Symantec by Broadcom

WHITE PAPER

Putting Adaptive Protection to the Test

Evaluations by MRG Effitas show how Adaptive Protection can block potential Living Off the Land attacks sooner

WHITE PAPER

Putting Adaptive Protection to the Test

Evaluations by MRG Effitas show how Adaptive Protection can block potential Living Off the Land

TABLE OF CONTENTS

For years, malicious actors were known for writing custom compiled programs to execute attacks. But lately, legitimate business applications are attracting threat actors looking to infiltrate corporate environments and levy serious blows while hiding in plain sight.

In The 2024 Ransomware Threat Landscape report, the Symantec Threat Hunter team noted a sizable presence of "living off the land" (LOTL) tools, where threat actors use either operating system features or legitimate tools to deliver sophisticated attacks. Because the Windows operating system provides a rich collection of powerful tools which obviate the need to deliver custom binaries, those LOTL tools have become a favorite target of attackers. PowerShell, WScript and CScript are fully functional languages that are equally capable of performing any actions that a compiled language can.

Recently, the definition of LOTL tools has expanded to include commercial third-party tools, such AnyDesk, Atera, Splashtop and ConnectWise, which are legitimate business applications but are also utilized by malicious actors. The fact that these otherwise legitimate tools are now being leveraged as vessels for cyber attacks presents a unique problem.

Of course, system administrators often require these tools to perform a host of activities to monitor and maintain systems. So how does an administrator figure out what "normal" behavior is—and which unusual behaviors could signal potentially malicious activity?

Now, Adaptive Protection, a powerful feature within Symantec Endpoint Protection, provides that capability. This white paper describes what Adaptive Protection does, how it works and how security teams benefit from it.

This paper also includes a report from independent testing firm MRG Effitas, whose tests reveal how this feature can reduce the attack surface of an organization and lessen the burden on security teams and the tools they use.

Introducing Adaptive Protection

No two enterprise environments are the same. In fact, no two divisions within the same organization are the same. Each company–or division–has its own way of doing things and its own combination of tools used to accomplish their goals. Those traits establish how that organization or division uses those tools within the scope of its business.

Simplifying a Complicated Problem

It would be extraordinarily difficult for most organizations to identify and block out-of-policy actions across their entire environment—while also allowing certain trusted groups or users to break policy if those actions fall within the scope of their roles. As a policy-driven technology, Adaptive Protection simplifies this complicated challenge. It allows or denies individual actions based on the policy that is in effect. It makes no judgments regarding the intent of any actions or actors which are blocked. However, by blocking out-of-policy behaviors, it can alert security teams to attempted actions that may indicate malicious activity—such as repeated attempts to run PowerShell when such an action falls outside policy. While it's up to security teams to make that determination, Adaptive Protection makes their job easier by shrinking their attack surface and surfacing anomalous behaviors that may indicate an LOTL attack.

Adaptive Protection learns those individual traits and blocks combinations which never occur. Moreover, it employs behavioral analytics to automatically learn and apply exceptions to cover the usage it does observe, allowing it to block items outside normal usage without impacting the organization's normal course of business.

Adaptive Protection employs extensive and constantly updated endpoint telemetry, granular policy controls and powerful global threat telemetry. This combination gives security teams greater network visibility, shrinks their attack surface, and enhances their ability to assess when unusual actions or behaviors occur that might signal malicious behavior (such as LOTL attacks) or otherwise unauthorized actions by end users.

Teams with this Symantec Endpoint Security feature can:

- Easily monitor and identify behaviors which are normal within the organization or subsets of the organization
- Craft policies that allow for these behaviors while preventing behaviors outside their norm—including unusual variations of normal behaviors
- Shrink their attack surface by automatically disallowing out-of-policy actions and by filtering out known malicious items via Symantec Endpoint Security's protections

How it Works

Adaptive Protection monitors a wide variety of behaviors and apps that are legitimately used to perform administrative tasks. However, most admins only use a fraction of these behaviors. Adaptive Protection allows an admin to perform the tasks they ordinarily do while preventing behaviors outside the scope of normal business. And by automatically allowing behaviors and apps that are observed to be part of the normal course of business for an organization, Adaptive Control is largely transparent to users—unless they attempt to break policy, either by mistake or on purpose.

Adaptive Protection can block more than 450 individual actions, such as Microsoft Word running PowerShell or CScript running WScript. Since Adaptive Protection begins by learning what "normal" behaviors are for the organization it is deployed to protect, these actions are placed into Monitor mode in an enterprise environment, where each time these actions are seen they will be flagged, but not blocked. After sufficient time—either 90, 180 or 365 days—an administrator can see if these actions have ever been seen in their company or within a group in their company. For actions that have never been seen across the Monitoring time period, they can be set to Deny without any fear of current business impact. Setting an action to Deny means it will be blocked from occurring. (For example, if Word has never been seen running PowerShell, it should not be allowed.)

Over time, a unique policy is created which, while allowing legitimate actions, also blocks legitimate actions that are *outside normal usage*. This reduces the attack surface, taking tools away from attackers while being invisible and undisruptive to end users.

About MRG Effitas

Among the reasons we chose MRG for this test is their proven track record of unparalleled expertise in finding real-world samples that thoroughly exercise security products. This, coupled with their highly advanced Tempus system, made them the obvious partner for this test. In the drive to protect businesses and home users from ever more advanced malicious threats, malware and viruses, MRG Effitas conducts innovative real-time testing and threat research to help security vendors and enterprises reduce their exposure to cyber risk and close the protection gap. A member of the Anti-Malware Testing Standards Organization (AMTSO), MRG Effitas is trusted by major vendors to assess their products and to provide the technical expertise and insight they need to keep improving cyber protection in real time.

Putting it to the Test with MRG Effitas

You may be thinking: "This sounds great in theory, but does it actually work? Can you prove it?" The answer to both questions is yes.

In conjunction with noted independent security testing house MRG Effitas, we devised a test to demonstrate the value of Adaptive Protection. We provisioned six machines.

- One was a control system, having all policies in Monitor mode, but running the Symantec Endpoint Security (SES) full behavioral protection stack with the static portions switched off as they were detecting all of the threats as they arrived.
- The other five machines were provisioned, each with a different policy used by our customers. These five policies are already deployed on hundreds of thousands of endpoints in the real world without business impact. Each system represented a unique customer environment.

A number of real-world threats were found and sent into the system that was set up to represent real customer Adaptive Protection environments. We were looking for threats that were difficult to keep up to date with, in which a generic policy would block every variance of that behavior—and thus block that attack.

MRG Effitas' Perspective: Report from the Testing Group

Outline of the Test

Broadcom approached us with a unique request, asking for the development of a tailored test to assess and compare their product using various configurations. The primary objective of this test is to emphasize the ability for Adaptive Protection to block attacks early on in the attack chain by identifying and blocking actions that fall outside the normal use of legitimate tools for each specific customer environment.

On the test systems, we installed Symantec Endpoint Protection 14.3.10148.8000. Six distinct environment setups were available: one exhibiting default settings for regular tests with all Adaptive Protection policies in Monitor mode, and five more featuring unique policy settings for Adaptive Protection.

In our In the Wild (ITW) test, we use our malware-hunting experience to select the freshest and most appropriate samples for testing. Sample sources include our own threat feeds and other sites and honeypots. We avoid using popular feeds and well-known threats. We use malicious URLs as a delivery method and follow the entire lifecycle of a malware attack, from initial infection to achieving its ultimate objective.

To initiate the test, we use valid in-the-wild URLs to deliver malware to the victim machine. We use hardened virtual machines to ensure the test environment is undetectable to the malware.

We simulate normal user behavior during our ITW test, opening the malicious link in Chrome, downloading the sample and running it. In parallel, we execute the sample in an unprotected system, and we compare the results with the protected systems. We count a Miss only if we see that the malware achieves its goal (e.g., the ransomware encrypts the files on the disk).

TO INITIATE THE TEST, WE USED VALID IN-THE-WILD URLS TO DELIVER MALWARE TO THE VICTIM MACHINE

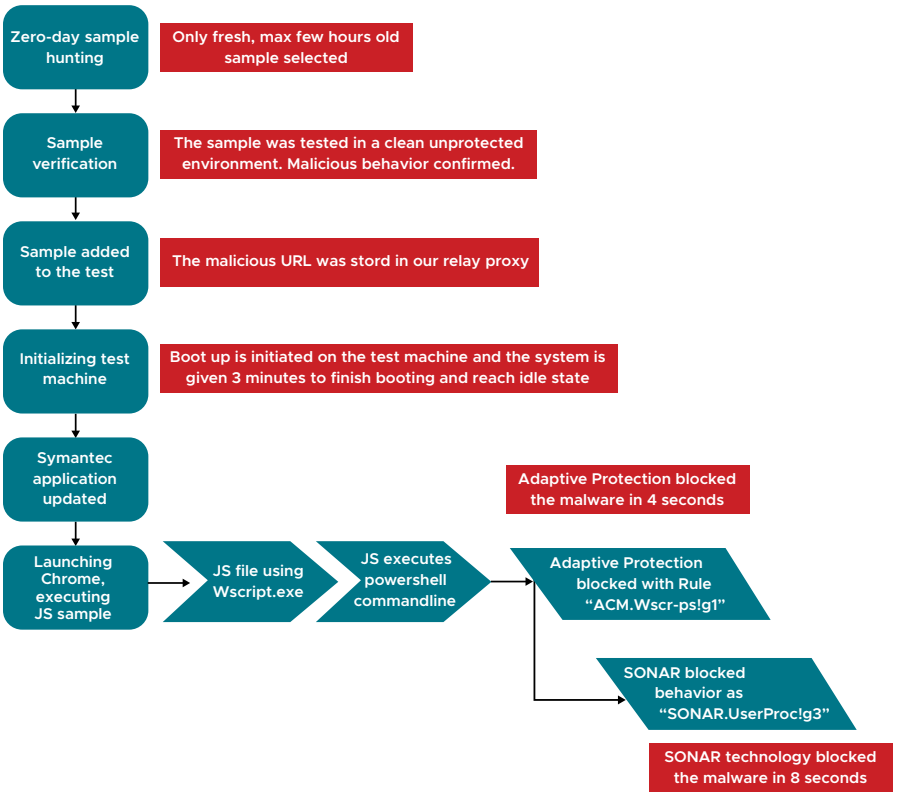

Detailed testing methodology

1. Windows 10 Enterprise 64-bit operating system is installed on a virtual machine, all updates are applied, and third-party applications installed and updated.
2. All possible virtual environment related artifacts are modified or hidden.
3. An image of the operating system is created.
4. A clone of the machine is made for each of the security applications used in the test.
5. An individual security application is installed using default settings or with the vendor provided, non-default settings.
6. A single binary executable or document, script, etc. is downloaded from its native URL and then executed in the clean, unprotected system. If the sample works, the sample is saved in a replay proxy to provide the same sample throughout the test.
7. The sample is selected for the test and tested in the systems where a security product is installed.
8. Boot up is initiated on the test machine and the system is given 3 minutes to finish booting and reach idle state.
9. Security product update is initiated via CLI command or the application's GUI.
10. Google Chrome is opened, and a sample is downloaded from its replay proxy URL.
11. Sample is started from Google Chrome's download bar.

TEMPUS, A REAL-TIME TEST ENVIRONMENT DEVELOPED BY MRG EFFITAS, EVALUATES EPP PRODUCTS AGAINST THE LATEST CYBER THREATS AND PROVIDES INSTANT ALERTS THROUGH EMAIL, SLACK, OR DISCORD WHEN MALWARE SAMPLES ARE MISSED BY EPP PRODUCTS

Based on different outcomes, each test case is marked either Blocked, Behavior blocked, AP Blocked or Missed.

- The test case is marked as *Blocked* if either the security application blocks the URL where the malicious binary was located, or the security application blocks the malicious binary whilst it is being downloaded to the machine.
- The test case is marked as *Behavior blocked* if the security application blocks the malicious binary when it is executed and either automatically blocks it or postpones its execution and warns the user that the file is malicious and awaiting user input.
- The test case is marked as *AP blocked* if the security application blocks the malicious binary after it is executed, and the unique Adaptive Protection type detection appears in the console.
- The test case is marked as *Missed* if the security application fails to block or behavior block the malicious sample during the test.

Tempus Real-Time Test Environment

Tempus, a real-time test environment developed by MRG Effitas, evaluates EPP products against the latest cyber threats and provides instant alerts through email, Slack or Discord when malware samples are missed by EPP products.

The online Tempus dashboard allows users to access sample data, including hashes, logs, screenshots and direct malware sample downloads. This facilitates quick and efficient updates to EPP, ensuring rapid protection against emerging threats.

In Tempus we run tests in parallel, testing the same protection product with the same threat samples but with different configurations at exactly the same time. This tests the efficacy of the product but also gives an accurate comparison and highlights differences between configurations.

In our testing environment, we use a hypervisor. Each virtual machine is configured with the same settings, running Windows 10 Professional, 64-bit edition. To maintain the integrity of our test results, we have modified all artifacts that could be attributed to the virtual environment.

The original ITW malware URLs usually live only a few hours. To minimize the risk that those URLs disappear during the test, we use our reply proxy, which stores the original IP address of the malicious URL.

The Results

In just a couple of days, over a half dozen cases were found where Adaptive Protection provided earlier protection against threats based solely on attack surface reduction. The systems with Adaptive Protection technology enabled detected the malware 4 seconds before the Default system. This highlights the efficiency of the Adaptive Protection technology, showcasing its ability to identify infections significantly earlier in the attack. Despite this time difference, it's crucial to underscore that both systems demonstrated robust protection, effectively shielding all users against potential data leaks.

Remember, the test used actual customer policies applied against in-the-wild threats. This approach was devised to show that an organization should be able to protect itself using a generic allow/deny approach to a list of behaviors. Let's delve in detail into one of the more interesting cases.

Detailed Case Study: AsyncRAT

Let's dive into a detailed analysis of one of the attacks executed by MRG: AsyncRAT. Symantec Adaptive Protection mitigation policies interrupted the attack at multiple levels as illustrated by the figure below. Despite the attack's complexity, Symantec Adaptive Protection policies made it impossible for this threat to execute. Instead, the attack was restricted at multiple steps early in the process.

Living Off the Land Attack Chain Disrupted By Adaptive Protection

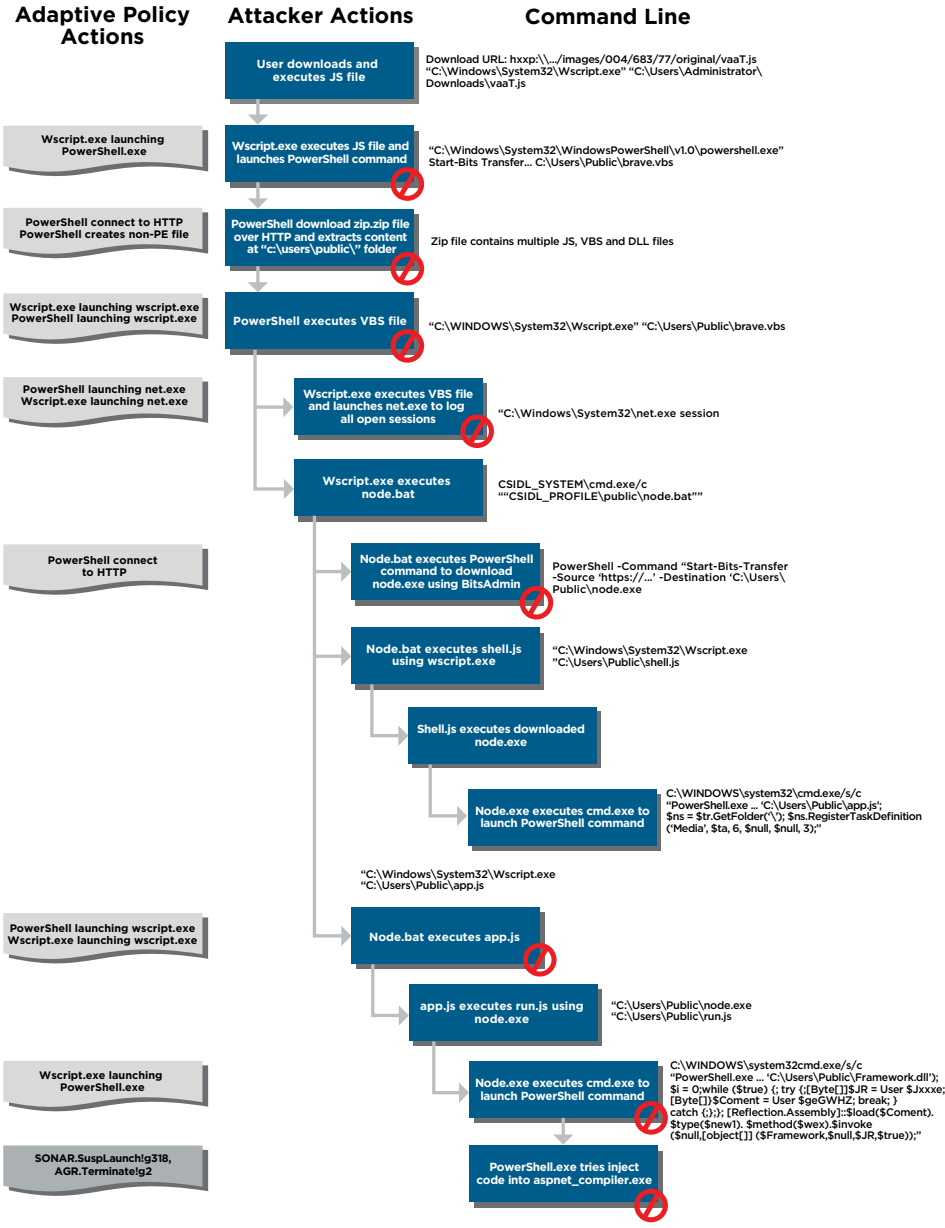

Test Summary

- Systems enabled with Adaptive Technology detected threats 4 seconds sooner than those without Adaptive Protection.
- Adaptive Protection blocked these attacks based on policy (and thus proactively), which is harder for attackers to circumvent than blocks based on behavioral signatures.
- Both systems—whether those with Adaptive Protection enabled or those with Symantec Endpoint Protection alone—demonstrated robust protection.

DESPITE THE ATTACK'S COMPLEXITY, SYMANTEC ADAPTIVE PROTECTION POLICIES MADE IT IMPOSSIBLE FOR THIS THREAT TO EXECUTE

Conclusion

While Symantec Endpoint Protection blocked all of the attacks, stopping those attacks happened sooner with Adaptive Protection. In fact, MRG's Tempus system measured that Adaptive Protection blocked the threats 4 seconds earlier than SEP's Behavioral system[1]—a significant timeframe in an environment where every second counts, and where stopping attacks earlier in the attack chain can prevent lateral movement, data loss and more. However, Adaptive Protection blocked these attacks via policy, and thus proactively. Moreover, these policies are significantly more difficult to circumvent than behavioral signatures.

The capabilities of MRG's Tempus system made the execution of this head-to-head test quite simple, and MRG's ability to acquire fresh, In the Wild samples to execute make the results highly relevant. Moreover, this also makes the results reflective of what a security analyst would expect from an actual deployment. Recall that the policies which were in effect are actual policies deployed to hundreds of thousands of real machines.

Adaptive Protection: Tested, Trusted and Proven

As attackers continue to ramp up their efforts, we at Symantec will continue to ramp up ours so we can equip organizations with the strongest solutions possible. These rigorous, independent tests by MRG Effitas demonstrate that Adaptive Protection works.

We proved that Adaptive Protection can:

- **Discover threats up to 4 seconds faster versus environments without Adaptive Protection enabled.** The lead time can mean the ability to shut down an attack before it reaches too far into an environment, protect mission-critical data or prevent a costly breach. When every second counts, this advantage is paramount.
- **Perform In the Wild with real-world threats.** LOTL techniques are designed to infiltrate environments and wreak havoc long before anyone notices. Adaptive Protection was tested with ITW threats using ITW customer policies, proving that it has real-world applications and viability for today's complex environments.
- **Interrupt attacks earlier in the attack chain, shifting the balance of power to organizations defending their data, users and assets.** Disarming attackers before they can burrow further into a system buys defenders precious time and more options.
- **Strengthen defenses in real time and reduce endpoint attack surfaces.** Easily managed, dynamic policies keep the focus on how legitimate tools are used within an organization—and blocks behaviors that may result from malicious intent.

[1] SEP's static technologies were disabled because they were detecting all of the threats as they arrived.

AS ATTACKERS CONTINUE TO RAMP UP THEIR EFFORTS, WE AT SYMANTEC WILL CONTINUE TO RAMP UP OURS SO WE CAN EQUIP ORGANIZATIONS WITH THE STRONGEST SOLUTIONS POSSIBLE. THESE RIGOROUS, INDEPENDENT TESTS BY MRG EFFITAS DEMONSTRATE THAT ADAPTIVE PROTECTION WORKS

For more on Symantec Adaptive Protection, explore these in-depth resources.

White Paper
Adaptive Protection, a Critical Capability in Disrupting Sophisticated Attacks

Product Brief
Symantec Endpoint Security: Adaptive Protection to Automatically Customize and Maximize Security

Video
Understanding Symantec's innovative Adaptive Protection technology

Symantec by Broadcom

For more information, visit our website at: www.broadcom.com

Symantec - White Paper - Broadcom & MRG Adaptive Protection September 30, 2024