ENERGYSEC
SERVING THE INDUSTRY'S SECURITY COMMUNITY SINCE 2004

Network Perimeter Defense

Analyzing the Data

September 2014

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

This material is provided for general information purposes only. You should make your own judgment as regards use of this material and seek independent professional advice no your particular circumstances. Neither the publisher, nor the author, nor any contributors assume any liability to anyone for any loss or damage caused by any error or omission in the work, whether such error or omission is the result of negligence or any other cause.

Executive Summary

An organization or security analyst can quickly become overwhelmed with monitoring logs and responding to security alerts. One problem is that, in order to be thorough, more logs will need to be collected than any team can realistically expect to check and respond to in a reasonable amount of time. Having a mature, effective security analysis operation, however, will allow important events to receive the attention they require. The first key to having a mature security analysis operation is to understand the normal, baseline operation of your networks. The second is to automate the collection and analysis of logs as much as possible, so that human time and resources are only being used to investigate events which require a human intellect to make sense of. The purpose of this paper is to provide tips for how to accomplish these two goals.

Establish a Baseline

In order to make the best use of logs, a baseline of what's "normal" should be established. The baseline can be fine-tuned by creating profiles for different time windows, devices, and network services.[1] There are several things which can be done to help ensure an accurate baseline is established:

Establishing a baseline of normal behavior is the first key to identifying abnormal, malicious behavior.

Ensure Logs are Consistent

An automatic, daily check can be performed on logs. The check is to ensure they are within expected parameters. For example, if the daily log for a given purpose has an average of 10,000 entries, and then on a certain day only has 50 entries, that would be a sign that there is something wrong with either the creation or collection of the logs, and the cause needs to be investigated. This check can also be used to ensure that the logs are formatted properly and have not been corrupted.

Firewall Rules

Review current firewall rules to ensure that only the approved ports and protocols are enabled. In some cases rules may have been approved for a specific business use, but they may no longer be required. In this case it is always a good idea to engage the business and determine if the given traffic is still required. Other times the review can be done by simply monitoring the firewall traffic for the rules in question, or in some cases working with the network team to get hit counts for the rules on the firewall.

To improve the efficacy of rule sets, tests can be arranged to sample traffic in order to determine how the firewall and rule sets are actually being utilized. For example, you can search for HTTP traffic which is neither sourced nor destined for a DMZ proxy. Ideally, all HTTP traffic would go through a proxy, with a “deny all” rule at the end of the ruleset, but it is good to actually test the performance in the event that configuration errors have found their way into firewall rules.

Minimize Noise

It’s not uncommon for servers and/or workstations to be configured improperly. For example, a proxy server could be missing or misconfigured, which may result in outbound denies at the perimeter. These log entries are an example of noise, or log entries which are insignificant from a security standpoint, but add to the data which must be filtered through in order to find actual threats. Collecting too much information adds costs due to the increased storage space required, but more importantly, it increases the time it takes to search and monitor.

Minimizing noise is the process of analyzing outbound denies and determining if there is a configuration issue, or something more nefarious.

Minimizing noise is a process by which one analyzes the outbound denies, and works with the system owners to determine if there’s a configuration issue, or potentially something more nefarious going on. This is a time-consuming and ongoing process, but will help to reduce the number of false positive alerts which are raised by a system. It also makes it easier to more rapidly identify anomalies, and helps to prevent an alert system from overwhelming analysts with insignificant alerts.

An example that we have seen where this was beneficial involved desktop wallpapers. Some users had attempted to configure their workstation to get a new wallpaper each day from a web service. These requests were denied at the firewall, but they created many false positive alerts. By correcting the configuration of the workstations, the organization was able to conserve network resources, while at the same time reduce the amount of noise in their alert system. By eliminating the false positives, it allowed the security analysts to focus on other, possibly more significant, alerts.

One way to begin to minimize noise is to set up a firewall inline, and use it for testing purposes.[2] If you set up this firewall with the rules you want to apply, plus a final rule to permit and log all traffic, then you will be able to see what traffic is hitting that final rule. This will tell you what traffic is on the network and not matching your current rules. This information can then be used to add to or refine the firewall rules, and repeated until you are able to take that ruleset and apply it to the actual perimeter firewall, with a “deny all” rule at the end.

Correlate Event Logs with Vulnerability Scans

A great way to identify problems, especially if you don't have access to firewall rules, is to use a vulnerability scan. Correlating your event logs with vulnerability scans which are performed provides several benefits. First of all, you want to ensure that the activity from the vulnerability scanning tools is logged. If that activity is not logged, then you may not be aware if an attacker is performing vulnerability scans on your network. By comparing your vulnerability scans to event logs, you are also able to check if an exploit, which you found with your vulnerability scan, had previously been used against a target. You can have a set source and time frame for vulnerability scans, and then if logs or alerts are raised by somebody scanning your network, you will be able to easily filter out the legitimate scans from a scan by an unknown person. This is a great way to put your logging and monitoring architecture through a real-world test to see if it is detecting threats, at the same time using compliance with NERC CIP-005-01 R4 (version 3), requiring vulnerability assessments, as a driving force to improving security, not simply marking a checkbox on an audit.

Once a baseline profile of “normal” traffic is established, it is time to examine things which should grab a security analyst's attention.

Identify Anomalies

Once a baseline profile of “normal” traffic is established, it is time to examine those things which should grab a security analyst's attention. You should regularly run reports, a minimum of biweekly, that identify anomalies in logs. Any anomaly should be actively reviewed, with a process in place for documenting any findings. It is important to not rely on filtering logs to identify events which require further investigation. Doing so would allow you to respond to known malicious traffic, but it would miss events which were not previously known as being bad traffic, and would not satisfy NERC CIP-005 R3 (version 3).[3] In this regard, it may be helpful to think of anomaly detection as “finding what you don't know to look for.”[4]

TCP Sessions

The first area to look for anomalies is in TCP sessions. There are several things which should raise red flags and draw an analyst's attention.

Examples of Anomalous TCP Sessions

Long-Lasting Sessions

Alerting on long-running TCP sessions is a great way to raise notifications of potential issues. For example, this could be a sign that a host on a network has been compromised and is communicating with a command and control (C&C) server.

High Bandwidth Sessions

Sessions which consume a lot of bandwidth, particularly outbound, should always draw the eye of the security analyst. This may be completely valid business traffic, or it could be an indication that data has been exfiltrated, or malware has been downloaded to a given host.

Unexpected Sessions

Look for sessions that aren't approved or typically allowed. For example, normal web traffic (port 80 and 443) is only allowed via the proxy server. If, for some reason, sessions have been discovered that weren't sourced from the proxy server(s), then this would highlight an unapproved rule or an undocumented change.

HTTP Traffic Which Bypasses Proxies

The purpose of having a web proxy is to have visibility and protection in place for both the network and end users. Some people, however, will intentionally bypass proxies. They may have an innocuous reason for doing this, such as to access a music streaming site they like to listen to, but these work-arounds create security vulnerabilities. Also, an attacker may attempt to bypass a proxy in order to connect to a C&C server without being detected. In order to combat this, packet sniffers should be deployed on demilitarized zones (DMZs) to look for HTTP traffic which bypasses HTTP proxies. Any traffic which avoids a web proxy should raise an alert. The packet sniffer in the DMZ would be able to do this. Another way to identify machines which are bypassing the proxy server uses the fact that a network should have a "deny all" rule at the end of the firewall rule sets. Logs generated by this rule should highlight machines that aren't configured to use the proxy or are trying to bypass it.

Perimeter Defense is not just about preventing access, but also about preventing data loss.

Monitor for Data Exfiltration

An important use of perimeter defense is not just to prevent access, but to prevent data loss. This is done by using egress filtering. You can monitor for keywords or sensitive information, such as personally identifiable information (PII), and use the network boundaries to block the loss of that information as well as raise an alert for security personnel. Along with this, you can monitor for unauthorized use of encryption, which can be used to hide data as it is being exfiltrated, or hide traffic going to an external C&C server. Commercial products are also available which protect against the exfiltration of data. It is important, even if the attempted exfiltration is successfully blocked, to follow up with an investigation of those attempts.

Proxy Server and Web Filter Logs

Having complete and accurate proxy server logs will aid immensely in an investigation of an infection. A Palo Alto Networks study found that 90% of unknown malware was delivered by web-browsing, rather than email.[5] This shows the necessity of being able to identify information related to web browsing. With the proper monitoring, you will have the capability of looking at the proxy logs and seeing where an infection had come from.

VPN Concentrator Logs

Virtual private network (VPN) concentrator logs allow you to see when and where users are accessing your network by using VPN services. One example of when an alert should be raised is if a user is accessing the VPN from multiple computers at the same time. This is a strong sign that their credentials have been compromised. Another example would be if a user uses their VPN credentials from a different address than they normally do. This could be benign, such as an employee getting some work done while they are on vacation, but it could also be a signal that the account has been compromised.

Mail and SMTP Logs

Mail and SMTP logs allow an analyst to investigate the source of an infection. They will allow an analyst to examine details after a malicious email. For example, was the email to one person, or to many people? Where was the suspicious email sourced from? That information can help in a forensic investigation.

IDS/IPS

Intrusion Detection Systems (IDS) can record a lot of information. For example, it can be configured to record all FTP traffic or all SMTP traffic. This information would allow you to be able to reconstruct information flows, which is invaluable in an incident response situation. It is important to configure an IDS or intrusion prevention system (IPS) to minimize excess noise, similar as with firewall logs. Having a flood of alerts, with a high incidence of false positives, creates a management nightmare for a security analyst. By using the systematic, albeit time-consuming, process of fine-tuning the alert system, you make it more likely that an important alert will be noticed and investigated.

Use your identified normal behavior to then identify anomalous behavior.

DNS logs

Having a history of DNS logs is very beneficial. You may have a workstation which is infected, and you find the infection immediately. Just as likely, however, is that some time will pass between the time of the infection and the time of discovery. Since DNS records change over time, it is important to have a record of how DNS requests have been resolved in the past, so that the source of the infection can be investigated.

Conclusion

Having an effective security analysis operation is one of the keys to having a secure network. This will help to ensure that cyber incidents are identified and remediated in a timely manner. One of the keys to having an effective analysis operation is to identify what the normal, baseline operation of the network looks like. This allows analysts to then be able to identify when anomalous behavior does happen. The next key is to automate the collection of data about the network as much as possible, so that analysts' time will be able to be spent efficiently, responding to any events that may happen.

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

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

[2] R. VandenBrink, “Egress Filtering? What - do we have a bird problem?,” InfoSec Handlers Diary Blog, 11-Jul-2014.

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

[4] T. Dunning and E. Friedman. Practical Machine Learning. Sebastopol, CA: O’Reilly Media, Inc., 2014.

[5] “The Modern Malware Review,” Palo Alto Networks, 1st Edition, Mar. 2013.