Input to the Commission on Enhancing National Cybersecurity

Applying Retrospective Network Analysis to Disrupt the Cyber Kill Chain

Comments for Improving Critical Infrastructure, State and Local, and Public Sector Cybersecurity

Executive Summary

99% of malicious cyber activity can be seen at the network. Monitoring networks eliminates the need for complex endpoint software, solves the issue of BYOD traffic and is generally the easiest deployment model. Cyber-attacks today are entering networks through a variety of means, and attempting to collect and exfiltrate sensitive information that includes intellectual property, personal identifiable information and cardholder data.

The attacks have become more sophisticated and elusive, bypassing preventative security defenses at the perimeter, the network and the endpoint. Attackers are hiding their activities and the data exfiltration in the noise of normal network traffic. This has created threat "dwell time," and a blind spot for security organizations that lasts on-average for more than 200 days.[1]

A simplified view of the cyber kill-chain illustrates the challenge:

Vulnerability weaponized -> malware discovered -> threat definition created -> organization protected.

We get smarter each day about new threats, but that knowledge is primarily helping us stop the "known" attacks going forward, and does not account for the unknown breaches that went undetected in the past.

The network is the common denominator and holds the key to truth about malicious activity.

By adopting a model of retrospective network visibility, where the latest threat intelligence is continuously applied to network history and user behavior, organizations can quickly uncover anomalous patterns on the network and find the advanced threats that are slipping past preventative security measures.

Challenges and Trends

When breaking down the challenges of critical infrastructure security in comparison to state/local or enterprise cybersecurity, there are going to be some obvious differences. For example, today's critical infrastructure relies on the use of Supervisory Control and Data Acquisition (SCADA) systems, which were designed for power distribution and measure frequency, voltage, and power at sensor locations.

However, there are commonalities in the challenges facing both critical infrastructure entities and public sector enterprises. The lack of network visibility and inability to continuously review that history stands-out as the most prominent. The "threat dwell time" is silently haunting today's security organizations. (Dwell time is characterized as the time that exists between a vulnerability becoming weaponized and a threat being detected).

[1] Verizon Data Breach report

Today's advanced cyber threats execute over long periods of time, hiding their communications in the normal flow of network traffic, and avoiding detection by preventative security tools. According to the Verizon Data Breach report, the dwell time is averaging more than 200 days. Most of the well-publicized data breach events illustrate this point:

Hackers breach Target and the company doesn't realize the full impact for two months

1,025 Wendy's locations hit by a malware-driven card breach that lasted six months

Home Depot admits hackers escaped detection with custom malware

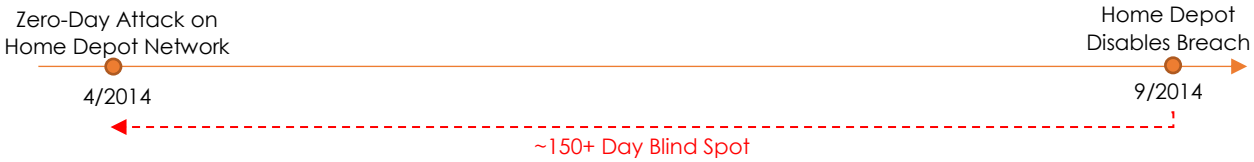

Hard Rock Hotel card processing network breached twice over two years

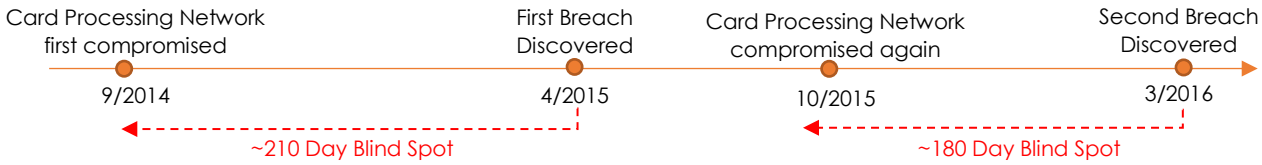

There are specific challenges related to gaining the appropriate level of network visibility for reducing or eliminating the threat dwell time:

1. *Lack of Application, Identity, Device, and Geolocation Information from Network Monitoring* – Tools, like NetFlow, cannot distinguish between multiple transactions such as those in a single email session, and at best will just provide a summary of the entire flow. This means you miss out on valuable information such as To, Cc, From, and Subject fields, as well as information about any potential malicious attachments. Certain obfuscated protocols like Tor can be difficult to detect on a network, but the ability to find and report these connections is critical to network security.

2. *Challenges Tied to Archiving and Network History Lookup* – While there are tools today that can store network log information for long periods of time, such as a security information and event management (SIEM) tool, there are challenges in inspecting that information quickly for the purpose of cyber investigations. The summarization of that data and the capabilities to tie the network data to an individual device or user becomes a challenge. There are also packet recording tools that have been available for some time. While these provide granular detail into network data, identifying what matters most for detecting a cyber event takes time, and the economics of storing full packets over an extended period of time becomes cost-prohibitive.

3. *Lack of Automated Workflows for Threat Detection* – The volume of new threat information coming in and the lack of skilled cybersecurity expertise available to process and understand that information is in high demand and difficult to acquire. "Log fatigue" quickly sets in with the amount of alerts and information that must be analyzed and processed. There is a general lack of automation to correlate that very latest in threat intelligence and tie that to actual events happening on the network.

   Currently, most cyber investigators still have to manually perform a series of complicated steps to generate useful investigative material from log reports and the limited history that full packet capture tools offer:

   a. Take in threat intelligence – The very first thing an analyst must do is scan various sources for threat intelligence information. These can come from third parties or from various security tools, and can be customized to an organization based on speed, origin, and a variety of other variables.
   b. Prioritize threat intelligence – Once threat intelligence has been gathered, it must be prioritized based on a variety of indicators, known dangerous behavior or potentially dangerous behavior.
   c. Comb through logs - After a threat feed has been prioritized, analysts must then go through logs to determine if the threat has been discovered on their network. If it hasn't, the process must be repeated with each threat feed, moving down through the list.
   d. If a breach has been discovered – further investigation is required to find out which machine (or, potentially, machines) have the indicators of compromise, and employ whatever steps are normally taken – either take the device offline while a deeper forensic investigation into its communications is performed, or keep an eye on traffic originating from and going to the DOI.
   e. Analyze the history of the DOI – To determine what information has been targeted on the DOI, an analyst would have to compile all logs from their organization's various security tools, such as firewalls, endpoint tools, and more, into a SIEM and run each threat feed against them. A workflow must then be built to identify suspicious protocols, analyze applications, isolate suspicious-looking flows, etc.

4. *Proposed Legislation to "Dumb it Down"* – Another recent trend has to do with recent introduced legislation that is intended to protect the nation's electric grid from cyberattacks by

"dumbing it down." The Securing Energy Infrastructure Act put forward by Senators Angus King, I-Maine, Jim Risch, R-Idaho, Martin Heinrich, D-N.M., and Susan Collins, R-Maine, would take what the bill's authors call a "retro approach" to critical infrastructure security by replacing vulnerable IT systems with unconnected, human-operated analog systems. The challenge and question here has to do with the cost and effort involved in f reengineering a power grid that is already operationally efficient. As the most technologically-advanced country in the world, this also raises a question about whether or not we are ready to take a step back, or look at alternatives that continue to advance our technology infrastructure.

Progress on Addressing the Challenges

The technology is improving, but adoption of new approaches is slow and the approaches vary. In terms of capturing higher-fidelity information from the network for faster detection and response, there are three separate camps of technologies: log capture, full packet capture and metadata extraction.

Log capture and aggregation tools benefit from collecting information from a wide variety of sources across the organization. The challenge remains about how to process the information quickly without overloading the security analyst. Full packet capture tools provide a high-level of detail for post-breach investigations, but storing that level of detail for long periods of time becomes cost prohibitive. And challenges remain around automated processing of the data for threat detection and investigation.

High-definition metadata generated from application-based deep packet inspection tools provides a great deal of promise, offering lengthy storage windows with reduced economics. The challenge with these tools is around detailed post-breach investigations.

The Most Promising Approach and What's Recommended

With the combined factors of threats executing over time, and with most cyber threats visible from the network, the ability to collect and store high-definition summaries from the network and have the ability to store that data for years, is essential. Let's break down the most-promising approach to solving this:

1. *Deep packet inspection (DPI)* capable of understanding network traffic in real time and at multi-gigabit speeds should be deployed at the Internet egress. This could be deployed using a TAP or SPAN port to passively scan every network flow, one packet at a time as they traverse the network.
    a. *Layered classification* of the traffic should be sequential and iterative, where each packet that is processed provides more information about what protocols are contained in a flow. This means as more packets are processed, the ability to classify the protocol is augmented.
    b. As classification is running, DPI should be performed simultaneously for all discovered layers. As new layers are classified, additional metadata is extracted at the new layer for all subsequent packets in the flow.
    c. *Heuristic pattern matching* should also be applied so that weak indicators like port numbers to classify protocols are not solely relied on. Heuristic pattern matching techniques should be used to classify traffic by inspecting flow content, which succeeds even if tunneling or obfuscation techniques are used.

2. Network history must be stored for more at least 12 months, and with reasonable economics. Compression should be applied to minimize the storage footprint. Once in storage, there should be capabilities to enrich the data with additional identity information from a directory store, as well as device and geolocation information to simplify hunting down a specific device of interest. There must also be accessibility to correlate a threat intelligence feed from any source with the network history.
3. A continuous and automated "rewind" process must exist, where the latest threat intelligence coming in is analyzed against the network history being stored and enriched. The benefit here is to uncover a previously unknown threat after a newly created threat definition or signature is published.
4. Some form of automated discovery must exist to alert a security analyst to potential threats and anomalous network behavior.

What can or should be done now or within the next 1-2 years to better address the challenges

Cyber threat activity is continuing at an alarming rate, and any organization from critical infrastructure to public sector enterprises are at risk of not knowing if the infrastructure has been, or is currently being breached.

The reality is that breaches begin with the standard process: beaconing, command and control communications, obfuscation of communications, obfuscation of content, and exfiltration. The issue is each of these individually are not noticeable, but when combined together they present a major problem.

What is needed today is the ability to wind the clock back by applying the latest threat intelligence to actual network history, and automating the discovery of any indicators-of-compromise and devices-of-interest. Our hope is that this provides a primer for a broader discussion among NIST and the Commission on Enhancing National Cybersecurity for placing greater emphasis and focus on network communications.

About SS8

SS8 is a time machine for breach detection. SS8 applies today's knowledge to history to find breaches now that you didn't know about before. By generating, storing and analyzing months, and even years, of enriched intelligence from all communications flows, SS8 customers benefit from unprecedented content- and context-aware insights that allows them to find the threats that matter most. SS8 is trusted by six of the world's largest intelligence agencies, five of the 10 largest communications providers and two of the world's largest critical infrastructure entities. Learn more at www.ss8.com.

About the author, Faizel Lakhani

Faizel Lakhani is president and COO of cybersecurity company SS8, which works with some of the world's largest intelligence agencies, telecommunications providers and critical infrastructure entities. Faizel has extensive experience in data security, network security, switching, routing, and Voice over IP technologies in both enterprise and service provider markets. Prior to SS8, Faizel was the vice president of Data Loss Prevention at McAfee, where he was responsible for the DLP business worldwide. Faizel is also credited with helping create Ontario Hydro's first SCADA system.