Network Perimeter Defense

What Information To Collect At The Network Perimeter

September 2014

EnergySec
8440 SE Sunnybrook Blvd., Suite 206
Clackamas, OR 97015
877-267-4732
energysec.org

Executive Summary

There is an almost never-ending supply of information that can be collected by network perimeter security devices. The purpose of this white paper is not to list all of those pieces of information. Rather, the purpose is to highlight some of the most important types of information which should be collected and spotlight where, how, and why to collect that information. After reading this paper, you should have an expanded view of what constitutes a network perimeter and what information should be collected from devices on that perimeter.

In order to know what information should be collected at your network perimeter, it is important to have a working definition of what that perimeter is. A perimeter can be defined as anywhere information enters or leaves a network. Firewalls are the first thing many people think of when the idea of perimeter security is discussed, and that is an important area to secure. By having a broader definition, though, we are able to encompass all that a perimeter does. For example, Internet traffic should enter and exit your network by using a proxy server. If we were to simply focus on firewalls, we would be missing this important source of information about what is happening on a network.

Verbosely Log All Traffic Arriving at Device

The first key to logging and monitoring at a network perimeter is that all traffic arriving at perimeter security devices should be logged. These logs serve multiple purposes. They allow you to have visibility into what is happening on the network. It is impossible to differentiate between normal traffic and potentially malicious traffic without such visibility. Another purpose for logging traffic arriving at perimeter security devices is that in the case of a data breach or cyber intrusion, it will allow forensic examiners to investigate what happened, when it happened, and the results of the incident. These logs should include date and timestamps, as well as source and destination addresses and ports. If your organization uses a consistent naming convention for its assets, you can use a Security Information and Event Management system (SIEM) to also log the location of the machine, and the asset owner or department.

Individual TCP Sessions

By having a complete record of TCP sessions, you can compare traffic to Indicators of Compromise as they become available.

As part of logging all traffic arriving at perimeter security devices, it is important to track individual TCP sessions. An example of the benefit of having logs of where connections came from, when they happened, and who was involved is from the recent Energetic Bear or Dragonfly attacks. Symantec released information about the command and control (C&C) servers which were used in the attacks.[1] By having a complete record of TCP sessions, you would be able to compare those C&C servers against traffic leaving your network and determine if you had possibly been a victim of the attack. These logs should include the source and destination IP addresses, timestamps, and can also include additional information such as the user who was logged in when the connection was made, if your infrastructure has such capabilities.

VPN Connections

All virtual private network (VPN) connections should be logged. It should include the source address for the VPN connection, as well as what user is logging in. This can help alert if user credentials have been stolen. For example, if a user normally logs in using their VPN connection from Salt Lake City and a connection comes in from Belarus, then that should raise an alert and be checked by security personnel. It's possible they are on vacation, but it could also mean their credentials have been compromised.

Configuration Changes

Configuration changes should be reviewed by security personnel since end users may not understand the security implications of the changes they make.

Any configuration changes made to perimeter devices should be logged and reported to security personnel. This protects against both malicious attackers as well as unintentional configuration errors. An authorized user could make a configuration change, and not realize that the changes they made also allowed a security vulnerability to occur. A change management program should include the location of the device which was changed, the department where it is located, and other details such as the user and time of the configuration change, as well as what changes were made.

Wireless Access Points

A network perimeter cannot be secure if the boundaries of the network are not known. If a rogue access point is on your network, then it is impossible to ensure that the network perimeter is secured. If you have an enterprise-grade router and access points (AP), then it likely has the capability to scan for rogue access points. For example, if you have Cisco products, they offer Infrastructure Rogue Discovery,[2] which embeds an authentication Information Element (IE) into beacons and probes from authorized access points. It then monitors for beacons that do not have that IE, and can create an alert if it finds a rogue AP.

If a rogue access point is on your network, then it is impossible to ensure that the network perimeter is secured.

If your infrastructure does not support rogue AP detection, you can discover them by using third-party software. For example, Kismet is a free wireless network detector, sniffer, and intrusion detection system.[3] It passively collects packets and detects named networks, and can even discover the name of hidden networks. The use of a tool such as Kismet would need to occur regularly, and would require additional effort to perform, but it would provide an added layer of security assurance.

Another possible option is to collect and monitor the Dynamic Host Configuration Protocol (DHCP) logs. Most, if not all, access points are configured to use DHCP by default. Monitoring the DHCP logs provides insight into what's being plugged into the network and requesting an IP address. The media access control (MAC) address of the client can be used to determine the vendor of the network interface card (NIC), and then be compared to a company-approved list of vendors. This approach can also be used to detect new equipment being connected to a network.

Whichever method is used to find access points, any APs found will need to be logged and checked against a list of authorized APs, so that the rogue APs can be discovered and neutralized if needed.

Two-Factor Authentication Services

Just as logging standard authentication events, any service used to perform two-factor authentication should be logged. This can be used to determine whether or not a user's token was stolen or potentially shared. For example, when the RSA was attacked back in 2011 and their SecurID tokens were compromised, companies could have detected if they were victims of an attacker using the compromised tokens.[4]

Excessive Login Attempts

Excessive login attempts are a tell-tale sign of a brute-force password attack. While it may sound more glamorous to defend against more advanced attacks, this kind of thing still happens and needs to be prepared for. For example, a brute-force attack was used to compromise a Supervisory Control and Data Acquisition (SCADA) system at a public utility just this year.[5] While having an Internet-connected portal with only a password protecting it was not a good idea, monitoring failed login attempts would facilitate detection and remediation of such an attack.

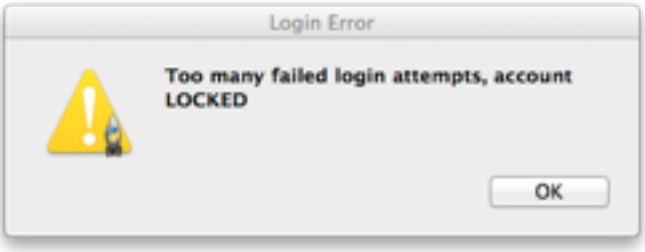

DNS Logs

All domain name service (DNS) queries should be logged. Since DNS records change over time, it is important to have a record of how DNS requests have been resolved in the past. Without this information, conducting a forensic investigation of a breach or security incident is immeasurably more difficult. For example, when the Havex vulnerability was disclosed, the disclosure included information on what domains had been used as command and control servers for the malware.[6] Having DNS logs would allow you to investigate whether any hosts from your network had been contacting those servers.

It is also important to log failed DNS queries. Since C&C servers oftentimes move to new locations, a machine which has been compromised as part of a botnet is likely to send out many failed DNS queries as it searches for its C&C server. By logging which machines are creating DNS queries that fail, you will have a clue that a machine may be infected.[7] Review of DNS query activity can also help identify potential DNS tunneling activity.

Summary

Start small, measure, and scale up in order to improve security monitoring and logging.

Logging and monitoring is necessary in order to make a perimeter defense solution effective. Logs must be thorough since they are likely the only evidence which will be available during and after a cyber attack. Logging information at scale can quickly create an overwhelming amount of information, however, as NERC says in their CIP-005 Compliance Analysis Report, “The bottom line is that entities should log as much as possible. While it could consume significant storage, detailed logs serve as a valuable source of forensic information.”[8] That collection of information is what will help to protect your organization. One way to begin to improve this process is to use the “start small, measure, and scale up” process.[9] You can make a small improvement, and measure the results. Those results can then be used to justify further investment, as well as providing tangible benefits themselves. Then, what you have learned can be applied to expanding your logging and monitoring program. To help analyze what could become a massive amount of information, our next paper in this series will examine how to make sense of all the information which has been generated from your monitoring and logging program.

EnergySec
8440 SE Sunnybrook Blvd., Suite 206
Clackamas, OR 97015
877-267-4732
energysec.org

[1] “Security Response,” Symantec, TLP: Amber, Jun. 2014.(Note: This document was restricted using TLP: Amber, and so was not publicly released.)

[2] http://www.cisco.com/c/en/us/support/docs/wireless-mobility/wireless-lan-wlan/70987-rogue-detect.html

[3] http://www.kismetwireless.net/

[4] P. Bright, “RSA finally comes clean: SecurID is compromised,” Ars Technica, 07-Jun-2011. [Online]. Available: http://arstechnica.com/security/news/2011/06/rsa-finally-comes-clean-securid-is-compromised.ars

[5] “ICS-CERT Monitor,” ICS-CERT, Apr-2014. [Online]. Available: http://ics-cert.us-cert.gov/sites/default/files/Monitors/ICS-CERT_Monitor_%20Jan-April2014.pdf

[6] “Security Response,” Symantec, TLP: Amber, Jun. 2014.

[7] “Using Log Correlation Engine to Monitor DNS (Revision 2),” tenable.com, 06-Sep-2013. [Online]. Available: http://static.tenable.com/prod_docs/LCE_DNS.pdf

[8] “CIP-005 Compliance Analysis Report: Electronic Security Perimeter(s),” NERC, May 2012.

[9] J. Gardiner, M. Cova, and S. Nagaraja, “Command & Control: Understanding, Denying and Detecting,” University of Birmingham, Feb. 2014.

