Continuously Hunt for Network Intrusions

A Question of Time

Today's network breaches often remain undetected for extended periods of time. A 2017 Ponemon Institute study found that the average time to identify adversarial presence within an enterprise was 191 days [1]. This is plenty of time for an attacker to wreak havoc on the enterprise network. A longer dwell time enables adversaries to establish persistence, perform network reconnaissance, and exfiltrate data. Additionally, remediation costs increase the further an attack progresses. This means that detecting network intrusions as early as possible is critical. Earlier detection minimizes damage, reducing remediation costs and speeding recovery efforts. Organizations that supplement automated preventive controls with a continuous and active hunt for unauthorized activity improve their ability to detect advanced threats sooner and reduce the time spent investigating and manually correlating network and host events [2].

The Hunting Process

Automated detection methods, such as Intrusion Detection Systems (IDS), Endpoint Detection and Response (EDR) capabilities, and Security Information and Event Management (SIEM) system alerts are useful, but they cannot detect all breaches. Advanced attackers obfuscate their actions to evade these common automated detection methods, ensuring longer dwell times. Organizations that continuously hunt for anomalous network activity do not rely exclusively on automated detection: they assume that malicious actors have by-passed automated detection and already reside in the network. They actively and continuously search for information about where the actors are, how they got in, and what they intend to do.

Identifying "normal"

Establishing a **baseline** is fundamental to devising a hunt strategy. A baseline is a representation of normal activity among an organization's network traffic, network performance, host and application activity, and user behavior. Having a baseline is necessary in order to detect anomalies. A baseline is collected over a period of time that is long enough to get an accurate snapshot of normal activity, but not so long as to risk accidentally incorporating anomalous activity.

In addition to baselining, **asset management**[1] practices ensure that all assets and network topology are recorded and that any changes are tracked. Be aware of all critical assets including software, hardware, network connections, and configurations. Knowing what assets and network topology to expect make it easier to identify suspicious changes to these assets. Information such as patching and asset vulnerabilities can allow a hunt team to focus efforts in the right areas. Actively managing assets will additionally make remediating issues more efficient by quickly providing information such as the asset's location, what the asset should be connected to, and how the asset should be configured.

Collecting data

An **Indicator of Compromise (IOC)** is a piece of data that a defender uses to detect malicious activity (e.g. in logs, files, and network traffic). Hunters detect IOCs or anomalies, pivot to gather additional information to understand the full extent of a compromise, and then evict the adversary from the network in a permanent way. The smallest missed indicator can often be an organization's downfall. In order for hunters to find malicious activity, organizations need to collect enough data to provide a detailed view of all network and host activity for comparison against IOCs and baseline behavior.

There are myriad specialized network inspection appliances on the market to collect different types of network data. Many expert network analysis shops collect long term network flow, short term packet capture, and network alerts. Such alerts can come from intrusion detection/prevention systems and gateways such as proxies and next-gen firewalls, while some go further to use technologies that extract and inspect files from email, web, and other traffic.

[1] For more information on asset management, please refer to "Actively Manage Systems and Configurations", part of the *NSA Cybersecurity Top 10 Mitigations*.

Additionally, end-user hosts, server equipment, network devices, and software services generate logs and store them locally by default. These logs contain relevant information such as successful and failed authentication, account creation and deletion, process creation and termination, driver loading, registry change, and file creation and overwriting events, along with alerts on identified malware, potential intrusions, or other suspicious behavior. A professional hunt team will learn the most critical logs to look for and where they reside, often then forwarding them to a SIEM solution for centralized analysis. Some SIEM solutions provide agents that can be installed on devices to forward logs and traffic from devices that do not natively provide forwarding capabilities.

Intrusion Detection Systems (IDS) monitor either network traffic (Network Intrusion Detection System or NIDS) or host activity (Host Intrusion Detection System or HIDS) and alert on suspicious behavior. A HIDS has agents installed on individual devices in a network to monitor each device's activities and state, while a NIDS is placed in one or more locations on a network and examines passing traffic. Intrusion Detection Systems (IDSs) are considered either rule-based/knowledge-based or anomaly-based depending on how they detect intrusions. A rule-based IDS examines network traffic or host activity for known attack patterns such as signatures or behaviors. Examples include a pattern of bits in traffic matching a regular expression (signature-based) or a process calling another process it should not normally call (behavior-based). An anomaly-based IDS establishes a baseline and then samples activity to detect deviations (e.g. a new traffic flow between endpoints or an unusual port). Defenders who understand their network and its baseline can make sense of such anomalies.

Endpoint Detection and Response (EDR) systems analyze behavior exhibited by users, systems, processes, and services to determine whether irregular activity indicates adversarial presence on an endpoint device. The ability to discover previously unknown tactics or seemingly-legitimate-but-actually-malicious techniques through the analysis of behavioral metadata makes EDR a valuable supplement to existing investments in defensive software. Important EDR capabilities that distinguish it from rule-based log analysis include instrumentation (which provides real-time visibility into security-relevant system functions), the ability to scan system memory with host agents, contextual awareness (which helps with incident timeline analyses), the potential for automated remediation (with severity estimation), and the ability to tune countermeasures in response to intrusion evidence. Often, EDR solutions employ machine learning techniques for anomalous behavior identification and third-party threat intelligence feeds for file reputation or IOCs.

Searching and analyzing

A SIEM system provides log aggregation and correlation, querying, visualization, and alerting, making it a necessary tool for hunters. The tool collects logs and traffic from across the enterprise and formats the data to allow for efficient searching and correlation. Hunters can leverage these features to search for IOCs and correlate data from different sources to follow the path an attacker took through the network. Logs and alerts from IDSs and EDRs can be forwarded to SIEM solutions. Analysis of logs in the SIEM system can uncover how attackers got in, how they traversed the network, and what activities they engaged in while in the system. SIEM tools often include built-in analytics and allow analysts to build their own analytics to develop IOCs. Third-party threat reputation services[2] can be integrated to enrich collected data or provide IOCs. When an IOC is found, an alert can be created to notify administrators in real-time if this behavior is seen in the future. Finally, SIEM tools usually include the ability to create dashboards with charts and visualizations to provide an overview of the network, for quick access to frequently needed or important information, or to dive deep into data and find relationships. When implementing a SIEM, key features to consider include data ingestion capability for the data collection products it must integrate with, data volume constraints, "canned" analytics offered, custom query development functionality, and third-party intelligence integration for enrichment.

Organizations can enhance their hunting capabilities by incorporating analytics and machine learning into their SIEM solution. **Analytics** turn a volume of data into meaningful information. An analytic will query data to make correlations and find patterns to describe an event. Analytics can drive dashboards, generate alerts, or trigger automated responses. Security analytics can be used in real-time for threat detection. Analytics can also be applied to a data set in hindsight for a reactive response (retrospective analysis). Some SIEM tools include built-in analytics or allow analysts to build their own analytics. Many third-party organizations sell pre-built analytics or create custom analytics as a service.

[2] For more information on threat reputation services, please refer to "Integrate Threat Reputation Services," part of the *NSA Cybersecurity Top 10 Mitigations*.

Many SIEM products are beginning to advertise **Machine Learning (ML)** analytics, which leverage artificial intelligence techniques to find patterns in data. Certain use cases for ML show immense promise, including the reduction of false positives in alerting. Before relying on ML techniques, users should understand what results they want to achieve, because it is difficult to objectively compare the claims of such products.

Frameworks that lay out the typical steps an attacker uses to execute an attack can be used as a guide while hunting for intrusions, helping network defenders identify potential vulnerabilities or IOCs. The **MITRE ATT&CK™ framework** [3] includes a matrix of attacker tactics and techniques and provides drilldowns on each technique including a description, examples, mitigation recommendations, and potential detection methods. Using this matrix, organizations can prioritize which tactics and techniques would be most harmful or likely to occur, and learn how to search for indications that the technique is being used on the network. NSA has also developed the NSA Technical Cyber Threat Framework (NTCTF), a lexicon to characterize and categorize malicious cyber activity across the adversary lifecycle [4].

Another useful guide while hunting is David Bianco's Pyramid of Pain [5]. The Pyramid of Pain categorizes IOCs by the amount of impact (pain) inflicted on the adversary if network defenders discover the indicator and use the information it provides to block the attacker's actions. The pyramid includes hash values, Internet Protocol (IP) addresses, domain names, network/host artifacts (e.g. registry keys, certain file activity, etc.), tools typically used by attackers (e.g. spear phishing software, software to establish command and control, password crackers, etc.), and tactics, techniques and procedures (TTP).

At the base of the pyramid are file hash values. Hash values are easy to detect and block because a computed hash value is almost unique to the original input. This means if a file hashes to a known-bad hash value, hunters can be sure the discovery is malware and can block it. However, this is not a huge inconvenience to attackers because it is easy to flip a bit, causing the hash value to change.

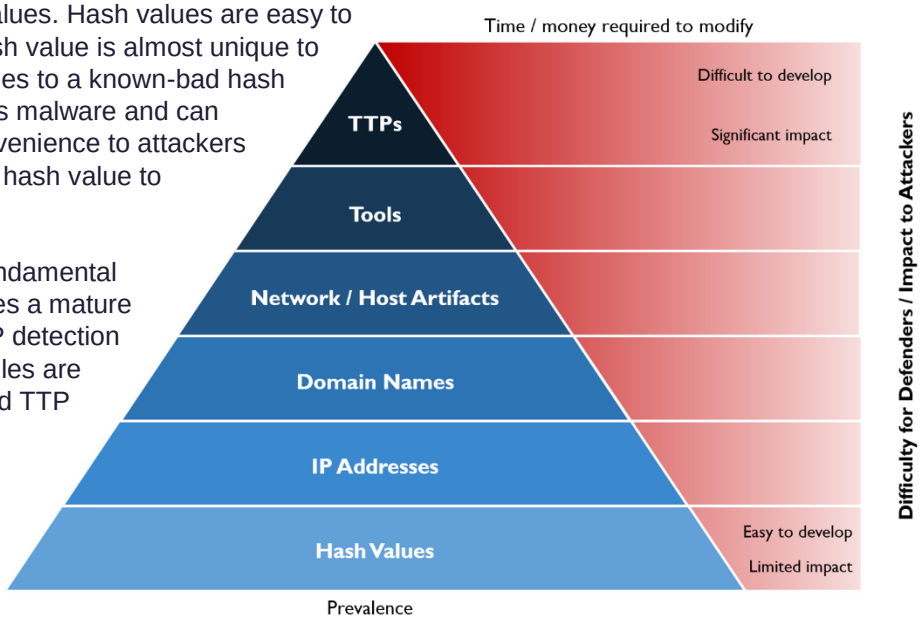

At the top of the pyramid are TTPs, the fundamental approaches for executing an attack. It takes a mature hunting operation to develop effective TTP detection behavioral rules from scratch, but these rules are more likely to thwart an attacker. A blocked TTP will force an attacker to have to start their attack over from scratch using a new method.

Detecting TTPs in a behavioral way is usually preferable for a hunter, but identifying specific attack infrastructure (e.g. by IPs and domain names) is generally still necessary, particularly in understanding the breadth and depth of an advanced compromise. When possible, hunters should try to pivot off of perishable IOCs to discover TTPs to establish a longer lasting mitigation.

Incident Response

When hunting leads to detection, Incident Response activities can help minimize data loss, mitigate vulnerabilities, and restore compromised services. To prepare for incidents, organizations create an incident response plan including procedures for handling an incident, important contact information, and network information. When an incident occurs, the incident response team assesses the target and scope of the attack, and any vulnerabilities enabling it. Once an incident has been detected and reported—and the evidence collected—it must be contained using the strategies and procedures outlined in the incident response plan. Sometimes carefully-orchestrated, extreme measures are required to evict

persistent actors. System recovery plans may need to be exercised.[3] After an incident, organizations should analyze how the incident was handled and make changes to the plan to better prepare for future incidents.[4]

Penetration Testing

Penetration testing, while not generally a hunting activity, is a good practice to assess network boundaries and defensive measures against realistic attack scenarios [7]. While hunting focuses on proactively searching for an attacker in the network, penetration testing focuses on proactively searching for vulnerabilities that attackers could leverage to gain access and attack an organization. Penetration testers can chain techniques together to form an attack path and emulate advanced or persistent adversaries, giving the hunting team a good training opportunity. Penetration testing provides valuable information regarding defensive gaps, procedural flaws, and configuration issues, identifying high-value targets for hunting. Additionally, penetration testing helps determine the sophistication an attacker would need to compromise a system, and helps identify additional countermeasures to reduce attack surfaces and inform incident response policies and procedures. Like hunting, penetration testing works best when repeated often, and shared openly with defensive organizations to stay current on evolving attack methods to improve protections and hunting processes.

Considerations

Hunting for intrusions requires the continuous collection of relevant data to examine. Permission and cooperation from the organization are required in order perform successful hunting operations. The cost of storing all relevant events, the amount of time and effort needed to continuously hunt, false positives and negatives, and difficulty inspecting encrypted traffic, and over-reliance on automation all pose challenges that can incur risk of missing malicious activities.

Be proactive

Hunting for intrusions involves proactively searching for evidence of cyber intrusions on host devices and in network traffic and then remediating any issues quickly and effectively. Continuously hunting on the network enables defenders to find attackers sooner and prevent damage to the organization.

Works Cited

[1] "2017 Cost of Data Breach Study." Ponemon Institute, 2017. Available: https://www.ponemon.org/library/2017-cost-of-data-breach-study-united-states
[2] "2018 Threat Hunting Report." Cybersecurity Insiders, 2018. Available: https://www.domaintools.com/content/2018-Threat-Hunting-Report.pdf
[3] "Att&ck Matrix for Enterprise." MITRE, 2018. Available: https://attack.mitre.org
[4] "NSA/CSS Technical Cyber Threat Framework v1." National Security Agency, 2018. Available: https://apps.nsa.gov/iaarchive/library/reports/nsa-css-technical-cyber-threat-framework-v1.cfm
[5] D. Bianco, "The Pyramid of Pain." 2014 January 17. [Blog] Available: https://detect-respond.blogspot.com/2013/03/the-pryamind-of-pain.html
[6] P. Cichonski, et al. "Computer Security Incident Handling Guide." NIST, SP 800-61 Rev.2, 2012 August. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf
[7] K. Scarfone, et al. "Technical Guide to Information Security Testing and Assessment." NIST, SP 800-115, 2008 September. Available: https://ws680.nist.gov/publication/get_pdf.cfm?pub_id=152164

Disclaimer of Endorsement

Contact

Client Requirements / General Cybersecurity Inquiries: Cybersecurity Requirements Center, 410-854-4200, Cybersecurity_Requests@nsa.gov
Media inquiries / Press Desk: 443-634-0721, MediaRelations@nsa.gov

[3] For more information on recovery plans, please refer to "Exercise a System Recovery Plan," part of the *NSA Cybersecurity Top 10 Mitigations.*

[4] Refer to "NIST SP 800-61 Rev.2: Computer Security Incident Handling Guide" for more guidance on incident response.