Cylance® vs. Traditional Security Approaches

Understanding Drives Informed Decisions

CYLANCE

Contents

Executive Summary – Cylance

Today's advanced cyber threats target every computer and mobile device, including enterprise endpoints, especially those that make up critical infrastructure like industrial control systems and embedded devices that control much of our physical world. The modern computing landscape consists of a complex array of physical, mobile, cloud, and virtual computing, creating a vast attack surface. Meanwhile, the cybersecurity industry is prolific with defense-in-depth security technologies, despite a threat landscape that remains highly dynamic, sophisticated, and automated.

Cylance, however, takes a unique and innovative approach of using real-time, mathematical, and machine learning threat analysis to solve this problem at the endpoint for organizations, governments, and end-users worldwide.

Cylance uses artificial intelligence (AI) to deliver security solutions that change how organizations, governments, and end-users approach endpoint security. Cylance's security solutions combine AI-driven predictive prevention with dynamic threat detection and response to deliver full spectrum threat prevention and threat visibility across the enterprise.

Cylance's next-generation antivirus product, CylancePROTECT®, delivers industry-leading malware prevention powered by AI, combined with application and script control, memory protection, and device policy enforcement in order to prevent successful cyber attacks.

Without the use of signatures or the need to stream data to the cloud, CylancePROTECT combats common threats such as malware, ransomware, fileless malware, malicious scripts, weaponized docs, and many other attack vectors, no matter where the endpoint resides. With unmatched effectiveness, ease of use, and minimal system impact, CylancePROTECT is the best way to prevent both known and unknown attacks before they can execute.

Augmenting CylancePROTECT prevention, CylanceOPTICS™ is an endpoint detection and response (EDR) component that enables easy root cause analysis, threat hunting, and automated threat detection and response. Unlike other EDR products that require organizations to make a significant investment in on-premises infrastructure and/or stream data to the cloud continuously, and employ highly-skilled security resources, CylanceOPTICS is designed to automate threat detection and response tasks using existing resources.

Cylance's Consulting Services provide pre-attack penetration and vulnerability testing, compromise assessments, and post-attack incident containment using AI-driven Cylance technology.

This powerful combination of predictive threat prevention, detection, response, and expert services allows Cylance to protect endpoints without requiring clients to increase their staff workload or costs.

Past	Present			Future
				
AV	Hips / Anti-Exploitation	Sandboxing	Isolation	AI
Humans Needed	Specialized Humans Needed Post-Execution: REACTIVE			No Humans Pre-Execution: PREDICTIVE

How Does Traditional AV Work?

> *"Antivirus software, or anti-virus software (abbreviated to AV software), also known as anti-malware, is a computer program used to prevent, detect, and remove malware. Antivirus software was originally developed to detect and remove computer viruses, hence the name. However, with the proliferation of other kinds of malware, antivirus software started to provide protection from other computer threats."* — *Wikipedia.*

Life, as the saying goes, is all about choices. Traditional antivirus (AV) products have claimed the lion's share of the security market for years. However, as the decades have rolled by, attackers' abilities to invent techniques, tactics, and procedures have improved exponentially. Threats like fileless malware, which writes nothing to disk, cannot be caught by signatures, so traditional AV is becoming less and less effective.

Even malware development itself has evolved. Attackers now run their own QA labs, use nation-state attack techniques (remember WannaCry?), commercial penetration tools, and validate their new malware samples using bootleg multi-engine scanning sites to see if they are detected. If so, they modify the code and try again until it passes under traditional AV products.

We need new ways of preventing the execution of malicious code – be it binaries, fileless, script-based, or whatever else is coming over the horizon. Let's look at how traditional AV products work in order to understand why it's so easy for the bad guys to bypass them.

How Traditional AV Works

1. Pattern Matching - Byte Matching

The first approach is pattern matching via signatures. Pattern matching is used to check a sequence of tokens for the presence of the constituents (parts) of a pattern. In contrast to the flexibility offered by pattern recognition, the match must be absolutely exact.

A signature is the digital fingerprint of a piece of malware. It's a unique string of bits, a binary pattern representing the malware. Each time a traditional AV product encounters a new file, the AV product looks through its signature list and asks, "does this byte in the signature match this byte in the file?" If it does, it moves on and checks the next byte. It continues through the whole file in this way. If every byte of the file matches every byte in one of its signatures, exactly, it flags the file as malware.

There are some optimizations to this process in place, and ways to only match certain parts of bytes, but for all intents and purposes, this is how traditional AV works when matching a signature.

However, *attackers easily bypass signatures by mutating, obfuscating, or otherwise changing up the code in their malware.* Herein lays the biggest weakness of signatures: if so much as a single byte is changed in any of the signature's important values, then the signature no longer matches the malware. It becomes toothless, to the extent that a single recompilation with different strings easily evades most signature detection algorithms.

All signature-based AV products operate pretty much the same way, and now this weakness is well known to the adversary.

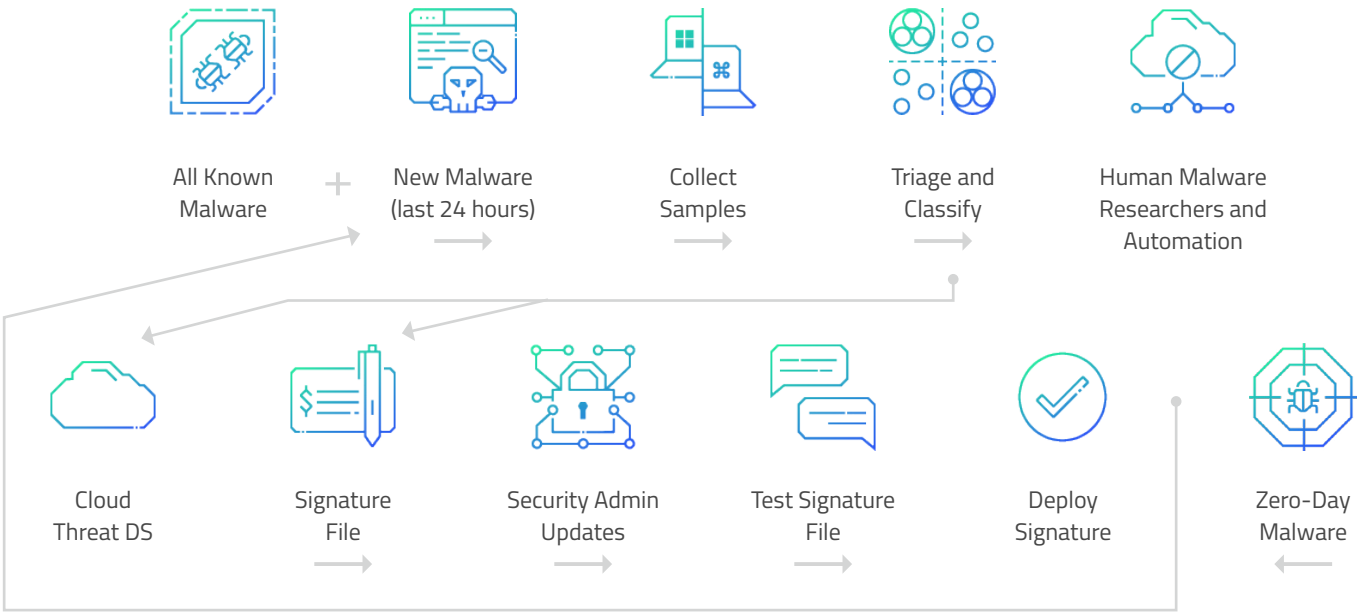

2. Heuristic Approaches

The second approach is heuristics. The AV looks at loose properties of the file, such as the file's size, whether it looks like it's using a set of dangerous functions, or whether it has abnormal permissions. With heuristic approaches, the AV matches things that aren't in the code directly. One example of how this might work is by asking the following questions of the file:

- Does the executable import VirtualAlloc?
- Is the executable greater than 30KB and less than 75KB?
- Does the executable have a section whose permissions are read, write, and execute?
  - If these things are true, then it is a virus.

With heuristic matches, there may be up to ten rules in place, but it's no more complicated than having more rules than the above illustrates in the real world. Traditional AV relies heavily on this set of rules to convict a sample. This is where the attackers get the last laugh. For an attacker, bypassing AV products that use heuristics requires knowledge of just a single feature; changing that one feature breaks the entire detection.

For example, adding random data to malware can easily bypass heuristic approaches. *An attacker only needs to change one tiny property of the file, and when heuristics can't match it, they win.*

3. Behavioural Analysis

A third approach is behavioural analysis, which is like heuristics and targets the actual behaviour exhibited by malware. Behavioural analysis looks at questions such as:

- What is the file doing on a file system level?
- What is the file doing on a registry level?
- What is the file doing on a process level?
- What is the file doing on a network level?

The trouble with this approach is that *the malware must run first before the AV product can detect it.*

4. Hash-Based Approaches

The fourth approach is hash matching. The AV calculates hashes over different parts of the file, then takes a hash over a certain area of the executables (MD5, SHA256, CRC32). The AV analyzes that hash to see if it matches the hash of a known virus. If it does, then the AV determines that the file is a virus.

Sometimes, engines will take many different hashes across the binary and see if any of them match. For instance, it may cut up the file into 1024-byte chunks and take the hashes of all of them and see if any of them match a virus.

The problem with hashes is, once again, if a single bit gets changed in any of the areas used to generate the hashes, the hashes produced are wildly different.

An attacker only needs to change one bit of the file, and it is game over for the AV.

Common Reasons People Buy Traditional AV

- Strong brand presence and market share.
- Maintaining existing security infrastructure is often seen as easier than removing and replacing with new technology.
- Traditional AV security companies often provide security components outside the endpoint space, allowing companies to buy all they need at one time, adding to the infrastructure layers.

The Weaknesses of Traditional AV

- Relies on signature files and requires daily updates.
- Performance impact on the endpoint is high, leading to higher help desk tickets and less productivity by users.
- Offline protection is limited due to the reliance on cloud-based protection capabilities.
- Leverages behavioral/heuristic rules in an attempt to take out zero-day malware — this is not effective because the malware has to run first.

Summary Table of Traditional Antivirus

Strengths	Weaknesses
Usually low false positives as signatures are very specific.	Specific signatures are easily bypassed.
Very strong when the threat is known to the vendor.	Human-driven approach to creating signature updates, too many unknowns to keep up to date.
	Cannot defend against threats that are still completely unknown today.
	High system resources utilization.
	Execution-based analysis carries a high risk of malware remaining undetected during execution and can only be seen as a means of reducing risk, not prevention.
	Demands frequent scans.
	Requires cloud connectivity all the time to do unknown hash lookups.

Cylance Is Different from Traditional AV

Cylance uses a fundamentally different, signatureless approach to traditional AV that leverages artificial intelligence and machine learning to prevent malicious code from executing. Instead of a simple, straightforward, step-based process, Cylance's algorithm is a deep neural network, a complex branched system that feeds back into itself and learns from the past to infer the future.

Here at Cylance, we have studied billions of files. In total, we're currently measuring 1.4 million features, which are extrapolated for analysis and used to train our machine learning models. Simple examples of these features could be the file length, the use of digital certificates (which are often legitimate but can be stolen), whether the file is using a packer, and the complexity or entropy of the file. But, instead of looking at five or ten features to make the decision about whether a file is good or bad, our machine learning algorithm looks at 1.4 million.

Malware vs. the Cylance Score

Each one of those features can be represented as a layer in our deep learning network. The presence or absence (and the weight) of a feature determines the path through the layers to reach a decision.

While we can make an analogy to an enormous, complex maze, the neural network we have designed is a deep, branched structure that outputs a confidence score. The higher the confidence score, the more certain we are that a sample is malicious – despite our model never having seen it before.

> *This is the basis for building a predictive model, learning from massive amounts of past data to identify malware and threats that may not even yet exist.*

To reverse-engineer a Cylance detection, the attacker would have to successfully backtrack through that entangled network of nodes processing features – a feat as impossible as trying to solve a maze with several million rows by making completely random turns.

How Cylance Uses Machine Learning Differently from Traditional AV Companies

When other vendors say they use machine learning (ML), what they typically mean is they are using it in one or more of the following ways:

- They use an ML algorithm to scan malicious software, generally from the cloud.
- They have the ML algorithm generate a signature, heuristic, or hash, as described above.
- They then have humans vet the resulting signature, heuristic, or hash to make sure that nothing non-malicious is blocked.

When Cylance says we use machine learning to detect and prevent malicious applications from executing, we mean that very literally.

> *Unlike other vendors who have added ML to the number of methods they use to identify malware (signatures, behaviors, sandboxes, etc.), Cylance uses ML exclusively for malware threat prevention. Our ML has been trained to make these decisions in milliseconds, locally on each endpoint.*

This allows organizations to eliminate the time-consuming and resource-intensive tasks associated with maintaining signature-based solutions.

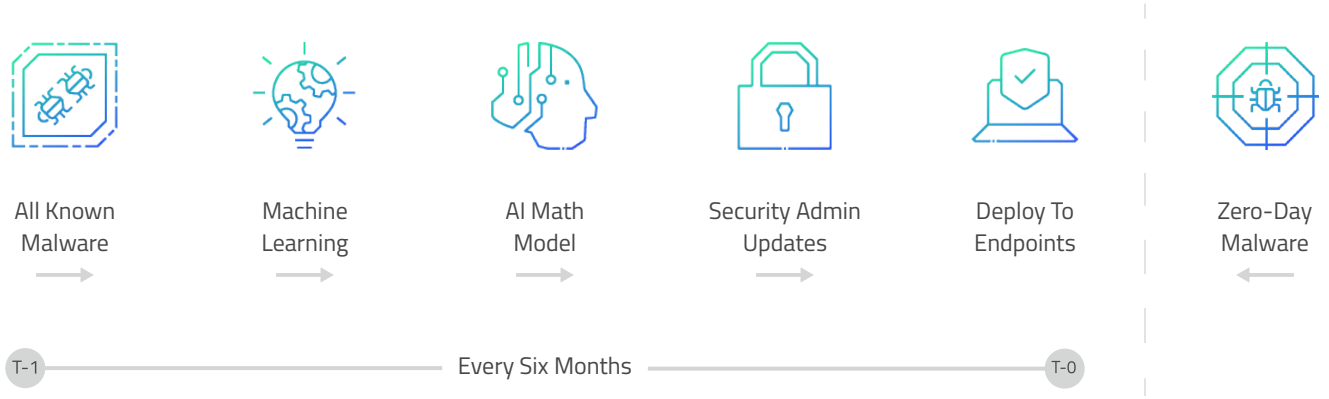

Why Do Companies Choose Cylance?

Cylance provides numerous advantages for customers seeking highly-accurate malware prevention, comprehensive attack protection, and robust EDR capabilities, all managed by a single management console.

Effectiveness

- Consistently prevents the execution of previously unknown, known, and custom-crafted malware and payloads without the need for signatures.
- Prevents execution of unauthorized scripts.
- Provides superior malware prevention accuracy whether online or offline.
- Leverages a combination of behavioral rules and AI-based ML models for EDR threat detection.

Simplicity

- Replaces, or if necessary, augments, existing anti-malware solutions (augmentation is only recommended for temporary/transition purposes for maximum solution value).
- CylancePROTECT is a Microsoft approved AV.
- Simple to deploy globally using GPO, login script, or third-party software management packages.
- Automate response actions to behavioral threats without human intervention.
- Updates are easy and infrequent (current model is 14 months old).

Performance

- Non-disruptive to the environment. No reboot required on workstations or servers.
- Improved end-user experience. A fully autonomous agent with a reasonable system resource footprint:
  - Eliminates the need for regular hard disk scans.
  - Reduces aggregate CPU and memory usage.
- Lowers network bandwidth usage by eliminating legacy solution DAT file distribution challenges.
- Returns performance to VDI infrastructure while providing a more complete guest OS-based anti-malware solution compared to hypervisor-level malware-only scanning.
- Enterprise-wide attack indicator queries returned in seconds.

Conclusion

An Ounce of Prevention Is Worth a Pound of Cure

The word prevention is used alongside words like detect and respond to compare the capabilities of a solution that stops a threat before it has impact, versus one that detects a threat post-impact, then mitigates the damage. But, not all prevention is equal. The goal of prevention is to stop a threat in its tracks, preferably pre-execution before it can cause damage. What is delivered by Cylance is predictive prevention — a technology that can stop unknown threats without signatures, heuristics, behavior, reputation, cloud look-ups, malware analysis, human interaction, or any time delay. This was demonstrated by an SE Labs report and by Cylance testing new malware attacks against older mathematical models.

> *"Not only does the data demonstrate that CylancePROTECT (agent v1300, model May 2015) was capable of preventing threats that did not exist at the time the AI model was 'trained', but it provides an insight into how far ahead in time it could be effective without new knowledge. In practical terms, this indicates that regular updates to the product are not always needed, although we would expect Cylance to develop and deploy newly-trained models over time, simply because product development is an ongoing process and machine learning continues to take into account new threats to predict future ones."*

Predictive Advantage Over Major Malware Campaigns (in Months)

Threat Family	Predictive Advantage
Bad Rabbit	29
Cerber	30
GhostAdmin	24
GoldenEye	23
Locky	20
NotPetya	25
Petya	26
Reyptson	27
WannaCry	24

Predictive Malware Response Test, SE Labs — March 2018

+1-844-CYLANCE
sales@cylance.com
www.cylance.com

20181130-1418