Global Information Assurance Certification Paper

Finding Bad with Splunk

GIAC (GCCC) Gold Certification

Author: David Brown, mrdavebrown@gmail.com
Advisor: Kees Leune
Accepted:

Template Version September 2014

Abstract

There is such a deluge of information that it can be hard for information security teams to know where to focus their time and energy. This paper will recommend common Linux and Windows tools to scan networks and systems, store results to local filesystems, analyze results, and pass any new data to Splunk. Splunk will then help security teams narrow in on what has changed within the networks and systems by alerting the security teams to any differences between old baselines and new scans. In addition, security teams may not even be paying attention to controls, like whitelisting blocks, that successfully prevent malicious activities. Monitoring failed application execution attempts can give security teams and administrators early warnings that someone may be trying to subvert a system. This paper will guide the security professional on setting up alerts to detect security events of interest like failed application executions due to whitelisting. To solve these problems, the paper will discuss the first five Critical Security Controls and explain what malicious behaviors can be uncovered as a result of alerting. As the paper progresses through the controls, the security professional is shown how to set up baseline analysis, how to configure the systems to pass the proper data to Splunk, and how to configure Splunk to alert on events of interest. The paper does not revolve around how to implement technical controls like whitelisting, but rather how to effectively monitor the controls once they have been implemented.

1. Introduction

Keep the bad actors out. This has been the battle cry of the information security profession. But, in an age where attackers have nation state resources at their disposal, and in an age where zero day exploits are discovered with sickening consistency, keeping the bad actors out is increasingly becoming an impossible charge. The National Security Agency has declared, "The question is not if a system will be compromised but when" (National Security Agency, 2015, p. 1). The castle wall (firewall) and mote (intrusion prevention system) approach to security is failing at a rapid pace.

To survive, the information security profession must turn its attention to finding these bad actors who have already breached. This new mantra of the information security profession can be summed up: "Prevention is ideal, but detection is a must." Early detection is essential to thwarting a skilled adversary. A Verizon study found that it took attackers mere minutes to compromise the vast majority of their targets and that they often achieved data exfiltration within days (Verizon, 2016, p. 10). A Mandiant Consulting study found that the median time to discover a compromise, whether by internal detection or external notification, was 146 days (Mandiant Consulting, 2016, p. 9). Somehow these bad actors are getting in, and somehow the information security professionals are failing to see them.

Trying to determine if there is an intruder on a network can appear to be an insurmountably daunting task. Turning to established security guidelines, such as the Center for Internet Security's Critical Security Controls, is a good place to start. The Australian Cyber Security Center estimates that 85% of targeted attacks could be thwarted by application whitelisting, patching, restricting administrative privileges, and creating defense in-depth (Australian Cyber Security Center, 2015, p. 18). The Critical Security Controls include the controls mentioned above as well as additional controls that can significantly enhance the security of a network. Considering the information that is generated from the Critical Security Controls can help the information security professional detect when the information security controls fail or are bypassed.

Take application whitelisting as an example of a security control that can help detect an intruder. Whitelisting is a highly effective control for stopping attacks dead in

their track by preventing unknown payloads from executing. Whitelisting cannot stop all attacks. Attackers who download their tools to a workstation and encounter whitelisting may turn to software that is already installed on the workstation, such as PowerShell, to advance their attack. Paying attention to the whitelisting control can help detect the initial, failed attack where the attacker tried to utilize his or her own tools. Assuming that the Critical Security Controls – like application whitelisting – have begun to be implemented, the key to detection lies in listening for the signals put off by the Critical Security Controls.

Take DHCP requests as another example of a security control that can provide early detection of an intruder. The first Critical Security Control directs that DHCP requests should be logged (Center for Internet Security, 2016, p. 6). This control accomplishes many purposes. For one, it helps identify new systems on a network, allowing those systems to be properly logged in inventory, verified, and securely maintained. If a system is not logged in inventory, chances are the system will not be securely maintained and will be sitting as an open invitation for any bad actor. Second, DHCP requests can immediately detect an unauthorized device that appears on a network. If a mysterious and unexpected DHCP request appears on a network, it is highly possible that someone has brought a device onto that network that does not belong.

Combing through and monitoring logs for metrics such as whitelisting blocks or DHCP requests is no trivial task. Logs are scattered across every machine on a domain and contain more noise than a 747 preparing for takeoff. This is where Splunk really excels with its ability to selectively collect targeted data from system logs. Splunk's Universal Forwarder makes it simple for any information security professional to gather log data that relates to the Critical Security Controls from across a network, consolidate it into one place, and leave the noise behind. Splunk can then be setup with dashboards, saved searches, and automatic alerts based on criteria such as application whitelisting blocks or new DHCP requests.

Collecting logs and effectively filtering through all the noise has become necessary to finding bad actors. The National Security Agency pointed out that this is becoming all the more essential as bad actors are becoming increasingly adept at hiding

their tracks and blending in with the deluge of network traffic (National Security Agency, 2013, p. 1). Trying to find bad actors without log aggregation and an effective means to filter through all the noise is akin to looking for a suspect by going door-to-door in a metropolis.

The Critical Security Controls are considered a baseline industry standard for cyber security by many. For example, in a recent publication, the California Department of Justice stated that the "… Critical Security Controls define a minimum level of information security that all organizations that collect or maintain personal information should meet. The failure to implement all the Controls that apply to an organization's environment constitutes a lack of reasonable security" (California Department of Justice, 2016, p. 30). The combination of the Critical Security Controls, built-in Linux and Windows tools, and Splunk's elegant data collection capabilities can truly help the information security professional cut through the noise and achieve the mantra: "Prevention is ideal, but detection is a must."

2. Initial Setup

Scanning networks and systems, analyzing changes from baselines, and alerting on those changes can largely be accomplished using common tools. Instead of purchasing more and more security tools and investing in the "fog of more," security professionals can take creative approaches that utilize built-in Linux and Windows tools and common network monitoring tools to accomplish the goals of the Critical Security Controls.

This whitepaper utilizes common tools such as Nmap and WMIC to scan networks and systems. Built-in tools such as diff and grep are then utilized to analyze and compare scan results and save the variances to local filesystems. Splunk is then used as the engine for alerting on network and system changes and for storing the history of these changes over time.

This section lays the groundwork for running the commands and setting up the alerts contained in this paper. The commands can be run manually at first, but should ultimately be automated using batch scripts with a task scheduler or shell scripts with cron jobs. This paper does not go into detail on how to incorporate the commands into

batch or shell scripts and assumes the security professional has the technical acumen to do so.

2.1. Tools

Most of the commands contained in this paper are designed to be run from Linux since Linux distributions come with text parsing tools such as diff and grep already loaded. Many of these commands could be adapted with minimal changes to run in a Windows environment. That said, the WMIC commands should be run from a Windows computer joined to the domain. This paper will not go into detail on how to set up the needed tools and will assume that the security professional has already installed the tools or has the technical acumen to set them up. The following tools are used: diff, grep, Nmap, Splunk Enterprise, WMIC, and Nessus. The free version of Splunk can also be used to follow the examples; however, alerts cannot be automatically generated with the free version. Dashboards can be setup instead of alerts in the free version.

2.2. Folder Structure

Some folders will need to be created to store and compare new scans with prior baselines. The security professional should choose a base location for the scan results to be stored and create three folders: current_scans, prior_scans, and results. The read and write privileges on these folders should be properly restricted. Splunk will need read access to the results folder and can monitor the results folder directly if the folder is on the Splunk server. Configuring Splunk to monitor a folder can be found within Splunk under Settings, Data Inputs, and then Files & Directories. If the results folder is on a different server, the Universal Forwarder should monitor the folder and pass the results to Splunk. The configuration file for the Universal Forwarder can be found at the location below:

```
splunk_installation_path\etc\system\local\inputs.conf
```

To monitor the results folder, modify the inputs.conf file to include the following:

```
[monitor://C:\path_to_results_folder\results\*]
sourcetype = syslog
```

David Brown, mrdavebrown@gmail.com

```
disabled = false
```

2.3. Establishing a Trusted Baseline

Many of the requirements for the first five Critical Security Controls that will be introduced in later sections pertain to detecting changes within a network. In order to detect changes, a trusted baseline must first be created. Initial scan results when creating a trusted baseline should be carefully reviewed for accuracy to ensure that malicious activity has not already taken place. For purposes of this paper, baseline scans will be held in the prior_scans folder and replaced with each subsequent scan. To establish a proper baseline scan for each section of this paper, the security professional should run the given command, but instead of saving the output of the command into the current_scans folder, the output of the command should be saved to the prior_scans folder. Then, the output in the prior_scans folder should be manually reviewed for accuracy and to detect any malicious activity that may have already taken place.

Once a trusted baseline has been established and reviewed, the commands in this paper will compare current_scans with prior_scans, output any identified changes to the results folder, move the contents of the current_scans folder to the prior_scans folder, and generate alerts based on what has changed.

2.4. Group Policy Settings

The audit policies, "audit account logon events" and "audit logon events," must be configured to audit "Success, Failure" for all domain-joined resources that are going to be monitored. These audit policies are located in a group policy object under Computer Configuration, Policies, Windows Settings, Security Settings, Local Policies, Audit Policy.

3. Critical Security Control #1: Inventory of Authorized and Unauthorized Devices

By implementing the first Critical Security Control and setting up alerting based on the control, the security professional can detect new systems that are plugged into a network.

David Brown, mrdavebrown@gmail.com

Critical Security Control (CSC) #1.1: "Deploy an automated *asset inventory discovery* [emphasis added] tool and use it to build a preliminary inventory of systems connected to an organization's public and private network(s). Both active tools that scan through IPv4 or IPv6 network address ranges and passive tools that identify hosts based on analyzing their traffic should be employed" (Center for Internet Security, 2016, p. 6).

CSC #1.2: "If the organization is dynamically assigning addresses using DHCP, then deploy dynamic host configuration protocol *(DHCP) server logging* [emphasis added], and use this information to improve the asset inventory and *help detect unknown systems* [emphasis added]" (Center for Internet Security, 2016, p. 6).

3.1. Know When a New Device Is Plugged In

Critical Security Control #1.1 deals with discovering assets that exist on a network. Nmap is a great tool for asset discovery. One of the strengths of Nmap is the ability to script out network scans and to export those scans into a grep-able format. The following command can be added to a batch file or shell script and can be used for quick device discovery. Before the command is run, a file called "network_targets.txt" should be created. The file should contain the subnets or IP address ranges to be scanned. The file should contain one subnet, IP address, IP address range, or URL per line. To perform the scan, run the following command:

```
nmap -sn -iL network_targets.txt -oG
     current_scans/nmap_sweep.txt
```

The "sn" flag tells Nmap to perform a ping scan. The "iL" flag specifies a file that contains the individual hosts and/or subnets to target. The "oG" flag tells Nmap to output the results of the scan to a grep-able file (Nmap, 2016). Once the scan is complete, diff and grep can be used to parse out the new hosts, as seen below:

```
diff prior_scans/nmap_sweep.txt
     current_scans/nmap_sweep.txt | grep '> Host' >
     results/nmap_sweep.txt
```

Diff will compare the two files and pass the output to grep. The "> Host" grep search string tells grep to find new host additions to the file only. The grep results are then saved as nmap_sweep.txt to the results folder. Assuming the steps were followed in the Initial Setup section, Splunk will automatically ingest the nmap_sweep.txt file from the results folder. To find any new hosts in Splunk, click "Search," and type the following command into the search box:

```
source="*nmap_sweep.txt" "Status: Up"
```

To begin alerting on any new hosts, select "Save As," "Alert," and create a "CSC New Device" alert. A sample of the alert being setup can be seen below:

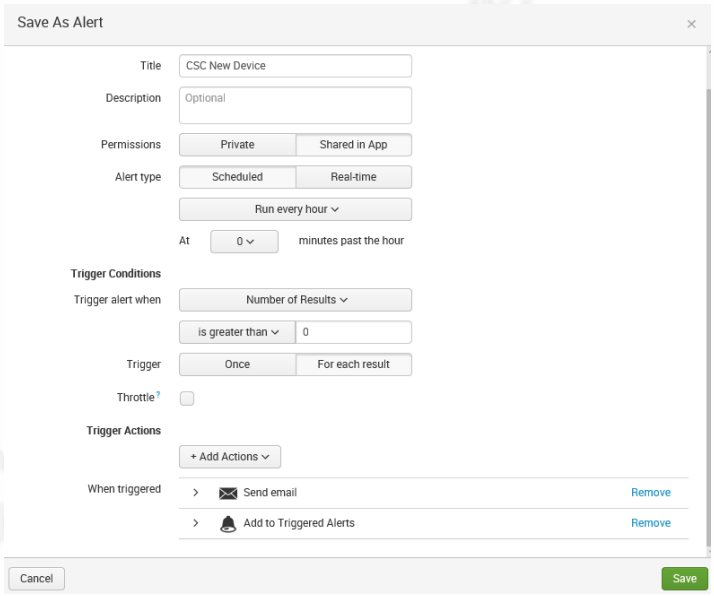

Splunk Alert. Screenshot Taken 08/15/2016.

After the results have been compared and passed to Splunk, the following command is run to prepare the system for future scans that will be run:

```
mv current_scans/nmap_sweep.txt prior_scans/
```

3.2. Mysterious DHCP Requests

Nmap scanning can successfully detect devices whenever the scan is run, but how often will the scan be run? Daily? Weekly? This is where DHCP requests come in. A

DHCP request can generate an alert the moment a new device is plugged into a network. Make sure that DHCP logging is enabled on each domain controller by opening the DHCP console, selecting the DHCP server then Action, Properties, General, and Enable DHCP audit logging. Then, to collect DHCP requests, install the Splunk Universal Forwarder on the Microsoft domain controllers and add the following configurations to the Universal Forwarder's inputs.conf file:

```
[hostname://C:\Windows\System32\dhcp]
disabled = 0
whitelist = DhcpSrvLog*
crcSalt = <SOURCE>
sourcetype = win_dhcp
```

As DHCP requests flow into Splunk, input the following search to find new DHCP assignments and create a "CSC DHCP Assignment" alert. Using the following search, alerts will be generated on initial DHCP assignments only and will not be generated on DHCP renewals so that known hosts that are renewing IP Addresses will not generate alerts:

```
sourcetype=win_dhcp Assign
```

4. Critical Security Control #2: Inventory of Authorized and Unauthorized Software

By implementing the second Critical Security Control and setting up alerting based on the control, the security professional can detect unauthorized software on systems.

CSC #2.2: "Deploy application whitelisting that allows systems to run software only if it is included on the whitelist and *prevents execution of all other software* [emphasis added] on the system. The whitelist may be very extensive (as is available from commercial whitelist vendors), so that users are not inconvenienced when using common software. Or, for some special-purpose systems (which require only a small

number of programs to achieve their needed business functionality), the whitelist may be quite narrow" (Center for Internet Security, 2016, p. 11).

CSC #2.3: "Deploy *software inventory tools* [emphasis added] throughout the organization covering each of the operating system types in use, including servers, workstations, and laptops. The software inventory system should track the version of the underlying operating system as well as the applications installed on it. The software inventory systems must be tied into the hardware asset inventory so all devices and associated software are tracked from a single location" (Center for Internet Security, 2016, p. 11).

4.1. Whitelist Monitoring

AppLocker is the built-in whitelisting solution for Microsoft Enterprise addition. The Splunk Universal Forwarder must be installed on every AppLocker protected workstation in order to track whitelisting blocks. Add the following to the Universal Forwarder's inputs.conf file to send AppLocker data to Splunk:

```
[WinEventLog://Microsoft-Windows-AppLocker/EXE and DLL]
disabled = 0
blacklist = 8001-8002
[WinEventLog://Microsoft-Windows-AppLocker/MSI and Script]
disabled = 0
blacklist = 8005
```

The AppLocker event logs are verbose. To avoid overwhelming the Splunk license with excessive data, successful program executions recorded in the event log are being blacklisted in the example above. Only application blocks or warnings are being passed to Splunk. Use the following search within Splunk to create alerts for AppLocker:

```
source=WinEventLog:Microsoft-Windows-AppLocker*
```

The "NOT" keyword can be used following the search if alerts are being generated on items that are not of interest, such as alerts on blocked Microsoft PowerShell scripts. Of course, it is better to disable unwanted PowerShell scripts in the registry or group policy instead of using AppLocker to block them, but either approach

can be effective. Once the search string is fine-tuned, set up a Splunk alert as "CSC Whitelist Block." The following is an example of a fine-tuned search:

```
source=WinEventLog:Microsoft-Windows-AppLocker* NOT
     "Message=*TS_UNUSEDDESKTOPICONS.PS1" NOT
     "Message=*TS_BROKENSHORTCUTS.PS1"
```

4.2. Software Inventory

WMIC is a helpful tool for monitoring software inventories. The "domain_targets.txt" file referenced in the command below should contain specific targets for the domain such as, "lab.example.com" and "mail.example.com," with one line per target. The target's file cannot contain IP address ranges or subnet expressions. In order to collect inventory, the commands must be run from a workstation joined to the domain and by a service account or user with administrative access to the target machines. The following commands can be used to identify running processes:

```
wmic /node:@domain_targets.txt
     /output:current_scans\tmp.txt process get
     name,executablepath /format:csv
type current_scans\tmp.txt > current_scans\process.txt
del current_scans\tmp.txt
```

The following command can be used to identify installed software:

```
wmic /node:@domain_targets.txt
     /output:current_scans\tmp.txt product get
     vendor,name,identifyingnumber /format:csv
type current_scans\tmp.txt > current_scans\product.txt
del current_scans\tmp.txt
```

This final command can be used to identify items in the startup folder:

```
wmic /node:@domain_targets.txt
     /output:current_scans\tmp.txt STARTUP GET Caption,
     Command, User /format:csv
type current_scans\tmp.txt > current_scans\startup.txt
del current_scans\tmp.txt
```

The following commands will compare the new scans with existing baselines and move the resulting comparison to the results folder for Splunk to integrate. These commands may be run from bash on Windows 10/Server 2016 or from a Linux machine:

```
diff prior_scans/process.txt current_scans/process.txt |
     grep '> ' > results/process.txt
mv current_scans/process.txt prior_scans/
diff prior_scans/product.txt current_scans/product.txt |
     grep '> ' > results/product.txt
mv current_scans/product.txt prior_scans/
diff prior_scans/startup.txt current_scans/startup.txt |
     grep '> ' > results/startup.txt
mv current_scans/startup.txt prior_scans/
```

Once software data begins flowing into Splunk, the following search commands can be used to setup alerts for "CSC New Process," "CSC New Product," and for "CSC New Startup Item":

```
source="*process.txt"
source="*product.txt"
source="*startup.txt"
```

4.3. Scheduled Tasks

Once a bad actor breaks into a network, the next step is to establish persistent access to that network. This is often accomplished through placing a program in the startup folder or by creating a scheduled task. Like whitelist monitoring, monitoring scheduled tasks is easy to accomplish with the Splunk Universal Forwarder. Add the following to the Universal Forwarder's inputs.conf to send event logs related to scheduled tasks:

```
[WinEventLog://Security]
disabled = 0
whitelist = 4698,4699,4700,4701,4702
```

Then, use the following Splunk search to find newly any scheduled tasks:

```
sourcetype="WinEventLog:Security" Message="A scheduled task
     was created.*"
```

As with whitelist monitoring, monitoring scheduled tasks can produce a lot of noise. Microsoft is constantly creating scheduled tasks in the background. To cut down on the noise, disable any unwanted Microsoft tasks in the registry or group policy, or fine-tune the Splunk search so the Microsoft tasks do not generate alerts. Once the Splunk search is refined, create the "CSC New Scheduled Task" alert. The following is an example of a refined search:

```
sourcetype="WinEventLog:Security" Message="A scheduled task
     was created.*" NOT
     Task_Name="\\Microsoft\\Windows\\Customer Experience
     Improvement Program\\Uploader"
```

5. Critical Security Control #3: Secure Configurations for Hardware and Software

By implementing the third Critical Security Control and setting up alerting based on the control, the security professional can identify new services being offered by a system.

CSC #3.6: "Implement and test an automated configuration monitoring system that verifies all remotely testable secure configuration elements, and alerts when unauthorized changes occur. This includes *detecting new listening ports* [emphasis added], new administrative users, changes to group and local policy objects (where applicable), and new services running on a system. Whenever possible use tools compliant with the Security Content Automation Protocol (SCAP) in order to streamline reporting and integration" (Center for Internet Security, 2016, p. 14).

5.1. Detecting New Listening Ports

It is critical to scan machines regularly to identify new listening ports, new administrative users, and new services. The first scan that is performed will serve as a baseline and will need to be carefully reviewed to make sure that the listening ports,

administrative users, and services are all authorized in a given environment. The following command can be used to detect new listening ports. This command may need to be run as sudo depending on the setup of nmap on the machine:

```
nmap -sS -iL network_targets.txt -oG
     current_scans/nmap_port_scan.txt
diff prior_scans/nmap_port_scan.txt
     current_scans/nmap_port_scan.txt | grep '> Host' >
     results/nmap_port_scan.txt
mv current_scans/nmap_port_scan.txt prior_scans/
```

The "sS" flag tells Nmap to do a TCP SYN instead of a normal SYN/ACK. This allows the scan to run faster (Nmap, 2016). Use the following Splunk search and create a "CSC Port Change" alert:

```
source="*nmap_port_scan.txt"
```

6. Critical Security Control #4: Continuous Vulnerability Assessment and Remediation

By implementing the fourth Critical Security Control and setting up alerting based on the control, the security professional can quickly identify new vulnerabilities within a network.

CSC 4.7: "*Compare the results from back-to-back vulnerability scans* [emphasis added] to verify that vulnerabilities were addressed, either by patching, implementing a compensating control, or documenting and accepting a reasonable business risk. Such acceptance of business risks for existing vulnerabilities should be periodically reviewed to determine if newer compensating controls or subsequent patches can address vulnerabilities that were previously accepted, or if conditions have changed, increasing the risk" (Center for Internet Security, 2016, p. 18).

6.1. Spotting New Vulnerabilities

Vulnerability scans are loaded with useful information about weaknesses that attackers may exploit within a network. However, since vulnerability scans can contain

thousands of data points on existing systems and vulnerabilities, spotting new vulnerabilities can be a very difficult task. It is important to have an automated way to spot new vulnerabilities as they creep into a network. Using Nessus, setup a "Basic Network Scan" of the target subnet(s). Schedule the scan to run each evening or on whatever schedule is appropriate. Each morning, after the scan is finished, log into the Nessus scanner and export the results of the evening's scan as a CSV to "current_scans\nessus.csv." While this section contains manual steps and cannot be automated as easily as other steps previously mentioned, it is essential for monitoring the overall health of a network. Once the nessus.csv file is placed in the current_scans folder, the following commands can be run to pass newly discovered vulnerabilities to Splunk:

```
diff prior_scans/nessus.csv current_scans/nessus.csv | grep
     -E '"Low"|"Medium"|"High"|"Critical"' >
     results/nessus.txt
mv current_scans/nessus.csv prior_scans/
```

Use the following search string to find the newly identified vulnerabilities in Splunk and setup a "CSC New Vulnerability" alert:

```
source="*nessus.txt"
```

7. Critical Security Control #5: Controlled Use of Administrative Privileges

By implementing the fifth Critical Security Control and setting up alerting based on the control, the security professional can track, in real-time, failed logins into administrative accounts and the creation of new administrative accounts.

CSC 5.4: "Configure systems to issue a log entry and *alert when an account is added to or removed* [emphasis added] from a domain administrators' group, or when a new local administrator account is added on a system." (Center for Internet Security, 2016, p. 22)

CSC 5.5: "Configure systems to issue a log entry and *alert on any unsuccessful login* [emphasis added] to an administrative account." (Center for Internet Security, 2016, p. 22)

7.1. New Administrative Accounts

Many times, the attacker will look to compromise existing account credentials. However, if the attacker takes the misstep of creating an administrative account, would the attacker be caught? If so, how fast? These are questions that should be considered when assessing detection capabilities. Use the following commands to identify newly created domain accounts. These commands must be run from a workstation joined to the domain:

```
wmic /output:current_scans\tmp.txt useraccount list full
     /format:csv
type current_scans\tmp.txt > current_scans\useraccounts.txt
del current_scans\tmp.txt
```

The following commands can be used to identify changes in group membership and can be used to identify if a user is added to an administrative group:

```
wmic /output:current_scans\tmp.txt path win32_groupuser
type current_scans\tmp.txt > current_scans\groups.txt
del current_scans\tmp.txt
```

Identifying local administrator accounts requires more effort on the part of the security professional. Setup a "workstation_targets.txt" file. It can be similar to the "domain_targets.txt" file discussed in prior sections, but only workstation names should be included in the file, such as "lab" and "mail," and not full domains. Then add the following commands to a batch file and run with a service or user account with administrative privileges on the target machines:

```
for /f %%x in (workstation_targets.txt) do wmic /Node:%%x
     path win32_groupuser where
     (groupcomponent="win32_group.name=\"Administrators\",d
     omain=\"%%x\"") >> current_scans\tmp.txt
type current_scans\tmp.txt > current_scans\localadmin.txt
```

David Brown, mrdavebrown@gmail.com

```
del current_scans\tmp.txt
```

Identifying changes to RDP users can also be useful for identifying potentially malicious users being set up for persistence through RDP by running the following commands:

```
for /f %%x in (workstation_targets.txt) do wmic /Node:%%x
     path win32_groupuser where
     (groupcomponent="win32_group.name=\"Remote Desktop
     Users\",domain=\"%%x\"") >> current_scans\tmp.txt
type current_scans\tmp.txt > current_scans\rdp.txt
del current_scans\tmp.txt
```

As mentioned in the Software Inventory section, the following commands will compare the files and move them to the results folder for Splunk to integrate. Run the following commands from bash on Windows 10/Server 2016 or from a Linux machine:

```
diff prior_scans/useraccounts.txt
     current_scans/useraccounts.txt | grep '> ' >
     results/useraccounts.txt
mv current_scans/useraccounts.txt prior_scans/
diff prior_scans/groups.txt current_scans/groups.txt | grep
     '> ' > results/groups.txt
mv current_scans/groups.txt prior_scans/
diff prior_scans/localadmin.txt
     current_scans/localadmin.txt | grep '> win32' >
     results/localadmin.txt
mv current_scans/localadmin.txt prior_scans/
diff prior_scans/rdp.txt current_scans/rdp.txt | grep '>
     win32' > results/rdp.txt
mv current_scans/rdp.txt prior_scans/
```

Once software data begins flowing into Splunk, the following search commands can be used to setup alerts for “CSC New Domain User Account,” “CSC New Domain Group Membership,” “CSC New Local Admin,” and for “CSC New Local RDP User”:

```
source="*useraccounts.txt"
source="*groups.txt"
source="*localadmin.txt"
```

David Brown, mrdavebrown@gmail.com

```
source="*rdp.txt"
```

7.2. Failed Logins

Administrative accounts, by default, are often set up to allow unlimited password attempts. This is by design so that administrative accounts do not get locked out. However, if countermeasures are not in place, this design can also allow attackers to brute force administrative accounts. It is essential to know if an attacker is trying to brute force an administrative account. To track failed login attempts, modify the Universal Forwarder's inputs.conf on the workstations to be monitored by employing:

```
[WinEventLog://Security]
disabled = 0
whitelist = 4625,4771
```

An additional benefit to tracking failed login attempts is that it will also generate alerts during vulnerability scans since part of vulnerability scanning is attempting to log in with default administrative credentials. The following Splunk search will show failed logins for all users and attempted logins for accounts that do not exist:

```
eventtype="wineventlog-security" (EventCode=4625 OR
     EventCode=4771)
```

To suppress alerts to failed logins on administrator accounts only, the Splunk search can be modified to explicitly look for specified accounts. Once the search is fine-tuned, create an alert called "CSC Failed Login." The following is an example of a fine-tuned search:

```
eventtype="wineventlog-security" (EventCode=4625 OR
     EventCode=4771) (administrator OR admin_jim OR
     admin_jane)
```

8. Conclusion

The information security professional can no longer afford to solely focus on prevention. A holistic approach to security – from prevention, to detection, to remediation – must be adopted. The Critical Security Controls and a data aggregation tool like Splunk can greatly assist the information security professional toward this holistic approach to cyber security.

It is also critical for the information security professional to begin cutting through the vast amount of noise that is generated on a given network. Bad actors are skilled at hiding in that noise, and there is a lot of noise. Setting up alerts based on the Critical Security Controls can inform the information security professional when a bad actor is attempting to bypass a control or when a potentially insecure and unauthorized device enters a network.

Cutting through the noise by focusing on the signals produced by the Critical Security Controls can be the difference between a massive and prolonged breach that makes newspaper headlines and a breach that is discovered and resolved in a single day. But the information security professional should always remember that detection without swift remediation is a losing strategy. Once Splunk has been fine-tuned to alert on the Critical Security Controls, the security professional should focus on that incident response plan and be prepared to remediate.

David Brown, mrdavebrown@gmail.com

9. References

Australian Cyber Security Center. (2015). *ACSC 2015 Threat Report*. Retrieved from https://www.acsc.gov.au/publications/ACSC_Threat_Report_2015.pdf

California Department of Justice. (2016). *California Data Breach Report*. Retrieved from https://oag.ca.gov/breachreport2016

Center for Internet Security. (2016). *The CIS Critical Security Controls for Effective Cyber Defense* (6.1). Retrieved from https://www.cisecurity.org/critical-controls/Library.cfm

National Security Agency. (2013). *Spotting the Adversary with Windows Event Log Monitoring* (2). Retrieved from https://www.iad.gov/iad/library/reports/spotting-the-adversary-with-windows-event-log-monitoring.cfm

National Security Agency. (2015). *NSA Methodology for Adversary Obstruction*. Retrieved from https://www.iad.gov/iad/library/reports/nsa-methodology-for-adversary-obstruction.cfm

Nmap. (2016). Nmap Reference Guide. Retrieved from https://nmap.org/book/man.html

Mandiant Consulting. (2016). *M-Trends 2016*. Retrieved from https://www.fireeye.com/current-threats/annual-threat-report/mtrends.html

Verizon. (2016). *Verizon 2016 Data Breach Investigation Report*. Retrieved from http://www.verizonenterprise.com/verizon-insights-lab/dbir/2016/