Security Monitoring State of the Union

Version 1.4
Released: July 10, 2018

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 McAfee.

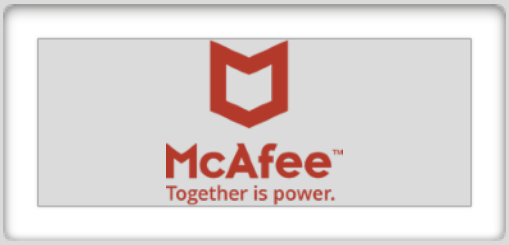

www.mcafee.com

McAfee is the device-to-cloud cybersecurity company. Inspired by the power of working together, McAfee creates business and consumer solutions that make our world a safer place. By building solutions that work with other companies' products, McAfee helps businesses orchestrate cyber environments that are truly integrated, in which protection, detection, and correction of threats happen simultaneously and collaboratively. By protecting consumers across all their devices, McAfee secures their digital lifestyle at home and away. By working with other security players, McAfee is leading the effort to unite against cybercriminals for the benefit of all.

Copyright

Security Monitoring State of the Union

Table of Contents

Revisiting the Team of Rivals

Things change. That is the only certainty in security today. Back when we wrote Security Analytics Team of Rivals, SIEM and Security Analytics offerings were different and did not really overlap. We were figuring out how they could coexist, instead of choosing one over the other. But nowadays the overlap is significant, so we see existing SIEM players bundling in security analytics capabilities, and security analytics players positioning their products as next-generation SIEM.

As usual customers are caught in the middle, trying to figure out what is truth and what is marketing puffery. So Securosis is here to help you sort it all out. This paper offers some perspective on which use cases make sense for SIEM, and where security analytics makes a difference.

Nowadays the overlap is significant, so we see existing SIEM players bundling in security analytics capabilities, and security analytics players positioning their products as next-generation SIEM.

Revisiting Security Analytics

Security analytics remains a fairly perplexing market, because almost every company providing security products and/or services claims to perform some kind of analytics. To level-set let's revisit how we defined Security Analytics (SA) in the Team of Rivals paper. A SA tool should offer:

- **Data Aggregation:** It's impossible to analyze without data. Of course there is debate whether a security analytics tool needs to gather its own data, or can just integrate with an existing security data repository such as your SIEM.
- **Math:** We joke a lot that math is the hottest thing in security lately, especially given how early SIEM correlation and IDS analysis were based on math too. But this new math is different, using advanced algorithms and modern data management to find patterns within data volumes which were unimaginable 15 years ago. The key difference is that you no longer need to know what you are looking for to find useful patterns, a critical limitation of current SIEM. Modern algorithms can help you spot unknown unknowns. Looking only for known and profiled attacks (signatures) is clearly a failed strategy.
- **Alerts:** These are the main output from security analytics, so they should be prioritized by importance to your business.

- **Drill down:** Once an alert fires an analyst needs to dig into the details, both for validation and to determine the appropriate response. So analytics tools must be able to drill down and provide additional detail to facilitate response.
- **Learn:** This is the tuning process, and any offering needs a strong feedback loop between responders and the folks running it. You need to refine analytics to minimize false positives and wasted time.
- **Evolve:** Finally the tool must improve because adversaries are not static. This requires a threat intelligence research team at your security analytics provider constantly looking for new categories of attacks and ways to identify them.

These are the requirements for an SA tool. But over the past year we have seen these capabilities show up not only in security analytics tools, but also in more traditional SIEM products. Though to be clear, "traditional SIEM" is really a misnomer — none of the current market leaders are built on 2003-era RDBMS technology, or sitting still waiting to be replaced by new entrants with advanced algorithms.

In this paper we discuss how well each tool matches up against the emerging use cases (many of which we discussed in Evolving to Security Decision Support), and how technologies such as the cloud and IoT impact your security monitoring strategy and toolset.

Wherefore Art Thou, Team of Rivals?

The lines between SIEM and security analytics blurred as we predicted, so what should we expect vendors to do? First, all vendor collaboration and agreements, such as between SIEM and security analytics vendors, are deals of convenience to solve short-term problems of a SIEM vendor not having a good analytics story or an analytics vendor not having enough market presence to maintain growth. The risk to customers is that buying a bundled SA solution with your SIEM can be problematic if your vendor acquires a different technology and eventually forces a migration to their in-house solution. This underlies the challenge of vendor selection as markets shift and collapse. At a minimum, make sure to get contractural assurances from the SIEM vendor about future support of the bundled solution.

First, all vendor collaboration and agreements, such as between SIEM and security analytics vendors, are deals of convenience to solve short-term problems

We are confident that the security monitoring market will play out as follows over the short term:

1. SIEM players will offer a basic level of security analytics.
2. Security analytics players will spend a bunch of time filling out SIEM reporting and visualization features sets to go after replacement deals.

3. Customers will be confused and unsure whether they need SIEM, security analytics, or both.

But that path ends with confused practitioners, and that's not where we want to be. So let's break the short-term reality down a couple different ways.

Short-Term Plan

The solution you choose for security monitoring should accommodate emerging use cases you will need to handle, along with questions you will need to answer about your security posture over time. But you probably already have security monitoring technology installed. Moving forward requires clear understanding of how your current environment influences your path forward.

SIEM-centric

If you are a large company or under any kind of compliance/regulatory oversight — or both — you should be familiar with SIEM products and services, because you've been using them for over a decade. Odds are you have selected and implemented multiple SIEM solutions, so you understand what SIEM does well… and less well. You have no choice but to compensate for its shortcomings because you aren't in a position to shut it off or move to a different platform.

Your objective is to get as much value out of your existing SIEM as you can. Your path is straightforward. First focus on refining the alerts coming out of the system to increase signal and decrease noise, and focus your team on triaging and investigating real attacks. Then integrate threat intelligence for a sense of the attacks happening to other organizations, which may enable you to respond faster to emerging threats as you face them.

But of course your SIEM vendor is certainly trying to add enough capability that you don't need to think about another platform, much less deploy a new tool.

Then add new capabilities the vendor is bundling into the system (either through in-house development or OEM), such as User Behavioral Analytics (UBA) and tracking insider threats. You are basically buying time by leveraging the platform you already have more effectively, letting the battle between SIEM and security analytics play out a bit before choosing a side. But of course your SIEM vendor is certainly trying to add enough capability that you don't need to think about another platform. There may be different underlying architectures to deliver both functions, but as long as you don't have to manage it that's not a real problem.

SA-centric

Perhaps you face sophisticated adversaries and/or have a mature security program, and have already decided security analytics is your future strategic platform, and current SIEM vendors not going to get there. So you already made your bet but your strategic platform isn't feature complete compared to your old SIEM. Then your focus needs to be on addressing shortcomings of your SA tool.

The first and most important gap tends to be compliance reporting. SIEM tools have ridden this use case to great success over the past decade, so many teams depend on these reports to prepare for increasingly frequent audits and assessments. Making sure you can generate the reports you need may thus be first on your list. If the reports aren't built in you are likely exporting data from the analytics tool and going back to the future. You know, the good old days when Excel was your compliance preparation tool. The good news is that reporting tools are much better and more intuitive today.

After making sure you address the compliance use case you can look at additional tools for response and forensics, because SA functionality in those areas is not yet as mature or complete as SIEM. There are a bunch of available options, specifically coming from next-generation endpoint protection vendors who provide far better response capabilities.

The best advice we can give is probably to get comfortable with "good enough". In the areas where SA is deficient compared to SIEM.

But the best advice we can give is probably to get comfortable with "good enough". In the areas where SA is deficient compared to SIEM think about the capabilities you need and then start pushing your SA vendor hard toward feature parity. All the standalone security analytics vendors are actively targeting the larger SIEM market and working to quickly address their shortcomings.

Both

Some organizations have the luxury of choosing. You've implemented both, continuing to rely on SIEM to detect standard attacks you know to look for and generate compliance reports. You deployed SA to profile activity in your environment and highlight potentially malicious actions which may represent adversary activity and warrant further investigation.

If this is you we recommend you continue holding with the tools you already have until you determine which will be your strategic platform. Then invest in your migration path.

Civil War

The State of the SecMon Union is: civil war is imminent. You need to pick a strategic platform because it doesn't really make sense to run both a SIEM and a security analytics platform from different vendors forever. You need staff trained on both, and the ability to work through contradictory info when tools reach different conclusions — even from the same data. And you need to pay for both tools. Details, details. As we saw in other security markets (including network and endpoint protection), the next generation and the incumbents all deliver a common set of capabilities and can't really sufficiently differentiate themselves from their competition, so the future of security monitoring is at stake.

Focus on Use Cases

The overlap between SIEM and security analytics has passed a point of no return. With a civil war brewing, our key goal is to determine which will be your strategic platform for security monitoring. This requires you to shut out the noise of fancy analytics and colorful visualizations, to focus on the problems you are trying to solve now, with an eye out for how they will evolve in the future. That means getting back to use cases. The use cases for security monitoring tend to fall into three major buckets:

1. Security Alerts
2. Forensics and Response
3. Compliance Reporting

With a civil war brewing, our key goal is to determine which will be your strategic platform for security monitoring.

Let's work through each of these to make sure you have a clear handle on success today, and how it will change in the future. After we work through the use cases we'll cover pros and cons of how each side addresses them.

Security Alerts

Traditional SIEM was based on looking for patterns you recognized as attacks. You couldn't detect things you didn't yet recognize, and keeping pace with dynamic attacks was always a challenge. So many customers didn't receive the value they needed. In response a new generation of security analytics products appeared — applying advanced mathematical techniques to security data, identifying and analyzing anomalous activity, giving customers hope that they would be able to detect attacks not covered by existing rules.

To have a handle on success today any security monitoring platform needs the ability to detect and alert on the following attacks:

- **Commodity Malware:** These are known attacks, often available from a Metasploit module, enabling even the least sophisticated attackers to use them. Although not sexy, these attacks are still prevalent because adversaries don't waste advanced attacks unless they need to.
- **Advanced Malware:** Advanced attacks are customized and effectively unique, so you are very unlikely to detect them using a known pattern already in your security monitoring platform.

- **User Behavior Analysis:** Another way to pinpoint attacks is to look for strange user activity. At some point in an attack a device will be compromised and act in an anomalous way, which is your opportunity to detect it.
- **Insider Threat Detection:** The last use case we'll describe overlaps with UBA because it is about figuring out if you have a malicious insider stealing data or causing damage. The insider tends to be a user, thus the overlap with UBA. Yet this use case is less about malware — the user is already within the perimeter — and more about profiling employee behavior and looking for signs of malicious intent, such as reconnaissance and exfiltration.

But the telemetry used to drive security monitoring today is much broader than in the past. The first generation technology — SIEM — was largely driven by log data, and sometimes network flows and vulnerability information. But with the disruption of the cloud and mobility, a much broader data set is needed. For instance you need to factor your SaaS applications into your security monitoring. You likely have IoT devices as well, whether work floor sensors or multi-function printers with vulnerable operating systems —they *all* need to be watched. Finally, mobile endpoints have become full participants in the technology ecosystem, so gathering their telemetry has become an important facet of monitoring.

The fact that corporate data lies both inside the perimeter and across a bunch of SaaS services and mobile devices makes it much harder to build a comprehensive security monitoring environment.

Aside from the main attack vectors, the fact that corporate data lies both inside the perimeter and across a bunch of SaaS services and mobile devices makes it much harder to build a comprehensive security monitoring environment. We described this need for enterprise visibility in our Security Decision Support paper.

Forensics and Response

This use case comes into play after an attack, when the organization is trying to figure out what happened and assess damage. The key functions required for response tend to be sophisticated search and the ability to drill down into an attack quickly and efficiently. Skilled responders are very scarce so they need to leverage available technology to streamline their efforts.

But given the scarcity of responders a heavy dose of enrichment (adding threat intel to case files) is needed, and even attack remediation must be increasingly automated. So it's not just about equipping responders — you need to help scale their activity.

Compliance Reporting

This use case is primarily focused on providing the information needed to make an auditor go away as quickly as possible, with minimal report customization and tuning. Every organization needs to deal with different compliance and regulatory hierarchies, as well as internal control reporting, so success entails having a tool map specific controls to regulations and substantiate that they are actually in place and operational.

Sounds simple, right? It is until you spend two days in Excel cleaning up the stuff that came from your tool. You could pay an assessor to go through all your stuff and make sense of things, but that might not be the best use of your or their time — and you could not ensure they reach the right conclusions.

As we look at the future, compliance reporting won't change that much. But the data you need to feed into a platform to generate your substantiation will expand substantially. It's all about visibility, as mentioned above. As your organization embraces cloud computing and mobility, you will need to make sure you have logs and appropriate telemetry from your controls to substantiate your security activity.

Assessing the Combatants

Given these use cases and what's needed for the future, we need a general assessment of SIEM and security analytics. As we've mentioned, there is already significant functional overlap between SIEM and security analytics products. The overlap will only grow, until there is no functional difference between SIEM and security analytics. It will just be security monitoring — whatever it's called.

Given this overlap we need some way to distinguish the players. The underlying means of analyzing data provides a useful way to distinguish an incumbent from a new entrant. To drill down, a system built on a rules engine is SIEM, while a tool based on a (Big Data-centric) analytical platform is a new entrant. But it gets a bit murky, because no current tool really only uses rules, and every analytics product has a rule option.

Nomenclature aside, returning to the attacks above, here is how each type of security monitor handles them:

- **Commodity Malware:** You know what these attacks are so traditional rules-based alerting works well. As long as you keep the rules current. A bit counter-intuitively, new entrants can have trouble with these attacks, because they don't always show up as anomalous behavior. This is analogous to why endpoint AV signatures are still useful: more sophisticated behavioral models can still miss recognizable old attacks.

- **Advanced Attacks:** The new entrants begin to shine when handling advanced attacks because they don't need to look for specific attacks as rules-based systems do. The behavioral and machine learning models underlying analytics engines do need to be updated periodically, but if you are worried about an advanced adversary a rules-based system definitely won't suffice.
- **User Behavioral Analysis:** UBA requires integration with the corporate identity store so you can associate specific devices with users, but leverages modern analytics to detect advanced attacks. This makes UBA difficult with a traditional SIEM.
- **Insider Threat Detection:** The major difference between UBA and insider threat detection is in specificity to an organization. UBA tends to look for generically anomalous behavior, while insider threat tools are tuned to the inner workings of a specific organization. This tends to involve integration of physical security and HR systems because there are many triggers for recognizing malicious employee intent.

To net it out, for the security alerting use case, beyond commodity malware a rules-based system is limited. Detecting advanced attacks and profiling users (either for UBA or an insider threat use case) requires higher-level analytics. So any tool you are considering from here on needs to provide broader analytics.

Thinking about forensics and response, the maturity of incumbents tends to make their tools more capable for the response use case because they collect broader security data and have better search and drill-down experiences. But the gap is closing rapidly as new entrants focus on feature parity with incumbent SIEM vendors.

Finally, for compliance reporting, incumbents have been cranking out auditor reports for well over a decade and have all those bases covered. But this is another area where the gap is closing. First because compliance reporting is fairly mechanical, so once the security control to mandate mapping is integrated into the product, the reports more or less pop out. No, it's not that simple, and there is still need for customers to customize reports, but it isn't rocket surgery.

Where does this leave us as the civil war continues to smolder? Right back where we started. The use cases continue to evolve, as do the tools.

Refreshing Requirements

Our next step is to translate use cases into requirements for your strategic security monitoring platform. Now that you have a handle on the problems you need to solve, what capabilities do you need to address them? Part of that discussion is inevitably about what you don't get from your existing security monitoring approach — this research would be pointless if existing tools were all sufficient.

Visibility

Maintaining adequate visibility across all the moving pieces in your environment is getting harder. When we boil it down to a set of requirements it looks like this:

- **Aggregate Existing Security Data:** We could have called this requirement "same as it ever was" — all your security controls *still* generate a bunch of data you need to collect. Kind of like the stuff you were gathering in the early days of SEM (Security Event Management) or log management 15 years ago. Given all the other things on your plate, what you don't want is to worry about integrating your security devices, or how to scale a solution up to your environment. To be clear, security data aggregation has commoditized, so this is table stakes for whatever solution you consider.

Amazingly enough, when you aggregate a bunch of security data, you need to manage it. So data management is still a thing.

- **Data Management:** Amazingly enough, when you aggregate a bunch of security data, you need to manage it. So data management is still a thing. We don't need to go back to SIEM 101 but aggregating, normalizing, reducing, and archiving security data is a core function of any security monitoring platform — regardless of whether it started life as SIEM or a security analytics product. One thing to consider (which we will dig into below under procurement) is the cost of storage, because emerging cloud-based pricing can be painful when you significantly increase the amount of security data.

- **Embracing New Data Sources:** In the old days the common complaint was that vendors did not support all the devices (security, networking, and computing) in organizations. As explained above, that's less of an issue now. But consuming and integrating cloud monitoring, threat intelligence, business context (such as asset information and user profiles), and non-syslog events all drive a clear need for streamlined integration to leverage additional data faster.

Seeing into the Cloud

When considering the future requirements of a security monitoring platform, you need to understand how it will track what's happening in the cloud, because the cloud seems to be here to stay. Start with API support, the *lingua franca* of the cloud. Any platform you choose must be able to make API calls to the services you use and/or pull information and alerts from a CASB (Cloud Access Security Broker) to track SaaS usage within your organization.

When considering the future requirements of a security monitoring platform, you need to understand how it will track what's happening in the cloud, because the cloud seems to be here to stay.

You'll also want to understand the architecture required to gather data from multiple cloud sources. You definitely use multiple SaaS services, and likely have many IaaS (Infrastructure as a Service) accounts to consider, possibly with multiple providers. All these environments generate data which must be analyzed for security impact, so you should define a standard cloud logging and monitoring approach, and likely centralize aggregation of cloud security data. You should also consider how cloud monitoring integrates with your on-premise solution. For more detail see our paper on Monitoring the Hybrid Cloud.

There are various considerations for different cloud environments.

- **Private Cloud/Virtualized Data Center:** There are differences between monitoring your existing data center and a highly virtualized environment. You can tap the physical network within your data center for visibility. But for the abstracted layer above that — which contains virtualized networks, servers, and storage — you need proper access and instrumentation in your cloud to observe what happens within virtual devices. You can also route network traffic within your private cloud through an inspection point, but the architectural flexibility cost is substantial. The good news is that security monitoring platforms can now generally monitor virtual environments by installing sensors within the private cloud.
- **IaaS:** The biggest and most obvious challenge in monitoring IaaS is reduced visibility because you don't control the physical stack. You are largely restricted to your cloud service provider's logs. IaaS vendors abstract the network, limiting your ability to see network traffic and capture network packets. You can run all network traffic through a cloud-based choke

point for collection, regaining a faint taste of the visibility available inside your own data center, but once again that sacrifices much of the architectural flexibility of the cloud. You also need to figure out where to aggregate and analyze logs collected from both the cloud service and individual instances. These decisions depend on a number of factors — including where your technology stacks run, the kinds of analyses you perform, and the expertise you have available on staff.

- **SaaS:** You see what your SaaS provider shows you, but not much more. Most SaaS vendors provide security logs. They don't provide visibility into the vendor's technology stack, but you are able to track your employees' activity within their service — including administrative changes, record modifications, login history, and increasingly application activity. You can also pull information from a CASB which polls SaaS APIs and analyzes egress logs for further detail.

Threat Detection

The key to threat detection in this new world is the ability to detect attacks you know about (rules-based), attacks you haven't seen yet but someone else has (via threat intelligence), and unknown attacks which cause anomalous user or device activity (security analytics). The patterns you are trying to detect can be pretty much anything — including command and control, fraud, system misuse, malicious insiders, reconnaissance, and even data exfiltration. So there is no lack of stuff to look for — the question is what you need to detect it.

- **Rules:** You can't ditch your rules — don't even think about it. Actually you *could* — but you'd likely miss a bunch of attacks you should catch because you already know their patterns. The behavioral models are focused on stuff you *don't* know about, not optimized to find known bad stuff. As with endpoint protection, rules (signatures) are not an either/or proposition. If you already know about an attack, shame on you if you miss it.

- **Threat Intelligence:** For attacks you hadn't seen yet, in the old days you were out of luck. But today there is a decent chance someone else has been attacked by them, which is where threat intelligence comes into play. You pump a threat feed into your security monitoring platform, in hopes of recognizing them when they show up on your doorstep. Make sure you can distinguish threat intelligence alerts so you can track the effectiveness of your feeds, to determine value and ensure you don't increase alert noise.

- **Security Analytics:** The final approach you need to consider is based on advanced math. You'll hear terms like security Big Data, machine learning, and the generic "It's fancy math, trust us!" to describe these techniques. Regardless of description, security analytics entails profiling devices, networks, and applications to baseline normal activity — then looking for deviations from that profile as indicators of malicious activity. It's very difficult to discern the differences between one analytics approach and another, so understanding what will work for your organization requires actually trying them out. We'll discuss procurement later in this paper.

After a few years using security monitoring technology, hopefully by this point you realize this isn't (and likely will never be) a set-and-forget scenario. You'll need to keep the system current and tune it.

After a few years using security monitoring technology, hopefully by this point you realize this isn't (and likely will never be) a set-and-forget scenario. You'll need to keep the system current and tune it, because not only are adversaries constantly changing and evolving their tactics, but your environment is constantly changing, requiring ongoing maintenance.

Build a learning and tuning step into your operational processes to ensure you improve detection based on false positives (alerts which weren't real attacks) and negatives (attacks you missed). A feedback loop is essential.

Forensics and Response

Obviously you cannot prevent every attack, and even if you do fire an alert about a specific attack your Security Ops team could miss it. So your security monitoring platform will also play a major role in your incident response process. The challenge is less gathering data or trying to link it together, and more how to make sense of the information at your disposal in a structured fashion — which is what you need to accelerate identification of the root cause of attacks. As we discussed in our Future of Security Operations paper, many aspects of the response process can be automated, so ensuring support for that is essential.

Key capabilities include:

- **Search:** Modern attacks do not limit themselves to a single machine, so you need to quickly figure out how many devices have been attacked as part of a broader campaign. Some of that takes place during validation/triage as the analyst pivots, but determining the breadth of an attack requires them to search the entire environment for attack indicators, typically via metadata.
    - **Natural Language/Cognitive Search:** An emerging capability is the use of natural language search terms instead of arcane Boolean operators. This helps less sophisticated analysts be more productive without having to learn a new language.
    - **Retrospective Search:** Responders often have a sense of what caused an attack, so they should be able to search through historical security data to find activity which did not trigger an alert because it wasn't recognized at the time.
- **Case Management:** The objective is to make each analyst as effective and efficient as possible, so you should have a place to store all information related to each incident. This includes enrichment data from threat intel and other artifacts gathered during validation. This should also feed into a broader incident response platform if the forensics/response team uses one.

- **Visualization:** To reliably and quickly validate an alert, it is very helpful to see a timeline of all activity related to an incident. That way you can see what actually happened across devices and get a broader understanding of the issue's impact. An analyst can take a quick look at the timeline to figure out what requires further investigation. Visualization can present all sorts of information, but be wary of overcomplicating the console. It is definitely possible to present too much.

- **Drill down:** Once an analyst has figured out which activity in the timeline raises concerns, they drill into it. At each stage of the attack they may find other things to investigate, so the ability to jump between events and devices helps identify the root causes of attacks quickly. There is also a decision to be made regarding how much data to collect and have at the ready. Obviously the more granular the available telemetry, the more accurate the validation and root cause analysis. But with increasingly granular metadata available you might not need full capture of networks or endpoints.

- **Workflows and Automation:** The more structured you can make your response function, the better a shot junior analysts have of finding the root cause of an attack, and figuring out how to contain and remediate it. Given the skills gap facing every organization, every bit of assistance helps. Response playbooks for a variety of different kinds of attacks within the environment can help standardize and structure response processes. Additionally, being able to integrate with automation platforms (now called SOAR — Security Orchestration, Automation, and Response) to streamline response — at least initial phases — dramatically improves effectiveness.

- **Integration with Malware Tools:** During validation you will also want to check whether an executed file is actually malware. A security monitoring platform can store executables and integrate with network-based sandboxes to explode and analyze files — to figure out both whether a file is malicious and what it does. This provides context for eventual containment and remediation. Ideally this integration will be native, enabling you to select an executable within the response console to send to the sandbox, with the verdict and report filed alongside the case.

- **Hunting:** Threat hunting has come into vogue over the past few years, as mature organizations decided they no longer wanted to be at the mercy of security monitoring, seeking a more active role finding attackers. So their more accomplished analysts started looking for trouble. They went hunting for adversaries rather than waiting for security monitors to report attacks in progress. Analysts need to figure out which behaviors and activities to hunt, then seek them out in the environment. The hunter starts with a hypothesis and runs through scenarios to either prove or disprove it. If the hunter finds suspicious activity then more traditional response functions — including searching, drilling down into available security data, and pivoting to other devices — all help follow the trail.

Compliance and Reporting

As we have mentioned, compliance reporting is extremely resource intensive and doesn't add much value to an organization. But if you screw up it can cost a lot of money in fines or other problems.

Compliance reporting is extremely resource intensive and doesn't add much value to an organization. But if you screw up it can cost a lot of money in fines or other problems.

The idea is to streamline the process of substantiating your controls as much as possible, so you can get the reports done quickly and back to real work.

This distinctly unsexy requirement is admittedly old hat, but the reports still need to be generated. You want your security monitoring tool to ship with dozens of reports to show your controls in place and map them to compliance requirements, so you don't need to do it manually.

You should have the ability to customize the reports which come with the tool, as well as develop your own reports when needed.

Scalability

Over the past few years, as you added mobile and cloud services and possibly endpoint data to your security data collection, you have been dealing with a lot more data — and there are no signs of abatement in the future. So you need to plan for scale.

- **Security Data:** Does your existing security monitoring platform keep pace with the increase in data and continue to perform smoothly? This is where the solution's underlying architecture shows through. Is the data aggregated on an appliance which gets bogged down at high insertion rates? Does the offering leverage a cloud-based architecture, which makes the particular scaling architecture less important, but leaves you still worrying about latency and response time. Does it combine technologies to support both on-premise assets and cloud-native technology stacks? The architecture you select must match your requirements — make sure any solution you consider can double in size within a reasonable timeframe without a forklift upgrade. Because the only surety in technology is that you will need to deal with more data sooner than you expect.
- **Pricing Scalability:** Security monitoring can be priced by events per second, a historical metric from when all data was collected by network sensors. We increasingly see pricing based on volume of data aggregated per day. Either way gives you a disincentive to collect more data, which is a problem when visibility into a sprawling IT environment is essential to your ability to detect attacks. So consider how the monitoring platform's pricing scales.

Intangibles

As much as we'd like to rely solely on technical requirements to select the best security monitoring platform, other factors always come into play.

- **Integration with Broader Product Line:** This is the age-old choice between big security company and focused upstart. We know smaller companies innovate faster, but many larger organizations are actively trying to reduce the number of vendors they deal with. A key question is: can you gain leverage by adopting a security monitoring platform from a big vendor which already provides various other solutions in your environment? One thing to check is that promised integration really exists. We don't say that facetiously — just because a vendor acquired a smaller company, or signed an OEM technology agreement, doesn't mean their solutions have been integrated beyond procurement. That's something to confirm during PoC.

- **Ease of Management:** How easy is it to manage the platform? To archive older data? To roll out new collectors, both on-premise and in the cloud? How about adding new use cases and customizing correlation rules? Are policy management screens easy to use, or do they consist of 500 checkboxes you don't fully understand? Make sure you have good answers to these questions during the PoC to ensure your new tool doesn't create more work.

- **Vendor Viability:** Have you ever bought a product from a smaller innovative company which didn't make it for whatever reason, and you are left holding the bag? We all have, so keep in mind that vendor fortunes can change dramatically and quickly. Your innovative small vendor might get acquired by a big IT shop and run into the ground. Conversely many larger security companies have struggled to scale (and show Wall Street growth and profits), forcing them to cut resources and miss huge innovations in the market. So buying from a big company isn't a safe bet either. Consider each vendor's ongoing viability and ability to deliver on its roadmap, to ensure it lines up with what you need going forward.

The Buying Process

Now we will work through a reasonably structured process to narrow down your short list and then test the surviving products. Once you've chosen the *technical* winner you need to make the business side of things work — and it turns out the technical winner is not always the solution you end up buying.

The first rule of buying anything is that *you* are in charge of the process. You'll have vendors who want you to use their process, their RFP/RFP language, their PoC Guide, and their contract language. All that is good and fine… if you want to by *their* product. But more likely you want the best product to solve **your** problems, which means **you** need to drive the process. Our procurement philosophy hinges on this.

But more likely you want the best product to solve your problems, which means you need to drive the process. Our procurement philosophy hinges on this.

The security monitoring market is very crowded and noisy. We have a set of incumbents from the SIEM space, and a set of new entrants wielding fancy math and analytics. Both groups share a set of base capabilities to address key use cases: threat detection, forensics and response, and compliance automation. But differentiation occurs at the margins of these use cases, so that's where you will be making your decision.

But no vendor is going to say, "We suck at X, but you should buy us because Y is what's most important to you." Even though they should. It is up to you to figure out each vendor's true strengths and weaknesses, and cross-reference them against your requirements. That's why it's critical to have a firm handle on your use cases and requirements *before* you start talking to vendors.

We divide vendor evaluation into two phases. First we will help you define a short list of potential replacements. Once you have the short list we recommend you test at least one platform from each category during a Proof of Concept (PoC) phase. It is time to do your homework. Even if you don't feel like it.

The Short List

The goal at this point is to whittle your list down to 3-5 vendors who appear to meet your needs, based on the results of a market analysis. That usually includes sending out RFIs, talking to analysts (egads!), or using a reseller or managed service provider to assist. The next step is to get a better sense of those 3-5 companies and their products. Your main tool at this stage is the vendor briefing.

The vendor brings in their sales reps and Sales Engineers (SEs) to tell you how their product is awesome and will solve every problem you have. And probably a bunch of problems you didn't know you had, too. But don't sit through their standard pitch — you know what is important to you.

You need detailed answers to objectively evaluate any new platform. You don't want a 30-slide PowerPoint walkthrough and generic demo. Make sure each challenger understands your expectations ahead of the meeting so they can bring the right folks. If they bring the wrong people cross them off. It's as simple as that — you don't have time to waste.

Based on the use cases you defined earlier in this process, have the vendor show you how their tool addresses each issue. This forces them to think about your problems rather than their scripted demo, and shows off capabilities which will be relevant to you. You don't want to buy from the best presenter — identify the product which best meets your needs.

Based on the use cases you defined earlier in this process, have the vendor show you how their tool addresses each issue. This forces them to think about your problems rather than their scripted demo!

This type of meeting could be considered cruel and unusual punishment. But you need this level of detail before you commit to actually testing a product or service. Shame on you if you don't ask every question to ensure you know everything you need. Don't worry about making the SE uncomfortable — this is their job.

And don't expect to get through a meeting like this in 30 minutes. You will likely need a half-day minimum to work through your key use cases. That's why you will probably only bring 3-5 vendors in for these meetings. You will be spending days with each product during proof of concept, so try to disqualify products which won't work before wasting even more time and effort on them. This initial meeting can be a painful investment of time — especially if you realize early that a vendor won't make the cut — but it is worth doing anyway. You can thank us later.

The PoC

After you finish the ritual humiliation of vendor sales teams, and have figured out which products can meet your requirements, it is time to get hands-on with the systems and run each through its paces for a couple days. This next step in the process, the Proof of Concept, is the most important — and vendors know that. This is where sales teams have a chance to win, so they tend to bring their best and brightest. They raise doubts about competitors and highlight their own successes. They have phone numbers for customer references handy. But for now forget all that. *You are running this show, and the PoC needs to follow your script — not theirs.*

Given the different approaches represented by SIEM and security analytics vendors, you are probably best served by testing at least one of each. As you read through our recommended process, it will be hard to find time for more than a couple, but given your specific environment and adversaries, seeing which type best meets your requirements will help you pick the best platform for your needs.

Preparation

Many security monitoring vendors have a standard testing process they run through, basically telling them what data to provide and what attacks to look for — sometimes even with their own resources running their product. It is like ordering off a *prix fixe* menu. You pick a few key use cases and the SE delivers what you ordered. If the vendor does it correctly it looks like a well-rehearsed ballet, where each participant precisely executes their assigned task. Everything quick and painless — just like security, right?

Wrong! Security is messy. Vendors design PoC processes to highlight their strengths and hide their weaknesses. We know this from first-hand experience — we have built them for vendors in past lives. We repeat this because it's that important. You need to work through *your* situation, not their scenario.

Wrong! Security is messy. Vendors design PoC processes to highlight their strengths and hide their weaknesses.

Before you start the PoC be clear about your evaluation criteria and which handful of use cases you want to conceptually test, based on your requirements from earlier in this process. Your requirements should spell out the key capabilities needed with a plan to further evaluate each challenger based on intangibles such as set-up/configuration, change management, customization, user experience/ease of use, etc.

We recommend investing in screen capture technology. It is hard to remember exactly what each tool did and how — especially after you have worked the tools through exactly the same paces. So capture as much video as you can of the user experience — it will come in very handy as you approach your decision point. Without further ado, let's jump in.

Testing

One advantage of testing security management products is that you can actually monitor production systems without worrying about blowing them up, taking them down, or adversely impacting anything. Pull the data you need to execute on the use case. The point is to run things according to your needs, with your data, alerting on your policies. You will also want to configure a custom data source or two and integrate with your directory store to see how that works.

If compliance is your key requirement use PCI as an example. Start pulling data from your protected network segment. Pump that data through the PCI reporting process. Is the data correct and useful to everybody interested? Are the reports comprehensive? Will you need to customize them for any reason? How easy is that? You need to answer these kinds of questions during PoC.

Pay attention to visualization and user interface. Security systems are not only used by security professionals. A configurable UI makes it easier for a wider audience to contribute to and benefit from the tool. Configure some dashboards and see the results. Mess around with reports a bit. Tighten alert thresholds. Does the notification system work? Will alerts work in a timely fashion at enterprise volumes? Is the information in the dashboards and reports useful? These are all things to check as part of the test.

Run a Red Team

The next step is to see how the tool runs under fire. This is particularly critical for analytics-based solutions, whose claim to fame is that they can find unknown attacks. Well, then, run an unknown attack against yourself. Clearly attacking production systems would make you unpopular with Ops, so set up a lab. Virtual environments are perfect for this — use the same base images for each vendor. The situation should be as realistic and consistent as possible.

Have attackers breach test systems with attack tools. Have your defenders try to figure out what is going on as it's happening. Does the system alert as it should? Will you need to heavily customize rules? Can you identify the nature of attacks quickly? Does their super-duper forensic drill-down give you the insight you need? The clock is ticking — how easy is it to use the system to search for clues?

The next step is to see how the tool runs under fire. This is particularly critical for analytics-based solutions, whose claim to fame is that they can find unknown attacks.

Obviously this isn't a real incident, but you need a feel for how the system performs in action. If an attacker is in your systems, will you find them? In time to stop or catch them? Once you know attackers are in, can you tell what they are doing? A red team exercise as part of the PoC will help determine that.

Knowing a tool will hold up in the heat of battle goes a long way toward giving security and operations teams confidence when they go live. Keep in mind that you cannot fully test scalability during PoC, so focus on what you can fully test. That's the user experience, and there is no better way to distill out the effectiveness of a tool than to use it during an attack.

The Postmortem

At the end of the test it's important that you evaluate both successes and failures of your PoC in terms of use cases and requirements. When you finish a red team exercise you should have a bunch of data which nicely illustrates what the attack team did — and perhaps what the defense team didn't do as well as they could have. This is a learning experience for everyone, and real attack scenarios nicely illustrate the particular value of a platform.

Your team should grade each candidate while memory is fresh and perceptions are raw. After a week or two with another product they won't remember what they liked and didn't about earlier ones as clearly — another reason screen grabs are handy.

Lather, Rinse, Repeat

You will probably test more than one product or service, so you get to do it all again, and make sure to use the same scenarios for each. Consistency helps make the testing process fair and comparisons more meaningful.

Now you have all the information you need to make a decision, so it is time to figure out what to do and substantiate your choice for your internal sales process. You can use the details of the PoC and screen capture videos you collected from each competitor.

Documentation

Your end product is a recommendation, so you need to document what you think and then present it to secure funding. You may not always be in the room when decisions are made, so your documentation must clearly articulate your reasons. We normally structure this artifact of the decision process as follows:

- **Requirements:** Tell them what you need and who said you need it. Compliance and security requirements come from different groups, so make sure to reference the folks driving those use cases.
- **Coverage:** What works and doesn't with the desired solution within the context of your requirements, both now and as they evolve. Make sure it's clear that your choice meets the requirements you just laid out.
- **Competition:** Which other vendors did you disqualify and why? What did you learn during Proof of Concept? Are any of the competitors viable? What compromises would need to be made if another product was selected?
- **Cost Estimate:** What will it cost to move to the new platform? How much is capital expense and what fraction is operational, and does that change the decision at all? What kind of investment in professional services will be required?
- **Migration Plan:** What will the migration look like? How long will it take? Will migration disrupt any services? Will you be more exposed to attack, and if so for how long? You need all these answers before you pitch the powers that be. Not a Gantt chart — that comes at the end — but enough to answer the tough questions.
- **Recommendation:** Your entire document should be building to this point, where you put the best path down on paper. If it comes as a surprise to your audience you did something wrong. This is about telling them what they already know, and making sure they have an opportunity to ask any remaining questions.

Now you have the thumbs-up from the internal team (we hope!). You need to negotiate with the vendor and get the deal done. We won't get into the specifics of negotiating — you likely have people to do that — but understand that you can use time-honored tactics such as waiting until the end of the quarter, playing one vendor against another (if either could meet your requirements), and possibly asking for non-cash add-ons (such as professional services or product modules).

Once all this is done you may need to go back for final sign-off. This is when your process is most vulnerable. Negotiations around price and services can be challenging. But at least in negotiating with a vendor you know who the adversary is. Inevitably, internal resistance will appear (or reappear),

and you may never see it coming. Especially if a losing vendor sells a lot of other stuff to your company and has friends in high places.

This entire process has prepared you to deal with these obstacles, so just work through the logic of your decision once more, making clear how your recommendation is best for the organization, and squash the resistance. Expect the losing vendor to go over your head — vendors don't go quietly into the long night when they lose a deal.

Migration

At this stage in the buying process it feels like you have gone through a lot of work but haven't even started the real security work yet. Let's step back a moment to focus on what's important: getting stuff done as simply and easily as possible. The migration process is not easy because you need to maintain service levels without exposing your organization to additional risk. This likely involves supporting two systems for a short while, or using two systems in a hybrid architecture — perhaps even indefinitely.

Either way, when you put your head on the block to choose a new platform, the migration needs to go smoothly. There is no such thing as a 'flash' cutover. We recommend you start deploying the new monitoring platform long before you get rid of the old. At best you will deprecate portions of the older system after newer replacement capabilities are online, but you will likely want the older system as a fallback until all new functions have been vetted and tuned. We learned the importance of this staging process the hard way. Ignore it at your peril — security monitoring supports several key business functions, so if things go haywire it's a bad day all around.

When you put your head on the block to choose a new platform, the migration needs to go smoothly. There is no such thing as a 'flash' cutover. We recommend you start deploying the new monitoring platform long before you get rid of the old.

Fast forward a few months; the challenging process of actually buying will be in the rearview and you will be getting to use your new tool. But don't get too comfortable. Just as we saw a new group of security analytics players challenge incumbent SIEM vendors, inevitably there will be something else newer and shinier in the market.

In 2-3 years, after that new-car smell wears off the platform you just bought, you'll be questioning whether today's choice remains correct. You may be building another Team of Rivals, which will eventually give way to the next strategic platform. Seems silly, but there is no use resisting — it's all part of the game.

Summary

Not that we spend a lot of time patting ourselves on the back, but our projection that SIEM and security analytics were converging was on the money. That and $4 will get you a coffee, so what does it mean for you? It means you need to think about what your strategic security monitoring platform will be, and how to get there.

Start with the use cases to determine your priorities. With a clear idea of what problems you are trying to solve and some time translating those needs into clear requirements, you are in a position to select a platform. Or at least have a good idea of when you'll be ready to start the move, if it makes more sense to leave things as they are for the time being. Remember there is no shame in deciding to stay put, so long as you understand the compromises inherent in that decision.

The Security Monitoring State of the Union is good. The technology continues to evolve and become more effective at detecting the attacks you care about. That platforms continue to scale and are even starting to embrace the cloud. That said, the technology isn't and hasn't really been the constraint. So remember to factor in the ability of your team to consume and leverage advanced technology. Structure your security monitoring program to give all parties the best chance of success.

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 publish the vast majority of our research for free through our blog, and package the research as papers that can be licensed for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements, and follow our Totally Transparent Research policy.
- **Cloud Security Project Accelerators**: Securosis Project Accelerators (SPA) are packaged consulting offerings to bring our applied research and battle-tested field experiences to your cloud deployments. These in-depth programs combine assessment, tailored workshops, and ongoing support to ensure you can secure your cloud projects better and faster. They are designed to cut months or years off your projects while integrating leading-edge cloud security practices into your existing operations.
- **Cloud Security Training:** We are the team that built the Cloud Security Alliance CCSK training class and our own Advanced Cloud Security and Applied SecDevOps program. Attend one of our public classes or bring us in for a private, customized experience.
- **Advisory services for vendors**: We offer a number of advisory services to help our vendor clients bring the right product/service to market in the right way to hit on critical market requirements. Securosis is known for telling our clients what they NEED to hear, not what they want to hear. Clients typically start with a strategy day engagement, and then can engage with us on a retainer basis for ongoing support. Services available as part of our advisory services include market and product analysis and strategy, technology roadmap guidance, competitive strategies, etc. Though keep in mind, we maintain our strict objectivity and confidentiality requirements on all engagements.
- **Custom Research, Speaking and Advisory**: Need a custom research report on a new technology or security issue? A highly-rated speaker for an internal or public security event? An outside expert for a merger or acquisition due diligence? An expert to evaluate your security strategy, identify gaps, and build a roadmap forward? These defined projects bridge the gap when you need more than a strategy day but less than a long-term consulting engagement.

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. For more information about Securosis, visit our website: <http://securosis.com/>.