Introduction to Threat Operations

Version 1.6
Released: April 7, 2017

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

Author's Note

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

Special thanks to Chris Pepper for editing and content support.

This report is licensed by ThreatQuotient.

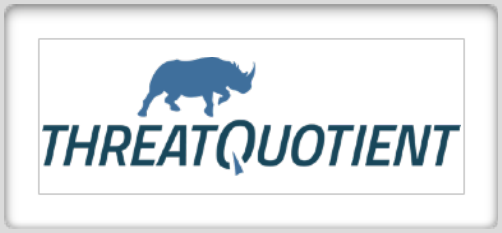

www.threatquotient.com

ThreatQuotient™ understands that the foundation of intelligence-driven security is people. The company's open and extensible threat intelligence platform, ThreatQ, empowers security teams with the context, customization and prioritization needed to make better decisions, accelerate detection and response and advance team collaboration. Leading global companies use ThreatQ as the cornerstone of their threat operations and management system, increasing security effectiveness and efficiency.

Copyright

Introduction to Threat Operations

Table of Contents

Thinking Differently

Let's start with a rhetorical question: Can you really 'manage' threats? Is that even a worthwhile goal? And how do you even define a threat? We have seen better descriptions of how adversaries operate by abstracting multiple attacks/threats into a *campaign*, capturing a set of interrelated attacks with a common mission. A campaign is a better way to think about how you are being attacked than the piecemeal approach of treating every attack as an independent event and defaulting to the traditional threat management cycle: Prevent (good luck!), Detect, Investigate, and Remediate.

The industry continues to be largely locked into this negative feedback loop: you are attacked, you respond, you clean up the mess, and then you start all over again.

That approach hasn't worked out well. The industry continues to be largely locked into this negative feedback loop: you are attacked, you respond, you clean up the mess, and then you start all over again. You never learn much from the last attack, which sentences you to continue running on the hamster wheel day after day. But this inability to learn isn't from lack of effort. Pretty much every practitioner we talk to wants better leverage, and to learn from attacks in the wild. The problem is that existing security controls and monitors don't really support that level of learning — not easily, anyway.

But inability to learn isn't our only challenge. Current threat management largely ignores the actual risk presented by an attack. Without some understanding of what an attacker is trying to do, you cannot prioritize effort intelligently. For example, if you look at threats independently, an advanced attack on your application may take priority due to its sophistication, which implies a capable attacker. So you take obviously capable attackers more seriously than simple phishing attacks.

But that can be a bad assumption. Advanced attackers seek the path of least resistance to compromise your environment. So if a phishing message will do the trick, they'll phish your folks. They won't waste a zero-day attack when a simple email will suffice. On the other hand you could be right that the phishing attempt just represents a kid in a basement trying to steal milk money. There is no way to know without a higher-level abstraction of attack activity, so current methods of prioritization are very haphazard.

Speaking of prioritization, we cannot afford hit-and-miss approaches any more. The perpetual (and worsening) security skills gap means we **must** make better use of limited resources. The main cost of false positives is time wasted by scarce and valuable folks. They need to work on the endless list of real attacks, not go on wild goose chases. And you probably don't have enough people to validate and triage all the alerts streaming out of the monitoring systems so things get missed, and it's all too easy to become a target of pissed-off customers, class-action lawyers, and regulators after a breach.

Clearly security hasn't been effective enough for a long time, and with the increasing complexity of technology infrastructure and high-profile security breaches, the status quo isn't acceptable any more.

But we aren't done. Once you figure out which attacks to focus on, current security programs generally remediate issues manually and serially. It's just another game of Whack-A-Mole, where you direct Operations to patch or reimage a machine, wait for the next user to click similar malware, and the next device gets compromised *exactly the same way*. Wash, rinse, repeat. That is no good either.

Clearly security hasn't been effective enough for a long time, and with the increasing complexity of technology infrastructure and high-profile of security breaches, the *status quo* isn't acceptable any more. Something needs to change, quickly.

Thinking Differently

Everybody loves people who think differently. Until they upset the apple cart and start agitating for massive change, upending the way things have always been done. As explained above, security has reached a point where we need to start going beyond what we've always done, because we cannot keep pace with attackers this way, or stem the flow of sensitive data being exfiltrated.

The move to cloud computing, succinctly described in our Tidal Forces posts (1, 2, 3), is going a long way toward destroying the security *status quo* because security is fundamentally different in the cloud. And if we could just do a flash cutover of all our systems onto well-architected cloud stacks, a lot of these issues would go away. Not all, but many.

Unfortunately a flash cutover isn't feasible. A huge amount of critical data still resides in corporate data centers, and will for the foreseeable future. So we need to maintain two realities in our minds for a while. First the older systems running in existing data centers, where we have to leverage traditional security controls and monitors. The other side is the new reality — enabled by cloud computing, mobility, and DevOps; architected for scale and security, presenting new governance and monitoring challenges.

It is tough to be a security professional, and it's getting harder. But senior management and boards of directors aren't interested in how hard your life is. They expect you to come up with answers.

So this “Introducing Threat Operations” paper is focused on addressing the issues which make dealing with attacks challenging:

- **Security Data Overload:** There is no lack of security data. Many organizations are flooded by it, without tools or expertise to manage it. These organizations often compound the issue by beginning to integrate external threat intelligence, magnifying the data overload.
- **Detection of Advanced and Dynamic Attacks:** But current security monitoring infrastructure is based on looking for attacks you have already seen. What happens with a custom attack built specifically for you? What if you want to actually *hunt* active threat actors in your environment? Either way you need to better utilize internal security data and intelligently leverage threat intelligence to look for attacks you haven’t seen yet.
- **Lack of Skilled Resources:** The sad fact is the industry can’t address the skills gap fast enough. We can and are focusing on education, but security requires a broad knowledge of technology and a lot of experience for effectiveness. So we need to make less experienced practitioners more effective, through smarter systems which guide them through their work. It’s not about replacing security analysts, but scaling up their impact.
- **Response and Remediation at Scale, Coordinated with Operations:** Finally, once you figure out what to fix, you face similar resource constraints dealing with Operations. The key is to figure out how to intelligently orchestrate and automate attack response and remediation.

It’s not just about managing threats any more. We need to build operational process to more effectively handle hostile campaigns.

We are talking about evolving how the industry deals with threats. It’s not just about managing threats any more. We need to build operational process to more effectively handle hostile campaigns. That requires leveraging security data through better analytics, magnifying the impact of the people we have by structuring and streamlining processes, and automating threat remediation wherever possible.

Accelerating the Human

With all the internal and external security data available, and the increasing sophistication of analytics, organizations should be doing a better job of handling threats. If what you are doing isn't working it is time to think differently about the problem and address the root causes underlying your inability to handle threats. We believe you need to focus on *accelerating the human*, making your practitioners better through training, process, and technology.

With all the focus on orchestration and automation in security, it would be easy to conclude that carbon-based entities (that's us) are on their way out as the key driver of security programs. That couldn't be further from reality.

With all the focus on orchestration and automation in security, it would be easy to conclude that carbon-based entities (that's us) are on their way out as the key driver of security programs. That couldn't be further from reality. If anything, as technology infrastructure continues to get more complicated and adversaries continue to improve, humans are *increasing* in importance due to limitations of analytics (particularly the potential for false positives) and downsides of automation run amok (taking down networks). That means your best investments will be in making your security team more effective and efficient at the ever-growing list of increasingly complex tasks they need to perform.

One of the keys in our Security Analytics Team of Rivals paper is the need to use the right tool for the job. That goes for humans too. Our security functions need to be delivered via a mix of technology and personnel, letting each do what it does best. The focus of our operational discipline is finding the proper mix.

Let's flesh out *Threat Operations* with more detail.

- **Harnessing Threat Intelligence:** Enterprises no longer have the luxury of time to learn from attacks they have seen, and then more time to adapt defenses accordingly. You need to learn from attacks on others using external threat intelligence. Make sure you can detect those attacks even if you've never seen them directly. Of course you can easily be overwhelmed with external threat data, so the key to successfully harnessing threat intel is focusing on attacks that pertain to you.

- **Enriching Alerts:** Once you have a general alert you need to add more information to eliminate as much as possible of the busywork many analysts need to perform just to figure out whether an alert is legitimate and critical. The data to enrich the alerts exists within your systems — it's just a matter of centralizing it where your analysts can use it.
- **Building Trustable Automation:** A set of attacks can be handled without human intervention. Admittedly that set is fairly small right now, but opportunity for automation is increasing quickly. As we have been saying for quite a while, the key to automation is *trust*, in the face of automation's dangers. Operations people need confidence that any changes you request won't crater your environment.
- **Workflow and Process Acceleration:** Finally, a Threat Operations mindset requires you to streamline how you handle threats, and apply structure as appropriate to provide leverage and consistency for staff. You need to find a balance between letting skilled practitioners do their thing, and providing the structure necessary to lead less sophisticated practitioners through a security process.

All these tasks focus on one objective: providing more context for analysts to accelerate efforts to detect and address threats in the organization — accelerating the human.

Harnessing Threat Intelligence

We have long believed threat intelligence can be a great equalizer, helping to restore some balance to the struggle between defender and attacker. For years the table has been slanted toward attackers, who can target a largely unbounded attack surface with increasingly sophisticated attack tools. Sharing data about attacks, enabling organizations to proactively look for new attacks *before* they get hit, can help alleviate this asymmetry.

But threat intel is an unwieldy beast, involving hundreds of potential data sources (some free, others paid) in a variety of data formats, which need to be aggregated and processed to be useful.

But threat intel is an unwieldy beast, involving hundreds of potential data sources (some free, others paid) in a variety of data formats, which need to be aggregated and processed to be useful.

To harness this data you need to:

1. **Integrate:** You need to centralize all your data. First the external data, because without eliminating duplicates and ensuring accuracy and relevance, your analysts will spend even more time spinning their wheels on false positives and useless alerts.

2. **Reduce Overlap and Normalize:** With all this data there is inevitable overlap among the attacks and adversaries tracked by different providers. Efficiency demands you address this duplication before putting analysts to work. This involves cleaning up your threat repository by looking for commonalities among indicators, and normalizing differences between data from different feeds.

3. **Prioritize:** Once you have all the threat intel in one place, you'll realize you have far too much data to address in any reasonable timeframe. This is where prioritization comes in: you need to address the *most likely* threats, generally sorted based on your industry and the types of data you are protecting. This requires some assumptions, which are likely to be wrong, so a tuning and feedback loop is essential.

4. **Drill down:** Sometimes analysts need to pull on threads within an attack report to find something useful in your environment. This is where human skills come into play. You want your analyst to drill into intel about the adversary and the specific threat, to give themselves the best opportunity to make non-obvious connections.

Ultimately threat intel, when fed into security monitors and controls, should provide an increasing number of the alerts your team handles. But an alert is only the beginning of the response process, and making each alert as detailed as possible saves analyst time. That's where enrichment enters the discussion.

Enriching Alerts

At this point you have an alert, and an analyst needs to validate it and assess it is criticality in your environment. They need context for these tasks. To streamline the processes of validation and assessment the analyst needs more information, faster, as they drill into the alert. For example, if they could see a set of attacks associated with unsophisticated attackers, that would facilitate prioritization.

On the other hand, it totally changes the analyst workflow to see indicators strongly linked to financial fraud, and recognize command and control traffic associated with a financial fraud botnet. If your organization is concerned about financial fraud (and who isn't?), this enrichment should bubble such alerts to the top of the list for priority investigation.

But that's just the automated stuff. If the alert is enriched with, say, a list of devices on your network which connected to the IP addresses associated with that botnet, you'd have a list of devices more likely to have been impacted by that attack. Or if you can pull from your change management system a list of devices which added a specific executable, which you recently learned was part of a broader attack, you are well ahead on figuring out where the attack spread.

Alert enrichment requires considerable foresight. You need to anticipate the kinds of questions analysts will ask based on what they find in alerts.

Alert enrichment requires considerable foresight. You need to anticipate the kinds of questions analysts will ask based on what they find in alerts. Obviously you run the risk of swamping them with too much data so you'll want proper instrumentation to know which data is useful and which isn't. It helps to have a feedback loop with analysts, to constantly track which data is helpful and which isn't.

With additional context on the severity of threats into your environment and attack types, you can figure out what can be remediated easily and what can't. Easy stuff is ripe for further automation.

Building Trustable Automation

Automation remains a divisive topic in security. Many operational folks are justifiably concerned about machines making changes which could cause downtime. But what choice do you have? You can't get all the work done with the resources you have, and bringing very expensive resources to bear on activities which don't add value is a waste. To bridge the gap we focus on building trust in automation.

How can you build this elusive trust? Start small. Select a set of very manageable use cases which make a difference operationally but have limited downside. A good place to look is outbound network connections. If you learn of a new phishing site, typically via threat intel, you can have your egress filters block connections. This is simple but very effective because many phishing sites are only up for a few hours. Having a staffer update the egress filter won't work — things move too fast for a human to keep up. But this is perfect for a machine. And downside is limited because if it is a false positive you'll hear about it quickly from an impacted employee. But that would likely be only a handful of folks — not the entire Finance department.

Implementing this kind of automation involves determining the trigger, designing the action, then deploying the controls. First you select the use case to look for. Maybe the trigger comes from your IDS or SIEM. Or perhaps it comes from a Threat Intelligence platform aggregating and analyzing external data for you. Once you know the data you are looking for, you set the alert in the system. When it fires, what happens? You need to design that set of actions, which usually involves changing something on an active control. Once that is designed you deploy the change to the security control.

Any automation that goes into production requires significant testing, but even well-tested automation isn't perfect. So to ensure trust in any kind of automation you'll want to be able to quickly walk back any changes you made, in case of unintended consequences. Operations will push back against things which can't be undone, so make sure your automation is designed with resilience in mind. Don't automate any changes which cannot be rolled back gracefully and quickly.

This is easy, right? In concept, especially for no-brainer use cases, it is. A new generation of orchestration and automation tools is emerging to handle this critical step. But as with most innovative security technology, this function is destined for eventual integration into broader security platforms.

With more contextual use of threat intelligence making alerts more impactful, and then automating changes which make sense, where do humans come into play? Unfortunately not everything fits into a nice clean set of circumstances you can model and tidily automate. Thus the importance of structuring operational process to be flexible yet consistent, allowing human analysts to diverge from standard procedures when necessary.

Structuring Activity

We understand process management is not nearly as sexy as hunting adversaries. But it's more important. Security teams face an unprecedented lack of talent to implement and run security programs. The secret to success has become rapidly improving the effectiveness of less sophisticated practitioners. That's a politically correct way to say making n00bs a bit less n00by.

Security teams face an unprecedented lack of talent to implement and run security programs. The secret to success has become rapidly improving the effectiveness of less sophisticated practitioners.

This requires looking at the process of how a responder works, and then implementing that process in a tool. Whether looking at an assembly line or a bank branch handling a deposit, a mature business function has a documented process representing best practices. So start by learning how your best folks do their jobs. Security folks mostly learned in the school of hard knocks, but they figure out the best way to get the job done.

Even if you have immature processes and don't know where to start, you can still use this approach to build these aspects of your security program. The good news is that many service organizations, as well as some product vendors, have already documented playbooks for many security processes and functions. As with automation, you start with a small set of playbooks, developing more and making them more sophisticated over time.

Look for opportunities to leverage other aspects of your threat operations process while implementing your playbooks. For example a quick response process could start with an alert from an endpoint security suite, enriched with information from a threat intel library, including a set of known malicious networks to automatically block at the outbound firewall. Then an analyst can start digging more deeply into the compromised device to figure out exactly what happened, and if any other devices in the environment have been impacted.

On the bright side, a very specific set of activities can be particularly useful for practitioners who aren't exactly sure what to do next. You already instrumented your platform to perform these functions, so folks can follow along to benefit from your best practice research.

You need to decide which playbooks to implement, and they don't help with issues which aren't in your playbooks. But innovation marches on, and you will soon be hearing about "cognitive security analytics," which can make some of these connections automatically using advanced analytics. It's very early for cognitive technology, but it looks promising, so that's something to keep an eye on.

The objective of threat operations is to a) stop thinking so tactically about dealing with only one attack at a time, and instead see the bigger-picture *threats* to your organization; and b) develop a consistent set of operational processes to make less-sophisticated practitioners more effective.

Threat Operations in Action

Let's bring these concepts into a scenario to make them more tangible. We'll tell the story of a high-tech component manufacturer named ComponentCo. Yes, we've been working overtime on creative naming. ComponentCo (CCo) products go into a leading smartphone platform, making their intellectual property a huge target of interest to a variety of adversaries with different motives.

- **Competitors:** Given CCo's presence inside a platform that sells hundreds of millions of units a year, their competition is keenly trying to close the technology gap. A design win is worth hundreds of millions in annual revenue, so these companies aren't above gaining parity any way they can.
- **Stock Manipulators:** Confidential information about new products and imminent design wins is gold to unscrupulous traders. But that is not the only interesting information. If they can see manufacturing plans or unit projections, they will gain insight into device sales, opening up another avenue to profit from non-public information.
- **Nation-states:** Many people claim nation-states hack to aid their companies. That is likely true, but just as attractive is the opportunity to backdoor hundreds of millions of devices by manipulating embedded components.

ComponentCo already invests heavily in security. They monitor critical network segments. They capture packets in the DMZ and data center. They have a solid incident response process. Given the money at stake, they have pretty much every new, shiny object that promises to detect advanced attackers. But they are not naive. They understand how vulnerable they are, mostly due to the sophistication of their various adversaries.

As with many organizations, fielding a talented team to execute on their security program is challenging. They have a high-level CISO, with funding to maintain a team of dozens of security practitioners. But it's not enough. So CCo is building a farm team. They recruit experienced professionals, but also high-potential system administrators from other parts of the business, who they train in security. Bringing on less experienced folks has produced mixed results — some of them were able to figure it out, but others weren't… about as expected when they started the farm team. They want to provide more consistent training and job experience to help junior folks come up to speed.

Given that backdrop, what should ComponentCo do? They understand the need to think differently about attacks, and how important it is to move past a tactical view to see threats more broadly. They want this wider perspective to help staff reach their full potential, and so they can protect information more effectively.

Harness Threat Intel

The first step in moving to a threat operations mindset is to make better use of threat intelligence, which starts with understanding adversaries. As described above, CCo contends with a variety of adversaries — including competitors, financially motivated hackers, and nation-states. That's a broad range of threats, so CCo decided to purchase a number of threat feeds, each specializing in a different aspect of adversary activities.

The first step in moving to a threat operations mindset is to make better use of threat intelligence, which starts with understanding adversaries.

To leverage external threat data they aggregate it all into a platform to reduce, normalize, and provide context. They looked at pumping the data directly into their SIEM, but the flood of external data would have overwhelmed their SIEM, which they didn't believe was a good idea.

They use their TI platform to alert based on knowledge of adversaries and likely attacks. But these alerts are not smoking guns — each is only the first step in a threat validation process which sends the alert back to the SIEM for supporting evidence of an actual attack. Alerts from these sources have higher priority because they match known real-world attacks.

Given what is at stake they don't want to miss anything. So they also integrate TI into some active controls — notably egress filters, IPS, and endpoint protection. They can quarantine devices communicating with known malicious sites or otherwise indicating compromise, hopefully *before* data is lost.

Enrich Alerts

We mentioned how an alert coming from the TI platform can be pushed to the SIEM for further investigation. But that's only part of the story. The connection between SIEM and TI platform should be bidirectional so when the SIEM fires an alert, information corresponding to the adversary and attack is pulled from the TI platform.

In case of an attack on CCo, an alert involving network reconnaissance, brute force password attacks, and finally privilege escalation would clearly indicate an active threat actor. So it would help the analyst performing initial validation to have access to all the IP addresses the potentially compromised device communicated with over the past week. These addresses may point to a particular bot network, providing a good clue to the adversary behind the attack. Of course it could be a false flag, but a hint still offers a head start for the analyst.

Additional information useful to an analyst includes known indicators used by this adversary. This helps understand how an actor typically operates, and the likely next step. You can save manual work by including network telemetry to/from the device for clues to whether they have moved deeper into the network. With destination network addresses you can have a vulnerability scanner assess other targets to help the analyst quickly determine if any other devices have been compromised.

Finally, given the indicators seen on the first detected device, internal security data could be mined to look for other instances of that attack, even on devices which have not yet exhibited suspicious network traffic. Then the analyst can tell whether the attacker has successfully used the same tactic to establish other footholds. This is critical when it is time to expel an adversary.

This is pretty simple stuff, which any semi-experienced analyst does as he/she validates an attack and assesses potential damage. The key is that all this data can be pulled automatically before an alert reaches an analyst.

This is pretty simple stuff, which any semi-experienced analyst does as he/she validates an attack and assesses potential damage. The key is that all this data can be pulled automatically *before* an alert reaches an analyst. When the analyst starts to dig in they shouldn't need to start with a bunch of manual digging to grab everything they need to investigate. Much better if they can start validation in a good position to quickly understand what happened and assess the blast radius.

Building Trustable Automation

Automation within threat operations can mean a lot of things. Assembling all the supporting information an analyst needs for threat validation prior to starting the process is clearly automation. But let's move a bit deeper into specific actions which can occur automatically. As described above, ComponentCo has a mature response capability, and typically removes all potentially compromised device from the network at the beginning of response to limit possible damage.

But this impacts response in a few ways. First, it may tip off the adversary to their discovery, prompting them to burrow deeper and find other points of entry. Additionally, CCo may prefer to monitor adversary activity to figure out what they were trying to do, and how.

Automation can help here. CCo automatically moves suspicious devices onto a VLAN where all network traffic is captured, which won't tip off adversaries to their discovery. They also start to pull EDR telemetry off the device at least every 30 minutes, to ensure data is captured even if the adversary is tampering with logs on the endpoint. This provides opportunity to see what adversaries are up to, and perhaps to establish preemptive workarounds in anticipation of their next move.

Given their sensitivity to exfiltration, another step CCo may add to their response playbook is to automatically update a network blacklist with any new external networks the compromised device has been communicating with, under the assumption they are likely botnets. To be clear, this approach may result in blocking traffic to legitimate unknown networks. So the organization needs to figure out whether possibly disrupting business outweighs the risk of missing exfiltration. They can also search network and device security data for other devices connecting to those networks to help identify additional compromised devices.

Workflow and Process Automation

Underlying all these functions is an "automation mentality," where the team builds playbooks which specify actions to respond to typical threats. This is valuable for several reasons, including consistent response and minimization of human error. But scaling the security team is the most important. CCo is a very desirable place to work, and doesn't generally have an issue finding talented folks, but skilled security staff are in high demand. By combining a threat operations mindset with a heavy dose of automation, CCo can make less sophisticated (and cheaper) analysts more productive.

CCo is a very desirable place to work, and doesn't generally have an issue finding talented folks, but skilled security staff are in high demand. By combining a threat operations mindset with a heavy dose of automation, CCo can make less sophisticated (and cheaper) analysts more productive.

Of course they still use Tier 3 analysts to handle tough and complicated incidents. But usually their playbooks can guide Tier 1 & 2 analysts. Let's consider a playbook response to phishing as the first phase of a targeted attack.

In our scenario a junior staffer in Finance received a phishing email claiming to come from his bank and require immediate attention. The employee fell for the ruse and clicked the link, which compromised his device. The compromised device began internal reconnaissance and connected to a known botnet. At that point an alert triggered and the automated playbook kicked in, putting the device in a fully logged VLAN and increasing its endpoint monitoring level, then updating egress filters and the IPS configuration to watch for indicators corresponding to the recognized attack. A full image of the device was taken prior to clean-up, and it was then quickly restored to normal operations without any real data loss or extensive manual effort.

But given the sophistication of its adversaries, CCo doesn't assume phishing attacks are nothing more. So they installed the device image in a sandbox and watched it. This secondary analysis showed the phishing attack was not the end. A secondary malware kit activated the next day, which had all the earmarks of far more sophisticated nation-state malware.

This was immediately escalated to Tier 3.

Handling a Targeted Threat

Escalation of an apparent nation-state-level attack started yet another playbook, which triggered threat intel assessment and alert enrichment specifically for that threat actor, because they didn't know it was a nation-state until that point. By the time the case reached a Tier 3 responder, they had all the information needed to quickly recognize the adversary, their tactics, and where else similar attacks have been seen — both inside and outside CCo.

At that point the response team knew they were under real attack by a sophisticated adversary, and automatically started capturing egress traffic from the DMZ and locking down their most critical assets as a precaution. Because related information had already been collected and associated with this case, the Tier 3 analyst could very quickly figure out the adversary's Tactics, Techniques, and Procedures (TTPs) and choose an appropriate response.

The key is that what looked like simple phishing, handled in a largely automated fashion, uncovered a sophisticated nation-state campaign. Their threat operations mindset enabled CCo to seamlessly escalate and provide a Tier 3 analyst with all available information to streamline attack and adversary research.

Obviously much more work and detail is required to actually eradicate a nation-state from CCo's systems, but they have a playbook for that too. The key is that what looked like simple phishing, handled in a largely automated fashion, uncovered a sophisticated nation-state campaign. Their threat operations mindset enabled CCo to seamlessly escalate and provide a Tier 3 analyst with all available information to streamline attack and adversary research. This accelerates both damage assessment and eventual expulsion of the adversary, and represents the innovative thinking needed in the security field to keep pace with increasingly sophisticated adversaries.

Summary

As we have explained through this paper, the types of adversaries you're facing require you to think differently about handling threats. Old-style threat management is giving way to structured and predictable threat operations. This evolution requires you to:

1. **Define Processes and Playbooks:** Consistency requires initial work to figure out appropriate responses for a number of different scenarios. This starts by defining how you want the team to behave, and then working to implement consistent processes.

2. **Implement External Threat Data Aggregation:** External threat data is key to understanding what adversary you are facing and what they are likely to do. Numerous feeds are available, but to avoid overload and effectively utilize the data, you'll need to aggregate and process it for better context.

3. **Integrate External and Internal Security Data with Analytics:** Once aggregated the external data needs to be analyzed alongside internal security data to pinpoint potential issues and identify patterns of malicious behavior to prioritize their efforts on the attacks that really matter. This provides more relevant alerts, enriched with supporting information about probable adversaries and indications of whether an attack proliferated in your environment.

4. **Orchestrate Existing Monitors and Controls:** The key to operationalizing a playbook is to have all the systems work together. So your TI aggregation platform (if a separate technology) needs a bidirectional connection to and from your SIEM. It can also send data to network security devices (IPS, egress filters, etc.) to block known bad sites. It can check with an advanced endpoint tool to confirm endpoint compromise.

5. **Automate First:** Finally, given all this analysis and integration, trusted automation can block traffic to known bad sites, move compromised devices into quarantine networks, and capture telemetry when you detect suspicious activity. Basically, if something can be documented in a playbook, you should be able to automate much of the process.

The end result is an orchestrated and automated ability to handle threats, equipping human analysts to do what they do best: pull on threads and make connections between isolated attacks which may actually represent sophisticated campaigns. Machines don't do this well or automatically. If it can be enumerated in a playbook, it should be automated. If not it remains the purview of the security team, and the immediate goal is to make them more productive by automatically aggregating the data they need to understand and address each situation.

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

About the Analyst

Mike Rothman, Analyst and President

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

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

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

About Securosis

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

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

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors. For more information about Securosis, visit our website: <http://securosis.com/>.