security data for that set of conditions to see if an attack is in progress.

Or a hunter might go to a security conference and learn about a new technique to overflow memory. After playing around with it in your lab, the hunter knows what to look for on each endpoint. Then a search can be initiated for that activity, even though there hasn't been evidence of that technique in the wild yet. Hunters have great leeway to follow their instincts, and SIEM tools need to offer flexible (and fast) enough search to find strings and pull them.

SIEM Kung Fu for hunters is about giving them a platform to do their job. They are skilled professionals, with their own tools for when they really want to dig into a device or attack. But a SIEM can help narrow their focus to devices that require more investigation, and provide a means to analyze activity patterns for clues to which threats are active.

Whether you are implementing a set of SIEM rules to react to attacks, or giving a set of hunters the ability to identify potential compromises in your environment, your security monitoring platform can be leveraged to enable faster detection and triage. And when you are racing an active adversary, time is not your friend.

Summary

Even though security monitoring (and SIEM specifically) has been much maligned, you don't really have a choice but to use technologies to detect attacks and provide a means for your security team members to find adversaries in your environment. SIEM Kung Fu is about making the most of a product or service you already likely have, by looking at advanced use cases to get more value, and improving processes to optimize your SIEM.

To summarize, here are our key tips:

- **Focus on the Adversary:** Your attack surface is infinite, so you need to figure out a way to narrow your aperture on alerts, because you can't handle everything. Understanding your most likely attackers and the typical attacks they use gives you a head start on optimizing your scarce SOC resources.

- **Eyes Wide Open:** Security monitoring is still work, and it always will be. Your infrastructure changes. Your applications change. Attackers change. You *can* just set up your SIEM and forget about it… if you want another war story about how an expensive SIEM project failed completely. Make sure you have the resources to both triage alerts and keep the platform current, or don't even waste your time and money on SIEM.

- **Master Response:** Alerts happen. The question is how efficiently and effectively you handle them. Make sure your front line can determine which alerts are real, and has the information required for effective escalation. And use information from every response to make the alert better. The best thing about failure (or partial success) is that you can feed it into an effective feedback loop.

- **A Hunting We Will Go:** If you have the maturity and resources, task some talented security analysts with hunting adversaries in your environment. They'll need tools to do so, and they can start with the SIEM to help figure out which devices should be scrutinized. This is an emerging but increasingly important use of security monitoring technology.

As Caine said in the legendary *Kung Fu* TV series: *"I do not seek answers, but rather to understand the question."* A SIEM won't give you definite answers about what is being attacked, but it can help make sure you understand what's happening in your environment, and give you information to ask the next round of questions when you are investigating an alert.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Mike Rothman, Analyst and President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security — such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who "knows where the bodies are buried" in the space.

Starting his career as a programmer and networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure. After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published The Pragmatic CSO <http://www.pragmaticcso.com/> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services. Our services include:

- **Primary research publishing**: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.

- **Research products and strategic advisory services for end users**: Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.

- **Retainer services for vendors**: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.

- **External speaking and editorial**: Securosis analysts frequently speak at industry events, give online presentations, and write and speak for a variety of publications and media.

- **Other expert services**: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <http://securosis.com/>.