

AI-Based Prevention and Detection

The Evolution of Endpoint Prevention and Detection

WHITE PAPER






Collectively, ransom demands exceeded \$1 billion. There were significant and costly business disruptions.

A Global Cyber Attack of Unprecedented Scale

In May 2017, a global wave of cyber attacks reached an unprecedented scale. Businesses, schools and universities, public health and transportation systems, and many other organizations in more than 150 countries became victims of the WannaCry Ransomware attack. Within 24 hours, more than 200,000 computers across the European Union were hit. Chinese state media reported nearly 40,000 public and private institutions in that country had fallen victim. Very large organizations were brought to their knees and unable to function normally during the height of the attacks.

Collectively, ransom demands exceeded \$1 billion. There were significant and costly business disruptions. For example, hospitals were forced to postpone patient operations, factories experienced manufacturing slowdowns and even shutdowns, and transportation systems were delayed. This wave of cyber attacks disrupted normal routines for millions of people around the world and will have cost many hundreds of millions of dollars to recover from the outages.

WannaCry, which contained both a ransomware application and a worm, exploited a vulnerability in Microsoft Windows to gain a foothold on endpoint computers, establish persistence, and then spread as far and wide as possible across enterprise networks. Cybersecurity experts say that organizations were vulnerable to WannaCry due to unpatched operating systems and inadequate defensive measures.



This attack, like many others, started at the points of highest vulnerability: the endpoint computers connected to business networks.

Too Many Endpoints Are Vulnerable

This attack, like many others, started at the points of highest vulnerability: the endpoint computers connected to business networks. In the 2016 SANS Institute “State of Endpoint Security” survey of IT and security professionals, 85% of the respondents reported compromises of desktop computers, and 68% reported compromises of laptop computers within the prior 24 months. Moreover, 13% of respondents considered their desktop compromise to be “widespread.”¹

This high state of compromised endpoints is just what cyber criminals need to pursue their goals. Malware lies at the very heart of cyber criminal operations. It’s a tool that allows organized gangs — or even nation states — to carry out espionage, sabotage, or theft. Major cyber criminal campaigns such as banking trojans, ransomware attacks, and even nation-state level cyber espionage begin with malware distribution to the target as the first stage of the attack.²

What’s good for the cyber criminal is very bad for the enterprise. Obviously, no organization wants to allow its endpoints or network to be compromised in any way, for any reason, and this is why endpoint protection is such a high priority for most companies. As network perimeter protection has all but dissolved, it’s critically important

to have effective security measures at every endpoint connecting to a network in order to block malicious access attempts and other risky activity at these points of entry. It all comes down to risk mitigation for the organization.

The technologies used to protect endpoint devices have changed significantly over the years. Here’s a quick look at where the industry has been, where it is today, and where it is quickly headed.

The Early Years: Prevention Based on Malware Signatures

Some of the early commercial software meant to protect PCs from viruses, worms, and other forms of malware became known as antivirus software, or simply AV. From the early 1990s until 2007 or so, AV software based on the principle of looking for the presence of unique file hashes (signatures) and then blocking the malicious programs was usually sufficient to protect the PC. When a new form of malware was discovered in the wild, AV researchers could deconstruct and analyze the code to understand its unique signature. Then they updated their AV program code to look for and block that signature to protect the computers running the AV software.

¹ G. W. Ray Davidson, PhD, SANS Institute, “Can We Say Next-Gen Yet? State of Endpoint Security,” March 2016

² Danny Palmer, ZDNet.com, “Why malware is still the beating heart of cybercrime,” May 4, 2017

Beginning in about 2007, however, the number of malware variants began to explode, making it impractical — and eventually totally impossible — to create a signature for every piece of malicious code in the wild. (See Figure 1.)

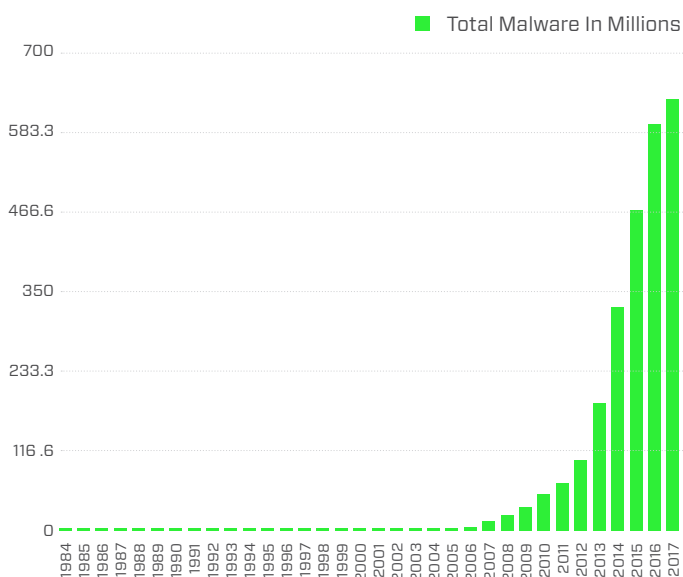


Figure 1: The exponential growth of malware since 2007

(source: www.av-test.org)

What happened to create such phenomenal growth in new malware? For one thing, the creation of polymorphic viruses, which can modify their source files automatically. If even one character in the code changes, the result is a totally new fingerprint being generated. Thus, one piece of malware can easily produce hundreds of variants with each variant requiring its own signature to be created by traditional AV software vendors. Considering that over 390,000 new malicious programs are registered every day, it's understandable that no one can keep up with creating the associated signature files.

In addition, the development and distribution of malware has become a commercial venture in its own right. Security blogger Brian Krebs described a process called crypting, in which malware developers run their software through a battery of tests to ensure that it is able to evade detection by commercial antivirus tools. The goal is to end up with malware that is completely undetectable by all of the AV tools on the market today.³

If Antivirus As We Know It Is Dead, What Fills the Void?

Security solution vendors practically abandoned the notion of preventing threats from executing at the endpoint and instead focused on detecting malicious activity as it happens or soon thereafter — hopefully in time to prevent, or at least minimize, the damage. All manner of new technologies have been developed for detecting activity that is suspected of being malicious after it has executed: log analysis and correlation, user behavioral analysis, retrospective searches, sandboxes, deception, and more.

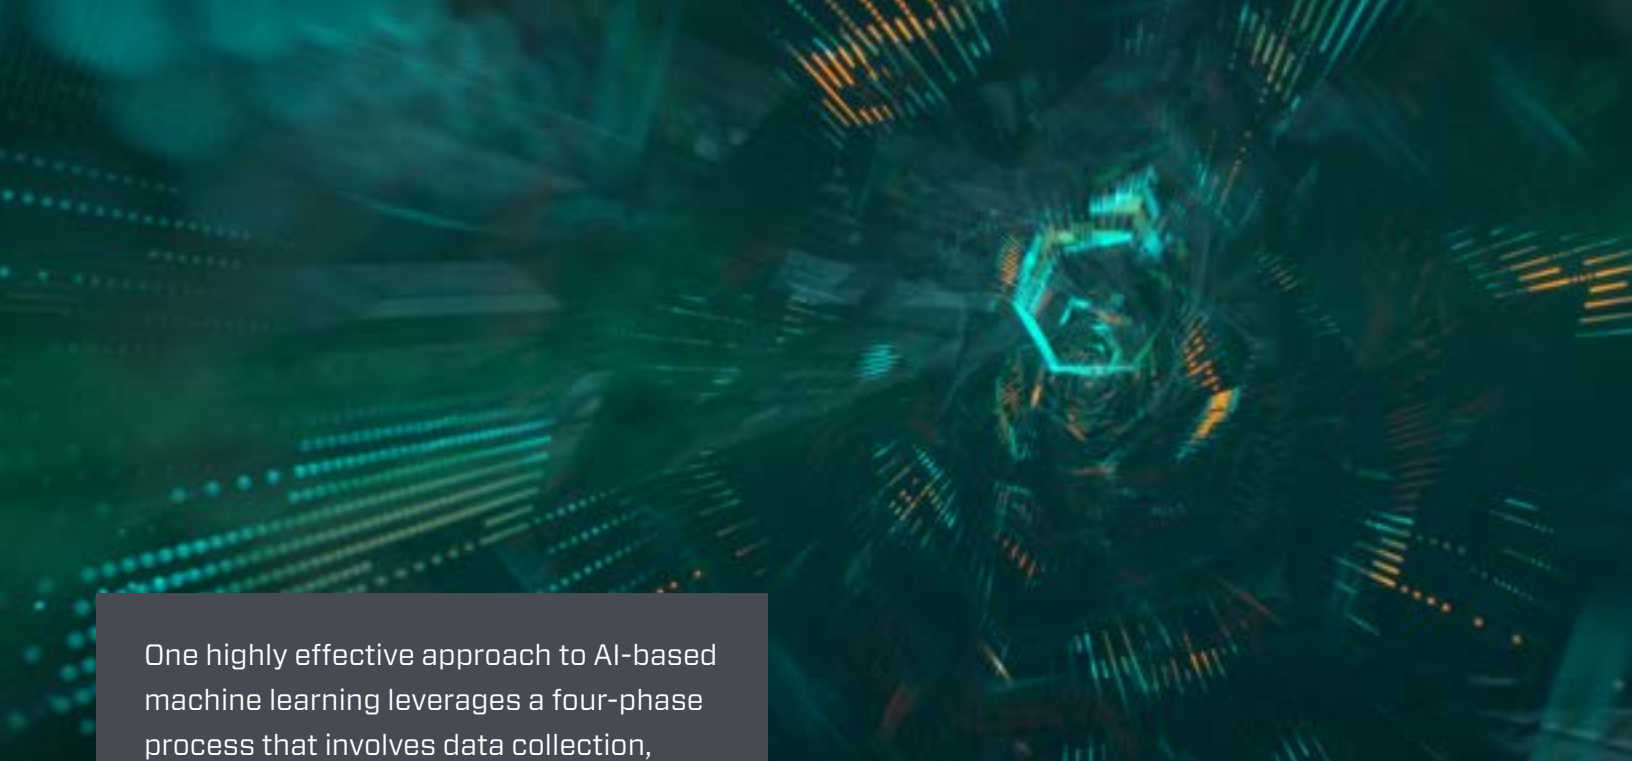
The idea behind these endpoint detection and response (EDR) tools is that files, programs, or activities that are suspected of having a malicious intent can be detected quickly enough so that a human can be alerted to the situation and take action in time to minimize the damage. And while many solutions have the ability to auto-remediate — for example, dropping a user's network connection if malware is suspected to be on the endpoint — organizations are often leery of taking action without human approval because of high false-positive rates. What's more, in situations such as with advanced persistent threats (APTs), detection and alerting can take place days, weeks, or even months after the initial infection occurs.

Detection-based EDR tools can have significant drawbacks. First, there is the obvious shortcoming of only discovering malicious activity after it has already gained a foothold. Attackers are developing very sophisticated code that is able to take evasive measures, so a failure to prevent it from gaining access to an endpoint in the first place is a very risky proposition.

Beyond that, most of these solutions require capturing, storing, and analyzing vast amounts of information pertaining to user and endpoint activities. Quite often this data is moved to the solution vendor's cloud for processing, which can create issues by using too much bandwidth to send so much data to the cloud, running up tremendous costs for data storage, and sending sensitive information to a third-party vendor. If the data isn't sent to the cloud, then it's kept on-premises, which requires the user organization to maintain the infrastructure of servers or appliances and storage devices to host the security application and data.

Detection solutions churn out a large number of alerts that must be prioritized for investigation. Some alerts get ignored if they are deemed low priority, or if the

³ Brian Krebs, Krebs on Security, "Antivirus is Dead: Long Live Antivirus!", May 7, 2014



One highly effective approach to AI-based machine learning leverages a four-phase process that involves data collection, extraction of attributes, learning, and file classification.

investigative team is stretched thin. Investigations into events require skilled security analysts, which are in short supply, and are expensive to employ. Moreover, large and complex rule bases are required to tune the detection solutions to create correlations among suspicious activities and to prevent false-negatives and minimize false-positives.

In short, endpoint protection based almost entirely on the detection of malicious activity is too complex due to all the rules to tune the system, generates too many alerts to realistically deal with, is resource-intensive in terms of bandwidth and storage and/or on-premises infrastructure, and requires too much effort by skilled security analysts. Even worse, detection and remediation might occur too late to prevent damage to the endpoint or the broader network.

Turning EDR on Its Head: The Rise of Prevention-Based EDR

The next generation of EDR solution is focused once again on prevention rather than detection; i.e., not even allowing malicious activity to execute on the endpoint at all, rather than trying to quickly detect when it does execute. The new approach does not use human-created file signatures at all. Instead, it uses artificial intelligence (AI) that is based on machine learning to automatically — without human intervention — distinguish good (benign) files

or activity from bad (malicious) files or activity based on mathematical risk factors. Once this good/bad classification is made, then it's possible to teach a machine to make the appropriate disposition decisions on these files in real time.

One highly effective approach to AI-based machine learning leverages a four-phase process that involves data collection, extraction of attributes, learning, and file classification. Here's a quick overview of how it can be applied to prevention-based EDR.

Building a New Type of AV Engine

The first phase starts with the collection of hundreds of millions of files of specific types, including executables, PDFs, Microsoft Word documents, Java, Flash, and so on. The files come from industry feeds, proprietary organizational repositories, live inputs from active computers, and various other sources. Once these files are collected, each one is reviewed and placed into one of three categories: known and verified as valid, known and verified as malicious, and unknown. It's important to the follow-on phases of the process that the categorization of the files be accurate.

The next phase is the extraction of file attributes. This process leverages the compute capacity of machines and data mining techniques to identify the broadest possible set of characteristics of a file based on the file type (.exe, .dll, .com, .pdf, .java, .doc, .xls, .ppt, etc.). The result of this attribute identification and extraction process is the creation of a file genome, very similar to that used by biologists to create a human genome. This genome is then used as the basis for which mathematical models can be created to determine expected characteristics

of files, much like human DNA analysis helps determine characteristics and behaviors of biological cells. The file attributes are converted to numerical values that can be used in statistical models.

Leveraging the millions of attributes of files identified in extraction, mathematicians then develop statistical models that accurately predict whether a file is benign or malicious. For each and every file, thousands of attributes are analyzed to differentiate between legitimate files and malware. The models can identify malware at an unprecedented level of accuracy, including exploits that target vulnerabilities that aren't yet publicly known (i.e., pre-zero-day vulnerabilities).

In the final phase, files which are unknown, such as files that have never been seen or analyzed before, are further classified to yield a confidence score which can be used to weigh decisions around what action to take on a specific file: block, quarantine, monitor, or analyze further.

This entire process is built into a self-contained engine that is small enough and fast enough to operate on individual endpoint computers, whether they are online or offline. Thus, all files and activities that attempt to execute on the endpoint are taken into the engine, evaluated, and deemed benign or malicious in mere milliseconds. If they are malicious, they are prevented from running, pure and simple.

Prevention Is More Than 99% Effective

“When artificial intelligence and machine learning are applied to the task of analyzing files on an endpoint, and files determined to be malicious are prevented from opening or running, the results are more than 99% effective in preventing malicious attacks at the endpoint.”⁴

In addition to the high efficacy of AI-based prevention, this approach reduces data collection and storage needs. The entire prevention process can be self-contained on the endpoint, so there is no need to send files to an appliance or to the cloud for analysis. In fact, the prevention process can work when the endpoint is offline because the process is totally local to the endpoint.

This approach reduces the burden on security analysts because the good/bad decision is automated, as is the course of action to take when a file is deemed to be malicious. With this type of solution having an accuracy

rate exceeding 99%⁵, organizations can feel comfortable allowing the automated disposition of malicious files. In addition, security analysts have fewer rules to develop, and fewer alerts to which they must respond.

AI-based prevention techniques are the future of endpoint protection. “Enterprises need a way to predict attacks and streamline the threat hunting and incident response workflows,” according to Doug Cahill, senior analyst, cybersecurity, Enterprise Strategy Group. “With the explosive rate of malware growth and other threat vectors, an AI-based solution that automates time-consuming parts of the threat hunting and incident response workflow is important.”

Prevention-Based EDR Makes Detection Much Better

Preventing 99.1%⁵ of malicious attacks at the endpoint is a tremendous feat, but a very small percentage of threats may still get by the prevention technology, hence the saying, “A cybersecurity team needs to be successful all the time, but an attacker only needs to be successful once.” Thus, even with a highly effective prevention strategy in place, organizations still need detection capabilities to clean up residual artifacts, or to find advanced threats that cannot be prevented pre-execution.

The most powerful solution, then, is to pair the new approach to prevention discussed above with an integrated detection and response capability. This one-two punch is the most effective way to get comprehensive endpoint protection.

As previously stated, detecting malicious activity can be complicated and costly, but if the total volume of what remains to be detected is vastly reduced through prevention measures first, then detection is much easier and far less expensive. The security analyst's job takes on a different mission when there are far fewer events to investigate. Analysts can transition from reacting to an event that has already happened to hunting for threats in the environment.

When the prevention and detection components are truly integrated, not simply complementary, but separate tools, they share underlying data that enables root cause analysis and threat hunting. This is best illustrated with a use case example.

⁴ NSS Labs Advanced Endpoint Protection: Cylance Security Value Map, April 2018

⁵ NSS Labs Advanced Endpoint Protection: Cylance Security Value Map, April 2018

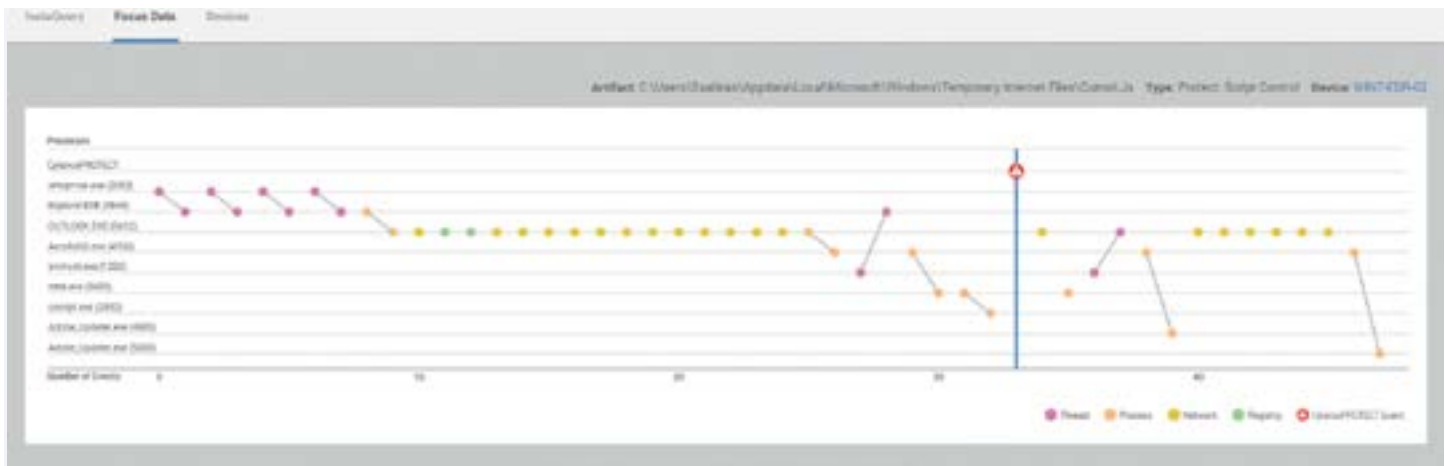


Figure 2: Tracking events for root cause analysis

Prevention and Detection: Two Sides of the Same Coin

The example in Figure 2 shows how the prevention solution CylancePROTECT® and the integrated detection component CylanceOPTICS™ work together to prevent an attack and allow an analyst to do root cause analysis to determine if any threat remains.

In this image, the symbol in the red circle at the right side represents the point at which CylancePROTECT interrogated an executable file on the endpoint and deemed it to be malicious. The executable was stopped — not permitted to run — so the risk was mitigated. However, a security analyst would want to get an understanding of where that threat came from, and whether any risk persists. This is when the detection capabilities of the integrated solution come into play. Since all the underlying data is readily shared between the prevention and detection components, the investigation can proceed seamlessly by simply pivoting to view what happened prior to prevention.

All the small dots and lines on the timeline in Figure 2 represent things that happened before the malicious executable was prevented from running. These artifacts are files, processes, network activities, registry changes, etc. The files and processes involved are listed vertically on the left side of the timeline. On their own, they weren't considered malicious and so they weren't stopped. However, together they tell a story of how the malicious executable got to the point of attempting to run. By analyzing these points, a security analyst can learn how to close vulnerabilities in the environment and improve the overall security posture.

On a live dashboard, the security analyst can hover over any of the dots to get more details. For example, the analyst might see that the malicious program opened a

browser session, ran Windows Task Manager, visited an IP address, used an old version of Adobe Reader, etc. This is a nice cache of information for the analyst to use to conduct further hunting, perhaps to see if any other machines in the environment have that same old and vulnerable version of Adobe Reader installed. The BlackBerry Cylance dashboard can pivot to search for and view forensic data in numerous ways.

Going from a Needle in the Haystack To a Needle in a Matchbox

Root cause analysis helps shore up defenses; however, another critical capability for organizations is the ability to hunt for threats that are lying in wait on their endpoints. Completion of this task requires the right combination of endpoint visibility, endpoint data, and search functionality.

With consistent visibility to the endpoints, analysts can hunt for threats using indicators of compromise as well as other relevant terms. When the search is initiated, the security solution communicates with the endpoints, searches through data stored on the endpoints, and returns any matching results. The analyst can then pinpoint the search and gather technical details based on a chosen type of artifact. This type of capability brings sophisticated threat hunting to the masses, providing instant access to the forensically relevant data collected from endpoints to identify potential security issues.

Quick Response Is Key

Going beyond identification, the organization needs the ability to respond quickly to protect the business from the fallout of a widespread successful attack. As an example, CylanceOPTICS has several built-in incident response options that enable an analyst to take action as soon as it's determined that a process, executable, file, or endpoint may be harmful to the environment. The security analyst can opt to download the suspicious file to complete a

deeper investigation with third-party tools; globally quarantine the suspicious item, restricting any endpoint in the environment from interacting with the item; or lockdown an endpoint that has been determined to be harmful to the environment for some reason.

Superior Prevention Makes Focused Detection Possible

It's the power of AI-based prevention, which stops the vast majority of malicious attacks before they can execute, that allows the integrated detection component to be an easy-to-use yet valuable forensic tool for the security analyst. Prevention removes the noise of too many alerts so that the security analyst can be very focused on hunting for and detecting residual threats.

Learn More About Deploying Prevention Plus Integrated Detection To Secure Your Environment

CylancePROTECT is an award-winning product that provides enterprise endpoint security by preventing advanced persistent threats and malware from executing. The product takes a radically different technological approach to cybersecurity, employing artificial intelligence to analyze the DNA of files before they execute. BlackBerry Cylance integrates a deep understanding of attackers and attack vectors with the most sophisticated AI and algorithmic science to neutralize virtually all malicious files before they execute. CylancePROTECT doesn't require a cloud connection or frequent updates and uses a fraction of the system resources associated with typical antivirus and endpoint security software.

CylanceOPTICS is an endpoint detection and response (EDR) component that augments the prevention provided by CylancePROTECT by enabling security analysts to:

- Perform root cause analysis for any threat blocked by CylancePROTECT or any other artifact found on an endpoint deemed important
- Proactively search endpoints for signs of threats such as threat hunting
- Take decisive action when a security incident, or potential incident, is identified

Unlike other EDR products that require significant investment in on-premises infrastructure, or that force an organization to stream data continuously to a cloud environment for storage and analysis, CylanceOPTICS is designed to run on the endpoint, using the existing CylancePROTECT agent for collection and a local database for storage.

CylanceOPTICS, in conjunction with CylancePROTECT, provides AI-based predictive threat detection and prevention solutions, enabling security analysts to take a more strategic approach to securing their business.

Visit the Cylance website at www.cylance.com to learn more.

About BlackBerry Cylance

BlackBerry Cylance develops artificial intelligence to deliver prevention-first, predictive security products and smart, simple, secure solutions that change how organizations approach endpoint security. BlackBerry Cylance provides full-spectrum predictive threat prevention and visibility across the enterprise to combat the most notorious and advanced cybersecurity attacks, fortifying endpoints to promote security hygiene in the security operations center, throughout global networks, and even on employees' home networks. With AI-based malware prevention, threat hunting, automated detection and response, and expert security services, BlackBerry Cylance protects the endpoint without increasing staff workload or costs.

 **BlackBerry**

CYLANCE

+1-844-CYLANCE

sales@cylance.com

www.cylance.com

