SentinelOne
The future of endpoint security

ENTERPRISE RISK INDEX

Risk assessment and control factors Q2 2017

Overview

In 2017, SentinelOne is releasing the Enterprise Risk Index, based on data obtained from enterprise environments from the second half of 2016. "The purpose of this report is to highlight the threats that are actually successful in reaching the endpoint."

We´re calling this the Enterprise Risk Index because, in order to achieve resilience from our security endeavors, we must accurately assess where the risk is coming from and apply satisfactory controls against those various factors.

This report is unique for two reasons:

First, our index is based on detections on the endpoint rather than detections from the gateway or statistical data from a cloud collection system; this may produce results that are out of sync with industry beliefs. We are not measuring the total encounters with malware, as the vast majority will be mitigated at the gateway or in the network. These threats are inherently blocked, and as such pose no risk to the environment. Rather, we are measuring the attacks that make it all the way to the endpoint and exhibit malicious behavior, unauthorized access and other nefarious activity. This is where risk is incurred and this is the real threat.

Second, the detections are based on machine learning systems focused on the behavioral characteristics of malicious activity. Detections are not specific to malware families or campaigns, and we don't natively attempt to identify what has gotten into the system, but we can tell you how they got in, when they got in and what they did. With this in mind, we won´t be announcing what the top malware family is - for example, Zeus, Diamond Fox or Upatre. However, we do build indicators of compromise to help with identification and response, and when we found a hash value, we submitted the hash to malware repositories to see what other submissions there have been for them.

What the report will tell you is the method with which malicious attacks are being deployed. We have classified these attack methods into three general risk categories:

Attacks detected from document based files; largely associated with Microsoft Word documents and Adobe PDF.

Attacks detected from executable files; Traditional Portable Executable based attacks.

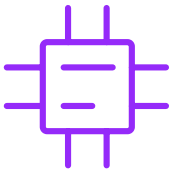

Attacks detected only in the memory of the system with no associated new artifacts on the system. The attack may exploit existing operating system resources, and run subsequent code or instructions directly from memory.

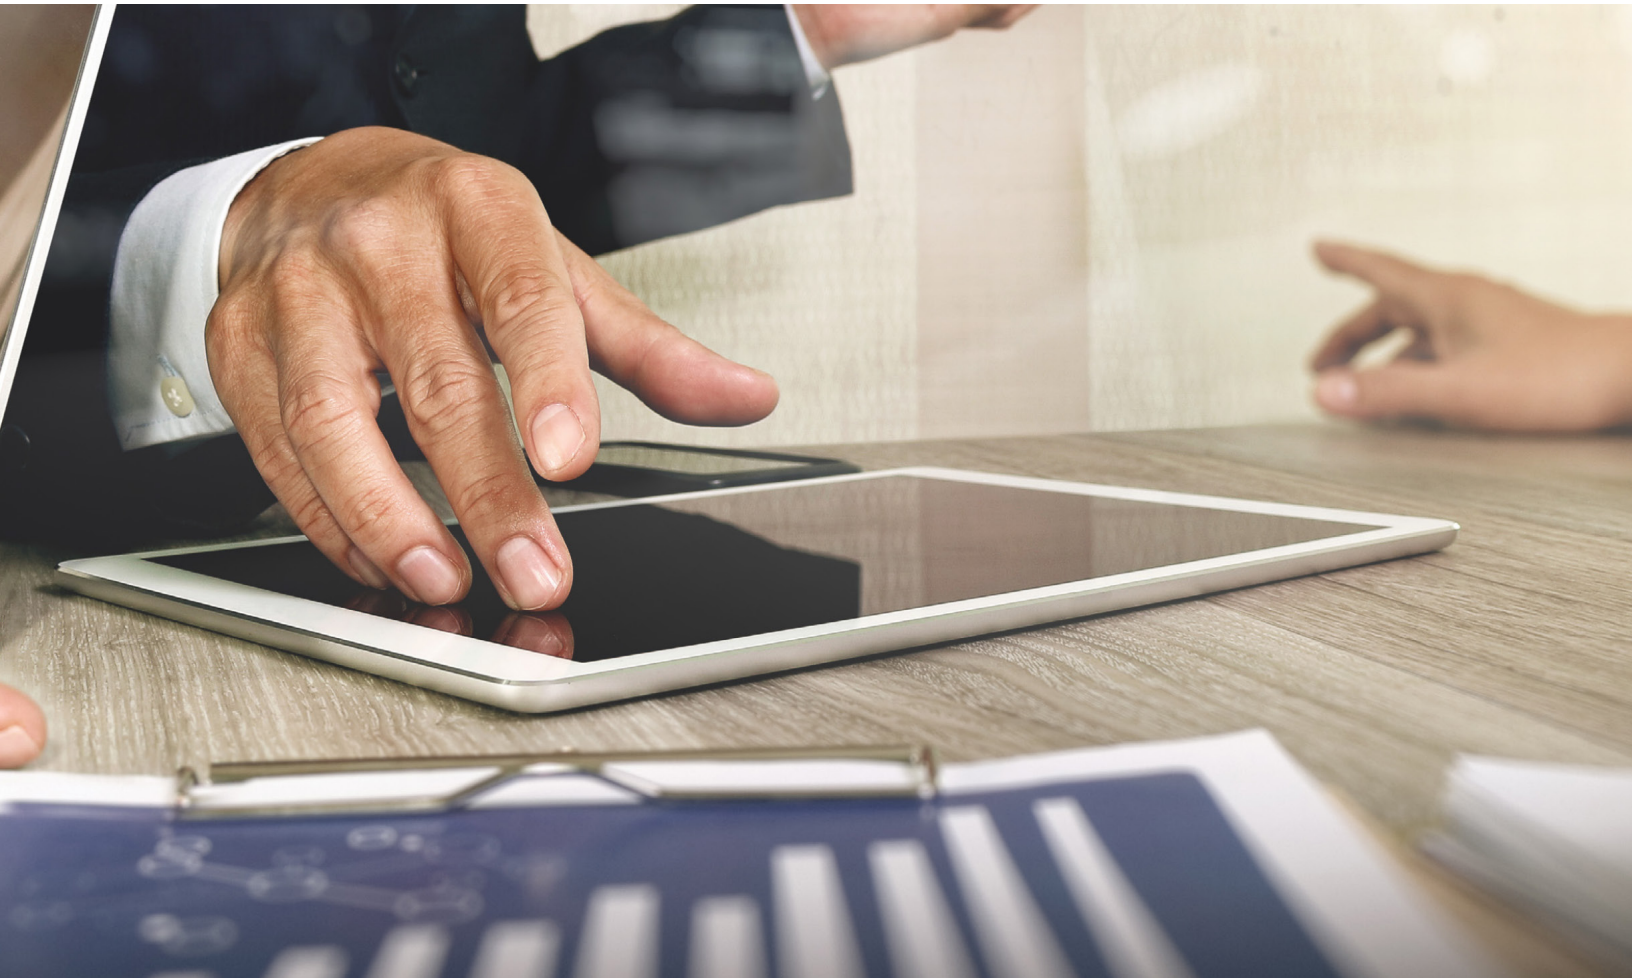

The SentinelOne Enterprise Risk Index detections are based on the SentinelOne EPP 1.6.2.5 agent and above. For Windows production environments, we did not include Proof of Concept or testing environments since there is a heavy bias towards testing against portable executables downloaded from a malware repository. The time frame represented in our results is August 2016 to November 2016 and is based on filtered data obtained from more than one million SentinelOne agents deployed worldwide. The vertical axis in the table below shows the percentage of methods successfully reaching the endpoint target.

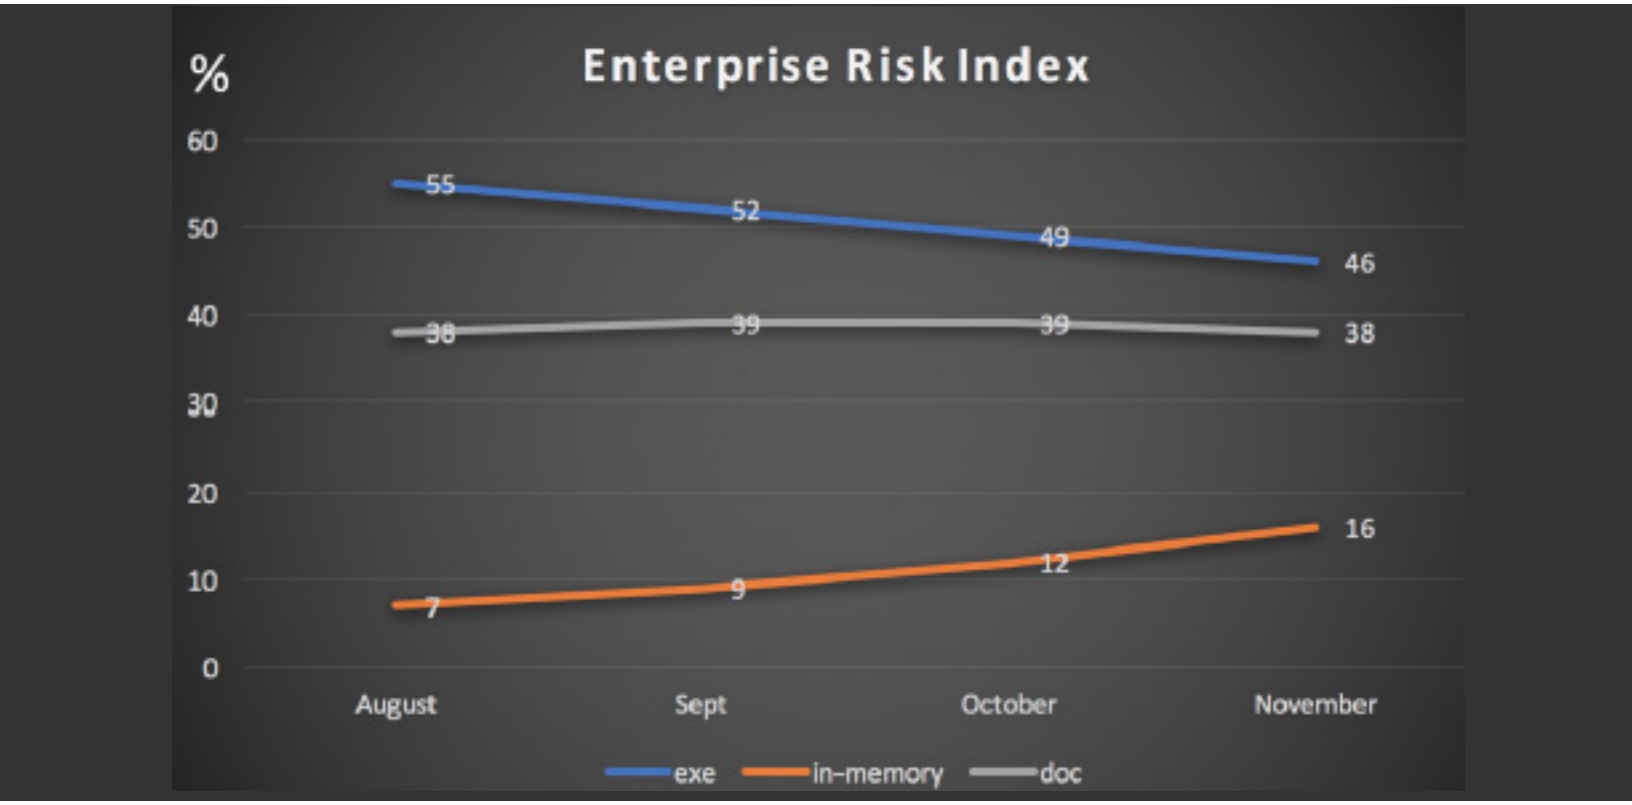

We will first address those threats that have an identifying hash value. Of the hashes obtained from the file-based attacks, which includes document-based attacks, only 50% had been previously submitted to malware repositories. And of that half submitted,

Only 20% had corresponding signatures from existing anti-virus engines.

We know from a file-based infection perspective our data reflects that, of other threat reports, in the latest Microsoft SIR21 report, the encounter and infection statistics mirror very closely what we are stopping in our customer base – that there is a large number of affiliate marketing syndicates offering pay-per-install revenues delivering types of Potentially Unwanted Programs, adware programs, browser modifiers, toolbars and installers.

In themselves they represent a nuisance, a potential for operating environment conflicts and, at worst, unauthorized usage and unauthorized activity from a risk assessment perspective. However, they can be conduits for more malicious deliveries over the duration of their existence in an internal environment leading to unauthorized access and a heightened level of risk.

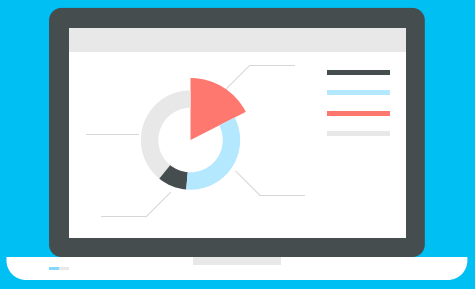

2016 was coined the YEAR OF RANSOMWARE,

and the number of encounters and the variety of malware families attempting to compromise our installed base certainly reflected this with many evolving payloads.

Global Development of New Ransomware Variants is Accelerating

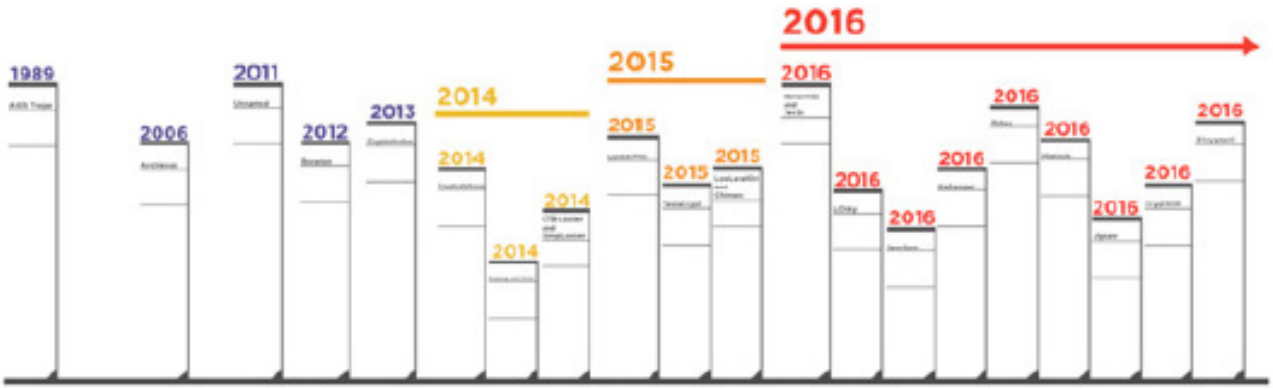

With a single bitcoin now approaching $1000 USD,

this type of threat is only going to get more painful for organizations. The cost from ransomware is not just in the money extorted, it is in the loss of operational capability evidenced by healthcare outages from around the world, which result in thousands of remediation dollars spent instigating a backup procedure of the infected devices , which takes an average of 33 hours to complete.

URSNIF

has been around since the late 2000s, and has been known by many names - Gozi, Vawtrak, Snifula. It has been a primary banking Trojan payload used by an organized crime gang (or gangs). They have also been identified as the Neverquest threat actor group, also known as TA530 or FIN6.

URSNIF & LONGEVITY

is an excellent example of the cat and mouse game played by the bad guys and the security industry vendors, as one jump in detection capability has led to an advancement in stealth techniques and the development of alternate infection methods. Today Ursnif is still active and now has strong evasion capabilities and a file or a file-less payload option.

NOTABLY,

one of the hashes picked up by the malware repositories that we submitted was identified as Zegost - Farfli and Gh0st by various AV vendors - it is a variant or descendant of the poison-ivy RAT and first came to prominence as a tool used in 2011 by the "Nitro" APT group targeting intellectual property leaders in the chemical industry. The sample detected by SentinelOne was previously submitted from a Korean entity and was seen calling back to a dynamic domain name system address. Interestingly, or coincidently, the SentinelOne detection also came from an organization in the chemical industry.

Of the malware payloads that have loftier goals - unauthorized access and spying capabilities as part of their modus operandi. We observed two types of threats in this category found in our timeframe.

… attacks making it to the endpoint that were detected during the memory resident phase of the attack have doubled.

Finally, we look at the risk from threats that reside in the memory of the target system. As we can see from the risk index, memory resident attacks have doubled during this time period and an equally important corresponding trend is the decline in traditional .exebased attacks, which are malicious campaigns that require live unauthorized access or interaction with the victim system and are opting for a memory resident tactic rather than a new payload on the filesystem. Often the originating object will be cmd.exe, powershell.exe or mshta.exe, as legitimate and essential operating systems resources that are subverted as the payload platform during the exploitation stage, instigated frequently either by a document received by email, malicious script or an active code component on a web page.

There are many different methods and tools that we detected trying to gain a foothold in memory; WMI persistence is one such tactic. This type of technique was first discovered during the investigation into Stuxnet and later also identified as a method used in the attack on the Democratic National Committee. The alleged threat actors behind these attacks are different groups, so this is an example of one group copying from another.

Another common attack pattern we see is a “live” or interactive attack, where the attacker delivers a weaponized document and is able to employ a meterpreter reverse shell, powersploit payload or red team testing frameworks. We often see hackers invoke reflective injection techniques to run late stage tools such as mimikatz, to gather credentials on the impacted system. We routinely spot the insertion of javascript into command line instructions and observed an increasing trend in exploits issuing malware payloads in shellcode rather than a file.

As with all industries, there are leaders & followers.

Below is the PE (file) or Shellcode compiling options for the notorious Poison-Ivy.

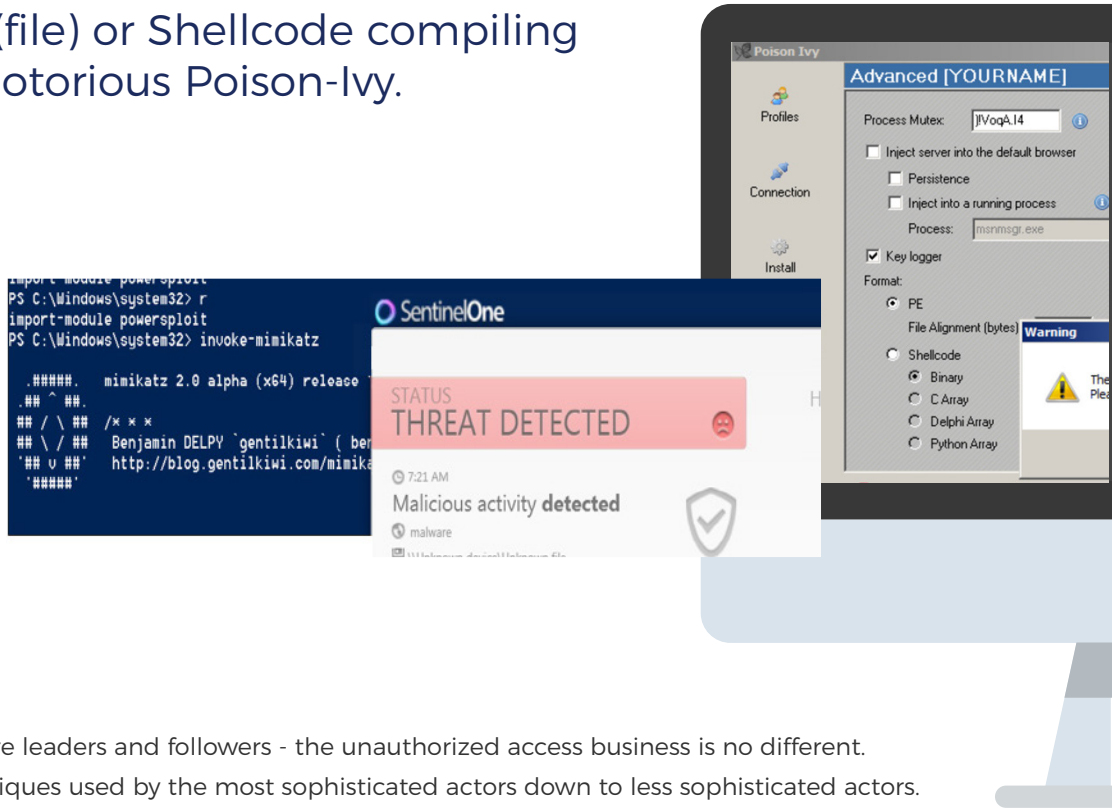

As with all industries, there are leaders and followers - the unauthorized access business is no different. There is a drip down of techniques used by the most sophisticated actors down to less sophisticated actors.

It´s not easy breaching an enterprise organization, the resilience rates are often in the thousands of encounters repelled before an infection takes place. The pattern adopted by Nation State Actors, is to place zero or as few new artifacts on the file system as possible to minimize the potential for detection by enterprise security controls, even if this means being ephemeral with the risk of having to re-infect the victim. Hackers prefer taking this risk rather than having a file based indicator of compromise detected and disseminated to the broader security community.

Now there appears to be a growing number of cybercrime authors copying these tactics. Angler EK had a file-less option, and Kovter, Phasebot, Powersniff and LatentBot are just some of the recent examples to employ in memory tactics.

Here come the TRIDENTS.

Our classifications in the SentinelOne Enterprise Risk Index represent the initiating object that caused the malicious activity to begin. Of course, in the real world, the line blurs. There are attacks that have all three aspects, which we call Tridents. They start as a document, exploit into memory, run additional shell code or execute instructions to whitelisted OS utilities, and if they want persistence they drop artifacts onto the file system or they simply drop additional payloads onto a disk. A hybrid multi-stage spread spectrum model.

AV is the Walking Dead, again.

The “AV is dead” slogan has been around for nearly 10 years, with the logic being that the bad guys are able to morph and encrypt their malware faster than the AV industry can write protections. We have statistics estimating around 390,000 new malware samples being uploaded every day to support this logic. “Traditional AV” simply cannot cope with developing corresponding protections with the current volume of uniquely hashed malware samples. Signatures only work if the exact malware stub is reused somewhere else.

However, what if that logic proves not to be the final nail in the AV coffin? What if it´s not the volume or snowflake uniqueness of files that need scanning that cause the ultimate death of AV? If the trend in memory resident attacks continue as a preferred option to a portable .exe, what if it´s the utter lack of files that render AV ultimately redundant? What if .exe and .dll files take a well-earned hiatus as a mainstream threat vector, just like macros in documents did in the past?

Plan for resilience.

Clearly, executable files are still a highly-encountered type of threat, however, this does not mean they are more likely to succeed at causing risk in the environment. Indeed in the Microsoft SIR 21 report of the file-based attack vectors, .doc extensions and JavaScript extensions were responsible for 70% of the successful attacks.

What´s omitted, ignored or simply not captured in many of the threat reports is the vector of in-memory or file-less threats. As stated in the opening paragraphs, SentinelOne is bringing visibility to the risk that organizations actually face across the board in a truly representative manner. This is done by providing actionable guidance of what is doing the damage and highlighting the behavioral trend in malware to operate their attacks in-memory, which has supplanted the payload delivery phase of file-based attacks as the more successful method because they have an easier time evading traditional and static file inspection dependent security models.

Rules of thumb.

Here are some takeaways from the Enterprise Risk Index data:

- For every handful of file-based infections you find, you need to hunt for the infections created without a file; measure and report on the things that matter. There is no risk from things that are blocked, only things that get in and pose the risk measure that share that data too.

- It´s not the breach that kills you; it´s how you deal with it that can. Reporting on levels of unauthorized access is not a weakness, it's a sign of resilience.

- The strongest probability in risk assessment is that, with almost certainty, you have it wrong.

Summary

Quite frankly, as a discipline of information security we are often flawed by the rigidity of policy advice. Our policies dictating best practices are, in some cases, actually bad practices.

When a user calls the help desk, saying,

"I may have been infected, I clicked a link and then a pop up appeared and my machine rebooted,"

our standard response is, “Can you make sure your anti-virus is up to date, and then do a full scan of the system?” The user calls back and says, “It didn’t find anything.” So, the help desk assumes a false positive, closes the case, until something else shows up at some point in the future.

Updating signatures and doing a scan simply does not achieve satisfactory levels of risk management. Our best practices and policies need updating frequently; 200 days of unauthorized access dwell time is common but unacceptable.

We continue to justify this help desk process as duly diligent and contextual in risk management best practice on the basis that our risk assessment was built on the fact that, in a testing environment, we downloaded 100 executable files from an online malware repository and detected them all. This naive view of security must end.

Enterprises

Cyber resilience is about having the right tools for the job. That starts with the right data, and we have attempted to present that here in the first threat agnostic SentinelOne Risk Index.

About SentinelOne

With intelligent automation becoming an obvious replacement for signature-based detection, SentinelOne offers a comprehensive solution for servers and endpoints. SentinelOne offers a lightweight solution secures endpoints and servers without compromising performance. Behavioral threat analysis that leverages machine learning to capture and neutralize both known and unknown threats, while providing a forensics package that allows administrators to visualize attack paths and remediate vulnerabilities.

In terms of compliance, behavioral threat analysis also removes some of the necessity of patching systems to their latest version. While this is best practice, oftentimes updating one system will break the dependencies of its connected subsystems—meaning that administrators must trade a functioning network infrastructure for security and compliance on the other. Organizations can rely on SentinelOne to monitor unpatched systems, meaning that even an out-of-date program retains its security.

In terms of mitigation, SentinelOne can block and identify malware, even if it hasn't been seen before in the wild. In Alert Mode, it can identify malware, such as ransomware, and detect malicious behavior, such as creating an executable file without permission. SentinelOne will display the entire attack path of malware—and then enable administrators to seamlessly rollback an infected machine.

With SentinelOne, IT teams finally have a viable path forward that allows them to stay ahead in the arms race against bad actors. Instead of spending limited time, money, and manpower remediating breaches that are already in progress, security practitioners can now usefully devote their time to reinforcing the solid foundation which SentinelOne provides.

SentinelOne
The future of endpoint security

For more information about SentinelOne Next-Generation Endpoint Protection Platform and the future of endpoint protection, please visit: sentinelone.com