ENDGAME.

Detecting Modern Adversaries: *Why Signatures Are Not Enough*

SERIES ONE, VOLUME ONE

MARK DUFRESNE

Cyber intrusions are continuing unabated with no end in sight.

Ransomware is on the rise, massive data breaches are announced with such regularity that the public is becoming numb to their significance, and APTs continue to burrow deep into networks going undetected for months or even years. At the same time, most organizations across all industries are increasing their cyber security budgets but usually fail to produce a meaningful increase in defensive effectiveness.

In short, the adversary continues to win. Fortunately, most security professionals and vendors are asking what must be done differently to increase defensive effectiveness.   We often hear that enhanced signature sharing is the primary solution.   From the other end, we hear that signatures are dead. The truth lies in between.

Signatures are effective in detecting a portion of what is already known and for hunting within your enterprise to understand the extent of a known intrusion. However, due to their brittleness and increasing specificity to only the targeted victim, signatures are an utterly insufficient foundation for the caliber of detection and prevention capability needed today to prevent compromise or detect and remediate compromise as rapidly as possible.

We need to do more. We need to add additional layers of detection around signature and IOC search, looking for indications of attacker techniques at low levels in the system while simultaneously hunting for higher-order patterns which could indicate maliciousness across large sets of monitored hosts. Moving from solely signature-based defenses to also including attacker techniques and patterns is the best way to maximize the defender's chance of success in minimizing damage and loss.

Why aren’t signatures enough?

For the purpose of this post, we use the terms signature and Indicator of Compromise (IOC) interchangeably. A good signature is a feature that, with a low false positive rate, uniquely corresponds to a known attack campaign or piece of malware. We can group these in two buckets: network signatures and endpoint signatures.

On the network

Network signatures usually come in the form of blacklisted domains, IP addresses, URI structure, or patterns in command and control or other communications. Two primary factors have massively reduced the effectiveness of network IOCs in recent years: attack infrastructure diversity and encryption.

First, adversaries know their infrastructure is a point of vulnerability in their campaigns and actively seek to diversify and blend in as much as possible. The ubiquity of cloud services has been a major enabler for adversaries, allowing them to rapidly stand up and tear down infrastructure for low cost. Others use legitimate cloud services for data exfiltration or command and control, bypassing a need for a dedicated C2 infrastructure. Adversaries also engage compromised, unwitting nodes as disposable hop points. Trying to keep up with every hop point to defend your network is not a winning strategy.

Adversaries would in the past often use the same infrastructure across many victims for long periods of time. This is much less common today. High caliber adversaries will usually use infrastructure across many victims for only very short-lived campaigns, sometimes going so far as to use entirely unique infrastructure for all phases of an operation targeting a specific victim. Today, signatures may only be useful retrospectively to identify whether a newly discovered campaign (which may have taken place weeks or months ago) targeted you. Signatures may actually prompt you to waste resources searching for something an adversary never would have used to target you in the first place.

Next, encryption has made it far more difficult to track patterns on the wire. Network-level pattern matching capabilities such as Snort or Bro signatures used to be relatively effective in detecting intrusions in your network. Malware authors need to design structured command and control communications to organize victim machines and direct victims to take certain actions. Analysts can often fingerprint these communications structures and detect them on the wire, even if unique or unknown infrastructure is in use. However, we are increasingly seeing malware communicate within end-to-end encrypted tunnels, usually using universal protocols such as SSL or TLS. When communications are encrypted, unless SSL proxying or other intrusive traffic inspection technology is put into place these patterns are not visible to network security appliances applying these signatures. Thus, the signatures for the malicious malware communication patterns will not fire and the intrusion will go unnoticed.

On the Endpoint

Evidence of an intrusion on workstations and servers can be found in numerous locations, including malware hashes, filenames, registry entries, and much more. As with network infrastructure, in the past, malware was regularly reused across many victims for long periods of times without diversifying these artifacts. Adversaries with any level of sophistication no longer make these mistakes. They have learned that it is important to avoid a detrimental (from their point of view) global impact from a single detection. Defenders need to understand this and pursue intrusions accordingly.

Malware is often polymorphic, changing itself to have a unique hash every time and automatically diversifying filenames, persistence mechanisms, and other features which can be signatured. In these cases, which are increasingly common, an artifact found in a single victim will not be effective as a global IOC. Strategies that focus on patterns within malicious binaries themselves (Yara signatures, for example) can at times be relatively effective in detecting new tools from a given known malware family, but these can be difficult to use across an enterprise and are very prone to false positives.

In addition, some adversaries are moving entirely away from malware as their default way of accessing and interacting a victim. Legitimate credentials and administrative tools like Powershell are often all that is needed to take desired actions on a network. Malware is often only used for persistence and sometimes not used at all. In these cases, the adversary does not leave behind a significant footprint to be used as the basis of IOCs. IOCs will be entirely ineffective and the problem turns into distinguishing malicious usage of tools and credentials from normal operations.

Do we still need signatures?

For the reasons described above, signatures are not a sufficient foundation for detection and prevention in your network. That said, they are still valuable. They are useful and effective in catching unsophisticated tools and actors. They can also help you determine if a given attack campaign has touched your systems.

Search functionality is very important to locate known IOCs on your systems and in your traffic. Signature search is also necessary to determine the extent of a given compromise in your environment. For example, if you find evidence that a certain registry key is being used for persistence on a compromised host, you need a way to look across your other systems to look for that same key. IOCs of this sort are useful much more often inside your network than they are to other possible victims of the same adversary. IOC searching is a part of threat hunting, but it's not enough.

So we need more. What should we do?

We need technologies to detect threats without relying on signatures. This takes two main forms: looking deep in the operating system for indications of malicious activity and hunting for suspicious patterns across key data from many systems. Basically, we must look a layer below and a layer above IOC search.

There are a few well established frameworks for understanding the sequencing and methodologies exhibited time and time again in cyber intrusions, such as Lockheed Martin's Kill Chain and Mitre's ATT&CK framework. While adversaries constantly change and adapt malware, they actually use the same techniques over and over – process injection, credential dumping, token stealing, host enumeration, and lateral movement being a few examples of many. An attacker can build a nearly infinite number of tools to do these things generating different IOCs, but they must go through the same choke points in the OS to execute these actions on the system. We can identify these key chokepoints, develop ways to detect and optionally automatically block the adversary, and alert the cyber security operations team that a malicious event has taken place. Effective tools can prevent malicious activity at the right chokepoints in real-time and alert the security team to a likely intrusion - all without signatures.

We also must look for suspicious activity and patterns across our endpoints. This is the core of effective threat hunting, improving from simply finding what's known to empowering security teams to find unknown and unique intrusions. This is possible because adversaries leave a trail which can be followed. Adversaries must operate on systems. They must execute code. They usually communicate on the network. They often read, create, or modify files. They do much more. All of these breadcrumbs can be followed by an astute hunter. The hunter can look at process activity information, network traffic, domain lookups, previously executed commands, persistence locations, and in other key areas. Suspicious activity can be flagged, investigated, and detections can occur. In this way, IOC search becomes a subset of hunting.

Hunting manually can be very difficult and will not scale. However, by combining hunt methodologies with automation, analytics, and machine learning, hunt operations can be scaled and optimized. Detections of unknown intrusions can be surfaced at speed and scale at this layer above traditional IOC search and then acted upon by the security team.

Conclusion

We still need to use signatures. It is important to have a capability to search for artifacts associated with known campaigns, to combat low caliber adversaries, and to pivot through your network once a unique adversary is discovered via other means.

Signatures are not enough to form the detection and prevention solution needed to defend against modern threats. They are neither effective on the host nor at the network level to detect advanced adversaries. Additional detection capabilities which look at low level chokepoints in the operating system are necessary, as are simultaneously executed hunt operations across systems for indications of suspicious or malicious activity. By combining hunting with automation, analytics, and machine learning, we can produce high quality detections which can be used by security operations teams in the same fashion as detections from chokepoint monitoring and signature monitoring.

Combining these three layers - low-level attacker techniques detections, signature-based detections, and detections from automated hunts - maximizes the chances of stopping adversaries before they succeed.

ENDGAME.