



SANS Institute

Information Security Reading Room

Automating the Hunt for Hidden Threats

Eric Cole, PhD

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.



Automating the Hunt for Hidden Threats



A SANS Whitepaper

Written by Eric Cole, PhD

Advisor: John Pescatore

October 2015

*Sponsored by
Endgame*

Introduction

Medicine is helpful only if you understand what is wrong with the person and provide the proper remedy. The wrong medicine does no good and, in some cases, can make things worse. In a similar vein, many organizations spend lots of money on security but fail to apply remedies that make it harder for adversaries and control the damage caused by attacks. Organizations need to recognize that their networks are constantly under attack and actively hunt for patterns and indicators of attack.

In other words, they must prepare their networks and systems for attacks, if they aren't already under attack. Start by asking this question: If attackers compromised your network, how would you know?

This is not a sign of a weakness. It merely acknowledges reality in order to plan for the inevitable and be ready to respond and remediate as quickly as possible. The focus should include robust scanning for potential attacks. Companies need to be proactive in preventing, detecting and responding in order to minimize the frequency and impact of attacks. This is the core goal and purpose of hunting for attackers and threats.

Threat hunting is the act of aggressively tracking and eliminating cyber adversaries as early as possible in what Lockheed Martin has dubbed the "Cyber Kill Chain."¹ The earlier in the chain responders can discover a threat, the less damage there is, the faster repairs take place and the sooner network operations get back to normal. Reactive incident response is no longer sufficient when dealing with advanced adversaries.

To provide improved security, the goals of hunting include the following:

- Gaining better visibility into the organization's weaknesses
- Providing early and accurate detection
- Controlling damage
- Improving defenses to make successful attacks increasingly difficult

Organizations make the news and incur significant fines because they do not hunt for breaches and cannot detect and contain them in a reasonable period, not because attackers succeed.

This paper defines the process of automating the hunt for threats, including how to deploy a continuous threat-hunting process and prepare a team to analyze threats to protect critical processes and data.

¹ Lockheed Martin website;
www.lockheedmartin.com/us/what-we-do/information-technology/cybersecurity/tradecraft/cyber-kill-chain.html



Why We Need to Hunt

Traditional security methods (such as antivirus, network IDSes and firewalls) can't catch today's advanced targeted threats because such threats work around security controls, hiding themselves in memory and other locations where malware defenses can't see. Companies continue to invest in traditional security measures, yet they remain vulnerable to breaches that bypass these defenses, as one incident after another proves with depressing regularity.

Organizations need to recognize that attackers will strike frequently and succeed occasionally, then focus on controlling and minimizing the damage as proactively as possible. Threat hunting, for the purposes of this discussion, is the proactive, aggressive and methodical discovery and pursuit of known threats based on indicators of compromise (IoCs) and detection of unknown malicious behaviors, disrupting and/or eradicating these threats and securing critical infrastructure. For example, threat hunting includes the following activities:

- Understanding the threats
- Mapping, geolocation and device understanding across the network
- Identifying critical data and business processes utilizing that data
- Distinguishing good from bad behavior
- Leveraging threat intelligence for discovery, detection and analysis
- Analyzing all this data, along with vulnerability data and other sources of network/endpoint behaviors, for anomalies that are both "known bad" and never before seen
- Looking for anomalies, learning abnormal behavior and understanding the network
- Tracing activities to IoCs
- Following IoCs to affected systems
- Taking appropriate action (e.g., removing bad actors from the environment, repairing systems, remediating vulnerabilities or training staff)

This list is not all-inclusive. The goal of hunting is to minimize business or mission impact and, ideally, to stop the attack before data exfiltration or other damage takes place.



Why We Need to Hunt (CONTINUED)

Automated or continuous hunts look for both known and unknown IoCs, such as anomalies, unusual connections, strange registry keys and anything else out of the ordinary. This is valuable and important, but it will not necessarily catch all attacks.

On-demand hunts look for a particular attack and/or a compromise within an organization. To do this, you need to know exactly what to look for. Threat intelligence and analytics play a key role here by telling you what type of activities, connections and behaviors to track and extinguish.

Focus the Hunt

Threat hunting is usually conducted two ways. One is continuous hunting, where the developers of security-monitoring systems augment their tools with intelligence and human analytics to continually search for and report on suspected, anomalous behaviors, known IoCs and malicious, previously unseen behaviors. The other method, on-demand hunting, focuses on a specific IoC or other event that sets off a search.

Hunt Cycle

Threat hunting starts with planning. This includes identifying processes and critical assets, as well as personnel responsible for threat hunting, response and remediation or follow-up. The cycle concludes with a reporting phase that enables responders to examine their actions and devise future defensive measures. Figure 1 shows a typical threat hunting cycle.

Figure 1 shows a typical threat hunting cycle.




Figure 1. Threat Hunting Cycle

Plan

The planning step of the hunt cycle includes assessing critical assets and learning their patterns. Here are some tips for getting started:

1. Start small by prioritizing assets and data of most interest to attackers. Identify the most critical systems and data, as well as all related assets. Ask yourselves: What are the most critical pieces of data that attackers would target and which servers contain them? Then ask what other systems, users and internal or external devices interact with these sensitive systems and data types.
2. Verify that user, file, process and network activities create detailed logs/activity data that can process quickly through the threat hunting system as required.



Why We Need to Hunt (CONTINUED)

In addition, be sure to monitor the proper network segments. Geolocation of assets is important here: physical and logical mapping of critical systems and data, as well as their related systems and users, provide big-picture and close-up views as needed. If you have trouble finding an adversary, it could be you do not have the right information about the network and systems and/or you lack the right intelligence to apply against the system information you are seeing. In addition to logs, you should collect detailed user, file, process and network activities to support the hunt.

Pursue: Search and Detect

Pursuing your adversary is critical to the hunt cycle. There are two approaches:

1. Searching for known threats by gathering existing IoCs or other tactical details, such as the signature of an attack. Implement techniques to harvest data from your critical assets (e.g., search for a specific malicious binary hash or for a command and control URL in a network flow database).
2. Detecting unknown threats. This type of advanced hunting is challenging due to a lack of intelligence to spark the investigation. To enable detection of unknown threats, leverage the power of data science (e.g., machine learning, clustering or statistical analysis). Confirm baselines of normal activity over time so you know what deviations from the norm look like. If you do not have a set of baselines, look for deviations from known or historic behavior.

In addition, be sure to monitor the proper network segments and collect detailed user, file, process and network activities to support the hunt.

Secure and Monitor

The hunt chain focuses on actively monitoring users, networks, servers, and other endpoints and systems, as well as egress points. By watching for signs of malicious activity and IoCs, organizations can catch attacks earlier, while categorizing lessons learned from the event and reusing them to detect and prevent future actions. Those interested in sensitive data and IP that they can monetize don't want to be caught, so they usually follow specific steps to avoid detection and dig deeper into their targets. These include assessing the enemy, exploiting a user or endpoint, hiding their existence and establishing a back door.

All of those leave IoCs that well-tuned threat hunting systems should pick up, such as the use of Remote Desktop Protocol (RDP) or other unusual or unapproved port connections, malformed DNS requests and internal server connections directly to the Web.

Locking down security and vulnerabilities on such systems reduces the likelihood of a successful attack, while monitoring for IoCs reduces the time to respond. Both are critical preparations that responders should include in any threat hunting cycle.



Respond

The active stage of a hunting program should include the means to detect and stop an attack at the reconnaissance, pivoting, attempted attack and command channel stages. Stopping the attack at any one of those stages would control and minimize the damage.

For example, while it may be hard to stop an attack in the reconnaissance stage, there are signs and indicators that attackers are scoping employees and systems for vulnerabilities. They include the following:

1. Running targeted scans against web and mail servers
2. DNS server tampering, registry lookups or registry changes
3. Targeting specific employees through social media and other web presence

A proactive response to any of those detected behaviors includes determining what value those systems offer to attackers, while also checking security and patch status the systems scanned by hackers, and educating the potential human targets.

Response would be different and much more intense to an attack. For example, an employee within the organization clicks an infected attachment or URL to a compromised site. As soon as the user clicks on the attachment, his now-infected system begins internal reconnaissance. The attack then pivots deep into the network, compromising a database system and creating an outbound command and control (C2) channel to the adversary.

Response in this case goes beyond determining motivation and targets and locking them down to prevent attack. Now response includes tracking all pathways and associations between the database and other systems, including like systems across the enterprise, determining the type and scope of data impacted, and determining egress points and data loss, all while still finding and kicking out the attackers and erasing all of their back doors and regenerating traces.



Metrics Matter

Event information should also be reutilized to repair impacted systems and associated assets, reduce risk and blacklist similar types of connected behaviors or IoCs in the future. To do that, organizations need to report in clear metrics a measurable reduction in risk that ties to their preparation, response and follow-up in the threat hunt cycle.

These metrics should consider the following:

- Fewer actual breaches
- Reduced attack surface/system hardening improvements
- Shorter dwell time (the time between when an attacker first gains unauthorized access and when the bad actor is removed from the network)
- Minimization and reduction of unauthorized lateral movement between internal systems
- Reduction of exposure by finding and stopping threats before they gain a foothold
- Speed and accuracy of response
- Measureable security of systems
- Fewer actual breaches
- Reduction in man hours and other expenses spent on response

When you see an organization in the news because of a data breach, the breach itself didn't make news; it was its degree of the breach: the amount of damage and the length of time the breach persisted. Organizations need to focus on controlling damage by catching attacks early (reducing dwell time) and controlling the damage (restricting lateral movement).

Don't Over-Rely on Intelligence

Mature threat hunting processes need some sort of outside intelligence against which to correlate discovered internal indicators of compromise. This threat intelligence will be valuable, but no single threat feed will contain every IoC. Further, the IoCs from very targeted attacks will never show up in external threat intelligence feeds. Importantly, overreliance on outside intelligence could consume resources by leading administrators to look for the wrong indicators in the wrong locations. Therefore, it is always important to properly validate the threat in the context of your own environment before pulling the fire alarm and sending all your resources at it.



Keys to a Successful Hunt

In many large organizations, hunting for breaches is like looking for a needle in a haystack. Such organizations are so large, with so much information and data to sort for actionable intelligence, that just knowing where to begin seems impossible.

The basic methodology of a successful hunting program includes the following:

- Augmenting humans with tools and automation across all areas of the hunt chain
- Segmenting and de-scoping the area of analysis
- Having focused goals
- Limiting the search (deeper is better than narrow)
- Improving processes by documenting what works (and what doesn't)
- Recording metrics that demonstrate business-relevant gains, such as reduced time to contain and mitigate

Those facilitate incident response not only by defining the target of the current investigation, but also by refining the process for future investigations and demonstrating gains in effectiveness and efficiency in the overall security program.

Threat hunting processes need to align and integrate with the three core areas of defense: prevention, detection and response.

Prevention

Over time, threat hunting can improve prevention through learning, tuning and constant improvement of preventative measures. Hunting should be part of a continuous loop. As hunts detect compromised systems, they are also used to determine how the compromise occurred. Then, as follow-up after remediation, this information is used to improve defenses, thereby preventing future attacks that operate in the same manner. Although all attacks are different, the behavioral characteristics of an attack are typically the same from one to the next. By understanding how attackers recon, pivot and create C2 channels, analysts can better detect future attacks.



Keys to a Successful Hunt (CONTINUED)

Understanding how to prevent one type of attack will not always prevent other similar attacks if attackers learn to further obfuscate their methods from detection systems. However, threat hunting will usually identify weaknesses in an organization's security efforts that need to be improved. For example, if hunters discover several compromised systems that were exploited because they were missing patches, the threat hunters just helped identify fundamental repairs that need to be made on the systems, as well as on their patch management or endpoint update processes. Implementing an effective patch management scheme is one of the most effective ways to prevent future attacks and reduce attack surfaces, according to the Critical Security Controls and other leading security frameworks and standards. So improving the patch status and program results in reduced attack surface and improved processes—all of which started with threat hunting.

Detection

Although manual hunting has some value and will always be used to some extent, it can be manpower-intensive or require skilled personnel that are either hard to hire or are fully committed to their areas. The real power of hunting comes to the fore when automation or advanced tools support these processes. Computers are good at repetitive tasks and humans are good at analytical tasks, so automating the former is likely to free up people for the latter. Upon detection, analyzing the anomaly can help determine whether it is an attack. One powerful way to do that is by using a scripting language to create an executable that can run automatically against traffic, looking for obvious anomalies.

There is enough data and information within even a small organization that it would be impossible for a human to process it in anything approaching a reasonable amount of time. By automating the hunt, organizations can benefit from almost real-time detection. This gives security analysts—and in larger entities, the security operations center (SOC)—more visibility into compromised systems, which in turn allows for more intelligent decision making. Any SOC needs tools that take all of the information from the data feeds, reduce false positives and prioritize high-confidence indications and, finally, enable analysts to review whether anomalies are actually compromised systems. Those tools are required as “force multipliers” to address the realities of limited security staff and the scarcity of deep hunting skills.



Keys to a Successful Hunt (CONTINUED)

Although breaches are damaging, the real business impact of an attack comes from attackers' post-breach activity. Once an adversary finds his way into a network and compromises a system, he will typically setup a "pivot point," then continue to move laterally across the network, gathering intelligence or doing more damage. A well-defined hunt catches this lateral movement by looking at network traffic and system logs and identifying anomalies in the data. These might include successful access attempts by users on systems they normally don't access, failed access attempts, registry changes, large data payloads or increased connections to a sensitive device or network.

A key component of any hunt is remembering you are dealing with an advanced adversary that does not want to be caught. Rather than looking for hostile activity, you want to identify normal behavior patterns and look for deviations from them, even subtle differences. You are hunting for a chameleon, so pattern matching plays to his strengths.

Figure 2 shows how threats constantly change their appearance and behavior to avoid detection.

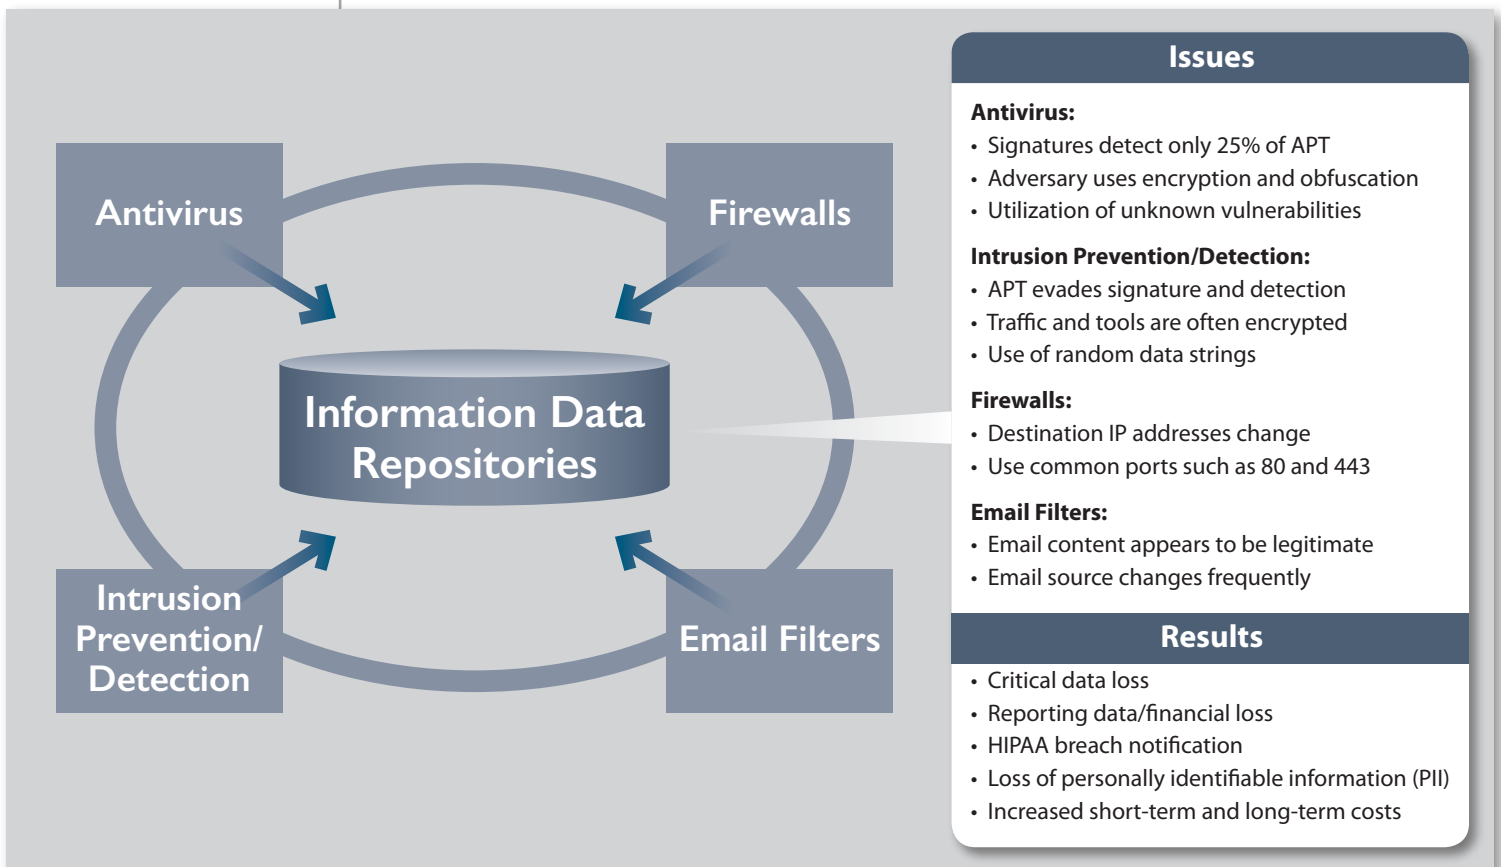


Figure 2. Why Hunting Matters: Advanced Evasion



Keys to a Successful Hunt (CONTINUED)

Even these evasion techniques create patterns. Because adversaries continue to change their patterns, the hunting process must do the following:

1. Adapt to changes in behaviors and learn how the adversary works. An advanced adversary will target a system, upload code, create a C2 channel and survive a reboot. Although the specific methods of doing this are always changing, the general behavioral patterns are the same.
2. Watch all behaviors of the adversary, including known good, known bad and unknown or unclassified behaviors. Looking for anomalies that deviate from normal behavior can help detect unknown or previously unseen hostile activity.
3. Identify adverse activity, track it and alert administrators to the suspicious activity.
4. Contain and control the damage by identifying attackers' lateral movements and removing infected systems from the network.

By constantly refining detection and remediation processes, future attackers will encounter a field-tested response and may look elsewhere for easier pickings. Once an attacker is detected, you should analyze behaviors of the adversary either in a safe contained environment that is undetectable to the adversary or through live activity on the network, which is usually more risky.

Through behavioral analysis, hunting can distinguish between real user activities and attacker activities. For example, although an attack may not have a specific signature, it will have unique behavioral characteristics that differentiate legitimate actions from hostile activity. These include uploading code, running processes, changing registry values, making a large number of connections at one time, multiple connections to odd services or at odd times of day or creating a C2 channel.

Hunting processes and tools need to scale to accept large amounts of hunting data, and they need to correlate and analyze the findings in a reasonable fashion. Instead of just using a database of behavior from one company, automated tools should enable exponential correlation based on the experiences of multiple organizations and industries, not just the perspective of a single company.



Response

As noted in the introduction, prevention is ideal, but detection is a must. However, detection without response has little, if any, value. The sooner an organization can respond, the less damage it will experience. If threats are caught early, response becomes more proactive and attacks can be cut off before significant damage occurs. For example, if an attack exploits an unpatched server or a system that is running an extraneous service, detecting the vulnerability, as well as the exploit, and fixing the vulnerability should be easy enough.

Proactive response should detect the compromise at the pivot point, taking action before the adversary begins his lateral movement and compromises more systems. In such situations, although the attack isn't prevented, the response occurs before the attacker steals sensitive information.

Even though we discuss detection and response as different activities, the closer we tie these together, the better. Instead of one team hunting to detect adverse activity and handing off to a separate team for response, integrating these functions allows for a more effective and seamless process. Remediation should begin upon detection. It should also provide a means of repair and workflow that can integrate with the efforts of SOCs and other parties to quickly repair the infected systems.

Part of response is reducing the time needed to fix systems and mitigate future threats. Hunting tools should pass on accurate and understandable vulnerability information, including physical system information and logical location information, to enable rapid mitigation.

Advice

Hunting is a key capability to enable cybersecurity programs to reduce business impact due to ever-evolving advanced targeted threats. An effective and mature threat hunting program starts with well-thought-out processes and/or playbooks, is sufficiently staffed and is integrated with other security operations processes. Tools can act as a force multiplier, enabling skilled staff to analyze higher volumes of security-relevant data, but they can be effective only when used on the proper foundation of basic security skills, knowledge and controls.

When looking for threat hunting solutions, make sure you test the system before purchasing it and verify that it brings together the highest volume of meaningful information (i.e., more needles, not more haystacks) with high-speed, correlated analytics. Select systems that analysts can update in an adaptive manner. In addition, choose tools that provide drill-down features that facilitate further analysis. Finally, integrate any such tool with the existing security dashboard to provide better visibility into attacks and remediation/repair.



Conclusion

Too many organizations continue to deploy traditional solutions and are frequently oblivious to attacks because they are not properly preventing, detecting or responding to the advanced persistent threats that exist today. Properly automated threat hunting could have kept many of the organizations that suffered widely publicized breaches out of the news by minimizing their exposure time.

A typical checklist that organizations can use to start an ongoing hunt includes the following:

- Identifying the data or information most critical to your organization
- Determining which business processes utilize or access this information
- Identifying all of the systems and networks that support key business processes
- Acquiring tools that can help with the correlation and analysis required for proper hunting
- Gathering information about the traffic flowing to the key systems and networks
- Gathering information about the operations of servers
- Utilizing threat intelligence to understand the threats and exposures to the organization
- Utilizing tools to perform automated analysis of normal behavior and attack behavior
- Filtering the output of the tools
- Responding appropriately to high-risk alerts

By utilizing automated tools, organizations can start to win the war against network-based attackers. It is important to remember that this is less about spotting malware and more about identifying hostile behavior and containing that behavior as quickly as possible.

By deploying hunting capabilities, you can start to control the damage and reduce the impact any attack can have on your organization.



About the Authoring Team

John Pescatore joined SANS as director of emerging technologies in January 2013 after more than 13 years as lead security analyst for Gartner, 11 years with GTE, and service with both the National Security Administration, where he designed secure voice systems, and the U.S. Secret Service, where he developed secure communications and voice systems “and the occasional ballistic armor installation.” John has testified before Congress about cybersecurity, was named one of the 15 most-influential people in security in 2008 and remains an NSA-certified cryptologic engineer.

Eric Cole, PhD, is a SANS faculty fellow, course author and instructor who has served as CTO of McAfee and chief scientist at Lockheed Martin. He is credited on more than 20 patents, sits on several executive advisory boards and is a member of the Center for Strategic and International Studies’ Commission on Cybersecurity for the 44th Presidency. Eric’s books include *Advanced Persistent Threat*, *Hackers Beware*, *Hiding in Plain Sight*, *Network Security Bible* and *Insider Threat*. As founder of Secure Anchor Consulting, Eric puts his 20-plus years of hands-on security experience to work helping customers build dynamic defenses against advanced threats.

Sponsor

SANS would like to thank this paper’s sponsor:

ENDGAME.

