savvius

WHITE PAPER

Automating Incident Response

The bad guys are winning! There are more of them, they are evolving faster than ever, and their weapons are highly automated. Enterprises that don't respond well are inviting effective attacks.

Automating aspects of incident response can level the playing field. Intelligently capturing the data in which the attacker is hiding enables rapid intrusion detection and response.

Need for Automation

Modern security professionals face a daunting task: malware attacks are becoming more frequent, more targeted, and more automated. Every day, new exploitable flaws[1] are discovered. Even dealing with known existing flaws is challenging. The number of products that have to be monitored for patch management is rapidly increasing. Compounding the problem are new attack surfaces such as cloud applications and personal smartphones.

Every level of the internet stack, especially the application layer, contains a large number of paths for an attacker to gain access to key corporate assets. The number and size of those targetable assets has grown; large databases of proprietary information may provide better services to customers and patients, but they also create attractive targets for cyberthieves.

Meanwhile, malware authors have continued to innovate. Drive-by downloads and spear-phishing tactics provide fast, low-overhead vehicles to deliver RATs (Remote Access Trojans). This automated attack malware embeds itself in a computer system and facilitates every aspect of the "Attack Chain" from reconnaissance to exfiltration.

And they are fast! According to Verizon[2] , the typical penetration attack takes only minutes — sometimes, only seconds — to succeed. While it might take somewhat longer to actually exfiltrate the asset the attacker seeks, it's still only a few hours.

According to Verizon , the typical penetration attack takes only minutes — sometimes, only seconds — to succeed

Security product vendors have responded by offering a wide variety of defensive tools. They focus on authentication (password logins, e.g.), databases, sales and customer service applications, web portals, smartphones, cloud storage, desktop and server hardware, and more. The old days of relying on just a firewall shell, a SPAM filter, and desktop antivirus software are well behind us.

For the security professional, this ever-expanding set of products is a mixed blessing. Each new defensive system generates more alerts which have to be investigated, often hundreds, even thousands of alerts per day. Most of them are false positives. Even the class of products intended to help manage this deluge of alerts, SIEMs (Security Incident and Event Managers), can themselves generate more alerts than the security staff can readily investigate.

What Can Be Done?

All of these factors highlight the fact that manual, sequential, plodding review of each and every alert generated by various defensive systems is unworkable today. It will get worse tomorrow.

Trained and experienced security professionals capable of the detailed investigations needed to find and stop attackers are rare. Given the speed with which an attacker can penetrate a system, even the most capable security professional needs to be able to react precisely and quickly to an assault.

The answer lies in network-focused automation of incident response.

While a thorough investigation of an alert and potential intrusion might eventually entail all three well-known computing domains — network information, host application, and system logs and computer memories — the fastest and most trustworthy approach is to focus on network information.

[1] According to the Symantec Internet Security Threat Report (ISTR), 400 million new variants of malware were created in 2011, which is an average of 33 million new variants of malware a month, or an average of one million new variants a day

[2] 2016 Verizon Data Breach Investigations Report

Why network information? The intruder must have traversed the network in his penetration. He or she may have touched the logs or memory of only a few individual computers or servers. Another factor is that modern attacker malware can automatically erase or, more insidiously, subvert logs making them untrustworthy.

By contrast, packets never lie!

If collected and sequestered on-the-fly, network packet information is a pristine record of the attacker's path and behavior.

The amount of information flowing through a modern enterprise network would be very expensive to capture and store. Searching for attacker behavior among that much data is tedious at best.

Studies have shown that alerts generated in the defensive systems can help filter which network data might be relevant. Even false positives don't add much to what will be stored — typically as little as 2-3%. This data reduction makes retention, even for long periods of time, economically feasible. It also reduces the volume of data that the investigator has to review so he or she can quickly isolate the attacker.

Despite the majority of alerts being false positives, most well-designed enterprise security defenses systems will capture some aspect of the attack. So, once an investigator is alerted to the presence of an attacker, he or she can be confident that an automated network information storage system has the data needed to find the assault.

False positives don't add much to what will be stored — typically as little as 2-3%.

One very important aspect of capturing the filtered network data is when and how much data is collected. By the time a defensive product triggers an alert, the actual attacker behavior has already occurred. Typically, the attacker will attack in multiple steps, and often, the first interaction isn't the one that triggered the alert.

For example, a DLP (Data Loss Prevention) system might alert because of an attempted access from another computer on the network. Since the attacker earlier corrupted the computer attempting the improper access, if we started collecting network data at the DLP alert, we wouldn't have seen the first computer getting compromised and wouldn't know how the attacker had gotten into the network in the first place.

The solution to this problem is to collect information from before the defensive systems produce the alert!

This isn't clairvoyance – the automated network information collection system maintains a rolling buffer of say, five minutes in length, for all network traffic. If an alert is detected, all of the network information for the previous 5 minutes is still available, so everything specifically related to that alert on the network can be sequestered and the rest discarded. Later, when the investigator looks at an alert, he or she can look back to the time before the alert to see what led to the trigger. Thus, all of the actual attacker behavior, including the initial penetration, will be evident.

Summary

Automation in the collection and storage of critical network security information can help the security professional deal with the challenges of protecting his or her corporation. By focusing on just the relevant data needed to find the attacker, the investigator can quickly identify and remediate the assault.