tenable network security

Discovering Malware by Looking for Abnormalities

November 2015

(Revision 1)

Table of Contents

Purpose

With modern malware threats, it's common to hear about a piece of malware going undiscovered for a year or more. As state sponsored actors become more prevalent in espionageware, this is expected to become even more common. In this paper, we will encourage you to turn the statistical hunt for malware on its head in order to identify abnormalities, and from those abnormalities, potentially discover malware that puts your organization at risk. This paper is intended to be of a higher level, and will not go into the technical details on implementation; existing Tenable resources should be utilized for specifics.

Introduction

Those engaged in network support have long known the adage: 80% of the problems are solved by addressing the top 10% (or 20%) of the problems. It's also taken for granted that the last 10% of problems (or systems) take much more effort to find and address than the other 90%. To that end, we're almost hardwired to look at the top 10 events in the intrusion, system, and network traffic logs. By remediating the top 10 events, we've expended minimum energy with maximum return on our efforts. Once the top 10 list has been addressed, we run the report again, getting a new top 10 list, lather, rinse, and repeat. It's to the point where we can almost do it in our sleep. Management is happy because we can show effective use of our time, and things are getting done.

While all of this is going on, many network teams don't take the time to do a normalization of their network. Since their time is spent addressing their top 10 issues, they can't articulate what is considered to be "normal" for their environment. What is the ebb and flow on the network? In neglecting this, people are surprised to learn they have been hosting malware, sometimes for years, due to the signal being hidden in the noise. Information security practitioners named 2001 "the year of the worm." Nearly every week, there was a high profile worm released that needed to be chased down and eradicated. What made this work easier was that there was no hiding the traffic; all of the worms were top traffic generators and very noisy by application standards. At that time, several malware researchers expressed concern that malware authors would learn their lessons and make their products more low profile and low noise. The news in 2014 and 2015 shows that the authors have learned that lesson, with malware such as Urberos, Chewbacca, Backoff, PoSeidon, and other state-sponsored espionageware having become the story leads.

Due to the low volume of network traffic created by modern malware, they will not be identified by looking at the top 10, 20 or even 100 identified issues. By knowing what is normal, these intrusions can be identified by being abnormal, and by patterning network activity. While this process will not tell you what is malicious, it will identify areas that need additional investigation, which can lead to previously unidentified malware. Examples of this could be noticing that a printer in the shipping department is connecting to an international system every evening, or systems are making web connections to the display terminal in the lobby at unusual times.

First Steps

Before we can define abnormal activity, we have to have a baseline from which to operate. Not only do we need to know what assets are in use on the network, but we also have to understand the flow and utilization of network traffic. Assets such as web servers, DNS servers, FTP servers, and file and print servers should be identified and categorized. Network traffic needs to be monitored and patterned. While in this mode, it will be common to identify hosts that come and go, as well as what seems to be odd peak utilizations. These can often be quickly discounted as transient mobile devices such as laptops or phones belonging to employees, or depending on your network policies, even guests with personally-owned assets, while traffic can be identified in reoccurring cycles (such as month-end billing cycles with increased traffic from the accounting department). Most business cycles can be normalized within a quarterly time period. While some cycles do occur once or twice a year, these are one off, limited situations and will be identified as part of the ongoing process.

Keep in mind that normalization is not a "once and done" task. Instead, it's a cycle. Once the baseline normal is defined, continuous network monitoring allows for refinement, and in some cases it may redefine what "normal" is. The longer "normal" is known, the more refined and granular it becomes. This process has allowed large Internet providers to adjust for seasonal peaks, and web hosts to adjust for high volume demands, but we can also use it to meet other corporate needs and overall awareness.

Some final words on normalization: just because you've identified a normal baseline, do not expect that single baseline to apply across the entire organization. While baselines create averages and set a level of expectations, there are and will be exceptions to every rule. Companies create standardized host images and policies, but failing to anticipate the need for exceptions has caused more issues than standardization has solved. Expecting to identify and deal with this type of abnormality in the early stages of normalization will cause the process to go smoother, and cause less hostilities from the supported environment. Scenarios such as anticipating engineering to temporarily stand up a web server while they try to troubleshoot a customer issue may be required and should be foreseen. Creating ways to notify the support staff of the exception and how to address the situation when procedures are not followed needs to be established early.

Moving Beyond "Normal"

The key to making all of this work is continuous network monitoring. Once the baseline is established, the cycle begins. While it's tempting to look only at the top 10 or 20 events, we still look at those, but we also look at the bottom 10 or 20 at the same time. While the bottom of the list will require more effort, the payoff will be larger. In order to be the most effective, we have to account for each deviation from the baseline. Using the earlier example of the printer in the shipping department, once this abnormality is identified, the network support staff can look closer at the traffic and determine if the printer is connecting to the manufacturer and checking for a firmware update. However, it could be that the display terminal in the lobby has had a web server installed, and compromised machines are performing HTTP POST commands, exporting sensitive data for the botnet herder to collect, and use the terminal to make further attacks or compromises.

Deviations from normal activity, such as the display terminal having a web server installed, could have brought a faster response. If malware is already hosted while the normalization period is occurring, these secondary identifiers will expose it early in the process.

While we've looked at network traffic, other normalizations should also be occurring. We already mentioned having standard desktop and server images. Monitoring hosts for deviations from image norms should also be leveraged. Look for things like unique run-on-starts, applications, or registry entries. While this may turn out to be limited use software, it's also an indicator of initial malware activities. In some cases, this may be the only indicator of malware, as malware authors continue to look for "fileless" infectors and embed JavaScript into the registry in order to infect hosts.

Completing the Cycle

Since the security or network support staff that monitor activity are often not the same individuals who investigate and remediate issues, it is critical that all responses and findings be closed out by reporting back to those who discovered them. Not only is it courteous, but it allows for the updating of the baseline and provides closure for the security team. In some cases, this can create actionable intelligence to support others in the company, or even in the information security community.

The results of investigation and remediation also feed into the corporate knowledge base. No matter if this is a formal database or just shared with the rest of the support team, knowing what the outcome is will shorten troubleshooting and remediation times in the future. This feedback also helps eliminate false positives, and allows the team to focus on known and previously unidentified abnormalities.

Understanding that the teams doing this work are always suffering under a time crunch, we do need to address the time required for these tasks. While it can be time-consuming to look at and chase after the abnormalities (or the bottom 10, as opposed to the top 10), it has long been said that it costs three times as much time to remediate an exploit as it does to prevent it. Let's also look from another point of view: had one of the victims of the latest high-profile espionageware infections engaged in this process, rather than finding the malware years after the infection, they could have found it in only months, if not weeks or even days. While we're all pressed for time and to "do more with less", the public relations fallout and loss of trust in the security of your own network would have been reduced along with the time to discovery. While it's not as high-profile as being the white knight that rode in saving and securing the network, there is no extra money spent on outside consultants and second guessing. What management team doesn't like to save money? This can be accomplished through an earlier discovery of threats.

The sooner that a baseline is established, the sooner patterns and abnormalities can be identified. Even though it's recommended to take a calendar quarter to establish a baseline, patterns and abnormalities will often start emerging during the process.

Conclusion

Not every abnormality will result in new malware. Quite the contrary, most often abnormalities will be misconfigurations or another explainable situation. Gaining an initial baseline and maintaining "normal" is a time-consuming process that can take as much time as you are willing to devote to the task. It can also be maintained with a couple of afternoons a week, depending on the amount of traffic and other items you are looking to monitor, as well as the tools used to complete the tasks.

While the effort may seem daunting, the payoff in early detection and identification of new malware or other attacks is well worth the expenditure. Through early detection, you reduce the cost of responding to the incident and prevent leakage of sensitive information from the company. This information isn't limited to customer credit card information from compromised POS devices, but can also be proprietary company data lost to industrial espionage, or a network topography that can be used for further exploitation, whatever an attacker finds value in.

Tenable Support for Abnormality Detection

Tenable has published multiple resources on utilizing our SecurityCenter and Nessus solutions to detect malware and abnormalities. These papers include:

- Comprehensive Malware Detection with SecurityCenter Continuous View and Nessus
- 24/7 Visibility into Advanced Malware on Networks and Endpoints
- Tenable Malware Detection: Keeping Up With An Increasingly Sophisticated Threat Environment

These resources are available for free download on the Tenable Network Security Resource Library, as well as having a discussion forum for Indicators of compromise and Malware.

About Tenable Network Security

Tenable Network Security provides continuous network monitoring to identify vulnerabilities, reduce risk, and ensure compliance. Our family of products includes SecurityCenter Continuous View™, which provides the most comprehensive and integrated view of network health, and Nessus®, the global standard in detecting and assessing network data. Tenable is relied upon by many of the world's largest corporations, not-for-profit organizations and public sector agencies, including the entire U.S. Department of Defense. For more information, visit tenable.com.