Global Information Assurance Certification Paper

Continuous Monitoring: Build A World Class Monitoring System for Enterprise, Small Office, or Home

GIAC (GMON) Gold Certification

Author: Austin Taylor, sans@austintaylor.io
Advisor: Hamed Khiabani, Ph.D.
Accepted: December 4th, 2016

Abstract

For organizations who wish to prevent data breaches, incident prevention is ideal, but detection of an attempted or successful breach is a must. This paper outlines guidance for network visibility, threat intelligence implementation and methods to reduce analyst alert fatigue. Additionally, this document includes a workflow for Security Operations Centers (SOC) to efficiently process events of interest thereby increasing the likelihood of detecting a breach. Methods include Intrusion Detection System (IDS) setup with tips on efficient data collection, sensor placement, identification of critical infrastructure along with network and metric visualization. These recommendations are useful for enterprises, small homes, or offices who wish to implement threat intelligence and network analysis.

1. Introduction

Imagine a large organization with millions of devices connected to its network and only a handful of security analysts to defend the network environment. This disproportionate work environment occurs more commonly in the information security field. Inadequately equipped settings prevent organizations from properly identifying and mitigating a threat. Analysts suffer from "alert fatigue" due to an overwhelming number of alerts to process, lack of event context, and a high false positive rate when analyzing threats. Inefficiently managed Security Operations Centers (SOC) exacerbate the problem by not aggregating and correlating data sources to help analysts determine if an incident occurred.

Many products offer visibility into only one or two layers of the seven-layer Open Systems Interconnection (OSI) model, which significantly limits the ability to detect, respond, and mitigate threats. This paper offers visibility recommendations into all seven layers of the OSI model, which include: methods to correlate data sources, storage, threat intelligence, Intrusion Detection System (IDS) sensor placement, network and metric visualization to aid analysts in detecting breaches, corporate policy violations, and misconfigurations.

2. Modern Cybersecurity Landscape

2.1. Product Visibility

Products on the market offer a broad range of insight into an enterprise network environment (Momentum Partners, 2015). It is important to focus on network solutions that are most applicable to an organization. For example, an organization that does not support Mobile devices may not need Mobile Security. Categories of products across the cybersecurity industry are captured in a Cyberscape diagram and are as follows:

- Information Security
- Endpoint Security
- Application Security
- Messaging Security
- Web Security
- Internet of Things Security (IoT)
- Security Operations and Incident Response

Austin Taylor, sans@austintaylor.io

- Threat Intelligence
- Mobile Security
- Data Security
- Transaction Security
- Risk and Compliance
- Specialize Threat Analysis and Protection
- Identity and Access Management
- Cloud Security

Implementing controls in each of these categories could be arduous, costly, and unnecessary as they provide varying levels of network visibility. When determining which controls to implement, organizations should evaluate which categories are most relevant to company operations and have the lowest cost with the most impact. For example, if an organization sells products online, web security would likely be a top priority. Within the web security category, an organization with limited funding might prioritize a web proxy over full Secure Sockets Layer (SSL) decryption. Web proxies provide a medium level of visibility and logs each visit to a website whereas SSL decryption would cost more due to storage and processing requirements and only provide visibility to a limited portion of the organization's web traffic.

Companies should map their key data, inventory of services, and hardware to each category. When companies understand what to protect, they can determine the required visibility level. If different products provide overlapping visibility it requires companies to use additional resources to store, correlate and process the information. The mapping process is important to provide the appropriate amount of visibility to devices critical to business operations while considering cost.

2.2. The Open Systems Interconnection (OSI) Model

The OSI Model characterizes the way network traffic communicates; each layer services the layer above and is served by the layer below. Understanding how applications follow the OSI Model's encapsulation process and the digital footprint it creates is important when designing an efficient monitoring system to identify the types of logs or network traffic generated by that application. This information will aid in sensor placement and collect the right data for the organization's environment. Thus, incident investigation processing time reduces significantly due to adequate network

visibility.

The visibility chart maps a few standard network devices, the Protocol Data Unit (PDU), and visibility sources to collect at each layer. Data taken from the visibility column could be ingested and correlated into a central management source for enriched analysis.

Visibility Chart

Product or Device	**OSI Layer Name**	**PDU**	**Protocols & Layer Description**	**Visibility Sources**
Proxy Firewalls	**Application**	Data	HTTP DNS HTTPS SMTP	Logs, Packet Capture
NexGen Firewalls	**Presentation**	Data	TLS/SSL, SMB	Logs, Packet Capture
Gateways *Proxy Servers* *Application Switches* *Content Filtering Firewalls*	**Session**	Data	NetBIOS, SOCKS,	Logs, Netflow Data, Packet Capture
Gateways *Proxy Servers* *Application Switches* *Content Filtering Firewalls*	**Transport**	Segment	TCP, UDP	Logs, Netflow Data, Packet Capture
Routers	**Network**	Packet	IPv4, IPv6, OSPF, ICMP	Netflow Data, Packet Capture
Switches / Bridges	**Data Link**	Frame	ARP, PPTP, L2TP, SNMP, WiFi	Logs, Netflow Data, Packet Capture
Hubs / Repeaters	**Physical**	Bits	Fiber Optic, Twisted Pair, & Connectors	Netflow Data, Packet Capture

2.3. Visibility Sources

Logs provide high-level visibility into events and are usually generated by endpoint devices. Correlation between multiple log sources is necessary to reconstruct a reliable event timeline and is a great starting point for incident responders and security analysts. Logs taken from key devices throughout the network will improve visibility and reduce the effort and time to reconstruct an event timeline. Figure 1 shows the three critical visibility sources to identify events of interest.

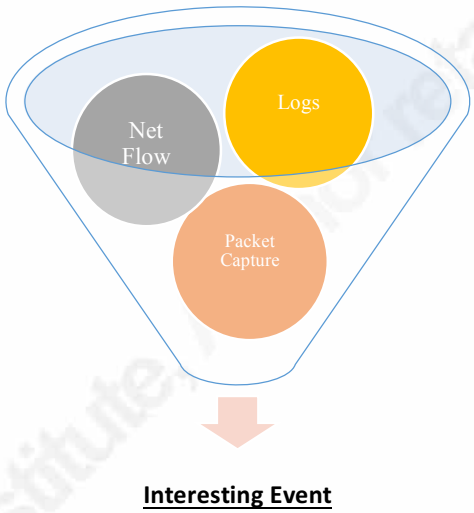

Figure 1 Critical Sources Funnel – Logs, Netflow, Packet Capture

Netflow captures packet metadata such as:

- Source IP & Port
- Destination IP & Port
- Bandwidth Consumption

Storing Netflow allows an analyst to view historical and real-time information regarding IP addresses communicating with their network. This will help an analyst establish a baseline of whether it is normal for IP A to talk to IP B based on the amount of connections from the enterprise to the destination. For example, an organization that

uses Google's DNS server might have a high degree[1] of source connections communicating to a valid IP address, whereas a small degree of connections would communicate to a potentially suspicious IP address.

Deep Packet Inspection (DPI) complements Netflow by providing artifacts such as Applications, Website Referrers, Protocols, Usernames, Cleartext Credentials, Hostnames, HTTP Headers, and File Hashes. Most of the time, by correlating logs, Netflow can help create a narrative for an incident by providing application level metadata; in most cases, an analyst will need packet capture data to validate his or her analysis.

Below is a DNS request from the perspective of each visibility source.

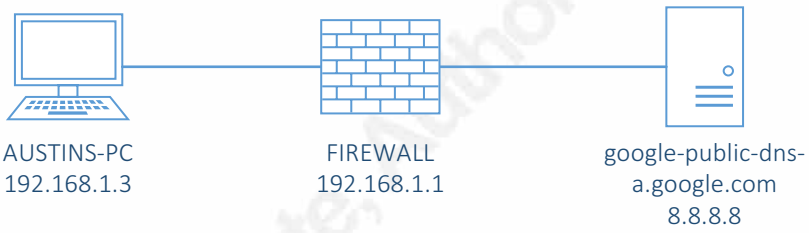

Log: 1475641856: A DNS request was made
Netflow: UDP: 192.168.1.3 → 8.8.8.8
Packet Capture: AUSTINS-PC made a DNS Request to www.google.com

Log information shows an event occurred; Netflow provides information on nodes involved with the event, and packet capture provides both log and Netflow information along with additional attributes such as hostname, website and much more. With appropriate visibility sources being collected and correlated it immediately provides answers to questions an analyst would otherwise have to seek, such as associating a hostname to an IP address. Thus, analysts can conclude their incident investigation in a much shorter time.

[1] https://en.wikipedia.org/wiki/Degree_distribution

Figure 2 illustrates types of information each source provides and serves as a reference when determining the level of visibility required.

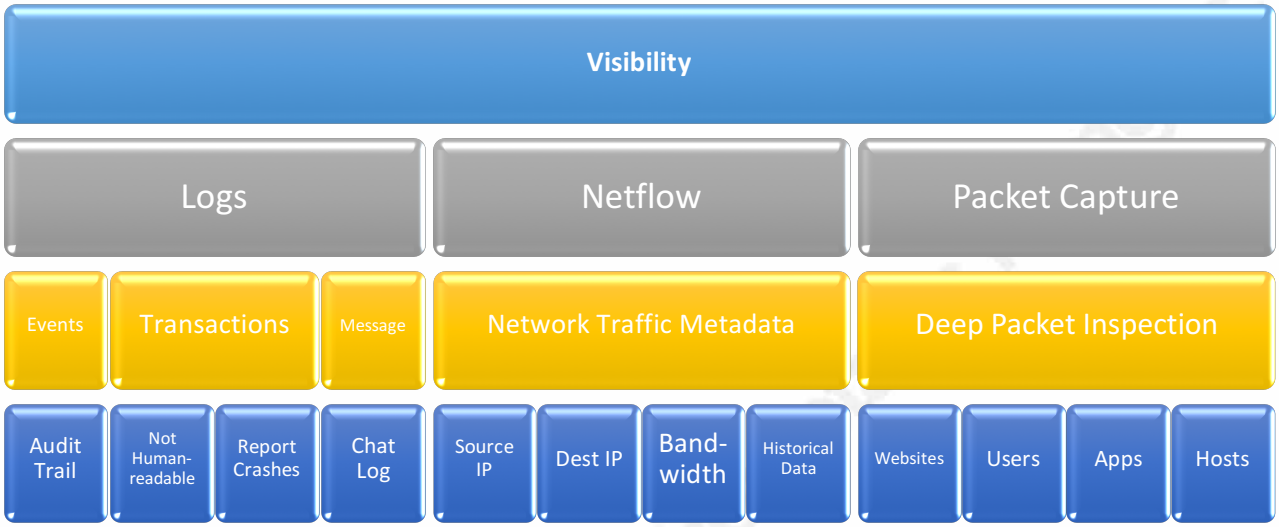

Figure 2 Relationship of visibility sources

For example, an organization that needs to identify which websites users are visiting can get that event information from logs. If the company wants to know how much network traffic or bandwidth is going to each site, Netflow is required. Additionally, packet capture provides details about which application is responsible for generating network traffic to each website. Organizations should identify which visibility sources are required to effectively monitor their network while considering the cost of each source.

2.4. Cost of visibility

While packet capture data provides additional benefits, for most network owners it is not feasible to capture and store data for extended periods of time due to cost and a lack of resources. Figure 3 illustrates the relationship between cost and visibility. The amount of information collected is directly proportional to the resources required for processing. Storage capacity makes log aggregation and Netflow a low barrier to entry for visibility. Full packet capture for large organizations requires more consideration to cost and resources.

Logs should be carefully collected to help identify threats involving critical processes. The National Security Agency's (NSA) article "Spotting the Adversary with Windows Event Log monitoring"[2] is an excellent resource to determine which Windows-based event types an organization should log. Also, system collection should be aligned to identify deviations from corporate policy. For instance, if an organization prohibits external media, section 4.13 in the article references how to detect external media in real-time.

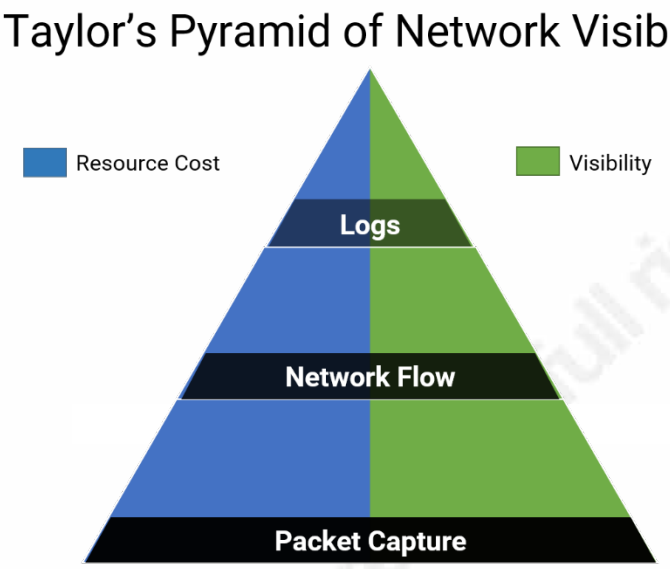

Figure 3

4.13 External Media Detection

Detection of USB device (e.g., mass storage devices) usage is important in some environments, such as air gapped networks. This section attempts to take the proactive avenue to detect USB insertion at real-time. Event ID 43 only appears under certain circumstances. The following events and event logs are only available in Windows 8 and above. Additional information can be found in the footnotes.

	ID	Level	Event Log	Event Source
New Device Information	43[54]	Informational	Microsoft-Windows-USB-USBHUB3-Analytic[55][56]	Microsoft-Windows-USB-USBHUB3
New Mass Storage Installation	400[57]	Informational	Microsoft-Windows-Kernel-PnP/Device Configuration	Microsoft-Windows-Kernel-PnP
New Mass Storage Installation	410[57]	Informational	Microsoft-Windows-Kernel-PnP/Device Configuration	Microsoft-Windows-Kernel-PnP

Table 14: External Media Detection Events

i

[2] https://www.iad.gov/iad/library/reports/spotting-the-adversary-with-windows-event-log-monitoring.cfm

2.5. Gain visibility

2.5.1 Logs

One way to gain network visibility is through log forwarding, which sends all results to a central location where they can be correlated. An essential role of log forwarding is to associate the data with NetFlow or packet capture and determine which process is responsible for generating that traffic. Most devices offer some form of log forwarding. Each operating system has a workflow for application log forwarding. For example, if an environment has Windows, the NSA has created a compilation of PowerShell scripts to help one create log subscribers with a minimum recommended set of Windows events to collect which provide value and insight into a system (Agency, 2016)

2.5.2 Netflow

The second form of network visibility is NetFlow capture or forwarding. Netflow provides metadata on packets traversing the network. Most Cisco routers and switches allow for Netflow forwarding. It can also be obtained passively from IDS devices such as Suricata or Snort. Solarwinds has produced a guide (Winds, 2016) for configuring various Cisco devices to forward NetFlow. The guide includes methods for creating flow records, configuring flow exporters, creating a flow monitor, and applying the flow monitor to an interface.

2.5.3 Packet Capture

The third and most complete form of network visibility is packet capture. If an analyst is not collecting NetFlow from routers or firewalls, it is recommended to install an inline Test Access Port (TAP). Network TAPs cost more than configuring a router or a switch with a span port. Figure 5 illustrates a basic sensor topology with a wireless access point.

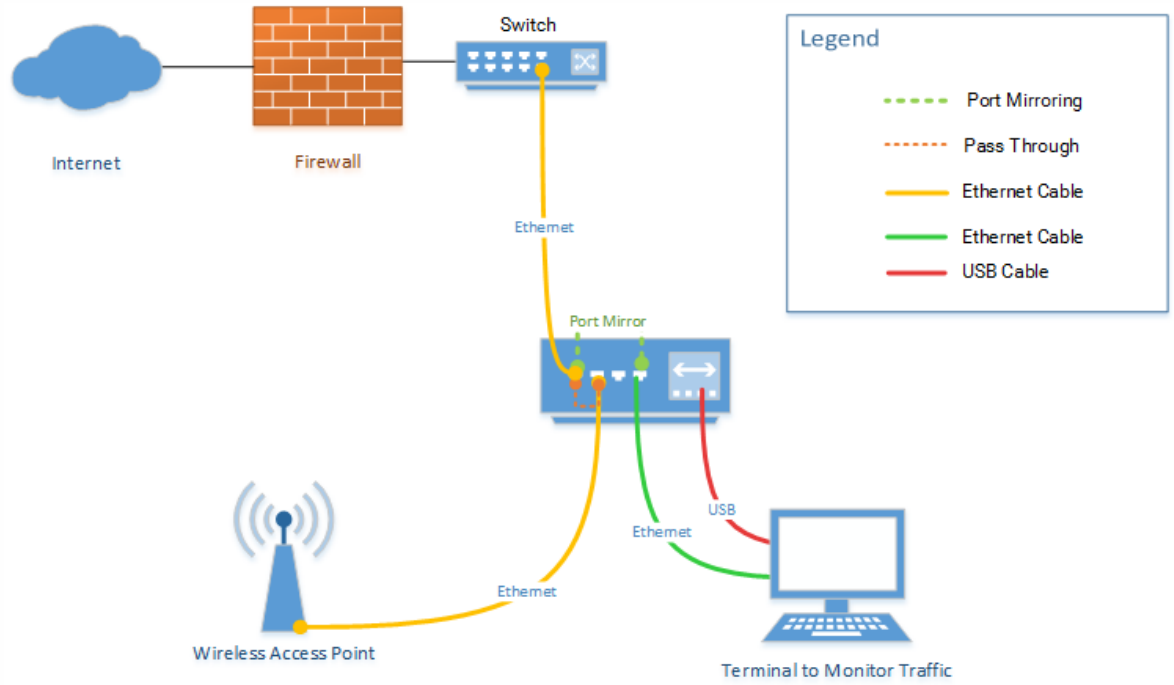

Figure 5 Sensor Topology Setup

Portable network visibility is helpful when performing incident response. It is ideal to have a network TAP or a sensor that can be powered by USB. If needed, it can be relocated to gain temporary visibility in different segments of your environment. In figure 5, each device plays a role:

- **Terminal or Managed Server**: Used to monitor traffic and power the TAP. Connected to port 4 and traffic is mirrored from port 1.
- **Wireless Access point** (or home router): The traffic you want to monitor.
- **Network TAP** – Port 1 and 2 are pass-through, and port 5 is a mirror.

An organization will want to monitor all traffic on a wireless access point, so the connection is mirrored going to the switch. This configuration would suit a home network, small office, or enterprise network segments.

2.6 Central Dashboard and SIEMS

After establishing visibility sources and sensor placement, the organization should consolidate its data. A central dashboard allows an analyst to search, analyze, correlate, and visualize data. In general, dashboard panels populate with aggregated values to help organizations identify events such as most commonly visited websites, most used applications, ports, or protocols. This information can be used to distinguish anomalous network behavior by looking at the tail end of the statistics; a method also known as “long tail analysis."

An example of a central dashboard is Security Information and Event Management Systems (SIEM), which provides a repository for data aggregation and correlation. It extracts data from multiple sources (web servers, IDS logs, proxy logs, databases) and correlates it in a central location. An effective SIEM will have ways to ingest data, allow for field extraction, and allow analysts to tailor the dashboard to their environment.

Each data source should represent different segments within the network, and be used as a data point, but may not accurately capture what is happening. For example, if there is an IP address that needs to be investigated, the next steps might be to:

1. Determine IP address hostname
2. Identify what process on the host is responsible for communicating with the IP in question.

Without SIEM correlation, answering these issues would be a manual process. If the SIEM is ingesting IP addresses from Netflow and passive DNS or DHCP logs, IP addresses should be enriched with the device hostname for that period. Figure 6 illustrates potential sources for a SIEM; an organization should aim to ingest sources from each section.

Figure 6 - SIEM Correlation

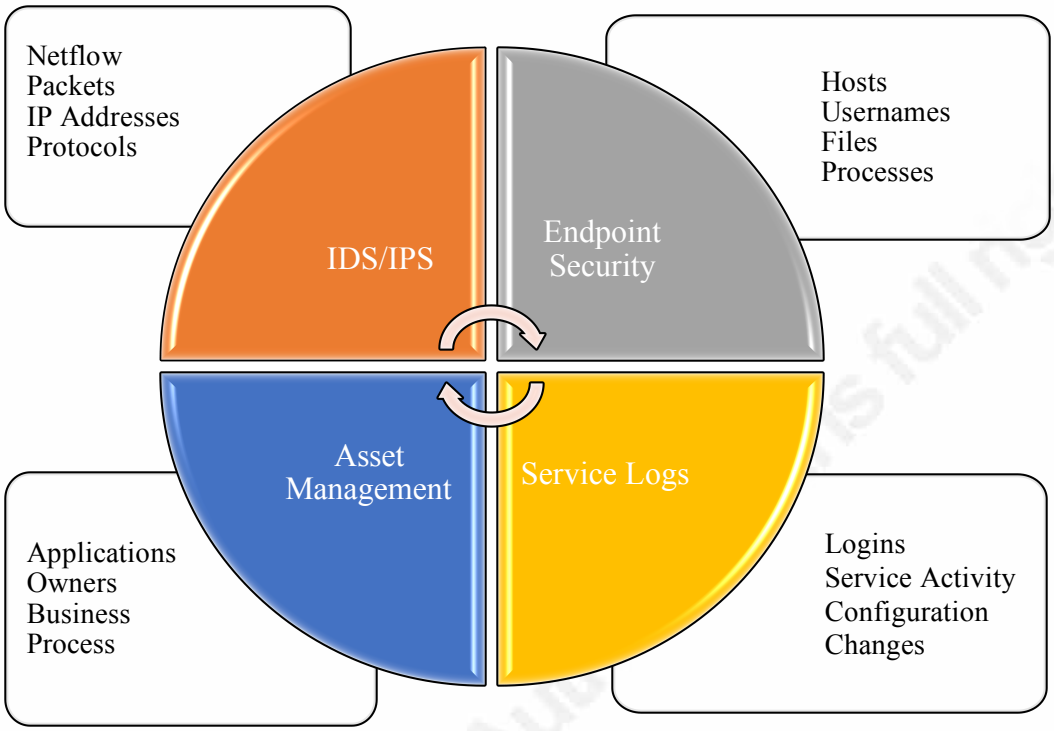

Many companies do not have the storage capacity to save every log source. If this is the case, critical infrastructure should be the priority. To identify critical infrastructure, ask the following question: *Which pieces of infrastructure are core to our company's objective*? Identifying critical infrastructure will help companies plan which logs to prioritize for collection.

Data forwarded to a SIEM allows an analyst to identify events of interest. A SIEM can provide answers to identify anomalous behavior such as:

- Which User-Agent strings are the least common in the environment?
- What Administrator accounts have the highest number of failed login attempts and with which devices does the account interact with ordinarily?

A correctly configured SIEM will provide an analyst with enriched data allowing for rapid analysis and well-informed decisions. Figure 7 provides recommended log sources, separated by category, for a SIEM[3].

[3] https://www.alienvault.com/blogs/security-essentials/what-kind-of-logs-for-effective-siem-implementation

Figure 7 - Recommended Log Sources by Category

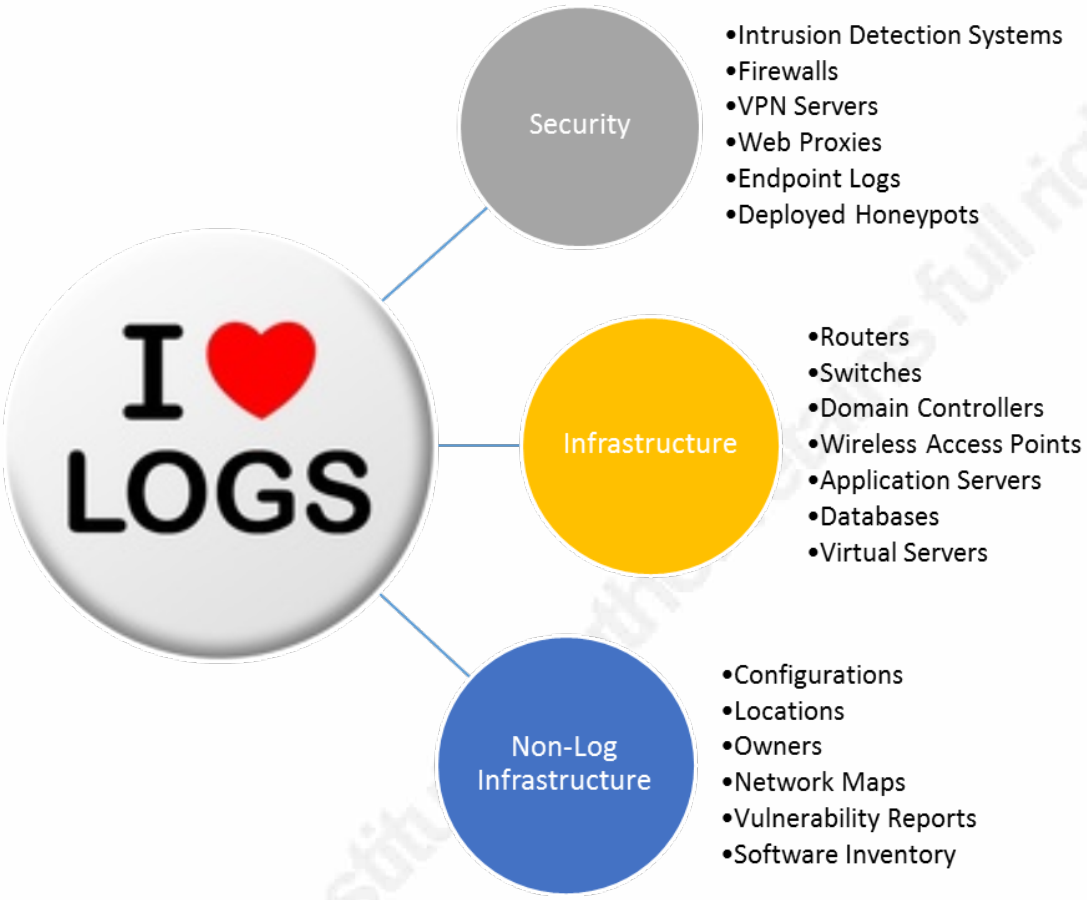

Consider these two examples:

1. 1475647717 User ATaylor AUTH SUCCESS 10.20.0.52 → 10.0.0.8
2. Austin Taylor remotely authenticated to a domain controller during off-duty hours from a workstation he does not regularly use on Wednesday, October 1st, 2016.

Example 1 is a log with no enrichment. Example 2 is a log that has been enriched by various sources and provides immediate context to an analyst, which saves the organization critical time and money.

A SIEM can be very effective when populated with the right data sources. However, it can also be a very arduous process to configure correctly. Companies can leverage pre-packaged, pre-configured solutions to reduce setup time.

Austin Taylor, sans@austintaylor.io

2.6.1 Suricata, Elasticsearch, Logstash, Kibana, Scirius (SELKS)

SELKS can complement a SIEM or serve as an alternative. Security analysts often find themselves playing system administrator. Projects like SELKS are helping lighten the workload by creating pre-bundled virtual machines that require minimal configuration. The recently released SELKS[4] [ii]is a collection of pre-configured tools allowing for easy deployment into an organization's network ecosystem. Figure 8 summarizes the SELKS workflow for each subsystem.

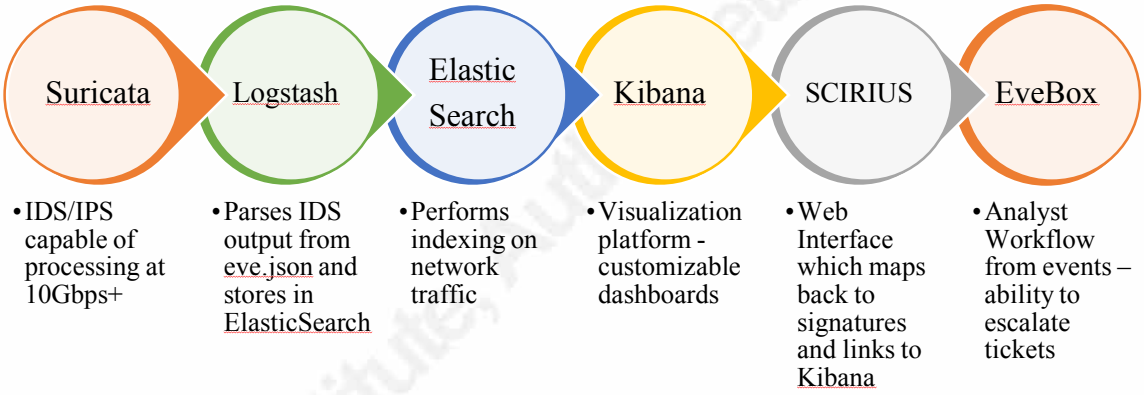

Figure 8 - SELKS workflow

SELKS is an installable Network Security Management (NSM) created primarily with Free Open Source Software (FOSS). A fresh install requires minimal tuning. SELKS allows organizations to monitor their network by leveraging Suricata, Logstash, Elasticsearch, Kibana, SCIRIUS, and EveBox. It is essential to understand the role each subsystem holds. The following is a summary of each subsystem:

Suricata - a high-performance Network IDS, IPS, and Network Security Monitoring engine. Capable of processing packets at 10+ Gbps. Suricata shares market space with Snort and Bro.

Logstash - Allows for central data processing; connects to a variety of sources and can stream to central analytics system. Logstash can take unstructured log data and transform it into structured data placing it into Elasticsearch.

[4] https://www.stamus-networks.com/open-source

Elasticsearch - A scalable framework which can perform very fast searches to support data discovery. Elasticsearch will take data in from Logstash and feed it to Kibana.

Kibana - Reads information from Elasticsearch and allows users to create rich dashboards to explore data.

Scirius - A web interface for Suricata's rule set management. Allows users to take in various sources, such as the Emerging Threat rule set or create rules unique to an organization.

EveBox - A web-based event viewer, which allows a user to generate reports on alerts, DNS, and Netflow. Also, EveBox includes a native workflow that allows an analyst to highlight events of interest by clicking the "Escalate" button.

The Scirius web interface has a suite of preconfigured dashboards geared toward common attack vector protocols such as HTTP, DNS, and TLS. Each panel presents summary statistics and answers questions an analyst may have of their environments, such as most or least popular file transactions. Scirius web interface also provides dashboards such as top alert signatures, TLS versions, and DNS resource records. Each panel inside the dashboard is also configurable and can be tailored to monitor threats specific to an organization.

In this scenario, an external source scanned my home network. Figure 9a is the beginning of the data normalization process SELKS uses to process Suricata's output.

Figure 9a - Output from eve.json

This output is in JavaScript Object Notation (JSON) format, but not ready for a user interface to query. Logstash will identify the structure, clean up the output, and forward the data to Elasticsearch for indexing. Ingested data is ready to be queried by user interfaces such as Kibana, SCIRIUS, and EveBox. After the data is indexed users can query Elasticsearch for the same information as illustrated in Figure 9b.

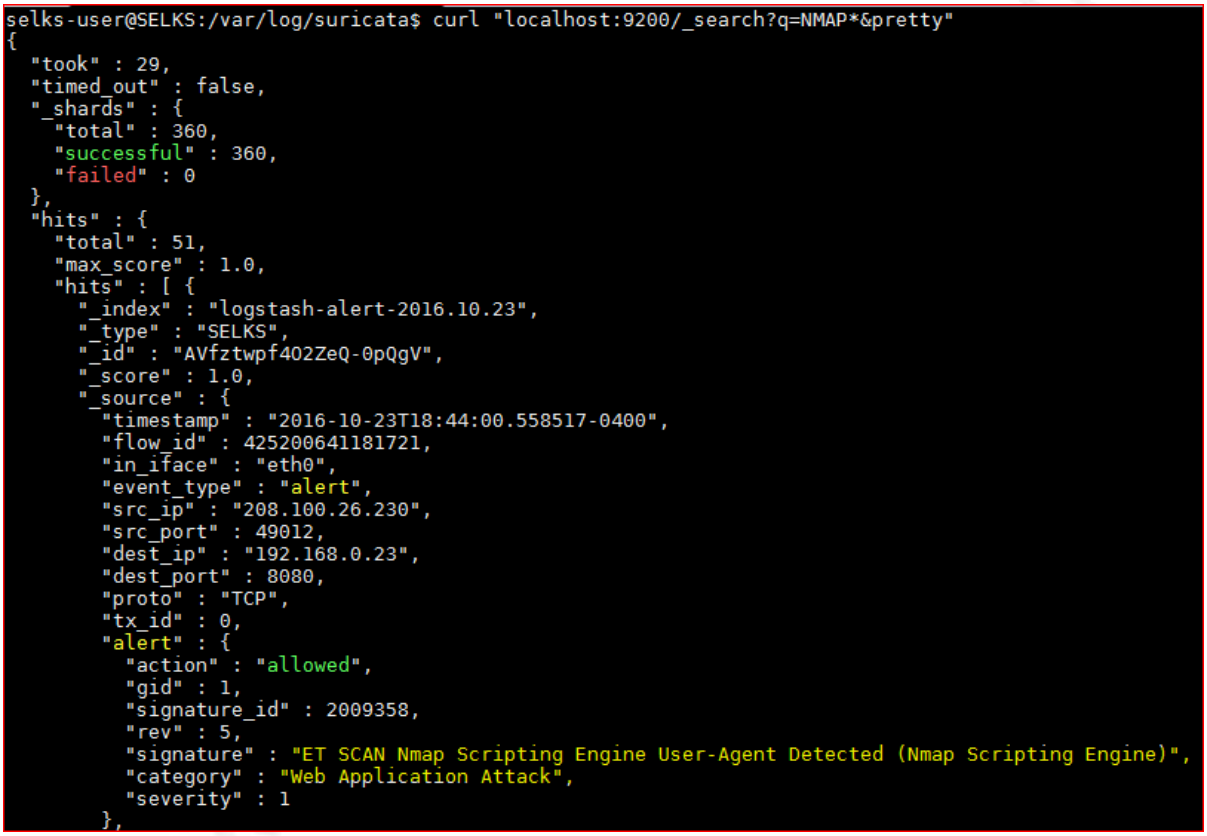

Figure 9b- Raw data from Suricata; normalized by Logstash; forwarded to Elasticsearch and consumed by Kibana/EveBox

The SELKS stack dashboard provides a high-level summary of alert activity. From the dashboard, a user can pivot to rules, Kibana dashboards, or the Event Viewer (EveBox). Figure 10 shows the primary interface that greets users after logging into SELKS.

Figure 10 Scirius Dashboard

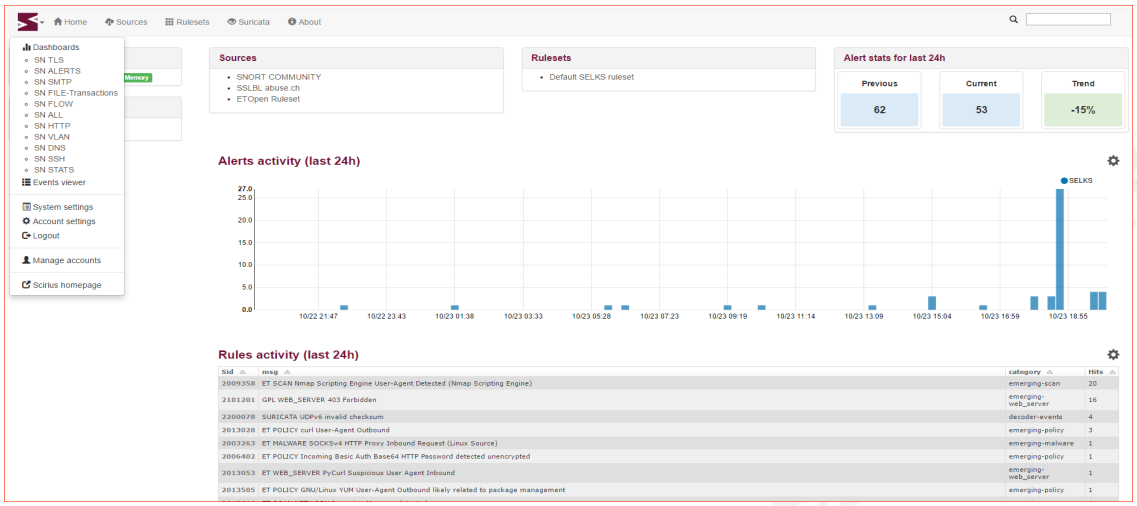

Continuing with the NMAP alert, a user can pivot to the alert dashboard. SELKS automatically enriches the IP address with coordinates (if applicable), subgroups, and plots them on a map. Figure 11 shows the NMAP alert plotted out geographically: dots represent alerts and a larger dot represents higher alert counts. This visual representation draws attention to anomalous network connections.

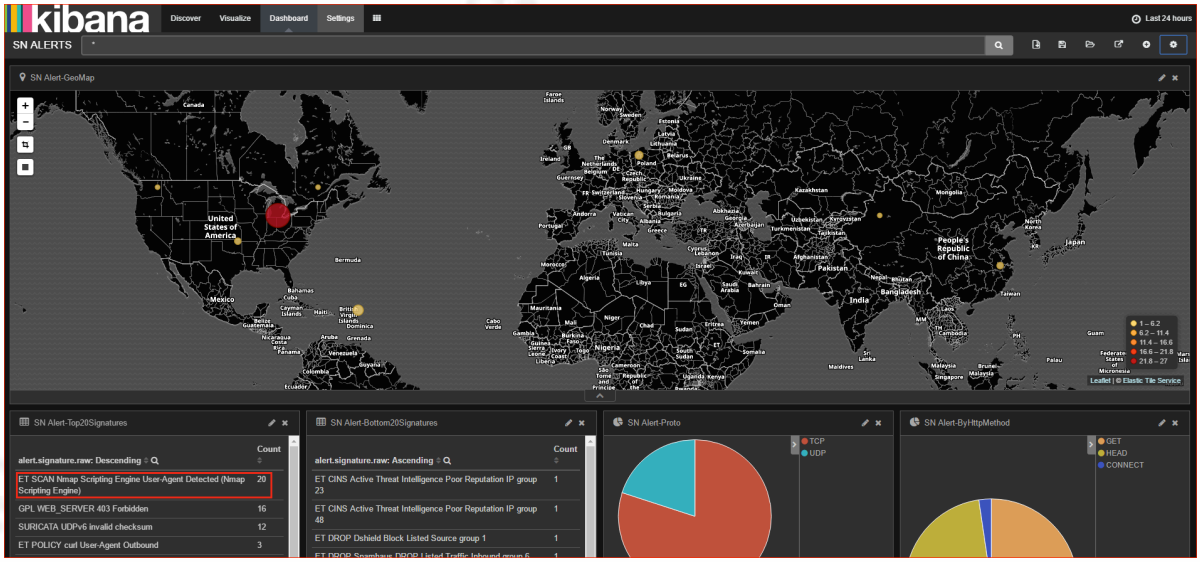

Figure 11 Alert Dashboard

Analysts have the option to filter results by selecting the alert; subsequently repopulating the entire dashboard. Figure 12 shows the option for users to pivot to EveBox.

Austin Taylor, sans@austintaylor.io

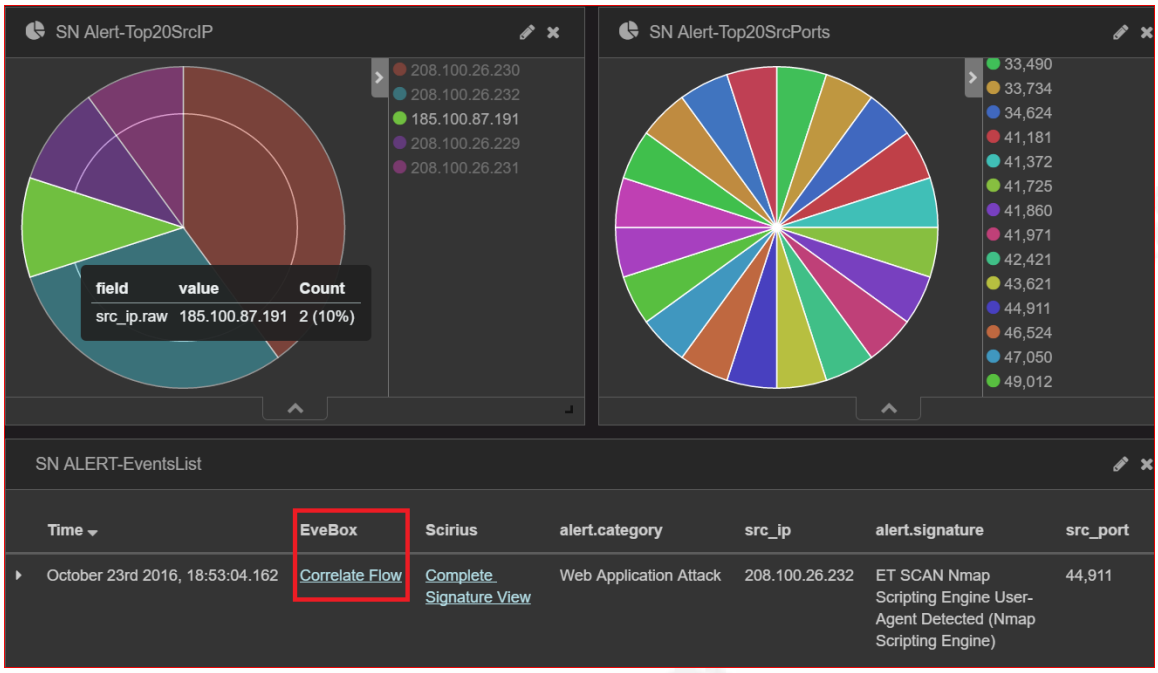

Figure 12 – Pivot to Event Box

After identifying the alert, an analyst can correlate the flow through EveBox. The event viewer will enrich the event with the triggered signature and relevant application protocol data. EveBox adds appropriate context to help quickly process the event. Figure 13 is a screenshot of EveBox populated with the NMAP scan.

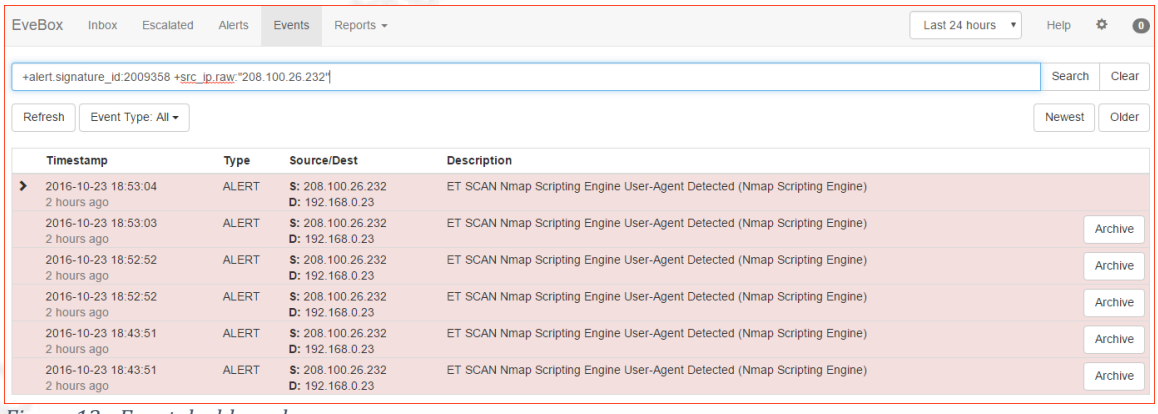

Figure 13 - Event dashboard

Finally, an analyst can archive or internally escalate the alert for additional assistance or queue for further processing. The ability to highlight events creates a workflow for SOCs with varying level of analysis. Junior analysts can escalate tickets to senior analysts for further review to validate their findings. Figure 14 shows the dashboard that gives users the ability to archive or escalate the alert.

Austin Taylor, sans@austintaylor.io

Figure 14 – Even Expansion with enriched data

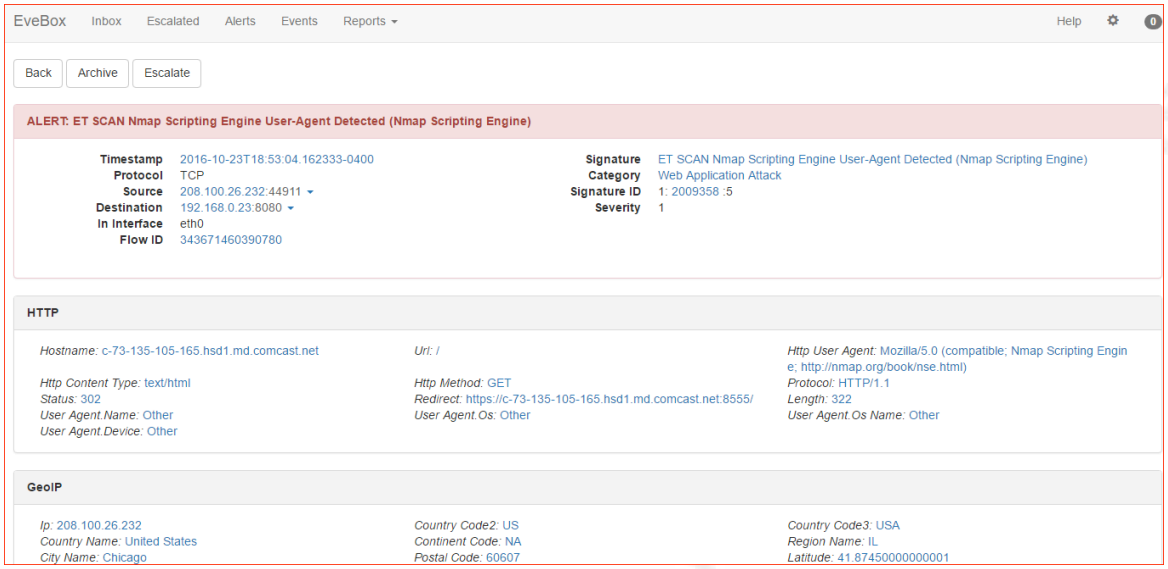

2.6.2 Tuning Recommendations

After the SELKS environment is setup, there are significant tuning considerations to implement for the sensor to perform reliably. SELKS has provided specific tuning recommendations at https://github.com/StamusNetworks/SELKS/wiki/Tuning-SELKS. Visit the tuning site and follow the recommendations presented by the SELKS team for optimal performance. The critical changes are as follows:

- **Suricata** – increase the default stream, memcap, and reassembly values to prevent packet loss.
- **ElasticSearch** – Increase the Heap Size to help to index; should be half of memory, but not over 32 GB.

ElasticSearch can be used to centralize all logs from systems throughout the organization, so more memory dedicated to ElasticSearch will increase speed and performance reliability. Proper tuning of Suricata and ElasticSearch will improve the speed, reliability, and integrity of the data.

2.7 Threat Intelligence

An effective threat intelligence program will help guide the analyst by generating actionable data to process. According to an article titled, "Definition: Threat Intelligence," threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications, and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard. (McMillan, 2013) The primary goal for threat intelligence data is to be actionable, therefore choosing a reputable source is important.

2.7.1 Open Threat Exchange (OTX)

The Open Threat Exchange (OTX) is a threat intelligence sharing platform used and populated by threat intelligence analysts, security researchers, and cyber analysts. Approximately 51% of the threat intelligence community is using OTX (Shackleford, 2015), which is one of the most widely known sources for threat intelligence. Users can upload data associated with a threat known on OTX as a "pulse." A pulse contains Indicators of Compromise (IOC) that are populated by the community. (Vault, 2016)

Indicators of Comprormise include:

- IP Addresses
- Domain Names
- Hostnames
- Emails
- Uniform Resource Identifier (URI)
- Uniform Resource Locator (URL)
- Filepaths
- File hashes (MD5, SHA1, SHA256, PEHash)
- CIDR or Netblock
- Mutex
- Common Vulnerabilities and Exposures (CVE)

Getting this information ingested in a relatively timely manner is important to detect recent threats inside the network. One way to do this is by converting IOCs into a rule format and keeping the IDS updated. Updated rules will continuously inspect network traffic for indicators of compromise. If rule conditions trigger, an alert is generated for an analyst to review.

Austin Taylor, sans@austintaylor.io

2.7.2 OTX Integration

AlienVault has created an OTX Suricata rule generator to retrieve, parse, and create Suricata rules from pulse feeds. To get started, the analyst needs an OTX API key and download the source code from https://github.com/AlienVault-Labs/OTX-Suricata. Figure 15 shows output from the rule generator. This process can be automated to populate Suricata rules with the latest threat intelligence from OTX. It currently supports file hashes and IP reputation.

Figure 15 - OTX Rule Generation (FILES)

```
[ataylor@smoke otx-suricata]$ head otx_file_rules.rules
alert http any any -> $HOME_NET any (msg:"OTX - FILE MD5 from pulse Flying
 Dragon Eye: Uyghur Themed Threat Activity";  filemd5:581916f4aa96ef71368c
8e47.txt; reference: url, otx.alienvault.com/pulse/581916f4aa96ef71368c8e4
7; sid:418010; rev:1;)
alert http any any -> $HOME_NET any (msg:"OTX - FILE MD5 from pulse Moonli
ght  Targeted attacks in the Middle East";  filemd5:5810d51fbe8776217ed00f
4a.txt; reference: url, otx.alienvault.com/pulse/5810d51fbe8776217ed00f4a;
 sid:416182; rev:1;)
```

2.7.3 Emerging Threat rule set (ET)

The ET rule set is comprehensive and creates rules for over 40 attack categories, including network behaviors, malware command and control, Denial of Service attacks, botnets, informational events, exploits, vulnerabilities, SCADA network protocols, and exploit kits. (Proofpoint, n.d.) Default with SELKS includes the ET rule set. Proofpoint populates the ET rule set and offers both a community and a pro version. The community version is the same rule set as the pro version but with a delayed release. The Scirius interface allows an analyst to maintain an updated set of rules by choosing a repository from which to download rules and determining the update frequency.

3. Visibility & Analysis

Figure 16 illustrates a small network with the two network TAPs in place. The TAP on the right will capture all traffic heading outbound and the TAP on the left will capture internal network traffic. The managed server, running SELKS, adjacent to each TAP will process flows, perform deep packet inspection, and apply threat intelligence

rules. Next, the central server will ingest metadata from the managed servers. Finally, the analyst will analyze the network through the Human Machine Interface (HMI) attached to the central server.

Figure 16 Network Topology with Log, Netflow, and Packet Capture

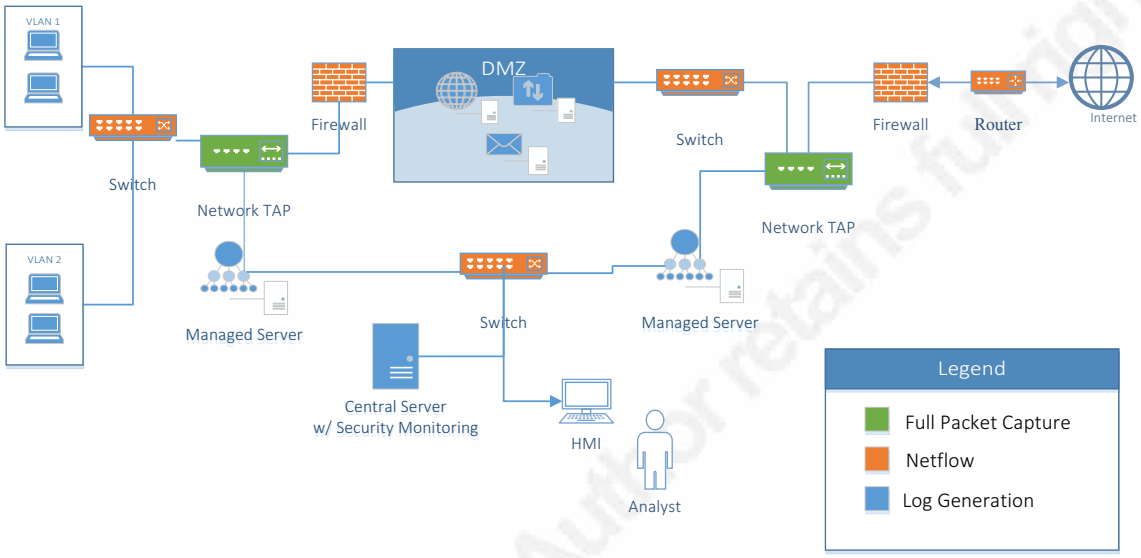

In this example, if both TAPs are not in place, it would complicate the job of an analyst. If the first firewall is providing Network Address Translation (NAT), the analyst would only see one IP address. With both TAPs in place, it offers the full story by showing the internal and public traffic. Thus, allowing an analyst to follow an incident from start to finish.

4. Conclusion

Network infrastructure is growing at a pace that can be difficult to monitor and control. Diagrams are outdated, inventory control is non-existent, and there are an insufficient number of analysts to monitor all the devices on the network. Some organizations have pieced together solutions that provide limited visibility into their network. This limited visibility makes it difficult for large Cyber Operations Centers and security analysts to detect, mitigate, and react to an incident and adequately defend the network.

Austin Taylor, sans@austintaylor.io

Organizations that do not accurately identify their critical infrastructure and services will not understand what they are trying to protect. Holistic environment monitoring is ideal. Logs, Netflow, and packet captures should be correlated to tell a complete story, which leads to shorter event processing times and increase in cost savings. Functional network visibility can posture an organization's security team to monitor and respond to cyber threats efficiently.

Austin Taylor, sans@austintaylor.io

References

(2016, October 01). Retrieved from Stamus Networks: https://www.stamus-networks.com/open-source/

Agency, N. S. (2016, 03 01). Event Forwarding Guidance. Retrieved from https://github.com/iadgov/Event-Forwarding-Guidance

Constantine, C. (2014, March 18). *Security Essentials.* Retrieved from Alien Vault: https://www.alienvault.com/blogs/security-essentials/what-kind-of-logs-for-effective-siem-implementation

McMillan, R. (2013, May 16). *Gartner*. Retrieved from https://www.gartner.com/doc/2487216/definition-threat-intelligence

Momentum Partners. (2015, October 05). *Momentum Partners Q4 Report.* Retrieved October 05, 2016, from Momentum Partners: http://momentum.partners/docs/Cybersecurity_Market_Review_Q4_2015.pdf

National Security Agency. (2015, August 07). *National Security Agency.* Retrieved from Spotting the Adversary with Event Log Monitoring: https://www.iad.gov/iad/library/reports/spotting-the-adversary-with-windows-event-log-monitoring.cfm

Proofpoint. (n.d.). *Emerging Threats Site*. Retrieved from Emerging Threats: https://www.proofpoint.com/us/products/et-pro-rule set

Shackleford, D. (2015, 02). *SANS Reading Room.* Retrieved from SANS: https://www.sans.org/reading-room/whitepapers/analyst/who-039-s-cyberthreat-intelligence-how-35767

Vault, A. (2016, January 20). *AlienVault.* Retrieved from AlienVault Open Threat Exchange User Guide: https://www.alienvault.com/doc-repo/OTX/user-guides/AlienVault-OTX-User-Guide.pdf

Winds, S. (2016, 03 01). Configure Devices for Flow Collection.

[i] Snippet from https://www.iad.gov/iad/library/reports/spotting-the-adversary-with-windows-event-log-monitoring.cfm

[ii]

Austin Taylor, sans@austintaylor.io