Evolving to Security Decision Support

Version 1.5
Released: April 9, 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 paper is licensed by Tenable, Inc.

tenable.com

Tenable®, Inc. is the Cyber Exposure company. Over 24,000 organizations around the globe rely on Tenable to understand and reduce cyber risk. As the creator of Nessus®, Tenable extended its expertise in vulnerabilities to deliver Tenable.io®, the world's first platform to see and secure any digital asset on any computing platform. Tenable customers include more than 50 percent of the Fortune 500, more than 20 percent of the Global 2000 and large government agencies.

Copyright

Evolving to Security Decision Support

Table of Contents

Visibility is Job #1

To demonstrate our mastery of the obvious, it is not getting easier to detect attacks. Not that it was ever really *easy*, but at least you used to know what tactics adversaries were using, and had a general idea of where they would end up, because you knew where your important data was, and which (single) type of device normally accessed it: the PC. It's hard to believe we now long for the days of early PCs and centralized data repositories.

You need a better way to assess your organization's security posture, detect attacks, and determine applicable methods to work around and eventually remediate exposures in your environment.

But that is not today's world. You face professional adversaries (and possibly nation-states) using agile techniques to develop and test attacks. They have ways to obfuscate who they are and what they are trying to do, further complicating detection. They prey on those ever-present gullible employees who click anything to gain a foothold in your environment. Further complicating matters is the inexorable march toward cloud services — which move unstructured content to cloud storage, outsource back-office functions to a variety of service providers, and move significant portions of the technology environment into the public cloud. And all these movements are accelerating — seemingly exponentially.

There has always been a playbook for dealing with attackers *when we knew what they were trying to do*. Whether or not you were able to effectively execute on that playbook, the fundamentals were fairly well understood. But as explained in our Future of Security Operations paper, the old ways don't work any more, which puts practitioners behind the 8 Ball. The rules have changed, and old security architectures are rapidly becoming obsolete. For instance it's increasingly difficult to insert inspection bottlenecks into your cloud environment without adversely impacting efficiency. And sophisticated adversaries can use exploits which aren't caught by traditional assessment and detection technologies — although they don't need such fancy tricks often.

So you need a better way to assess your organization's security posture, detect attacks, and determine applicable methods to work around and eventually remediate exposures in your environment. As much as the security industry whinges about adversary innovation, at least it has made progress at improving your ability to assess and detect attacks. We have written a lot on threat intelligence and security analytics over the past few years. Those are the cornerstone technologies for dealing with modern adversaries' improved capabilities.

But these technologies and capabilities cannot stand alone. Just pumping some threat intel into your SIEM won't help you understand the context and relevance of the data you have. And advanced analytics on the firehose of security data you collect is not enough either, because you might be missing a totally new attack vector.

What you need is a better way to assess your organizational security posture, determine when you are under attack, and figure out how to make the pain stop. This requires a combination of technology, process changes, and clear understanding of how your technology infrastructure is evolving toward the cloud. This is no longer just *assessment* or *analytics* — you need something bigger and better. It's what we now call **Security Decision Support** (SDS). Snazzy, huh?

This paper will delve into these concepts to show how to gain both visibility and context — so you can understand both *what* you have to do and *why*. Security Decision Support provides a way to prioritize the thousands of things you *can* do, enabling you to zero in on the few you **must**.

Visibility in the Olden Days

Securing pretty much anything starts with visibility. *You can't manage what you can't see* — and a zillion other overused adages all illustrate the same point. If you don't know what's on your network and where your critical data is, you don't have much chance of protecting it.

In the *olden days* — you know, way back in the early 2000s — visibility was fairly straightforward. First you had data on mainframes in the data center. Even when we started using LANs to connect everything, data still lived on a raised floor, or in a fairly simple email system. Early client/server systems started complicating things a bit, but everything was still on networks you controlled in data centers you had keys to. You could scan your address space to figure out where everything was and what vulnerabilities needed to be dealt with.

That worked pretty well for a long time. There were scaling issues, and a need (desire) to scan higher in the technology stack, so we started seeing first stand-alone and then integrated application scanners. Once rogue devices started appearing on your network it was no longer sufficient to scan your address space every couple weeks, so passive network monitoring enabled you to watch traffic and flag (and assess) unknown devices.

Those were the good old days, when things were relatively simple. Okay — maybe not really *simple,* but you could size the problem. That is no longer the case.

Visibility Challenged

One of our favorite memes is a man from the 1870s, blissfully remembering the good old days when he knew where his data was. The image always gets a lot of laughs from audiences when we present. But it's brought on by pain, because everyone in the room knows it illustrates a very real problem. Nowadays you don't really know where your data is, which seriously compromises your ability to determine the security of the systems with access to it.

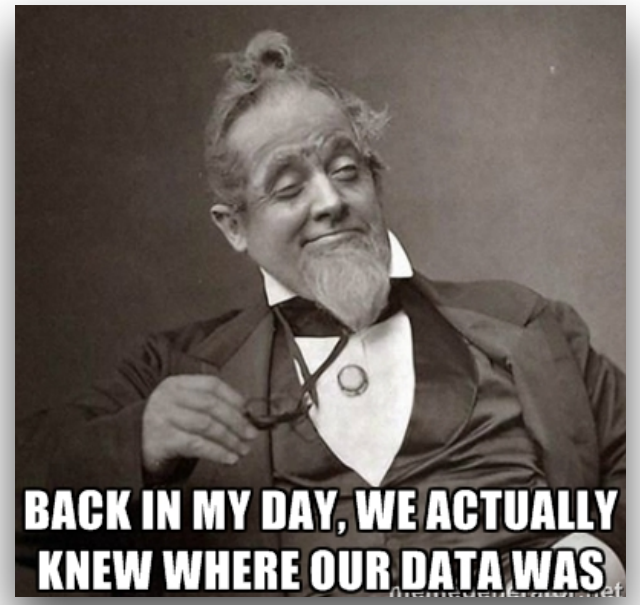

These challenges result directly from a number of key technology innovations:

- **SaaS:** Securosis talks about how SaaS is the New Back Office, which has drastic ramifications for visibility. Many organizations deploy CASB (Cloud Access Security Broker) just to figure out which SaaS services they are using, because they cannot depend on business folks to ask permission to use a business-oriented service. This problem isn't going away — if anything more business processes are moving to SaaS.

- **IaaS:** Speaking of cloudy stuff, you have teams using Infrastructure as a Service (IaaS) — either moving existing systems out of your data centers or building new systems in the cloud. IaaS changes how you assess your environment, breaking most old techniques. Scanning becomes a lot harder and some of the 'servers' (now called instances) live only a few hours. Network addressing is different, and you cannot really implement taps to see all traffic. It's a different world, where you are pretty much blind until you come up to speed with new techniques to replace tricks the cloud broke.

- **Containers:** Another new foundational technology, containers bring much better portability and flexibility to building and deployment of application components. Without going into detail about why they are cool, suffice it to say your developers are likely working extensively with containers as they architect new applications — especially in the cloud. But containers raise new visibility and security challenges, in part because they are short-lived (they spin up and down automatically, responding to load and other triggers), self-contained (usually not externally addressable) and don't provide access for traditional scans. They also break existing discovery and assessment processes.

- **Mobility:** It seems a bit old hat to even mention that you have critical data on smart devices (phones and tablets), but they expand attack surface and make it hard to understand where your data is — as well as how devices are configured.

- **IoT:** A little further out toward the horizon is the Internet of Things (IoT). Some argue it's here today, and with the number of sensors being deployed and *smart* systems already network connected, they are arguably right. But either way, if you look even just a year or two out into the future, you can bet there will be a lot more network connected devices accessing your data and expanding your attack surface. So you need to find and assess them.

And we're just getting started. It won't be long before the *next* discontinuous innovation makes it harder to figure out where critical data resides and what's happening with it. To put a bow on the challenges facing you, we'll talk about some reasonable bets to make. We are confident there will be more cloud tomorrow than today. And equally confident more devices will be accessing your stuff tomorrow. That's pretty much all you need to know to understand the extent and magnitude of the problem.

It won't be long before the next discontinuous innovation makes it harder to figure out where critical data resides and what's happening with it.

Challenge Accepted

To once again state the obvious, it's hard to be a security professional nowadays. We get it. But curling up into the fetal position on your data center floor isn't an option. First of all, you may not even have a data center any more. And if you do, it might have been repurposed as warehouse space or sold off to a cloud provider. But even if you have a cozy spot, curling up won't actually solve any problems.

So what can you do? Remember you cannot manage or protect what you cannot see, so you need to focus on visibility as the first step toward Security Decision Support. Visibility across the enterprise, wherever your data resides, on whatever platform. That means discovery and assessment of all your stuff.

We're pretty sure you haven't been able to totally shut off your data centers and move *everything* to SaaS and IaaS yet — even though you might want to — so you need to make sure you aren't missing anything within traditional infrastructure. You need to continue your existing vulnerability management program.

- **Network, security, databases, and systems:** You already scan your network and security devices, all the servers you control, and probably your databases as well (thanks, compliance mandates). You should keep doing all that. Hopefully you have been evolving your vulnerability management systems, and have some way to prioritize all the stuff in your environment.

- **Applications:** You are likely scanning your web applications as well. That's a good thing — keep doing it. And working with developers to ensure they are fixing the issues you find before deploying them to millions of customers. Obviously as developers continue to adapt agile methods of building software you will still need to evangelize finding issues with your application stacks and — given the velocity of software changes — fixing them faster.

That's the stuff you should already be doing. It might not be going as well as you want (there is always room for improvement, right?), but at least for compliance you are probably already doing something. It gets interesting when discovery and assessment meet the new environments and innovations you need to grapple with. Let's look at the innovations above for a sense of how they change things in the new world.

- **SaaS:** Many of you have deployed a CASB to monitor your network egress traffic and figure out which SaaS services you are actually using. Customers tend to be shocked when they are using 10 times the number of SaaS services they thought. To be clear, you don't need a purpose-built device or service to find SaaS in use — many secure web gateways offer this kind of visibility, as do DLP solutions to control exfiltration. So there are many alternatives for discovery, all of which examine egress traffic. And don't forget you also have some options to leverage a SaaS provider's API (Application Programming Interface) to access reasonably granular usage and activity metrics via API.

To be clear, you don't need a purpose-built device or service to find SaaS in use — many secure web gateways offer this kind of visibility, as do DLP solutions to control exfiltration.

- **IaaS:** Unlike SaaS, an egress filter cannot provide much detail about what's running within or going into an Infrastructure as a Service (IaaS) public cloud. In this case the API really is your friend. Visibility tools need to poll the cloud provider's API to learn what systems are running in their environment and assess them. Although keep in mind that API limitations of cloud providers may drive you to an multiple account per application deployment architecture, which also helps to minimize the blast radius of attacks and failures with stronger functional isolation between applications.

- **Follow the Money:** For both SaaS and IaaS, someone is getting paid for any services you use. Whoever pays the bills should be able to let you know — at least at a gross level — which services are in use, with pointers to who can tell you more. So make sure you are friendly with Accounting. Of course an Accounting report is no replacement for pulling information from API or monitoring egress traffic, especially given the lag between when something is consumed and when you pay for it -- which may be 4 - 6 weeks.

- **Containers:** Containers encapsulate microservices — which are often not persistent, and cannot really be accessed or scanned by external entities such as vulnerability scanners, so you need to build discovery and assessment into your container system. First make sure any containers are built using an image which is not vulnerable. Then track container usage to make sure nothing drifts by instrumenting the container with an agent or API call as containers spin up to track each container through its lifecycle, report back to your central repository, and watch for signs of attack. Like most of security you can't effectively bolt this on later, so make sure you are friendly with the developers too. Without their participation you'll have little visibility into your container environment.

Containers cannot really be accessed or scanned by external entities such as vulnerability scanners, so you need to build discovery and assessment into your container system.

- **Mobility:** With mobile devices now full participants in the IT environment, Mobile Device Management (MDM) is pretty well established technology to address common problems like managing and securing thousands of devices. These platforms can provide an inventory not just of devices, but also what is installed on each. There is still work to be done to integrate that information into the rest of the Security Decision Support stack. You should be able to use telemetry from your MDM environment in your security analytics strategy. For example a person's mobile device accessing cloud data stores they aren't authorized to look at, or their desktop performing reconnaissance on the finance network — or even both! — may indicate a compromise. Your analytics should detect and connect both events across the enterprise.

- **IoT:** The problem with most IoT devices don't have an API you can poll to assessing their security and watching for potential misuse, nor can you install an agent to monitor activity. These devices often appear on network segments which are less monitored and protected, such as the shop floor or the security video network. IoT requires a different approach — largely passive monitoring to profile the devices on each network, baseline typical traffic patterns, and then watch for devices acting unusually. Yet, it is more challenging than simply collecting NetFlow records from the shop floor network. IoT devices often use non-standard proprietary protocols which complicate discovery and assessment.

(Re)Visitation

Of course we need a few caveats around these concepts. First, emerging technologies are moving targets. Let's take IaaS as an example. Like other technology providers, cloud providers are rapidly introducing API and other mechanisms to provide visibility into their environments. Device makers across all device types realize customers want to manage their technology as part of a larger system, so many (alas, not all) are providing better access to their innards in more flexible ways. That's the kind of progress we like to see.

But we cannot wait for tomorrow to solve today's problems. You need to build a process and implement tooling with what's available today.

But we cannot wait for tomorrow to solve today's problems. You need to build a process and implement tooling with what's available today. So build periodic reassessment into your SDS process, similar to how you likely revisit malware detection periodically.

We know reviewing your enterprise visibility approaches can be time-consuming, and expensive when something needs to change. Reversing course on decisions you made over the past year can be frustrating. But that's the world we live in and resisting will just cost more in the long run. If you expect from the get-go to revisit all these decisions, and at times to toss some tools and embrace others, that makes it easier to take. Even more important, managing the expectations of senior management that this might (quite likely) happen will go a long way to maintaining your employment.

Data to Intelligence

Given all the moving pieces in your environment — including various clouds (SaaS and IaaS), mobile devices, containers, and eventually IoT devices — it's increasingly hard to know where all your critical data is and how it's being used.

Enterprise visibility is necessary but not sufficient. You still need to figure out whether and how you are being attacked, as well as whether and how data and/or apps are being misused. Nobody gets credit just for knowing where you can be attacked. You need to *stop* attacks and protect critical data. Ultimately that's what matters. The good news is that many organizations already collect extensive security data (thanks, compliance!), so you have a base to work with. It's just a matter of turning all that security data into actual intelligence you can use for Security Decision Support.

The History of Security Monitoring

Let's start with some historical perspective on how we got here, and why many organizations already perform extensive security data collection. It all started in the early 2000s with deployment of the first SIEM, deployed to make sense of the avalanche of alerts coming from firewalls and intrusion detection gear. You remember those days, right?

SIEM evolution was driven by the need to gather logs and generate reports to substantiate controls (thanks again, compliance!). So SIEM products focused more on storing and gathering data than actually making sense of it. You could generate alerts on things you knew to look for, which typically meant you got pretty good at finding attacks you had already seen. But they were very limited for detecting attacks you hadn't seen.

SIEM technology continues to evolve, but mostly to add scale and data sources to keep up with the number of devices and amount of activity to be monitored. Unfortunately that doesn't address the fact that many organizations don't want *more* alerts — they want *better* alerts. To provide them, two separate capabilities have come together in an interesting way:

1. **Threat Intelligence:** SIEM rules were based on scanning for what you had seen before, so you were limited in what you could look for. What if you could leverage attacks *other companies* have seen and look for *those* attacks so you could anticipate what's coming? That promise is driving organizations to embrace external threat intelligence.

2. **Security Analytics:** The other capability isn't exactly new — it's using advanced math to look at the security data you've already collected to profile normal behaviors, and then look for stuff that isn't normal and might be malicious. Call it anomaly detection, machine

learning, whatever — the concept is the same. Gather a bunch of security data, build mathematical profiles of normal activity, then look for activity that *isn't* normal.

Let's consider both these capabilities to gain better understanding of how they work, and then we'll be able to show how powerful the combination can be for improving alerts.

Threat Intelligence Identifies What Could Happen

Culturally, over the past 20 years, security folks were generally the kids who didn't play well in the sandbox. Nobody wanted to appear vulnerable, so data breaches and successful attacks were their dirty little secret. Sure, they happen, but not to *us*. Yeah, right. There were occasional high-profile issues (like SQL*Slammer) which couldn't be swept under the rug, but they hit everyone so weren't as big a deal personally.

Security practitioners realized no one is perfect, and we can collectively improve our ability to defend ourselves by sharing information about adversary tactics and specific indicators from those attacks.

But over the past 5 years a shift has occurred in security circles, borne out of necessity as most such things are. Security practitioners realized no one is perfect, and we can collectively improve our ability to defend ourselves by sharing information about adversary tactics and specific indicators from those attacks. This is something we dubbed "benefiting from the misfortune of others" a few years ago. Everyone benefits because once one of us is attacked, we all learn about the attack and can look for it. The modern threat intelligence market emerged.

In terms of the current state of threat intel, we typically see the following types of data shared via commercial services, industry groups/ISACs, and open source communities:

- **Bad IP Addresses:** IP addresses which behave badly, for instance by participating in botnets or acting as spam relays, should probably be blocked at your egress filter because you know no good will come from communicating with those networks. You can buy a blacklist of bad IP addresses, the low-hanging fruit of the threat intel world.
- **Malware Indicators:** Next-generation attack signatures can be gathered and shared to help look for activity representative of typical attacks. You know these indicate an attack, so being able to look for them within your security monitors helps keep your defenses current.

The key value of threat intel is to *accelerate the human,* as described in our Introduction to Threat Operations research. But what does that even mean? To illustrate a bit, let's consider *retrospective search*. This involves being notified of a new attack via a threat intel feed, and using those indicators to mine existing security data to determine whether you experienced this attack *before* you knew to look for it. Of course it would be better to detect every attack as it happens, but the ability to go back and search old security data for new indicators shortens the detection window.

Another use of threat intel is to refine your hunting process. A hunter learns about a specific adversary's tactics and then hunts for that adversary. It's not like the adversary is going to send out a memo detailing its primary TTPs, so threat intel is the way to figure out what they are likely to do. This makes the hunter *much* more efficient ("accelerates the human") by focusing on typical tactics used by likely adversaries.

Much of the threat intel available today is focused on data to be pumped into traditional controls, such as SIEM and egress filters. There is an emerging need for intel on new areas of exposure including the cloud, IoT, and mobility. As more attacks leverage new attack vectors, more data becomes available, making us all better. But in the meantime there is a clear gap in the data available on these emerging technologies.

Yet effectively leveraging threat intel alone cannot realize the full potential of Security Decision Support. Knowing what could happen is very helpful. But you still end up with a long list of stuff to triage and potentially remediate, and little real context to prioritize efforts. That's where analytics comes in.

Analytics Identifies What Is Happening

Security analytics is conceptually simple. Use advanced math to establish a baseline of activity in the mass of collected security data. Then look for situations which could indicate malicious activity or misuse of critical systems or data. In reality, of course, it's anything but simple.

Security analytics is conceptually simple. Use advanced math to establish a baseline of activity in the mass of collected security data. Then look for situations which could indicate malicious activity or misuse of critical systems or data.

The basic technology underlying security analytics is anomaly detection. You remember that, right? Security analytics vendors come up with fancy terms, but at its core this isn't new. We have been looking for anomalies in security data for over a decade. Remember Network-Based Anomaly Detection (NBAD)? We certainly do as security historians — it was the first "security analytics" offering we recall.

To be clear, NBAD worked and still does. The technology has been wrapped up in a variety of different offerings so there aren't really any standalone NBAD companies any more, but the approach has evolved and morphed into what we now call security analytics. So what's different now? First, you can analyze a lot more data, much more efficiently. Instead of just looking at network flow records (like in the NBAD days), you can now look at full packet captures, endpoint telemetry, log activity from pretty much all devices and applications in use, and even potentially transaction data — and then build a baseline to learn what is *normal* in your environment.

Most enterprise networks are pretty complicated (and only getting more so), so building a baseline across a number of seemingly unrelated security data sources is challenging. To simplify things a bit, you can start by looking at different use cases to chunk up the universe of activity into a manageable set for a reasonable place to start.

For example you probably want to find compromised devices. One way is to look for devices doing strange things which could indicate misuse. Typically someone in Finance shouldn't be reconnoitering devices in Engineering. Or vice-versa. *That's not normal,* and as such should probably be investigated. So prioritizing device behavior to detect malware is a common use for security analytics.

Along with device behavior, you could also expand the scope of analysis to include broader insider threats by building a baseline for how specific employees use their devices (User Behavior Analytics: UBA). Then when an employee does something funky, such as connecting to the finance system from a tablet at home, you can flag that for review.

You could extend that use case to broader analysis by adding data from physical access systems (to see when an employee is in the office), as well as HR records (listing employees under investigation or considered flight risks). Or you could look at employee usage of specific applications, especially the ones accessing critical proprietary corporate data. With all this data you can profile employee usage and transaction patterns. That provides a baseline to help identify activity which should be investigated.

Security analytics provides you with a list of things not to look for (like threat intel or threat modeling), but *happening* in your environment — providing more actionable alerts which likely warrant investigation.

Yet security analytics on its own also doesn't rise to the standard of Security Decision Support. Analytics can identify things to look for, but doesn't provide a sense of importance or prioritization in terms of other things the platform is alerting on. So we need to ask a few more questions of the system to make better decisions.

Driving Security Decisions

Now that you have an understanding of which attacks are being used in the wild (via threat intel) and which systems/users/applications are potentially being misused (via security analytics), the next step is to use that context to drive action. But what action? How can you prioritize all the things (both internal and external) which should be looked at? You need to balance the following to accurately decide which actions will be most impactful in your environment:

- **Asset/Data Value:** There are corporate systems and data which are important to your company. When this kind of stuff is compromised, heads roll — probably yours. So obviously you prioritize potential situations involving these systems or users above others. This is a subjective measure of value, so you will need discussions with senior management to understand the value of various systems. But if you want to stay employed you need to factor asset value into the mix — nobody has the resources or time to do *everything*.

- **Confidence:** False positives hurt you by wasting time on stuff which turns out to be nothing. Reducing this wasted time is possible by tracking your *confidence* in data/intel sources and analytical techniques. Obviously if a specific intel source sends you crap, you should not jump when it alerts. Likewise if a bunch of alerts based on a user's mobile activity turn out to be much ado about nothing, you should de-emphasize that source over time. But if you find something which indicates an attack via retrospective search, that is a high priority — you know that attack is real, and happening in your environment.
- **Internal Skills/Resources:** The industry is making progress at automating some security activities, but there still seems to be an infinite amount of work. You also need to consider the skills at your disposal when prioritizing. If you are weak at Tier 1 response because your front-line staffers keep getting poached by consulting firms, you may want to send alerts that might be urgent directly to Tier 2. Similarly, if you know your Tier 3 folks find file-less attacks to be heavy sledding, you may want to just quarantine those devices until your external forensics team can take a look.

This kind of process requires the ability to measure effectiveness of both threat intel and analytics over time. A *gut feel* isn't the ideal tool to determine which sources and methods work best. The sooner you start quantifying value the better. Not just for prioritization but also to save money. If you are spending money on sources or analytics platforms which don't provide value, stop.

This kind of process requires the ability to measure effectiveness of both threat intel and analytics over time. A gut feel isn't the ideal tool to determine which sources and methods work best.

By leveraging threat intel and more advanced security analytics, you can narrow the aperture of all the things you *could* look at to help identify what you *need* to look at. We don't claim this makes your to-do list manageable, but every little bit helps. By focusing on likely attacks (according to threat intel) on devices which are acting abnormally, considering the value of the asset being attacked and your confidence in the data, you have a much better chance focus on attacks which matter.

Bringing It All Together

As we have mentioned, the first step in making better security decisions is ensuring you have full visibility into your enterprise assets, because if you don't know assets exist you cannot make intelligent decision about protecting them. Next you focus on how threat intelligence and security analytics can be brought to bear to get both internal and external views of your attack environment, again with the goal of turning data into information useful for prioritizing your efforts.

Once you reach this stage you have the basic capabilities to make better security decisions. The key is to **integrate** these practices into your day-to-day activities. This requires process changes and a focus on instrumentation within your security program to track effectiveness, in order to constantly improve performance.

Implementing SDS

To implement Security Decision Support you need a dashboard of sorts to help track all the information coming into your environment, to help decide what to do and when. You need a common interface to visualize alerts and determine their relative priority. This entails tuning your monitors to your particular environment so prioritization improves over time.

To implement Security Decision Support you need a dashboard of sorts to help track all the information coming into your environment, to help decide what to do and when.

We know — the *last* thing you want is another dashboard to deal with. Yet another place to collect security data, which you need to keep current and tuned. But this need not be a *new* system. You have a bunch of tools in place which certainly could provide these capabilities. An existing SIEM, security analytics product, and vulnerability management service, to name a few. So the SDS platform you ultimately pick may already be deployed, but these advanced capabilities are probably not yet fully implemented and utilized. That's where process changes come into play.

Before you worry about what tool will to do this work, let's work through the capabilities required to implement this vision.

The first thing you need in a Decision Support platform to visualize security issues is data. So what will feed this system? You need to understand your technology environment, so integration with your organizational asset inventory (usually a CMDB) provides devices and IP addresses. You'll also want

information from your enterprise directory, which provides people and can be used to understand each user's role and entitlements. Finally you need data from security monitors — including any SIEM, analytics, vulnerability management, EDR, IPS/IDS, etc.

You'll also need to categorize both devices and users based on their importance and risk to the organization. Not that some employees are more important than others as humans (everyone is important — how's that for political correctness?). But some employees pose more risk to the organization than others. That's what you need to understand, because attacks against high-risk employees and systems should be dealt with first.

We opt for simplicity here, recommending 3-4 different categories with very original names:

1. **High:** These are the folks and systems which, if compromised, would cause a bad day for pretty much everyone. Senior management fits into this category, as well as resources and systems with access to the most sensitive data in your enterprise. This category poses risk to the entire enterprise.

2. **Medium:** These employees and systems will cause problems if stolen or compromised. The difference is that the damage would be *contained*. These folks can only access data for a business unit or location, not the entire enterprise.

3. **Low:** These people and systems don't really have access to much of import. Of course there is enterprise risk associated with this category, but it's *indirect*. An adversary could use a low-risk device or system to gain a foothold in your organization, then attack stuff in a higher-risk category.

We recommend categorizing adversaries and attack types as well. Threat intelligence can help you determine which tactics are associated with which adversaries, and perhaps prioritize specific attackers (and tactics) by how motivated they are to attack your environment.

Once this is implemented you will have a clear sense of what needs to happen first, based on type of attack and adversary — and of the importance of each device, user, and system. It's a *priority score*, although security marketeers call it a risk score. This is analogous to a quantitative financial trading system. You want to take most of the emotion out of decisions, so you can get down to what is best for the organization. Many experienced practitioners push back on this concept, preferring to make decisions based on their *gut* — or even worse, a FIFO (First In, First Out) model.

We'll just point out that pretty much every major breach over the last 5 years produced multiple alerts of the attack in progress, along with opportunities to deal with it, before it became a catastrophe. But for whatever reason those attacks weren't dealt with.

The final output of a Security Decision Support process is a *decision* about what needs to happen — meaning you still need to actually do the work. Integration with a security orchestration and automation platform can help make changes faster and more reliable. You will probably want to send tasks to a work management system (trouble ticketing, etc.) to route to Operations, and to track remediation.

Feedback Loop

We talk about Security Decision Support as a process, which means it needs to adapt and evolve to both the changing environment and new attacks and adversaries. You want a feedback loop integrated with your operational platform, learning over time. As with tuning any other system, you should pay attention to:

1. **False Negatives:** Where did the system miss? Why? A false negative is very serious because it means you missed a legitimate attack. Unfortunately you might not know about it until you get a breach notification. Many organizations start threat hunting to find active adversaries their security monitoring system miss.

2. **False Positives:** A bit more visible, and sources of much annoyance, are false positives. These are generated by the system but turn out to not be real security issues — although they are certainly workflow issues. They crop up particularly when you try to tighten detection. Work on false positives to ensure you have proper thresholds.

3. **Time to Respond:** You are trying to improve operations so you need to track duration of the incident. How long did it take to detect an issue? To remediate and close out the problem? This is an area where more sophisticated organizations can start setting service levels — within both security (how quickly you detect) and operations (how quickly you remediate), especially for higher risk attacks.

4. **Data Source Effectiveness:** As much as we like to collect and analyze more data, rather than less, at some point you will reach a point of diminishing returns from adding more analytic data. This is especially important for external threat intel, which typically carries a financial cost. Of course false positives from unreliable data sources also carry opportunity cost and distract from real attacks.

A key requirement for better decisions is quantification. Always provide sufficient instrumentation for any new control or tactic, ensuring that you can substantiate whether to do more or less, addressing the need to more effectively allocate resources.

A key requirement for better decisions is *quantification*. Always provide sufficient instrumentation for any new control or tactic, ensuring that you can substantiate whether to do more or less, addressing the need to more effectively allocate resources.

Over time your ability to benchmark performance against similar organizations provides yet another way to gauge the effectiveness of your operations. This is aspirational at this point, because what to track is still an open

question, and how to anonymize and share that data hasn't yet been defined. But there is clear value in being able to pinpoint areas of under and over performance to target continuing investment.

Picking the System

Now that you know what your decision support system needs to do, the question is how you get one. Can you go to the local computer store to pick one up? Can you just check out your favorite analyst's quadrant chart and send over a PO? Alas, it's not that easy — yet. You already have pieces of the puzzle, but not all in one place.

As we hinted earlier, four technologies are the most likely candidates to evolve into the Security Decision Support platform:

1. **SIEM:** The SIEM is already the aggregation point for security data, and can integrate threat intelligence for simple alerting to current attacks. Some tools offer some visualization, as well as minimal asset tracking and risk scoring. The SIEM has key pieces of the puzzle and is therefore a possibility, but you'd need to bolt on the analytics, and perhaps even enlist an enhanced visualization mechanism. Additionally, the data management requirements of SDS require substantial advancement over existing SIEM offerings.

2. **Security Analytics:** By definition analytics tools have the advanced math capabilities, and can ingest a subset of security data and threat intel. We say 'subset' due to performance limitations of their data models. Advanced analytics does not lend itself to comprehensive security data collection, so the main concern for analytics platforms is how to scale internal and external data collection and analysis. Additionally, analytics tools typically include decent visualization to drill down into security data after an alert.

3. **Vulnerability Management:** These tools aren't really limited to their original scope any more. The vendors have been actively broadening their offerings over the past few years — focusing on analytics, systems management, reporting, and simple automation. The tools are already asset-focused with prioritization baked in.

4. **GRC (whatever that means):** Governance, Risk, and Compliance tools include central aggregation with visualization and dashboards, offering a basis for managing a security program. But security programs aren't generic and interchangeable, so these tools require significant customization to fit any organization's program. Customizability includes some scope for adding capabilities, such as more sophisticated analytics and threat intelligence integration.

With a few types of options to consider, how do you decide? It depends on which aspect of SDS is most interesting to you, and how much customization you want to perform. If you are focused on better alerts you will lean towards analytics and SIEM offerings. If you want more effective visualization and dashboards, vulnerability management may be a better initial choice.

Summary

This paper discusses how to go beyond security monitoring, to integrate external threat intelligence and more advanced security analytics, to evolve towards Security Decision Support. We envision a day when a more scientific methodology — based on behaviors and activities within your extended technology environment, including the cloud and IoT — drives security decisions.

We envision a day when a more scientific methodology — based on behaviors and activities within your extended technology environment, including the cloud and IoT — drives security decisions.

You have a few technologies already in-house which could become your Security Decision Support platform — including SIEM, security analytics, and vulnerability management. Regardless of which direction you choose, you face compromises. None of the available tools satisfies all the requirements *today*. Yet if there is one thing we have learned over the past 20 years in security, it's that sooner rather than later, tools mature and add capabilities to meet more generalized needs.

Within a few years these types of tools will almost totally overlap into a more general *security management* umbrella embracing these broader concepts. To survive today's threats and adversaries you need a process to make best use of available resources by working smarter, leveraging technology and external resources to improve the effectiveness of your security team. That's what Security Decision Support is all about.

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 elQnetworks 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/>.