



SOC 2025: The Future of Security Operations Centers

Version 1.2
Released: May 24, 2022

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](#) and has been enhanced, reviewed, and professionally edited.

This report is licensed by Splunk.



www.splunk.com

The Splunk platform removes the barriers between data and action, empowering observability, IT and security teams to ensure their organizations are secure, resilient and innovative.

Founded in 2003, Splunk is a global company — with over 7,500 employees, Splunkers have received over 1,020 patents to date and availability in 21 regions around the world — and offers an open, extensible data platform that supports shared data across any environment so that all teams in an organization can get end-to-end visibility, with context, for every interaction and business process. Build a strong data foundation with Splunk.

Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.

<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>



SOC 2025

Table of Contents

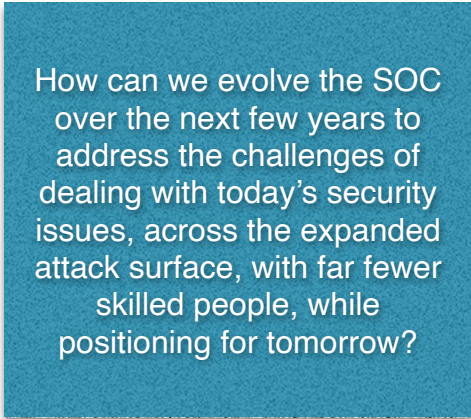
The Coming SOC Evolution	4
Making Sense of Security Data	7
Detection/Analytics	12
Operationalizing the SOC	18
About the Analyst	23
About Securosis	24

The Coming SOC Evolution

It's brutal running a security operations center (SOC) today. The attack surface expands exponentially as data moves to SaaS, applications move to containers, and infrastructure migrates to the cloud. The tools used by the SOC analysts are improving, but not fast enough. It seems adversaries remain one (or more) steps ahead. There aren't enough people to get the job done. Those you can hire typically need significant training, and retaining them is problematic. As soon as they are decent, they head off to their next gig for a huge bump in pay.

At the same time, security is under the spotlight like never before. Remember the old days when no one knew about security? Those days are long gone, and they aren't coming back. Thus, many organizations embrace managed services for detection and response, mainly because they have to.

Something has to change. Actually, a lot has to change. That's what this paper is about. How can we evolve the SOC over the next few years to address the challenges of dealing with today's security issues, across the expanded attack surface, with far fewer skilled people, while positioning for tomorrow?



How can we evolve the SOC over the next few years to address the challenges of dealing with today's security issues, across the expanded attack surface, with far fewer skilled people, while positioning for tomorrow?

SOC, what's it for?

We see two prominent use cases for the SOC. Detecting, investigating, and remediating attacks and substantiating the controls for audit/compliance purposes. We are not going to cover the compliance use case in this paper. Not because it isn't important, audits are still a thing, and audit preparation should still be done in as efficient and effective a manner as possible. Here we're going to tackle the evolution of the Security *OPERATIONS* Center, focusing on the detection, investigation, and remediation aspects of the SOC's job.

You can't say (for most organizations anyway) that there hasn't been significant investment in security tooling over the past five years. Or ten years. Whatever your timeframe, security budgets have increased dramatically. Of course, there was no choice given the expansion of the attack surface and the complexity of the technology environment. But if the finance people objectively look at the spending on security, they can (and should) ask some tough questions about the value the organization receives from those significant investments.

And there is the rub. We, as security professionals, know that 100% security is impossible. No matter how much you spend, you can (and will) be breached. We can throw out platitudes about reducing the dwell time or make the case that the attack would have been much worse without the investment. And you are probably right. But as my driver's education teacher told me over 35 years ago, *"you may be right, but you'll still be dead."*

What we haven't done very well is manage to *Security Outcomes* and communicate those achievements. What do we need the outcome to be for our security efforts? Our mindset needs to shift from activity to outcomes. So what is the outcome we need from the SOC? We need to find and fix security issues before data loss. That means we have to sharpen our detection capabilities and dramatically improve and streamline our operational motions. There is no prize for finding all the vulnerabilities. Like there are no penalties for missing them. The SOC needs to master detecting, investigating, and turning that information into effective remediation **before** data is lost.

Improved Tooling

Once we've gotten our arms around the mindset shift in focusing on security outcomes, we can focus on the how. How will the SOC get better at detecting, investigating, and remediating attacks? That's where better tooling comes into play. The good news is that SOC tools are much better than even five years ago. Innovations like improved analytics and security automation give SOC's far better capabilities. *But only if the SOC uses them.*

The good news is that SOC tools are much better than even five years ago. Innovations like improved analytics and security automation give SOC's far better capabilities. *But only if the SOC uses them.*

What SOC leader in their right mind *wouldn't* take advantage of these new capabilities? In concept, they all would and should. In reality, far too many haven't and can't. The problem is one of culture and evolution. The security team can handle detection and even investigation. But remediation is a cross-functional effort. And what do security outcomes depend on? You guessed it – remediation. So at its root, security is a team sport, and the SOC is one part of the team.

This means addressing security issues needs to fit into the operational motions of the rest of the organization.

The SOC can and should automate, especially the things within their control. But most automation requires buy-in from the other operational teams. Ultimately if the information doesn't consistently and effectively turn into action, the SOC fails in its mission.

Focused Evolution

We will deal with both internal and external evolution in this paper. We'll start by turning inward and spending time understanding the evolution of how the SOC collects security telemetry from both internal and external sources. Given the sheer number of new data sources that must be considered (IaaS, PaaS, SaaS, containers, DevOps, etc.), aggregating the right data is the first step in the battle.

Next, we'll tackle detection and analytics since that is the lifeblood of the SOC. Again, you get no points for detecting things, but you have no chance of achieving desired security outcomes if you miss attacks. The analytics area is where the most innovation has happened over the past few years, so we'll dig into some use cases and help you understand how frameworks like ATT&CK and buzzy marketing terms like eXtended Detection and Response (XDR) should influence your SOC plans.

Finally, we'll wrap up the paper by taking the what (accurate detections) and turning them into the how (effective remediation), resulting in positive security outcomes. Operationalizing is a critical concept in that context.

Making Sense of Security Data

Intelligence comes from data. And there is no lack of security data, that's for sure. Everything generates data. Servers, endpoints, networks, applications, databases, SaaS services, clouds, containers, and anything else that does anything in your technology environment. Just as there is no award for finding every vulnerability, there is no award for collecting all the security data. You want to collect the right data to make sure you can detect an attack before it becomes a breach.

As we consider what the SOC will look like in 2025, given the changing attack surface and available skills base, we've got to face reality. The sad truth is that TBs of security data sit underutilized in various data stores throughout the enterprise. It's not because security analysts don't want to use the data. They don't have a consistent process to evaluate ingested data and analyze it continuously. But let's not get the cart before the proverbial horse. First, let's figure out what data will drive the SOC of the Future.

The sad truth is that TBs of security data sit underutilized in various data stores throughout the enterprise. Security analysts don't have a consistent process to evaluate ingested data and analyze it constantly.

Security Data Foundation

The foundational sources of your security data haven't changed much over the past decade. You start with the data from your security controls because 1) the controls are presumably detecting or blocking attacks, and 2) you still have to substantiate the controls in place for your friendly (or not so friendly) auditors. These sources include logs and alerts from your firewalls, IPSs, web proxies, email gateways, DLP systems, identity stores, etc. You may also collect network traffic, including flows and even packets.

What about endpoint telemetry from your EDR or next-gen EPP product? Endpoint data has a renewed interest because remote employees don't always traverse the corporate network, resulting in a blind spot regarding their activity and security posture. On the downside, endpoint data is plentiful and can create issues in scale and cost. The same considerations must be weighed regarding network packets as well.

But let's table that discussion for a couple of sections since there is more context to discuss before truly determining whether you need to push all of the data into the security data store.

Use Cases

Once you get the obvious stuff in there, you need to go broader and deeper to provide the data required to evolve the SOC with advanced use cases. That means (selectively) pulling in application and database logs. You probably had an unpleasant flashback to when you tried that in the past. Your RDBMS-based SIEM fell over, and it took you three days to generate a report with all the needed data. But hear us out; you don't need to get all the application logs, just the relevant ones.

This brings us to the importance of threat models when planning use cases. That's right, old-school threat models. You figure out what is most likely to be attacked in your environment (think high-value information assets) and then work backward. How would the attacker compromise the data or the

This brings us to the importance of threat models when planning use cases. You figure out what is most likely to be attacked in your environment (think high-value information assets) and then work backward.

device? What data would you need to detect that attack? Do you have that data? If not, how do you get it? Aggregate and then tune. Wash, rinse, and repeat for additional use cases.

We know this doesn't seem like an evolution; it's the same stuff we've been doing for over a decade, right? Not exactly; the analytics you have at your disposal are much improved yet continue to be constrained by the availability of security data. Yet you can't capture all the data, so focus on the threat models and use cases that can answer the questions you need to know.

Cloud Sources

Given the cloudification of seemingly everything, we need to mention two (relatively) new sources of security data: your IaaS (infrastructure as a service) providers and SaaS applications. Given the sensitivity of the data going into the cloud, over the seemingly dead bodies of the security folks that would never let that happen, you're going to need some telemetry from these environments to figure out what's happening, if those environments are at risk, and ultimately to be able to respond to potential issues. Additionally, you want to pay attention to the data moving to/from the cloud, as detecting when an adversary can pivot between your environments is critical.

Is this radically different from the application and database telemetry discussed above? Not so much in content, but absolutely in location. The question then becomes what and how much of the cloud security data do you centralize?

What About External Data?

Nowadays, you don't just use your data to find attackers. You use other people's data, or in other words, threat intelligence, which gives you the ability to look for attacks that you haven't seen before. Threat intel isn't new either, and threat intel platforms (TIP) are being subsumed into broader SOC platforms or evolving to focus more on security operations or analysts. There are still many sources of threat intel, some commercial and some open source. The magic is understanding which sources will be helpful to you. That involves curation and evaluating the relevance of the third-party data. As we contemplate the security data that will drive the SOC, effectively leveraging threat intel is a cornerstone of the strategy.

Chilling by the (Security Data) Lake

In the early days of SIEM, there wasn't a choice of where or how you would store your security data. You selected a SIEM, put the data in there, started with the rules and policies provided by the vendor, tuned the rules and added some more, generated the reports from the system, and hopefully found some attacks. As security tooling has evolved, now you've got options for how you build your security monitoring environment.

Let's start with aggregation. Or what's now called a *security data lake*. This new terminology indicates that it's not your grandad's SIEM. Rather it's a place to store significantly more telemetry and use it better. This new-fangled data lake doesn't have to be new at all or even a data lake. You still have the option to buy a SIEM, have it ingest and process your security data, and generate alerts. Same as it ever was, but with a shiny new name.

Alternatively, you can use a vendor's multi-tenant cloud-based aggregation service, collecting your telemetry and doing the analytics within their cloud estate. Like other SaaS services, you get out of operating the infrastructure. You also use the vendor's analytics and other ancillary services, like SOAR, because it's a closed environment.

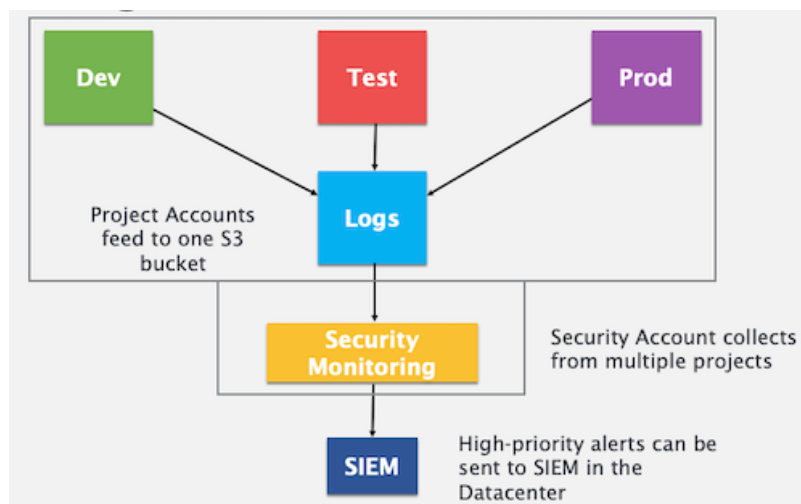
Finally, you can do it yourself (DIY). The DIY option involves using a modern data store (like Mongo or Snowflake) to store your data. You find or build the connectors for ingestion. You use the platform's (or a third party) analytics capabilities to design detections and generate your reports. It seems like a lot of work, but it's an option.

The next decision is how many lakes do you need? Traditionally, you'd centralize your data because it made no sense to operate two on-prem SIEMs. But given the plethora of options in terms of on-prem vs. cloud vs. managed services, it's feasible to have different security monitoring environments covering other areas of your infrastructure. This decision comes down to the operational motions occurring when an alert fires.

If you decide to centralize the security data, should it be added to your SIEM, or do you migrate to a cloud-based environment? This depends on your future platforming strategy. If the stated direction is to be cloud-first and migrate existing data and applications to the cloud, then your decision on location is easy – choose the cloud. Then do you leave your existing SIEM in place to handle the on-prem systems? Again, that'll depend on the operational motions.

And security isn't the only consideration. It's out of the scope of this research to consider application observability, but if you have moved to a model where the DevOps team operates the application and takes on some security responsibility, looking at an integrated platform that monitors for both security and performance makes sense.

We generally recommend a cascading log architecture for modern infrastructures (read cloud and DevOps), where all of the data is aggregated within the application stack. Then security-relevant data is moved to a separate security repository, which only the security team can access. Security analytics is done here, and then alerts (and other relevant context) go to the security operations group, be it an on-prem or cloud-based environment.



This aligns with the move to decentralize technology efforts, but will it meet the security requirements? Is it possible to do XDR (extended detection and response) if you don't capture and centralize all security data, including network, cloud, and endpoint data? That will be one of the critical discussions later in the paper, so let's defer that question for the time being.

That's the balancing act. You want to leverage data sources of all types to identify complex attacks. But it needs to be done in the most cost-efficient way.

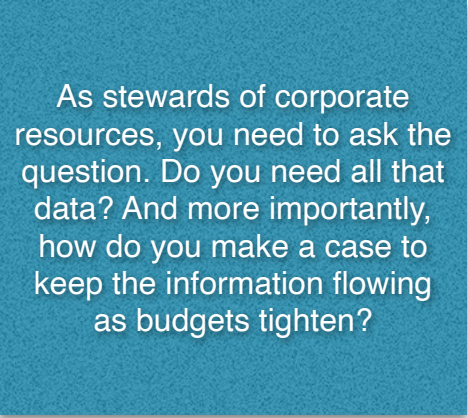
Stack Ranking

Let's wrap up with data retention strategies. It seems security data sources never go away. And you'll have more data tomorrow than you have today. That's great for a security vendor, who gets to expand the system in perpetuity, but less so for your CFO, who's trying to figure out how to find eight figures to pay to a SIEM vendor.

As security purists, we want to keep security data as long as possible. That makes sense if you're finding low and slow attacks or determining the proliferation of an attack. However, there are logical constraints on storage availability and system performance (on-prem) or cost (cloud-based).

As stewards of corporate resources, you need to ask the question. Do you need all that data? And more importantly, how do you make a case to keep the information flowing as budgets tighten? You need a consistent approach to evaluate the usefulness of both internal and external security data. We're going to take a page from sales and marketing to track the impact of these data sources. Highly functioning go-to-market teams can quantify the value of their campaigns, events, and other marketing spend based on what contributes to closed sales. We want to take a similar approach with incidents, figuring out which security data source contributes to impactful detections.

Adding a "data source" attribute to each incident (or even alert) gives you a sense of which security data sources provide value. You could get very granular in identifying where the information source provides value (earlier is better) in the detection process and possibly determine whether it's a primary or secondary detection source. Since most SOC's don't know which incidents are triggered from which data sources, starting with anything will help.



As stewards of corporate resources, you need to ask the question. Do you need all that data? And more importantly, how do you make a case to keep the information flowing as budgets tighten?

Once you have the incidents instrumented, you evaluate which sources are most effective and stack rank them. We recommend you run this exercise periodically (quarterly or so) and scrutinize which sources provide value, giving you the option to double down on the stuff that works and move away from the stuff that doesn't.

And you can feel good about that decision because it's based on data.

Detection/Analytics

Unfortunately, a lake of security data doesn't find attackers, so now we have to use it. Security analytics has been all the rage for the past ten years, resulting in many security analytics companies having emerged promising to make sense of all of this security data.

It turns out analytics aren't separate; they are part of every security thing. That's right, analytics drive endpoint security offerings. Cloud security products? Yup. Network security detection? Those too. It's hard to envision a security company of scale without analytics playing a central role in providing value to its customers.

As a security leader, what do you have to know about analytics and detection as you figure out how the SOC should evolve? First, it's not about [analytics technique A] vs. [analytics technique B]. It's about security outcomes, and to get there, you'll need to start thinking about the *SOC platform*.

Defining the SOC “Platform”

The initial stab at the SOC platform already exists with some overlapping capabilities. You already have a security monitoring capability, maybe an on-prem SIEM. As discussed earlier, the SOC platform should include threat intelligence. Currently, some organizations use a separate threat intel platform (TIP) to curate and prioritize the incoming external data. The third leg of the SOC platform is operations, where validating, verifying, and ultimately addressing any alerts happens. We'll have a lot to say about security operations later in the paper.

Though the evolved security operations platform may seem to be bolting together a bunch of stuff you already have, we are advocating for an evolutionary approach in the SOC. You certainly could ditch the existing toolset and start from scratch, and as liberating as that may be, it's not practical for most organizations. For instance, you've spent years tuning your on-prem SIEM to handle existing infrastructure, yet you have to keep the SOC operating. It's not like the attackers will give you a break to accommodate your platform migration. Thus, it may not make sense to scrap it. Yet.

Although you do have to decide where the SOC platform will run, here are some considerations:

- **Data Location:** It's better to aggregate data as close to the originating platform as possible. You keep cloud-based security data in the cloud, and on-prem systems go into an on-prem repository. That minimizes latency and cost. In addition, you can centralize alerts and context if your operational motions dictate.

- **Operations Approach:** Once the alert fires, what then? You'll need to centralize if you have an operations team that handles both cloud and on-prem issues. The next question becomes do you consolidate the raw security data, or just the alerts and context?
- **Care and Feeding:** How much time and resources do you want to spend keeping the monitoring system up and running? There are advantages to using a cloud-based, managed platform that gets you out of the business of scaling and operating the infrastructure.

The long-term trend is towards a managed offering in the cloud, but your migration strategy depends on how quickly you need to get there. If you've decided that your existing SIEM is not salvageable, you pick a new platform for everything and migrate as quickly as possible. But we see many organizations taking a more measured approach, focusing on building the foundation of a new platform that can handle the distributed and hybrid nature of computing in the cloud age while continuing to use the legacy platform during the migration.

We see many organizations taking a more measured approach, focusing on building the foundation of a new platform that can handle the distributed and hybrid nature of computing in the cloud age while continuing to use the legacy platform during the migration.

Analysis

Once you have internal and external data collected and aggregated, you analyze the data to identify the attacks. Easy, right? Unfortunately, there is a lot of noise and vendor puffery about how the analytics work, making it confusing to figure out the best approach. Let's work through the different types of techniques used by SOC tools.

1. **Rules and Reputation:** Let's start with signature-based controls, the old standard. You know, the type of correlation your RDBMS-based SIEM performed for decades. Adding patterns enumerated in the ATT&CK framework helps narrow the scope of what you need to look for, but you still need to recognize the attack. You'll need to know what you are looking for.
2. **Machine Learning:** The significant evolution from simple correlation is the ability to detect an attack you haven't seen. Advanced analytics can be used to define an activity baseline, and with that baseline defining normal behavior within your environment, your detection engine can look for anomalies.

Digging into the grungy math of different machine learning models and cluster analyses probably won't help you find attackers faster and more effectively. Continue to focus on the security outcomes during your evaluation. Does it find attacks you are likely to see? How much time and effort will it take to isolate the most impactful alerts? What's involved in keeping the platform current? And ultimately, how will the platform's analytics make the team more efficient? Stay focused on ensuring any new platform makes the team better, not on who's math is better.

Use Cases

You may be bored (and maybe frustrated) with our constant harping on the importance of use cases in detecting attacks. There is a method to our madness in that use cases make a pretty nebulous concept more tangible. So let's dig into a handful of use cases to understand how a SOC platform will favorably impact your detection efforts.

Ransomware

Ransomware doesn't seem to get as many headlines nowadays, but don't be fooled by the media's short attention span. Ransomware continues to be a scourge, and every company remains vulnerable. Let's examine how an evolved SOC handles ransomware? First, ransomware isn't new, particularly not the attacks — it typically uses commodity malware for the initial compromise. Attackers are more organized and proficient — once they have a foothold within a victim's network, they perform extensive reconnaissance to find and destroy backups, increasing pressure to pay the ransom.

Many of your controls will be lighting up during a ransomware outbreak, and the SOC tooling can take those alerts and pinpoint the type and extent of the outbreak, minimizing noise and focusing efforts on the root cause of the attack.

The evolved SOC platform integrates telemetry and data from external and internal sources, using analytics to identify malicious intent. Many of your controls will be lighting up during a ransomware outbreak, and that's good. The SOC tooling can take the alerts from each of the controls and pinpoint the type and extent of the outbreak, minimizing noise and focusing efforts on the root cause of the attack.

For example, you'll first see the ransomware outbreak on the network, along with ongoing reconnaissance and attempts to compromise additional devices. Then you'll see privilege escalation on a compromised device and likely some more recon to identify the backup systems and other vital assets. The SOC platform can weave these data sources together to provide invaluable context to help responders understand the situation faster, making a huge difference in containing the damage.

Threat Hunting

Threat hunting is proactively looking for attackers in your environment before getting an alert from another detection method. Hunting involves more art than science as hunters start with a plan to look for certain activities indicating active adversaries. Then they mine security data to find and follow an attacker's trail to identify what the attackers have done and project what they'll do next. A modern SOC platform provides broad and deep security data collection and the ability to pivot through data effectively.

No security tool will turn an entry-level analyst into a world-class hunter. But the SOC platform can accelerate and improve any reasonably capable security professional's hunt. Further helping the hunter are common queries, typically pre-loaded into detection tools to kickstart hunting efforts. These rules don't make the hunt, but they can codify common searches likely to uncover malicious activity — including drive-by attacks, spear phishing, privilege escalation, credential stuffing, and lateral movement.

Insider Threat

The classic “inside job” typically involves an employee acting maliciously to steal data or sabotage systems. The insider knows the company's defense's weak points, so this is a particularly impactful and damaging attack. We have a broader definition of insiders, as external actors with control of a device inside the network are technically insiders because they have access and privileges to internal resources.

You use the collection and analysis capabilities of the SOC platform to identify anomalous activity from employees (and their devices). Given that insiders can be anywhere, you'll need a broad collection effort, including telemetry from all remote employees and cloud resources. In terms of analysis, deviations from the typical activity baselines are the strongest indicators of malicious activity. This capability was called UEBA (user and entity behavioral analytics), but nowadays it's considered just another use case for a SOC platform.

ATT&CK

We learned in Security 101 class that understanding TTP (Tactics, Techniques, and Procedures) is key to detecting attacks. These help piece together attack timelines, accelerating the assessment of attack damage and proliferation. But the challenge of an attack timeline is that it occurs after the attack has succeeded. It's helpful to ensure protection from that same attack vector in the future, but you are still blind to attacks you don't know.


What if you could get a list of most of the attack techniques you're likely to see in your environment? You could use that list to tune your detection techniques to what you'll likely see coming over the transom. You could identify gaps in your data sources to eliminate blind spots to detection. You could even test your controls against these attacks to determine holes in your defenses. MITRE and other industry contributors have given us this kind of list. It's called the [ATT&CK framework](#). Of course, you don't find attackers by running around with a list of TTPs. But ATT&CK can help focus your efforts.

The ATT&CK framework enumerates hundreds of attack techniques. And it's increasingly expected that SOC tooling will map data into ATT&CK, so you can visualize the gaps in your defenses and determine how to address any deficiencies most effectively.

The ATT&CK framework enumerates hundreds of attack techniques. And it's increasingly expected that SOC tooling will map data into ATT&CK, so you can visualize the gaps in your defenses and determine how to address any deficiencies most effectively.

XDR

Let's start with a definition: "Extended Detection and Response (XDR) collects and correlates data across multiple security layers - network, servers, endpoint, and cloud." So how is that different than what we've been doing with SIEM for two decades? Yeah, it's not significantly different. But all the same, XDR is very shiny, and any company doing security detection is planting a flag on Mt. XDR. Since detection is one of the critical functions of the SOC, that means we need to discuss where XDR fits into your SOC tooling.



Suppose you consider XDR an uber-detector assuming the responsibilities of SIEM, EDR, NDR, and even some cloud detection leveraging a common datastore and performing advanced analytics to find anomalies and attack indicators for broader, multi-faceted campaigns. In that case, that sounds pretty good, no? Who doesn't want to collapse various monitoring tools into one platform that handles everything more effectively?

Here's the problem, at least in XDR's current incarnation. It doesn't speak to the reality of the installed base within most enterprises. You already have security monitoring technology, which you probably don't want to (or can't) forklift. You have a compliance requirement, so you've got to keep your SIEM around. While we believe the idea of a standard, open way to integrate the telemetry from all of your sensors and do advanced analysis is great, *it's all in how you get there*.

You have options to achieve the promise of XDR. If you need to overhaul your SOC tooling and have limited compliance requirements, considering an XDR would make sense. But it also may make sense to embrace an “Open XDR” concept that allows you to leverage existing sensors and analytics with supplemental solutions to address weaknesses in the monitoring environment. You can also consider how your existing monitoring vendor(s) are moving towards their vision of XDR, which may satisfy your requirements. Both revolution and evolution are reasonable paths.

Operationalizing the SOC

Yet, there is an inconvenient fact that warrants discussion. Unless someone does something with the information, the best data and analytics don't end in a positive security outcome.

Security success depends on consistent and effective operational motions. Sadly, this remains a commonly overlooked aspect of building the SOC. As we wrap up the paper, we're going to go from alert to action and do it effectively and efficiently, every time (consistently), which we'll call the 3 E's. The goal is to automate everything that can be automated, enabling the carbon (you know, humans) to focus on the things that suit them best. Will we get there by 2025? That depends on you, as the technology is available, it's a matter of whether you use it.

The 3 E's

First, let's be clear on the objective of security operations, which is to facilitate positive security outcomes. Ensuring these outcomes is to focus on the 3 E's.

- **Effectiveness:** With what's at stake for security, you need to be right because security is asymmetric. The attackers only need to be right once, and defenders need to defeat them every time. In reality, it's not that simple, as attackers do need to string together multiple successful attacks to achieve their mission, but that's beside the point. A SOC that only finds stuff sometimes is not successful. You want to minimize false positives and eliminate false negatives. If an alert fires, it should identify an area of interest with sufficient context to facilitate verification and investigation.
- **Efficiency:** You also need to do things as quickly as possible, consuming a minimum of resources due to limited available resources and the significant damage (especially against an attack like ransomware) that can happen in minutes. You need tooling that makes the analyst's job easier, *not harder*. You also need to facilitate the communication and collaboration between teams to ensure escalation happens cleanly and quickly. Breaking down the barriers between traditional operational silos becomes a critical path to streamlining operations.
- **Every Time (Consistency):** Finally, you need the operational motions to be designed and executed the same way, every time. But aren't there many ways to solve a problem? Maybe. But as you scale up your security team, having specific playbooks to address issues makes it easy to onboard new personnel and ensure they achieve the first two goals: Effectiveness and Efficiency. Strive to streamline the operational motions (as associated playbooks) over time, as things change and as you learn what works in your environment.

Do you get to the 3 E's overnight? Or course not. It takes years and much effort to get there. But we can tell you that you never get there unless you start the journey.

Defining Playbooks

The first step to a highly functioning SOC is being *intentional*. You want to determine the proper operational motions for categories of attacks **before** you have to address them. The more granular the playbook, the less variance you'll get in the response and the more consistent your operations. Building the playbooks iteratively allows you to learn what works and doesn't, tuning and refining the playbook every time you use it. These are living documents and should be treated as such.

So how many playbooks should you define? As a matter of practice, the more playbooks, the better, but you can't boil the ocean, especially as you get started. Begin by enumerating the situations you see most frequently. These typically include phishing, malware attacks/compromised devices, ransomware, DDoS, unauthorized account creation, and network security rule changes. To be clear, pretty much any alert could trigger a playbook, so ultimately you may get to dozens, if not hundreds. But start with maybe the top 5 alerts detected in your environment and start with those.

What goes into a playbook? Let's look at the components of the playbook:

- **Trigger:** Start with the trigger, which will be an alert and have some specific contextual information to guide the next steps.
- **Enrichment:** Based on the type of alert, there will be additional context and information helpful to understanding the situation and streamlining the analyst's work handling the issue. Maybe it's DNS reputation on a suspicious IP address or an adversary profile based on the command and control traffic. You want to ensure the analyst has sufficient information to dig into the alert immediately.
- **Verification:** At this point, a determination needs to be made as to whether the issue warrants further investigation. What's required to make that call? For a malware attack, maybe it's checking the email gateway for a phishing email that arrived in the user's inbox. Or a notification from the egress filter that a device contacted a suspicious IP address. You want to list the facts that will lead you to conclude that this is a real issue and assess the severity of each trigger.
- **Action:** Upon verification, what actions need to be taken? Should the device be quarantined and a forensic image of the device be captured? Should an escalation of privileges or firewall rule change get rolled back? You'll want to determine what needs to be done and document that motion in granular detail, so there are no questions about what should be done. You'll also look for automation opportunities.
- **Confirmation:** Was the action step(s) successful? Next, confirm whether the actions dictated in the playbook happened successfully. This may involve scanning the device (or service) to ensure the change was rolled back or making sure the device is not accessible anymore to an attacker.

- **Escalation:** What's next? Does it get routed to a 2nd tier for further verification and research? Is it sent directly to an operations team to be fixed if it can't be automated? Can the issue be closed out because you've gotten the confirmation that the issue was handled? Be specific in where the information goes, what format the supporting documentation needs to be delivered, and how you will follow up to ensure the issue has been addressed to completion.

Building playbooks is a skill, which means you'll be pretty bad when you start. The first playbook will take you a while. The next will go a bit faster, and with each subsequent playbook, you'll get the hang of it. By the time you've built the 10th, you'll start cranking them out. Also, factor in a feedback loop, ensuring you capture what works and what doesn't work every time the playbook runs. This practice of constant improvement is critical, given the dynamic nature of technology.

In terms of playbook design, modularity is your friend. There will be commonalities in terms of how you handle parts of the playbooks; for instance, connecting to devices or services can be standardized (via standard APIs), as can remediation actions (block this or roll back that), as well as escalations. If anything needs to be done across multiple playbooks, look to build a common module. This becomes even more important when automating the operational motions, where you can create scripts/code and reuse them across multiple playbooks.

You may find experienced security practitioners pushing back on strict adherence to the playbook approach, and they have a point. Your rock stars don't need that level of guidance, but it's not about them. Achieving a consistent response requires concrete actions to be defined and executed consistently, making it more likely that your less-experienced staffers will be able to complete the playbook successfully.

Automating (What Makes Sense)

Once the playbooks are stable, the operational motion will be effective (the first E). Next is to improve efficiency, which means figuring out which actions within the playbook can be automated. A modular playbook approach allows you to introduce automation where appropriate without reinventing the wheel for common actions.

How do we start this automation journey? First, you need to orchestrate between the different systems and then develop and deploy the automation.

- **Orchestrate:** Start with the playbook and define the devices and systems to be managed. Then you determine how to connect to and manage the devices. Do you need to use an API, build a script, or develop a home-grown integration? An advantage of using a commercial SOAR platform is that pre-built connectors already exist for most, if not all, end devices. These platforms also have a scripting language or visual studio to develop the connectors and automations. And let's not forget about security, given that you are managing these devices. How do you ensure proper authentication and authorization of any commands sent to the devices/services?

- **Develop (and Maintain):** Once you have connectivity to the device, you need to build, test and deploy the automation. Keep in mind that you are in the software business once you build automations. Maybe you can use a low-code environment, but it's still code, so you've got to decide who will maintain it. Also, consider who monitors for changes in the end device and updates the automation when new capabilities are available or needed and how you will fix defects, especially if they break the automation.
- **Deploy:** It's too bad the acronym is SOAR and not SO+B+D=R because you can't respond to anything until you deploy the automation. You'll need to select an execution environment and define the process to test and iterate until the automation is ready for production. That involves formal functionality testing (with actual, documented unit tests) and a burn-in period where the automation runs in monitor or debug mode to ensure it works as intended. Burn-in is critical because nothing will set back an automation program faster than bringing down systems. So the automations need to be READY before being deployed to production.

We don't bring up all of these sticky issues (like maintaining automation code) to scare you away from embracing automation. More to make the point that security teams need additional skills in the SOC of 2025. It's not that the days of the console expert have passed, but you'll also need staffers with specialized coding chops. And as more and more of security becomes code, more and more of the Ops skills you'll need will skew towards development.

So which categories of automations make sense to start? Job #1 is to build credibility, so start with functions that won't bring down systems or otherwise cause damage. Think about alert enrichment, quarantining compromised devices, and maybe blocking egress on known-bad IP addresses. As you get some quick wins and build your credibility, you can look at more sophisticated operational motions. By developing the automations modularly, you can string them together to implement advanced, multi-step processes.

Beyond 2025

Although the technology to get to SOC 2025 is here today, most organizations will take the next 2-3 years to accept this new approach culturally. We fully expect most organizations to adopt a more flexible data collection and aggregation approach and introduce more sophisticated analytics in this timeframe. Automation for alert enrichment, policy changes (block known bad IP addresses), and quarantine will become commonplace. We also expect to see some aspects of security automation built directly into application stacks, especially as organizations increasingly move to the cloud and build all code using CI/CD pipelines.

But what happens then? Let's look beyond the mid-term planning horizon to what's in store for SOCs.

1. **Security Data Lakes:** Many SOC's send telemetry to different places. The first is the SIEM for short-term correlation, alerting, and reporting. But many SIEMs can only maintain 60-90 days of data without killing performance or breaking the bank. Thus, telemetry is also sent to object storage (typically in the cloud) for longer-term storage, forensics, and cheap archival. Why not handle both objectives on one platform? An emerging approach called the security data lake indexes telemetry in object storage quickly and cost-effectively. This separates the analytics plane from the data plane, providing more scale for a lower price. Some recent entrants in the SIEM market use this model within their platforms, offering a buy (versus build) option for those interested in the approach.
2. **Security Justification:** If you stick around long enough, you'll see the security cost pendulum swing back and forth. First, the budget is there, and then it's not. You'll need to justify security expenditures in belt-tightening times, especially for threat intel and specific controls. The best way to justify spending is to substantiate value by instrumenting your SOC processes to attribute alerts and remediated issues to specific intel sources and controls. It's like how sales teams use their CRM systems to track lead sources to determine the effectiveness of marketing programs and campaigns.

But first things first, there is a lot to do before we get to SOC 2025. Start by making sense of your security data and more effectively analyzing it. Being intentional about your SOC motions and systematically focusing on effectiveness and efficiency will get you to the promised land of consistency. Before long, you'll be executing security operations flawlessly every time.

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. 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 was an analyst at META Group prior to founding SHYM Technology and then held executive roles at CipherTrust and TruSecure. Mike then 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, Mike joined Securosis with a rejuvenated cynicism about the state of security.

Mike published [The Pragmatic CSO](http://www.pragmaticcco.com/) <http://www.pragmaticcco.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.

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 several 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 tells 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. Keep in mind that we maintain our strict objectivity and confidentiality requirements on all engagements.
- **Custom Research, Speaking, and Advisory:** Do you 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/>.