SANS Institute
Information Security Reading Room

Intelligent Network Defense

Jake Williams

Copyright SANS Institute 2020. Author Retains Full Rights.

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

Intelligent Network Defense

A SANS Spotlight Whitepaper

Written by Jake Williams

August 2016

Sponsored by
ThreatSTOP

©2016 SANS™ Institute

When an army invades a sovereign nation, one of the defenders' first goals is to disrupt the invader's command and control (C2) operations. The same is true when cyber attackers invade your network. Network defenders must prevent adversary communication, stopping the attack in its tracks while alerting the incident response (IR) team to the point of compromise and nature of the attack.

In the past decade, there has been a shift in focus from network to endpoint when considering how to best defend the enterprise, but smart defenders are turning that tide back to the network. While it's true that endpoint instrumentation is part of a complete defense-in-depth program, all adversaries must communicate on the network to accomplish their goals.

Consider a hypothetical customer environment with 100,000 endpoints and five Internet egress points. Which are easier to instrument and monitor: endpoint or network protections? The answer is the small number of Internet egress points, which accounts for 100 percent of attacker C2 traffic. Using this model, the idea of deriving useful threat intelligence and then pushing threat intelligence out to security at the egress points doesn't seem so daunting.

TAKEAWAY:

Any organization considering large-scale DNS monitoring must carefully consider how it will separate benign traffic from malicious traffic. Monitoring solutions must operate at the enterprise scale.

Considerations for Network Defense Solutions

To ensure maximum survivability and avoid detection, most attackers' C2 traffic connects to domain names instead of connecting directly to IP addresses. In these cases, attacker C2 actions result in DNS lookups, which are easily detectable with the correct instrumentation, but false positives and finding the needle in the haystack are problematic for large environments, some of which experience more than a million DNS requests per hour.

Firewall traffic logs are also useful in detecting attacker C2 traffic, but the scale of firewall log data is orders of magnitude greater than the data provided by DNS queries, which can also provide more definitive results of threats in action.

Utilize Existing Infrastructure

Ideal monitoring solutions should require minimal changes to the network and use existing infrastructure as much as possible. Network defenders shouldn't require new communication paths (e.g. holes in the firewall that an attacker will eventually find) to obtain updates for new defensive solutions. Similarly, installing one more network appliance to perform detection is not ideal because all new appliances will consume untold person-hours for patching, troubleshooting and maintenance.

Maintaining and parsing lists of domain names and IP addresses, manually applying whitelists and blacklists as well as configuration rules to network device configurations, for example, is a suboptimal use of existing network support staff. The administrative overhead associated with creating and applying new device configurations to take advantage of new and existing threat data, particularly in a heterogeneous network, is unacceptable. Threat data provided in a form that can automatically be applied to routers, firewalls, DNS servers and IDS systems is critically important.

When Deep Content Inspection Isn't an Option

Making effective use of network threat intelligence data frequently requires deep content inspection, also known as deep packet inspection (DPI). Many feeds include hashes of known malicious files, but use of encryption in the network may render DPI ineffective. More important, full content inspection is not always possible. Some organizations operate in regulatory environments, for example, where full content inspection of network traffic is simply not allowed. Other organizations choose not to perform deep content inspection of network traffic, citing privacy or efficiency reasons. Therefore, threat intelligence solutions should be maximally effective even when deep content inspection is not employed.

Threat Data Should Provide Actionable Intelligence

The days of supporting a single vendor's product are long gone in most enterprise networks. Especially in the culture of mergers and acquisitions, enterprise networks now consist of products from a multitude of vendors, few of which operate on compatible standards. Actionable threat data feeds must support the widest possible set of vendor product lines while including fallback options that are widely used in the industry.

In particular, this data needs to feed seamlessly into the protection devices that block malicious activity on the network. An ideal fallback solution for obtaining network threat data is iptables. Most network devices run some form of Linux under the hood, and iptables is built into all recent Linux kernels. Threat data vendors that supply data in the iptables format effectively include an "everything else" option.

Future-Proof Your Program

Network defense architectures should also be future-proof. Too many solutions deployed just a few years ago have been abandoned either by the vendors that produce them or the organizations that have deployed them. By integrating solutions that require no dedicated network hardware or appliances, organizations effectively future-proof their defensive strategies.

A final hallmark of a strong, network-based threat data solution is extremely low false positive rates. If blocking actions are taken on a false positive detection, the organization suffers needless reduction in service.

TAKEAWAY:

Ease of integration must not come at the expense of precision in detection.

Selecting the Right Tool for the Job: A Hypothetical Healthcare Case Study

A large healthcare organization wants to measurably increase its security. Given the recent high-profile compromises in the healthcare and insurance industries, the organization wants to protect its reputation. Because seconds save lives, it prioritizes the availability of the information; false positive detections and firewall blocks are unacceptable in this environment. The organization has three Internet points of presence, each with its own firewall. The organization's networking team manages the firewalls centrally.

The organization has more than 25,000 endpoints running various versions of Windows, Linux and OS X. It also has critical systems running on more obscure operating systems, such as Solaris, HP-UX and IRIX—platforms for which most vendors don't supply endpoint agents.

Most of the data has little context, and using the data for hunting has produced occasional false positives, leading the team to believe that much of the "intelligence" in these feeds has not been vetted or is stale.

The organization examines requirements for centrally managing endpoints and is astounded at the hardware and software costs just to run a pilot program. After carefully studying the problem space (small hard drives, limited CPU cycles, FDA certification and excessive management costs), the organization ultimately decides that endpoint detection solutions are not ideal for its needs. The security team looks to network-based detection as the answer.

The threat intelligence team subscribes to several freely available and paid threat intelligence data feeds. The feeds send data in the form of lists of domain names, IP addresses, file names and file hashes that are suspected to be malicious.

The security team has experienced significant challenges in operationalizing its threat intelligence data. First, the security team must convert the data into formats that can be integrated into firewall configurations. This is no small challenge. Two different vendors' boundary firewalls are used to protect the Internet points of presence. Additionally, the organization employs internal firewalls to isolate its research, PCI and patient care environments. Finally, the research team uses a custom-built firewall running on a Linux OS.

These firewalls were deployed without standardization in mind, meaning the IT security team must account for multiple vendors' configuration standards to process threat intelligence data. Due to the heavy focus on availability, the network team is rightfully reluctant to deploy any configuration that might contain unvetted data. Additionally, such a configuration requires outage windows to be scheduled in order to make weekly updates. The organization operates at human speed, substantially slower than its adversaries who are attacking at machine speed.

The team briefly considered a full packet capture solution but did not deploy it because of the volume of confidential information, including patient data, transiting the network. The security team ultimately discounted the idea of full packet capture due to privacy concerns and the high cost of network taps, software and storage.

The security architecture team recognizes that the organization needs a security solution that operates on the network and allows it to operationalize threat intelligence data rapidly. The team recognizes the organizational structure will be a consistent challenge in getting firewall updates processed and seeks an autonomous solution. For this goal to become a reality, the organization will need threat intelligence data that is vetted, is fresh and has an extremely low false positive rate. Installing new equipment or standardizing on a single firewall vendor is a non-starter for this organization.

The organization ultimately employs a threat intelligence data service that makes its data actionable by proactively blocking known malicious traffic both inbound and outbound at the network boundary locations. A large percentage of malware infecting the endpoints was, thus, never allowed to communicate with C2 servers or to exfiltrate data due to this proactive network solution. The solution provides many unique attack detections—detections that were only made possible through the operationalized threat intelligence service. The service operates in a continuous feedback loop that utilizes new threat information in detection and response. This way, the organization achieves a measurable increase in security with a relatively small investment in technology.

Blocking inbound traffic also reduces unwanted traffic, such as spam and known port-scanning hosts. This both minimizes the adverse effects on network devices and prioritizes the traffic the security team must sift through.

Conclusion

For the reasons outlined throughout this paper, detecting threats at the endpoint is becoming increasingly complex and costly for most organizations. These organizations can implement network detection for a fraction of the cost and improve their visibility, detection and response—but only if the team chooses solutions that can rapidly operationalize threat intelligence using existing heterogeneous network infrastructures. After all, threat data that's not rapidly actionable is likely stale data.

About the Author

Jake Williams is a SANS analyst, certified SANS instructor, course author and designer of several NetWars challenges for use in SANS' popular, "gamified" information security training suite. Jake spent more than a decade in information security roles at several government agencies, developing specialties in offensive forensics, malware development and digital counterespionage. Jake is the founder of Rendition InfoSec, which provides penetration testing, digital forensics and incident response, expertise in cloud data exfiltration, and the tools and guidance to secure client data against sophisticated, persistent attack on-premises and in the cloud.

Sponsor

SANS would like to thank this survey's sponsor:

ThreatSTOP®

Last Updated: September 23rd, 2020

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS October Singapore 2020	Singapore, SG	Oct 12, 2020 - Oct 24, 2020	Live Event
SANS Community CTF	,	Oct 15, 2020 - Oct 16, 2020	Self Paced
SANS SEC504 Rennes 2020 (In French)	Rennes, FR	Oct 19, 2020 - Oct 24, 2020	Live Event
SANS SEC560 Lille 2020 (In French)	Lille, FR	Oct 26, 2020 - Oct 31, 2020	Live Event
SANS Tel Aviv November 2020	Tel Aviv, IL	Nov 01, 2020 - Nov 06, 2020	Live Event
SANS Sydney 2020	Sydney, AU	Nov 02, 2020 - Nov 14, 2020	Live Event
SANS Secure Thailand	Bangkok, TH	Nov 09, 2020 - Nov 14, 2020	Live Event
APAC ICS Summit & Training 2020	Singapore, SG	Nov 13, 2020 - Nov 21, 2020	Live Event
SANS FOR508 Rome 2020 (in Italian)	Rome, IT	Nov 16, 2020 - Nov 21, 2020	Live Event
SANS Community CTF	,	Nov 19, 2020 - Nov 20, 2020	Self Paced
SANS Local: Oslo November 2020	Oslo, NO	Nov 23, 2020 - Nov 28, 2020	Live Event
SANS Wellington 2020	Wellington, NZ	Nov 30, 2020 - Dec 12, 2020	Live Event
SANS OnDemand	OnlineUS	Anytime	Self Paced
SANS SelfStudy	Books & MP3s OnlyUS	Anytime	Self Paced