SANS

A SANS Whitepaper

The Need for Speed: Integrated Threat Response

Written by **Matt Bromiley**

September 2018

Sponsored by:
LookingGlass Cyber Solutions

Introduction

Let's face it: Defending an organization is no easy task. A world of attackers wants to steal nearly every data point that an organization can have. Some attackers don't even want data—they want an organization's excess computational cycles or bandwidth. And attackers don't stand still. Lower-level malware and opportunistic attackers now employ techniques that were once regarded as advanced adversary capabilities.

And any organization can be a target. So leaders and stakeholders should be asking themselves one question: Are we providing our security team with the tools, resources and capabilities they need to successfully protect the organization's data and other resources? Without visibility and the right processes in place, attacks can be almost impossible to defend against. But it's more than just the right processes; organizations need to have the right systems in place that allow defenders to respond faster than the attackers can change. The only way to be faster is to use technology to the advantage of the defenders when and where possible.

In this paper, we discuss the concepts of security automation and integration and provide recommendations on how to use technology to make your team faster and more efficient. It's time that defenders started looking at technology as an asset, not a liability they constantly have to shield from the bad guys.

Specifically, we examine the following topics:

- The current threat landscape, and whether organizations can keep up with modern attackers.
- The need for automation and integration within organizations and security teams.
- Automation and integration in practice, with a focus on preparing to combat threats as they arise, also known as active threat response. Specifically, we'll discuss:
  - Perimeter defense
  - Deception
  - Host-based active response

We hope that by the end of this paper you will not only have a better understanding of the need for security automation and integration, but also see that automation and integration are enhancements to, rather than replacements for, a security program. Organizations should still have good security policies that protect their assets and users. Automation simply helps make those goals easier to achieve.

Let's get started.

Keeping Pace with the Attackers

Today's threat landscape is constantly changing and presenting new challenges. Gone are the days when security teams were responsible for one operating system, a certain build of hardware and a manageable number of geographic locations. Instead, they are facing complexities that appear to grow exponentially daily. Today, most organizations face:

- A wide range of operating systems and sub-versions
- Mobile devices that are just as accessible and capable as laptop/desktop systems
- A workforce that expects to have remote working options
- Cloud-based applications to help facilitate the above

Those things make it difficult to rely on one management tool or enforce the same policies and limitations on all systems. With organizations struggling to effectively manage so many variables, it would be nice if all that complexity made things more difficult for attackers.

Unfortunately, attackers have had no problem keeping up with those changes and exploiting them to their advantage. Despite advancements in enterprise security, including network- and host-based technologies, organizations are still being hit by many security incidents. In our 2017 SANS Incident Response survey,[1] 87 percent of respondents reported having at least one incident in the past year. A fifth of all

[1] "The Show Must Go On! The 2017 SANS Incident Response Survey," June 2017, www.sans.org/reading-room/whitepapers/incident/show-on-2017-incident-response-survey-37815

respondents reported at least 100 incidents within a 12-month period—100 incidents that, regardless of whether they led to data exfiltration, required time and resources to investigate, contain and remediate. How easily can your team keep up with a new incident every 3.5 days?

TAKEAWAY

While organizations grow in size, complexity and geographies, attackers have no problem overcoming each to consistently achieve their objectives and steal data.

Attacker Success

One factor that contributes to attackers' success is automation.

Consider one of the more significant malware attacks during the past year, NotPetya. Masquerading as the Petya ransomware variant, NotPetya was destructive malware that made use of leaked, APT-level exploits and credential harvesting to cause significant damage to its infected victims. It inflicted more than $300 million worth of damage to a major shipping company[2] that had to rebuild 49,000 systems within a 10-day period. How did the attackers manage such a prolific, wide-scale attack? Via automation.

Many exploitation tools used by attackers these days include a significant amount of automation. APT-style tools will look for vulnerabilities, steal credentials and establish additional back doors within seconds. Ransomware, one of the largest annoyances in today's attack landscape, will automatically delete volume shadow copies, destroy system backups and prevent eradication techniques from removing it.

Defender Success

How, in the face of capable, automated attackers, do security teams turn the tables? The answer is to play the same game—but play it better. Enterprises should be arming their defenders to automatically respond to threats with accuracy and precision based on a clear understanding of the threats and their relevant severity. Where attackers can quickly pivot to different strategies, including the ability to move through networks, defenders should be able to close those gates just as quickly, if not faster. Additionally, there should be areas where automatic defense mechanisms are enabled in the event of an imminent threat. We will discuss specific points of automation and integration in later sections, but the theme of each is the same: Automatically defend the environment from attacks.

Another way to effectively arm defenders is to ensure that they have access to reliable, high-fidelity threat intelligence sources. Luckily, there is an abundance of threat intelligence sources available to most organizations, both free and paid. Your enterprise does not need to have deep pockets to join Information Sharing and Analysis Centers, or ISACs, which are typically industry-focused and have a wealth of data about cyber attacks. Additional resources available include law enforcement-sponsored threat sharing groups and paid intelligence feeds from private companies. Note that no matter what sources you choose, your team should vet threat intelligence before integrating it within your environment.

[2] "IT 'heroes' saved Maersk from NotPetya with ten-day reinstallation blitz," The Register, Jan. 25, 2018, www.theregister.co.uk/2018/01/25/after_notpetya_maersk_replaced_everything

The Need for Automation and Integration

The fact that attackers are using highly automated infrastructures to perform their attacks, maintain presence and monitor networks only strengthens the argument that defenders should be using automation as well. Again, there's no reason to replace human beings—they are still critical to any security program. Organizations rely on their teams to train staff, understand their infrastructure and its quirks, and know the organization inside and out. Because many entry vectors take advantage of human nature (both spearphishing and malicious insiders are examples), it's important that the human element in information security is not lost. Additionally, and perhaps most importantly, humans can focus on the problems and decisions that automated systems cannot solve.

Our focus is on sliding scales of automation that allow defenders to work more efficiently and tackle the tough problems that require human analytics. A balance must be struck to determine what can be automated versus what still requires humans to make decisions. To determine the best methods to implement automation, we recommend considering the following:

- Can the organization increase efficiency with the use of automated systems (and free up humans from doing manual tasks)?
- Is the team drowning in data, with no hope of coming up for air?
- Does the team have the right type and number of human resources, and are they ready for automation?

Let's examine each in detail and discuss how they may apply to your enterprise.

TAKEAWAY

Automation does not replace the human. It makes parts of the job more efficient, so the human can focus on the tough problems.

Increase Efficiency with the Use of Automated Systems

Many defenders still have to work through tasks that, quite frankly, should be automated. Modern networks have dozens, if not hundreds, of network systems that can assist in the defense of the perimeter and internal networks—and that's not even including the cloud! There are a multitude of ways to interact with network systems, but in many cases the story is still the same: Analysts have to manually block external threats *when they get around to it*, and *only after* they have learned about them.

Automation allows for the automatic deployment and mitigation of potential threats to the enterprise. However, not all organizations are ready to automate part of their workflows. It's worth asking a series of questions to see if automating a task is the right step to take. Figure 1 provides some of the relevant questions that your organization might ask when it decides to look at automation as an option.

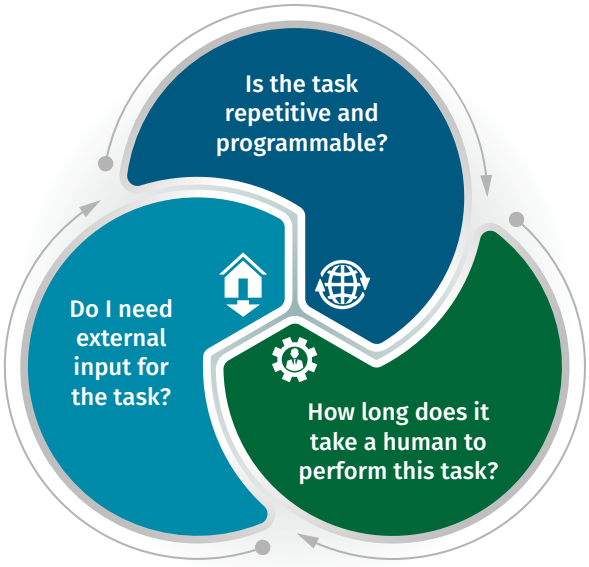

Figure 1. Questions to Ask When Considering Automation

Drowning in Data

Wading through the sheer amount of data that organizations generate and capture these days prevents team success. Not only quantity, but also the complexities of connections can be daunting. What with endpoint logs, third-party applications, NetFlow, DNS, full packets and everything in between, many defenders are drowning in data. This inundation precludes a complete (and necessary) understanding of the enterprise environment, which may lead to additional issues in the future.

Just to be clear, we advocate capturing as much data as possible. But there's a caveat. Captured data must be correctly cataloged and organized, and, most importantly, it must be *accessible*. Unfortunately, many organizations have followed only the first part of this classic advice: They implemented SIEMs and turned on data capture at as many points as they could. But they are failing to complete an effective catalog and organization of the data, which prevents their defenders from ever coming up for air. This influx of data can be lessened with automation and integration.

Understaffed and Overwhelmed

Not only are humans left to perform tasks that are (and should be) easily automated, but many organizations are also finding themselves without the requisite number of humans to achieve the bare minimum. In the 2017 SANS Incident Response Survey cited earlier, approximately 49 percent of respondents indicated that lack of resources was the top impediment to incident response teams, with "staffing and skills shortage" right behind that, at 47 percent. It's clear that information security teams need help where they can get it. Automation and integration bring help without the extra overhead and allow teams to get back to the business of investigating.

Active Threat Response in Practice

We've talked about the need for automation and how defenders can benefit by leveraging it within their environments. To some, this may seem like an easier-said-than-done situation. Luckily, there are already scenarios where automation and integration are in place and assisting organizations in active threat response.

The concept of active threat response relies on automation and integration to help organizations react and adapt to threats as they arise. It's important to note that organizations that rely on automated and integrated systems do not wait for a human to consume threat intelligence and then perform actions based on what was consumed. Simply blocking or logging based on intelligence is not enough. Active response is a sophisticated order of defense that goes beyond the traditional "stop" reactionary techniques, with connected systems that automatically deflect threats, obfuscate data, deceive attackers and deploy other deception techniques. Figure 2 depicts this contrast between reactive and proactive defense graphically.

Let's examine three areas where automation and integration, when implemented properly, can assist organizations in defending against various threats.

Figure 2. Reactive (Human) vs. Proactive (Automated) Defense Against Threats

Perimeter Active Response

The first place to think about automation and integration is the external surface of the organization, at the perimeter. Here's why:

1) Every organization has an external surface, so this common denominator allows for widespread development and implementation.
2) Because most attacks originate from *outside* the organization, it's a natural point that attackers *must* cross to breach the environment.
3) Defensive mechanisms at the perimeter may be easier to implement than internal, due to network devices and infrastructure.

The first place to think about automation and integration is the external surface of the organization, at the perimeter.

Automation of perimeter defenses starts with high-fidelity sources. As discussed earlier, the abundance of external threat intelligence means that organizations can typically get reliable data about imminent threats. Unfortunately, many organizations purchase threat intelligence subscriptions but do not act upon the data. Even worse, upon being breached, some organizations realize that they had the related intelligence days, weeks or even months prior to the breach.

The blocking of known-bad IP addresses and domains is an easy mechanism that organizations can employ to assist in automated threat prevention. When threat intelligence feeds provide new insight into attacker campaigns or imminent threats, the external perimeter devices should be able to adjust accordingly. IP addresses and malicious domains can be blocked. Vulnerable ports can be closed. By the time defenders have read up on the latest threat, the organization's automated tools have already adjusted to the possibility of a breach.

Perimeter defenses can also assist in limiting the spread of internal malware or already-infected systems. These days, many malware variants, especially ransomware, require Internet connectivity to function. Investigative findings may lead to suspicious IPs or domains that can be blocked, which will prevent additional infections from reaching out.

TAKEAWAY

An organization's perimeter is its most exposed attack surface. Automation helps keep the perimeter healthy and secure with up-to-date threat data.

Deception

The concept of deception is a tried-and-true technique that has been used in many forms for decades, and not just within interconnected networks. Within enterprises, deception can be applied in multiple ways:

- As a network
- As a server/host
- As a user
- As a set of files or system configurations

Using a deceptive technique within your organization means you may stage data, a handful of systems or even an entire subnet within the network. Monitoring and logging are typically wrapped tightly around the deception, and if an attacker enters the space or touches the file(s), alarms are tripped and the organization is alerted.

One of the more common deception techniques is the use of a honeypot, but honeypots are not always the most successful internal network detection mechanisms. For starters, attackers know to look for them. Once inside a network, attackers will often perform extensive internal reconnaissance and become intimately familiar with the network architecture. They know what belongs and what doesn't and can usually detect when an "inactive" subnet is live, or when documents are placed in "convenient" locations.

Additionally, attackers will find the honeypot only if a route to it is available. If an attacker entry vector is to spearphish the human resources department, but the honeypot is accessible only from engineering, the organization could suffer a breach and data loss, while the honeypot never gets noticed. Even worse, an internal honeypot assumes the worst: Attackers are already inside the network, have moved laterally and are enumerating the domain. It's better to detect those activities earlier.

Active response and network deception involve much more than simple honeypots. Internal network deception focuses on understanding the goals of attackers during each step of the breach and placing dynamic detectors throughout the environment. These detectors, while unseen by endpoint users, will align with attacker objectives—objectives such as gaining privileged credentials or moving laterally between systems.

Network deception focuses on understanding the goals that attackers will have during each step of the breach and placing dynamic detectors throughout the environment.

Consider an example of attackers successfully spearphishing a user with a least-privileged account. The attackers will want to elevate to a user with a more privileged account and may make use of a password dumper or credential harvester to gain additional account credentials. An organization's endpoint monitoring tools may or may not detect the attackers' password-stealing tool; endpoint monitoring does not have a perfect record of stopping all attacker tools.

If the organization in this scenario used internal network deception, it might focus on planting legitimate—but heavily monitored and alert-tied—account credentials around the environment. The attackers would then stumble upon a seemingly privileged account to use, but the organization would be alerted that a deception had been triggered. Every action taken by that account would then be monitored. At this point the organization has the upper hand, and it faces two choices:

1) It can monitor the account activity to gain intelligence about the attackers and understand their tools, tactics and procedures (TTPs).
2) It can immediately cut off the activity and/or remove the attackers from the network.

Each approach has its benefits, and organizations should assess which one they are comfortable with. Watching an attacker to gain intelligence is an important step in a well-structured incident response plan but can be risky if the attacker suddenly changes accounts or subverts the monitoring processes.

So where do automation and integration tie in? Throughout the whole process. The deploying of dynamic indicators around the network should be handled by automated systems. This approach allows for the credentials to change dynamically to align with current attacker techniques. Realistically, it's tough to expect that defense teams, regardless of size, can effectively manage an entire enterprise as well as develop and maintain an active, internal deception network that aligns with attacker techniques.

The trigger of deceptions should be integrated with security operations, such as a SIEM and/or an alerting platform. Because a tripped network deception is a solid indication that an attacker is inside the network, such alerts should have the highest priority. With alerts and depiction triggers combined, an automated internal deception program can help organizations detect attackers faster.

Host-Based Active Response

The result of even the best-laid deceptions is an alert; the organization is still in a reactive mode. The next step is to set up endpoints so that they can automatically defend themselves against attacks, both reactively and proactively. First, let's discuss reactive host-based response.

Typically, host-based response has been associated with IDS and host IDS, in which IDS is deployed on various endpoints. An IDS monitors for malicious activity and reports to a central repository or a SIEM. But humans have to wade through the alerts generated by hundreds, thousands or tens of thousands of hosts, and determine which to act upon (spinning up more resources), and which to ignore (false positives that can likely be tuned out).

Active response with automation and integration asks where we can remove the humans from this equation—and again, this is to make their lives easier and the organization more efficient.

A proactive host-based active response focuses on preparing hosts for threats that might enter the environment. As we discussed with internal network deception, attackers typically take predictable steps once they gain entry into an environment.

Conclusion

Automation and integration can help security teams work faster and smarter—but only if the right workflows and mindset are in place. Teams need to understand where automation can lead to success, and ultimately a higher sense of security for the enterprise. Remember, automation is not a way to replace the analysts but rather a means of detecting evil faster and giving the analysts an advantage over attackers. When teams don't have to worry about mundane tasks, they can focus on the complex problems of detecting advanced attackers and permanently removing or preventing them from entering the environment. Figure 3 presents a high-level checklist for getting started with automation and integration.

How to Introduce Automation and Integration

- ☑ Integrate high-fidelity threat intelligence sources into the enterprise.
- ☑ Reserve human expertise for the most complex threats.
- ☑ Use internal defenses to detect threats early.

Figure 3. Quick Checklist for Introducing Automation and Integration into an Organization

Automated defense has a valuable allied resource in high-fidelity threat intelligence sources. When integrated properly within the enterprise, threat intelligence can be the difference between preventing an attack and suffering a breach. Defense teams should internally rank the stability of threat intelligence sources, relying in part on the sources' external reputations, and they can filter feeds so that automated responses are limited to certain trusted types of results.

Organizations should also look for opportunities to implement active threat response. This can include perimeter and host-based active response and network deception. Catching an attacker via deceptive techniques can lead to invaluable intelligence gathering. With active response in place, defenders can begin to box attackers in, degrading their mission, all the while studying their methods and techniques.

Some organizations may resist automation and integration. Some have systems in place that require more manual, intervened work. Others don't trust external sources enough to automate security based on the data of others. While those are valid stances, neither one does anything to solve the problem—and in fact they introduce problems and take analysts away from solving the problems that truly matter.

But any organization looking for a way to optimize team workflows and save time, money and resources should investigate how automating active response can be integrated into its environment. Its security team will thank it for the support, and its data will be better protected.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor and a GIAC Advisory Board member. He is also an incident response consultant at a major incident response company, bringing together experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools. Matt teaches the SANS DFIR courses FOR508 and FOR572.

Sponsor

SANS would like to thank this paper's sponsor: