Taming the Endpoint Chaos Within: A Review of Panda Security Adaptive Defense 360

Justin Henderson

A SANS Product Review

Taming the Endpoint Chaos Within: A Review of Panda Security Adaptive Defense 360

Written by **Justin Henderson**

March 2019

Sponsored by:

Panda Security

Introduction

To survive in a world dealing with automated malware and targeted adversarial attacks, organizations must move past traditional mindsets and implementations. Antivirus is insufficient, but so is focusing heavily on preventive controls. At the same time, too much focus on detection also introduces security risks. Organizations must establish a balance between protection mechanisms and detection. A successful implementation requires automated solutions that scale while simultaneously providing ease of administration for both preventive and detection capabilities.

A significant problem regarding endpoint suites is the failure to deploy and tune such solutions. In the 2018 SANS "Endpoint Protection and Response" survey, 50 percent of organizations had acquired next-gen antivirus, but 37 percent had not implemented the capabilities. A similar finding shows that 49 percent had malwareless attack detection, but 38 percent had not implemented the capabilities.[1] Standalone solutions such as endpoint protection platforms (EPPs) and endpoint detection and response (EDR) fail to be implemented fully. EPP solutions focus on preventive controls that ultimately

[1] "Endpoint Protection and Response: A SANS Survey," June 2018, www.sans.org/reading-room/whitepapers/analyst/endpoint-protection-response-survey-38460 [Registration required for access.]

are bypassed, whereas EDR focuses on detection controls that are for an after-the-fact response. Neither solution is comprehensive enough on its own. Gartner states that a convergence of the EPP and EDR is necessary and includes them as key parts of Gartner's Adaptive Security Architecture.[2]

The solution seems as if it requires more full-time hires and disparate point products. In reality, the answer lies in choosing an appropriate solution for the environment. An effective endpoint platform integrates into an organization to provide quick, meaningful results by:

- Preventing **known and unknown** malicious executables by systematically identifying the context, behaviors and attributes of the executable
- Allowing teams to **respond to meaningful alerts** by automatically handling malware and identifying more sophisticated attacks
- Helping security analysts, incident handlers and forensic teams identify targeted attacks by **providing high-fidelity information**
- Automating complex tasks and simplifying management to **maximize time spent** so the team can achieve higher effectiveness
- **Providing scalability** without the cost of complexity and labor

Within the search to identify a better endpoint solution, SANS had the opportunity to review a product from Panda Security—a business with a unique approach that scales easily while providing the benefits previously mentioned and solving the hurdles of standalone EPP or EDR solutions. In this paper, we examine:

- Panda Security's approach to endpoint protection, detection and response
- How to use the Panda Adaptive Defense 360 platform to prevent malware from executing and causing damage
- How information security teams can use Panda Adaptive Defense 360 to deploy preventive technologies while retaining insight into their environments
- Why prevention controls can and should combine detective controls and how Panda Adaptive Defense 360 applies each in conjunction with one another

SANS found Panda Adaptive Defense 360 to be easily deployable, with instant results in preventing malware and identifying targeted attacks. Within the platform, we found that tasks associated with large amounts of labor investment, such as tuning and patching, instead are automated or minimal. The solution brings synergy and success with groundbreaking preventive and detective capabilities.

Traditional endpoint protection solutions focused on preventing known malware and known malwareless or in-memory attack techniques still work today, but they are not effective enough for today's threats.

[2] "Build Adaptive Security Architecture into Your Organization," June 30, 2017, Rob van der Meulen, www.gartner.com/smarterwithgartner/build-adaptive-security-architecture-into-your-organization

Introducing the Panda Adaptive Defense 360 Platform

Before we can test any endpoint protection process, it is important to understand today's evolving malware landscape. See the sidebar, "Understanding the Endpoint Chaos," for a glimpse at where we are.

Understanding the Endpoint Chaos

Today's evolving malware landscape ranges from traditional worms and viruses to more modern phishing attacks and ransomware. See Figure 1 for a timeline of the changing malware landscape.

1980s	Viruses and worms spread across the internet.
1986	Obfuscation and encryption techniques are found in malware.
1990s	Polymorphic and oligomorphic viruses occur in the wild.
1999	Melissa virus introduces mass spreading via email.
2000	Malware heavily uses social engineering.
2001	Code Red worm spreads, and fileless malware attacks become better known.
2008	Conficker worm infects 15 million machines.
2010	Organization-targeted malware becomes more common.
2013	Ransomware makes an appearance.
2017	WannaCry ransomware massive attack causes $1 billion in direct and associated costs.
2017	74 percent of threats are entered as an email attachment or link; 48 percent via drive-by web or download.
2018	New variants of ransomware occur almost daily.

Figure 1. Evolving Malware Landscape[3]

Evolving malware requires better solutions, not necessarily more of them. Too many products and organizations have interoperability, performance and complexity issues. Instead, an integrated solution that adapts to malware as it evolves is desirable. The Panda Adaptive Defense 360 platform is a single solution that combines revolutionary techniques focusing on stopping attacks up front while providing fine-grained analytics for identifying even the most advanced attacks.

Organizations need to evaluate endpoint solutions on their merits, requiring time and testing. For this review, we were fortunate to have a month to test the Panda Adaptive Defense 360 platform. Panda Security realizes the many challenges present when dealing with malware and malwareless attacks, with malwareless referring to targeted attacks such as phishing with PowerShell payloads. Panda Adaptive Defense 360 is a platform design encompassing protection from all directions. The Panda Adaptive Defense 360 platform:

- Contains advanced next-generation prevention controls to stop malware before running, thanks to the automated classification of any binary
- Employs monitoring of processes, registry keys, network connections and files
- Provides incident handling information, workflows and response capabilities

So, is Panda Adaptive Defense 360 an EPP plus EDR? Yes, but with more bells and whistles. Outside of prevention and detection controls, the platform includes the following:

- Built-in managed services: 100% Attestation Service and Threat Hunting and Investigative Service (THIS)
- Patch management for major applications and operating systems
- Device control
- Web URL filtering
- Asset identification
- Firewall, IDP/IPS

[3] Adapted from the timeline defined in www.darkreading.com/risk/the-evolution-of-malware/a/d-id/1322461 and results of the "2017 Threat Landscape Survey: Users on the Front Line," August 2017, www.sans.org/reading-room/whitepapers/threats/paper/37910 [Registration required for access.]

The 100% Attestation Service claims to prevent 100 percent of malware by classifying and controlling all executables and is a core benefit of the overall platform. THIS identifies new malwareless attack techniques that allow security teams to detect pure malwareless attacks or stop them in early steps of the Cyber Kill Chain®, before any malicious binary is deployed or run.[4] Both services provide information for hunt team exercises, forensic investigations and alerting about odd behaviors.

Based on the preceding features, we set on a journey to test the Panda Adaptive Defense 360 platform, seeking to answer the following questions:

- Does Panda Adaptive Defense 360 provide an easy-to-use interface that is welcoming to information security teams?
- Can the 100% Attestation Service correctly classify and prevent all malware?
- Can Panda Adaptive Defense 360 simplify investigations and identify malwareless attacks?
- How easily can the Panda Adaptive Defense 360 platform contain an infected asset?

> At its roots, Panda Adaptive Defense 360 is a full-stack EPP and EDR but with two key additions: 100% Attestation Service and Threat Hunting and Investigative Service (THIS).

Our review began with access to the Panda Adaptive Defense 360 web console, the design of which security practitioners and system administrators alike will find simple and easy to use.

Panda Adaptive Defense 360 Web Console

The web user interface comes with full documentation that is largely unnecessary due to the well-thought-out design and layout. As shown in Figure 2, the Panda Adaptive Defense 360 platform is broken down into four main sections: Status, Computers, Settings and Tasks.

The dashboard immediately presents an overview of an organization's deployment. From the main dashboard, users can quickly identify newly discovered assets not managed by an agent and outdated agents requiring updates, as well as any possible systems that malware or potentially unwanted programs (PUPs) attempted to run on. On the left side of the page is an easy-to-use set of dashboards users can select from to drill down into specific information.

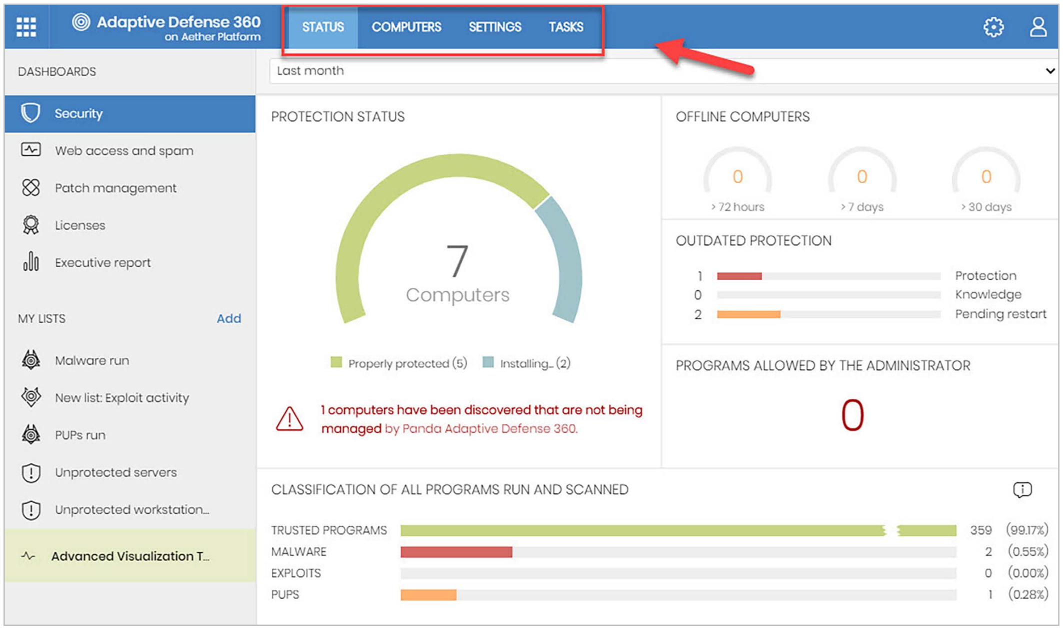

1 computers have been discovered that are not being managed by Panda Adaptive Defense 360.

CLASSIFICATION OF ALL PROGRAMS RUN AND SCANNED

TRUSTED PROGRAMS	359	(99.17%)
MALWARE	2	(0.55%)
EXPLOITS	0	(0.00%)
PUPS	1	(0.28%)

Figure 2. Panda Adaptive Defense 360 Dashboard

[4] Pure malwareless attacks are the ones that never deploy and run malware. However, attackers almost always need, at some point of the Cyber Kill Chain®, to drop a malicious binary to communicate with the command and control. In all these cases, the 100% Attestation Service will detect the malicious program and block the attacker before it causes any critical damage.

Agent Deployment and Management Capabilities

Before analyzing the security controls of the solution, we first assessed the agent deployment and management capabilities. Agent management within the Panda Adaptive Defense 360 platform is as simple as it can get. In our case, we followed a simple three-step process. We switched to the *Computers* menu, clicked on *Add computers* and selected the agent type to deploy (see Figure 3).

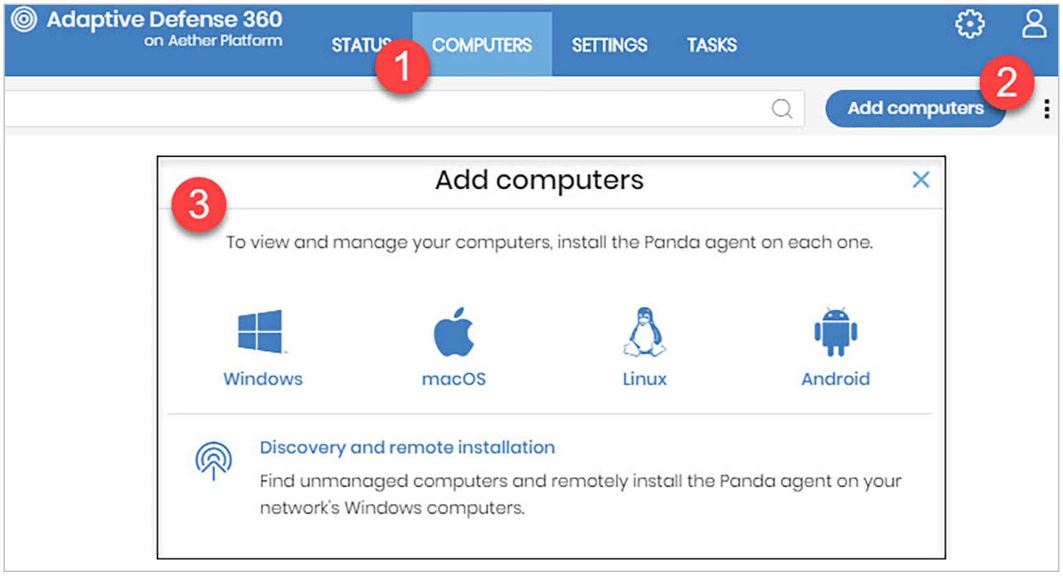

Add computers

To view and manage your computers, install the Panda agent on each one.

Discovery and remote installation

Find unmanaged computers and remotely install the Panda agent on your network's Windows computers.

Figure 3. Adding Agent Options

Organizations should spend a minimal amount of time deploying agents. Unfortunately, this is often not the case. However, we found deploying agents was a simple task of pulling down an agent and deploying it. On Windows, the file is an **`MSI`** file. For Mac, it is a **`dmg`**. On Linux, it is either a **`deb`** or **`rpm`**. Please note that no complex switches, group policies or Windows transform files are required. The downloaded agents contain everything that is necessary, and asset settings change when the agent first checks into the organization.

Organizations can achieve auto-deployment of agents using any asset management software or with a second option we tested, which is letting Panda Security identify assets and push agents from an existing agent. Why spend time identifying assets and installing agents if existing agents can do it for you? As shown in Figure 4, one or more agents can easily be configured to scan subnets or an Active Directory domain to identify assets and push agents to them.

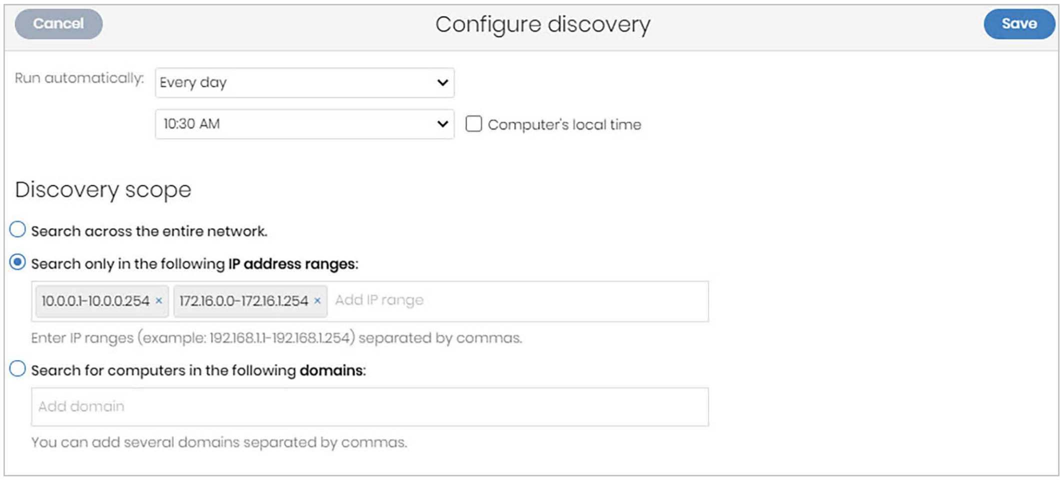

Configure discovery

Run automatically: Every day

10:30 AM ☐ Computer's local time

Discovery scope

○ Search across the entire network.

● Search only in the following **IP address ranges**:

10.0.0.1-10.0.0.254 × 172.16.0.0-172.16.1.254 × Add IP range

Enter IP ranges (example: 192.168.1.1-192.168.1.254) separated by commas.

○ Search for computers in the following **domains**:

Add domain

You can add several domains separated by commas.

Figure 4. Automatic Agent Deployment

Panda takes a novel approach to maintain agent configuration. It automatically integrates with Active Directory (AD) because the agent retrieves the path of the asset as defined by AD. Organizations can use this method to apply settings to computers via AD organizational units (OUs) or with dynamic filters, all without having custom hooks between Panda Adaptive Defense 360 and an organization's domain controllers. So on Day One, we were able to start applying key settings to all systems or subsets of systems because the agent is aware of AD, subnets and hostnames. Given these options, we wanted to determine whether Panda Adaptive Defense 360 would easily apply different security configurations and patches to various OUs.

To do this, we moved into the *Settings* menu. Here we were able to distinguish which security controls are enabled and how they are tuned. We had expected that there would be a tremendous number of settings, areas in which to add exceptions and overwhelming amounts of documentation to identify what everything does. Instead, we found each of the options to be easy to understand and change with nominal tuning requirements. Figure 5 shows an example of turning on the Advanced Protection service, which is part of the 100% Attestation Service.

Panda takes a novel approach to maintain agent configuration. It automatically integrates with Active Directory (AD) because the agent retrieves the path of the asset as defined by AD.

The configuration of advanced protection requires enabling it and setting one of three modes:

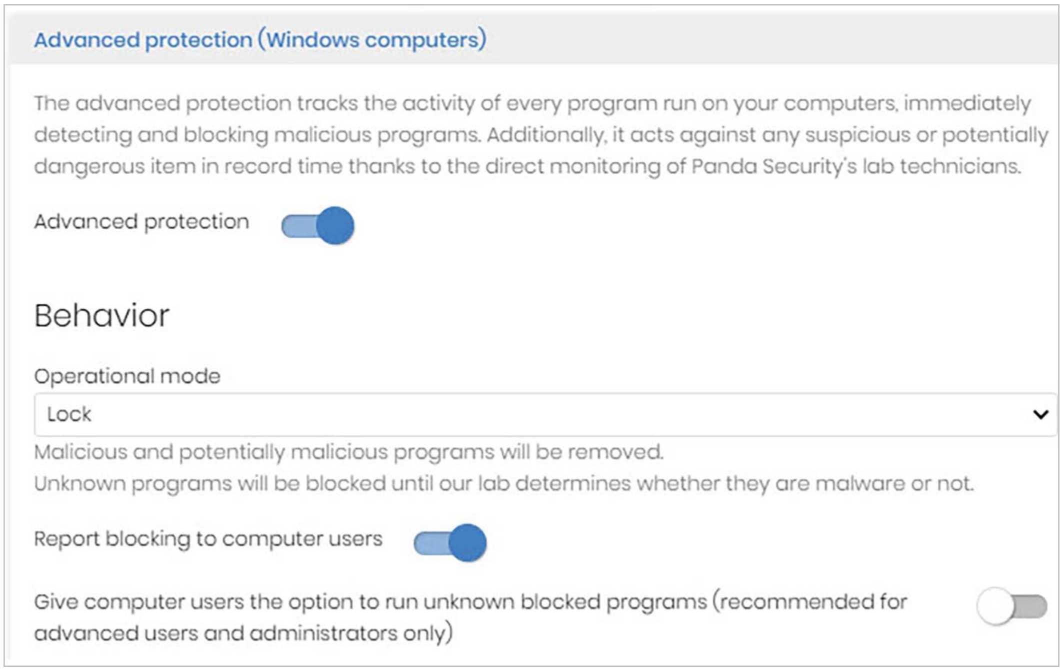

Advanced protection (Windows computers)

The advanced protection tracks the activity of every program run on your computers, immediately detecting and blocking malicious programs. Additionally, it acts against any suspicious or potentially dangerous item in record time thanks to the direct monitoring of Panda Security's lab technicians.

Advanced protection

Behavior

Operational mode

Lock

Malicious and potentially malicious programs will be removed.
Unknown programs will be blocked until our lab determines whether they are malware or not.

Report blocking to computer users

Give computer users the option to run unknown blocked programs (recommended for advanced users and administrators only)

Figure 5. Enabling Advanced Protection

- **Audit—**Monitors system activity and tracks every binary but does not take any action.
- **Hardening—**Malicious content is removed, indicators of fileless and malwareless attacks (IoAs) are detected, and unknown binaries from the internet or external storage drives are not allowed to execute until classified as trustable by Panda Security. Other unknown binaries, previously present in the endpoints, will run while being continually analyzed and classified by the 100% Attestation Service in the background.
- **Lock—**Malicious content is removed, IoAs are detected, and unknown binaries are not allowed to execute until classified by Panda Security.

The recommendation and steps we took during testing are to place assets into Hardening mode for two weeks or more and then consider switching to Lock mode. It is easy to forget that Panda Adaptive Defense 360 platform is designed to be like a managed service. The goal is for the platform to work for an organization rather than the organization constantly working to adjust a product. Panda Adaptive Defense 360 platform did a great job of this because we experienced no false positives, which are common within traditional endpoint protection solutions or whitelisting solutions. The best part is that this experience occurred with zero exceptions being added.

To be clear, organizations can add exceptions, but we did not experience performance issues or legitimate content being removed or blocked without them. With less tuning required and no false positives on legitimate software, we began to wonder whether the Panda Adaptive Defense 360 platform could truly be competitive at both stopping malware and providing data for detection. We'll cover this line of thinking in the "Endpoint Prevention Capabilities" and "Visibility and Detection at the Endpoint" sections later in this paper.

Within *Tasks*, the last main menu, we were able to establish scheduled antivirus scans and—more importantly—schedule patch installations. Endpoint suites that do not include patch management forget that a key control for the prevention of malware execution is to patch software vulnerabilities. The Panda Adaptive Defense 360 platform accounts for this by providing the means to patch operating system files as well as major third-party applications such as Java, Adobe Flash Player, Adobe Reader and many others. As shown in Figure 6, patch scheduling is as simple as establishing reoccurrence, selecting a group or individual systems, identifying what patches to install, and finally deciding whether a restart is necessary.

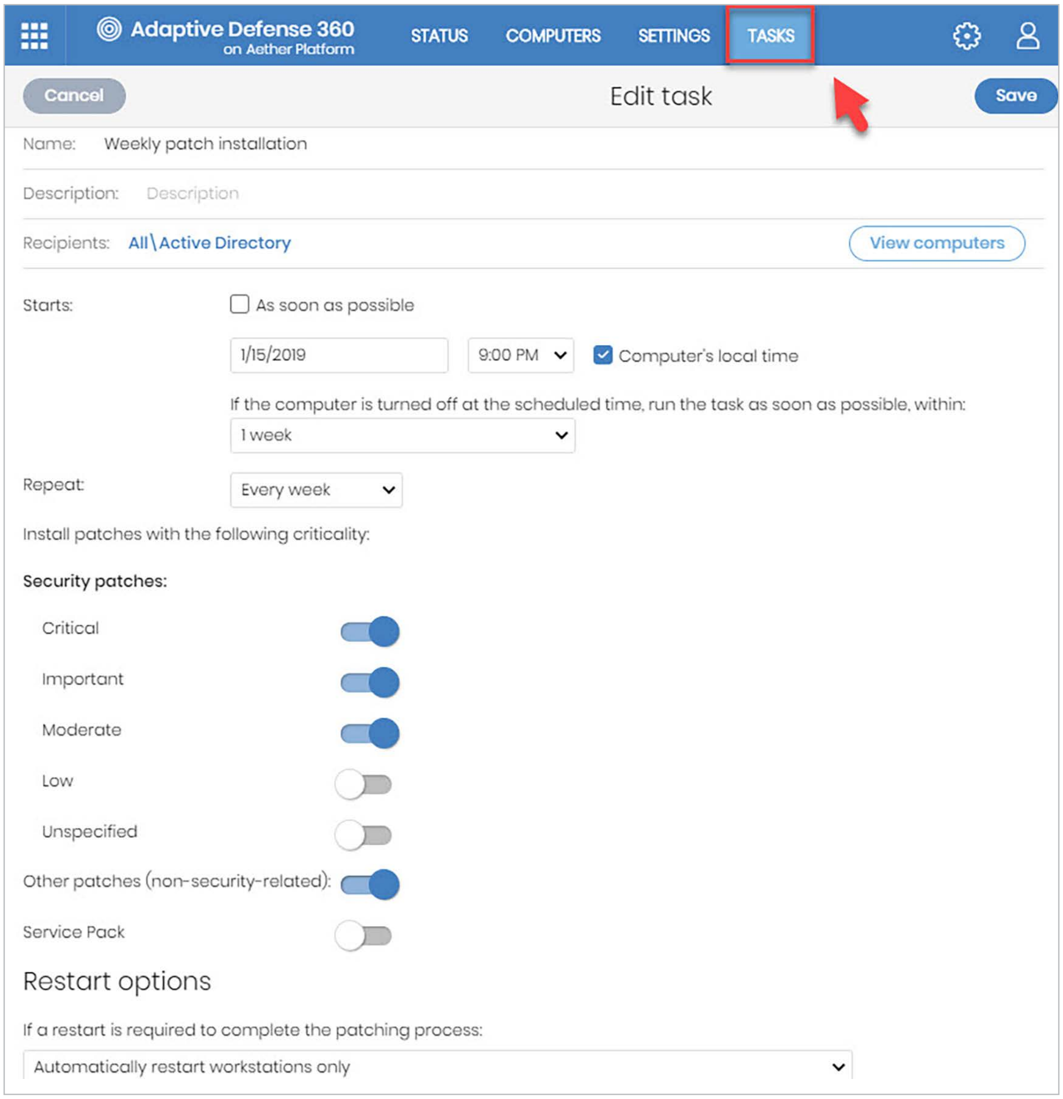

Edit task

Name: Weekly patch installation

Description: Description

Recipients: All\Active Directory

Starts: ☐ As soon as possible

1/15/2019 | 9:00 PM | ☑ Computer's local time

If the computer is turned off at the scheduled time, run the task as soon as possible, within: 1 week

Repeat: Every week

Install patches with the following criticality:

Security patches:

Critical ☑

Important ☑

Moderate ☑

Low ☐

Unspecified ☐

Other patches (non-security-related): ☑

Service Pack ☐

Restart options

If a restart is required to complete the patching process: Automatically restart workstations only

Figure 6. Patch Installation Task

Endpoint Prevention Capabilities

After installing agents and configuring their settings, the next step was to assess the 100% Attestation Service, which states it blocks 100 percent of malware. This capability is at the core of the Panda Adaptive Defense 360 prevention capabilities. Effectively, every binary is classified and based on those classifications, the binary is either authorized to execute or not. (Note that if your agents are in Hardening mode, which was previously discussed, binaries already present in the endpoint may be able to run during the classification process, while under Lock mode they would not.) The word *classification* needs clarification because it oversimplifies what the 100% Attestation Service is doing. Essentially, the service is applying machine learning and statistical modeling to more than a **thousand flags** varying from static comparisons to **behavioral analysis and context of execution**. Ignore the data science talk for a moment. Instead, think about how effective it would be if you could have an immediate answer to the main question, "Is this an unknown malicious binary never seen before?" even before letting it execute and without expending the efforts needed to validate and triage any single alert.

If the malicious binary were able to execute, Panda Security would answer the following questions:

- Is a binary making network connections to the internet or to internal machines?
- What are the path, parent path, process and user underlying the creation of the binary?
- What time is the binary being requested to execute compared to the creation time?
- What process wrote the binary to disk compared to what process is trying to execute it?
- Is this the first time a machine is attempting to execute a file, or has it previously run on this machine or others within the organization?
- What are key attributes about the binary? (For example, is it a PDF, a JAR, an EXE or things metadata about the inner workings of the binary?)
  - Is there shellcode?
  - Is code obfuscation identified?

As a rule of thumb, we recommend that no security control be taken at face value. Even third-party studies are not validation that a solution is ideal for your environment. The best test is to audit the security controls against multiple solutions and evaluate the results.

In truth, the preceding description is an example (if not an oversimplification) of what the 100% Attestation Service is automatically checking. The 100% Attestation Service is designed to collect more than a thousand data entry points and then apply data science to classify files automatically. On the off chance that an algorithm is unable to classify a file, the Panda Adaptive Defense 360 system assigns analysis to staff at Panda Security for manual classification. In our experience, the classification process was quick and led to immediate results. (Again, in Hardening mode, the executable—if it is not from an external device or the internet—would run until classified. In Lock mode, employees are prevented from executing the binary until classification is completed.) Because of how quickly classification of unknown files occurred, we found this to be an effective solution with an acceptable time trade-off.

Ultimately, we believe prevention technologies need to be assessed. We began testing the prevention capabilities of the Panda Adaptive Defense 360 platform. For this, we tried placing and executing malware samples ranging from modern-day ransomware to rootkits to traditional viruses. During our tests, malware met with repeated removal or failure to execute. A clear majority was prevented before execution. One of the malware samples tested was the infamous Petya ransomware. This malware never made it to execution and was properly logged. See Figure 7 for an example of the front page, showing that one of the test assets had malware activity.

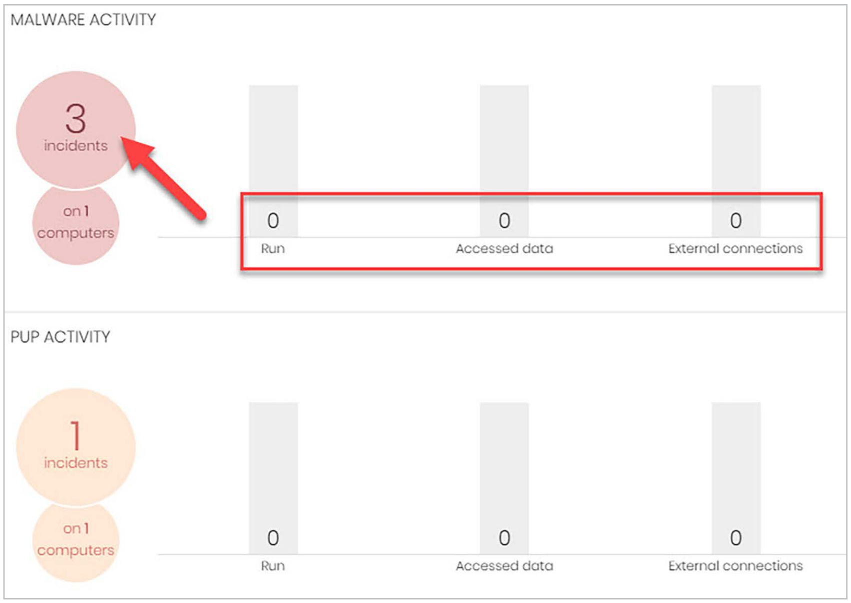

Figure 7. Front Page of the Malware Activity Section on the Dashboard

We found one of the best features is that everything is click-through. For example, clicking on the circle labeled “1 incidents” or the circle that states “on 1 computers” drills down into the hosts on which malware was blocked, as shown in Figure 8.

Note that the screen also includes additional information analysts need to make decisions, such as threat category and path, as well as other key information, such as whether the binary executed on any systems in the organization, accessed other files or made external connections. The immediate presence of actionable information made it simple to build a hypothesis on what may or may not occur, yet it was largely unnecessary because of the number of malware samples the 100% Attestation Service stopped before execution. However, further drilldown is possible by clicking on the malware sample or computer name, as shown in Figure 8. Figure 9 demonstrates the complete drilldown into the forensic information for malware samples.

Notice, with just three clicks from the front page, how much data an analyst can easily access. The overall theme is to provide not only information but also color highlighting and emphasize asking further questions. We found the native capability to see where malware occurs on other organization systems and its dwell time to be great additions that often require more work to identify. Perhaps one of the more helpful capabilities is the ability to view malware activity on a graph using the View activity graph button. Figure 10 illustrates the activity discovered in our testing.

Panda Security realizes that catching malware requires a combination of signatures, behavioral monitoring, context identification, reputation scoring and file/memory analysis. The problem is not the application of each of these methods individually but the combination of them, in the right order, and in a scalable fashion. Doing so requires technologies such as machine learning and other data science best practices.

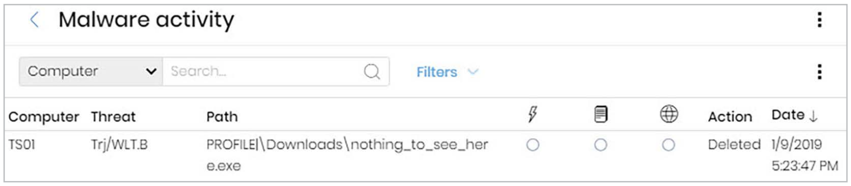

Computer	Threat	Path				Action	Date ↓
TS01	Trj/WLT.B	PROFILE|\Downloads\nothing_to_see_here.exe	○	○	○	Deleted	1/9/2019 5:23:47 PM

Figure 8. Drilldown on Malware Activity

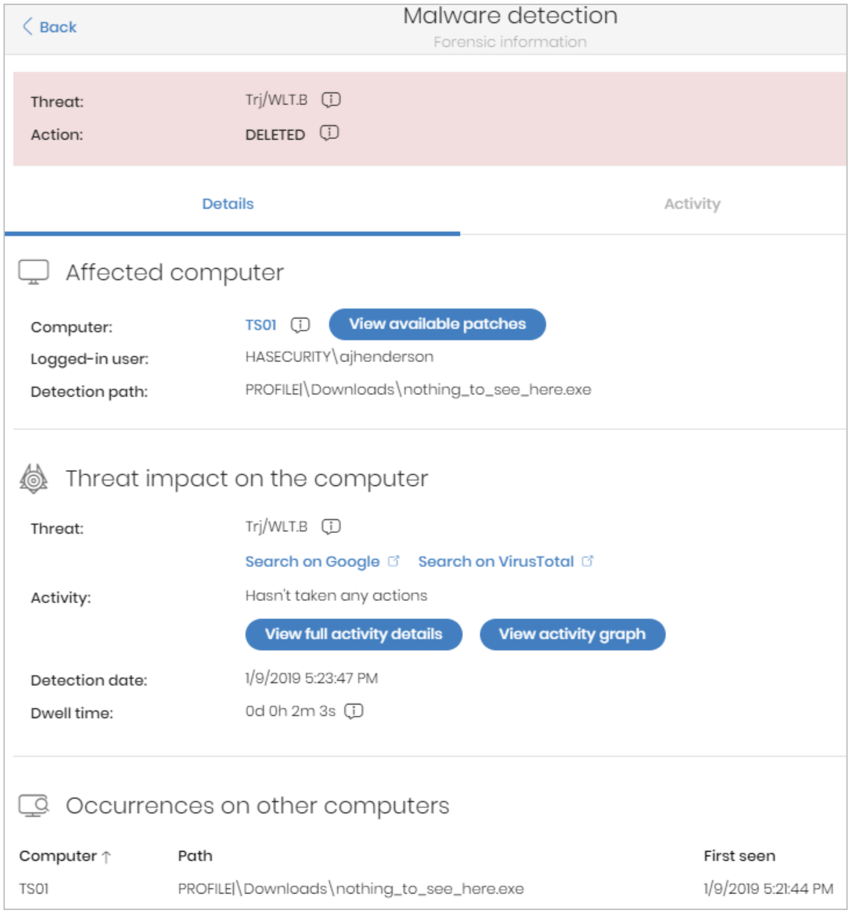

Threat: Trj/WLT.B
Action: DELETED

Affected computer

Computer: TS01 View available patches
Logged-in user: HASECURITY\ajhenderson
Detection path: PROFILE|\Downloads\nothing_to_see_here.exe

Threat impact on the computer

Threat: Trj/WLT.B
Search on Google Search on VirusTotal
Activity: Hasn't taken any actions
View full activity details View activity graph
Detection date: 1/9/2019 5:23:47 PM
Dwell time: 0d 0h 2m 3s

Occurrences on other computers

Computer ↑	Path	First seen
TS01	PROFILE|\Downloads\nothing_to_see_here.exe	1/9/2019 5:21:44 PM

Figure 9. Forensic View of Malware

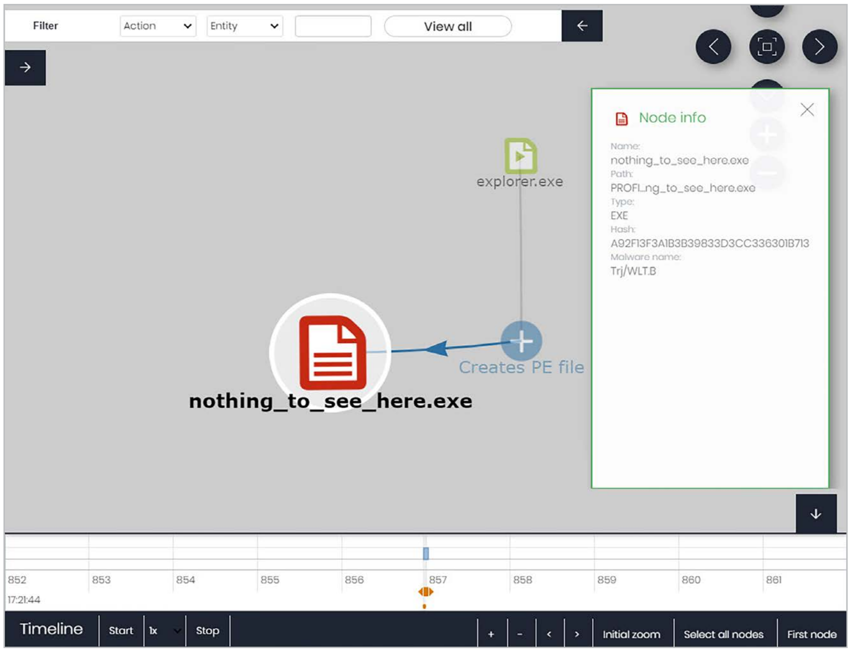

Name:
nothing_to_see_here.exe
Path:
PROFI...ng_to_see_here.exe
Type:
EXE
Hash:
A92F13F3A1B3B39833D3CC336301B713
Malware name:
Trj/WLT.B

Figure 10. View Activity Graph

Visibility and Detection at the Endpoint

There is a significant difference in dealing with malware-based attacks versus malwareless attacks. The 100% Attestation Service focuses on catching malware. In our tests, it worked 100 percent of the time. Organizations, however, are plagued with attacks that are both direct malware and attacks that are memory resident, such as scripting frameworks and built-in operating system binaries. Panda Security understands this, and as a result, endpoint, detection and response capabilities are built into the Panda Adaptive Defense 360 framework.

Today, malware uses some of the most evasive, destructive and sophisticated techniques available: an organization's software. Scripting languages such as Python and PowerShell, macros within Microsoft documents, or even macroless attacks such as Microsoft Dynamic Data Exchange (DDE) attacks spread rampantly across the internet. These attacks live off the land and achieve fileless malware.

As seen in previous images, drilldown reporting is a click away within the platform. Figure 11 shows the output of an attempt we made using Veil-Evasion. Veil-Evasion is routinely used to generate code, such as PowerShell code in our attempt, which evades antivirus products. In this case, it was blocked and logged.

Visibility on blocked or quarantined items is important, but even more important is the ability to find potentially undetected attacks. This ability requires access to data about an organization's environment. To achieve this, the Panda Adaptive Defense 360 platform has multiple add-on modules to provide access to this data. Two modules that we reviewed are the Advanced Visualization Tool and the Panda SIEMFeeder.[5] Both allow an organization to perform manual analysis and hunt team exercises based on a wealth of endpoint data collected. Panda SIEMFeeder enables an organization to export the endpoint data collected inside the Adaptive Defense 360 platform into the SIEM of your choice. The Advanced Visualization Tool allows users to directly query the data without the need to bring any data on-premises. The default view in the Advanced Visualization Tool is shown in Figure 12.

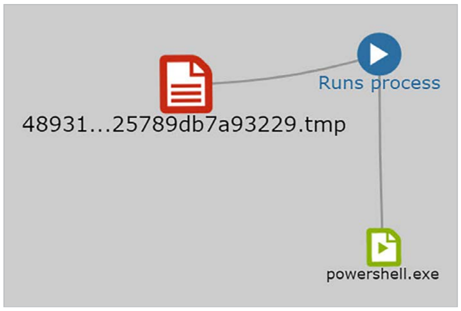

Figure 11. Veil-Evasion Attempt

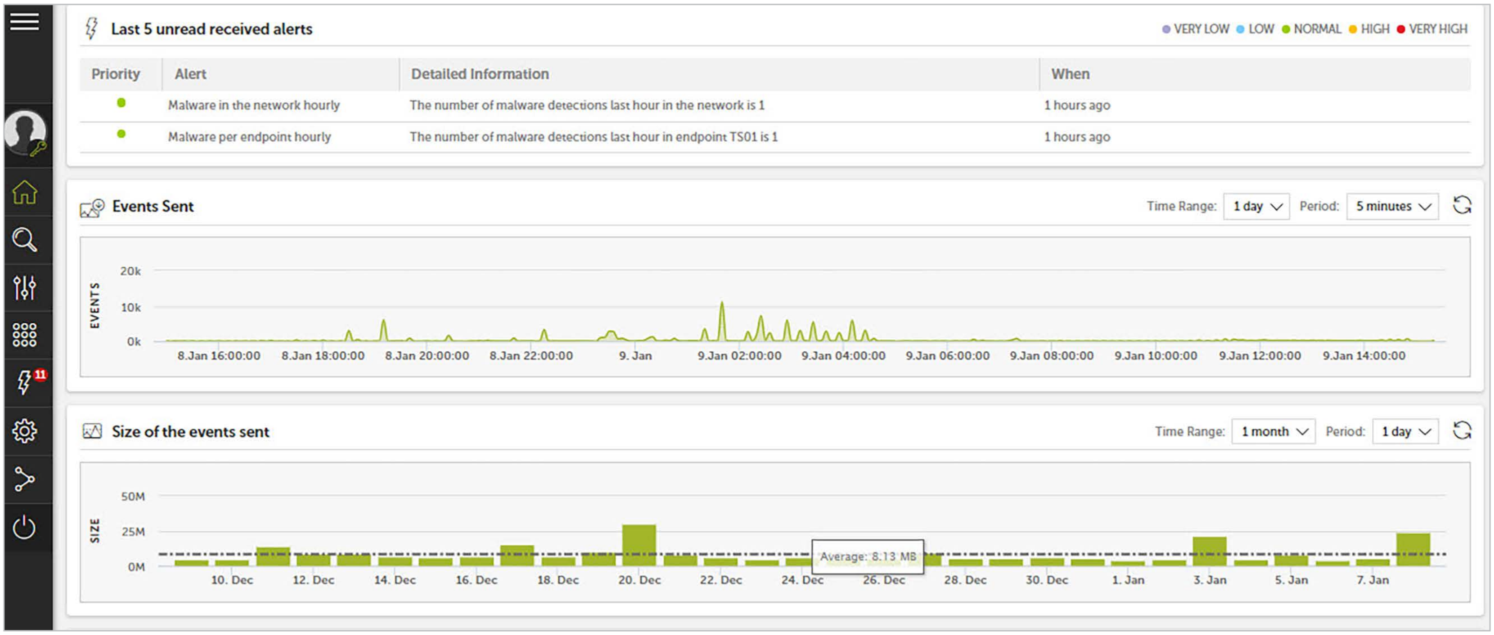

Last 5 unread received alerts

● VERY LOW ● LOW ● NORMAL ● HIGH ● VERY HIGH

Priority	Alert	Detailed Information	When
●	Malware in the network hourly	The number of malware detections last hour in the network is 1	1 hours ago
●	Malware per endpoint hourly	The number of malware detections last hour in endpoint TS01 is 1	1 hours ago

Events Sent

Time Range: 1 day Period: 5 minutes

Size of the events sent

Time Range: 1 month Period: 1 day

Figure 12. Advanced Visualization Tool

[5] In addition to these tools, Panda Security offers a Threat Hunting Platform to managed security service/managed detection and response (MSSP/MDR) providers and to any mature organization with an advanced security operations center (SOC). The platform and Panda's tools are designed to cover and automate the whole process of hunting for threats that were able to bypass any other security controls, accelerate the investigation and respond to attackers.

Analysis tools make or break an analyst. Opening a console and seeing hundreds of millions—if not billions—of events is overwhelming even to seasoned professionals. When reviewing the Advanced Visualization Tool, we looked for easy ways to find items of suspicion, and we were not disappointed. The reporting capabilities make it easy to search across processes and make it possible for users to find things such as suspicious PowerShell use. Figure 13 shows a query for abnormal PowerShell use by looking for encoding as well as unusual parameter use.

Analysis tools make or break an analyst. Opening a console and seeing hundreds of millions—if not billions—of events is overwhelming even to seasoned professionals.

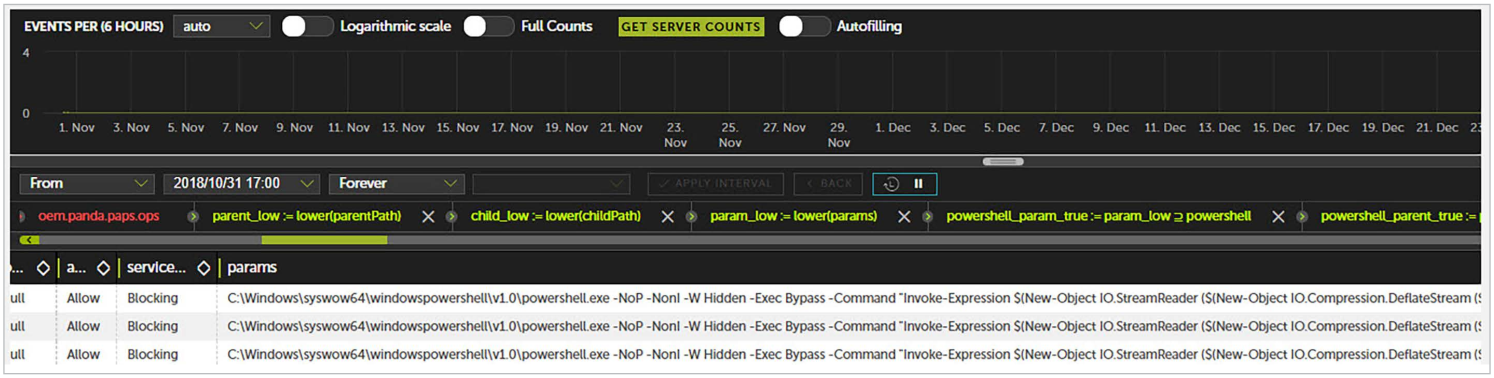

2018/10/31 17:00

oem.panda.paps.ops › parent_low := lower(parentPath) › child_low := lower(childPath) › param_low := lower(params) › powershell_param_true := param_low ⊇ powershell › powershell_parent_true :=

...	a...	service...	params
ull	Allow	Blocking	C:\Windows\syswow64\windowspowershell\v1.0\powershell.exe -NoP -NonI -W Hidden -Exec Bypass -Command "Invoke-Expression $(New-Object IO.StreamReader ($(New-Object IO.Compression.DeflateStream ($
ull	Allow	Blocking	C:\Windows\syswow64\windowspowershell\v1.0\powershell.exe -NoP -NonI -W Hidden -Exec Bypass -Command "Invoke-Expression $(New-Object IO.StreamReader ($(New-Object IO.Compression.DeflateStream ($
ull	Allow	Blocking	C:\Windows\syswow64\windowspowershell\v1.0\powershell.exe -NoP -NonI -W Hidden -Exec Bypass -Command "Invoke-Expression $(New-Object IO.StreamReader ($(New-Object IO.Compression.DeflateStream ($

Figure 13. Abnormal PowerShell Discovery

We found the interface supported hunt team investigations and indicator-of-compromise searches with relative ease. If you know what to look for, the Advanced Visualization Tool will let you discover it.

However, Panda Adaptive Defense 360 can also be extended with Panda Security's Threat Hunting Platform. While we did not review this platform, it is designed for corporate internal security operations centers (SOCs) and external MSSP/MDR providers that must actively look for adversaries on the protected networks and endpoints to avoid breaches from happening by detecting hackers and insiders as early as possible. It is designed to enable security practitioners to discover advanced threats and insiders faster, to accelerate the investigation and response, and to act immediately on compromised endpoints from a single cloud-based console without deploying additional agents to endpoints.

The Threat Hunting Platform integrates key capabilities for efficient hunting and incident response, such as:

- Machine learning algorithms for automatic profiling
- Anomaly and suspicious activity automatic detonation
- Big data-scaled querying for alerts triage
- Threat hunting hypothesis validation
- Scan of indicators of compromise (IoCs) derived from threat intelligence or other sources, such as past investigations
- Malicious activity containment (remote shell, isolation and so on) on endpoints
- Tools to automate routine and repetitive tasks during the hunting, investigation containment and reporting phases

A new, cloud-based version of Panda Security Threat Hunting Platform is under development. Plans for this upgrade include advanced functionality for automating the investigation and threat hunting process.

Figure 14 shows the Threat Hunting Web Console for incident investigation.

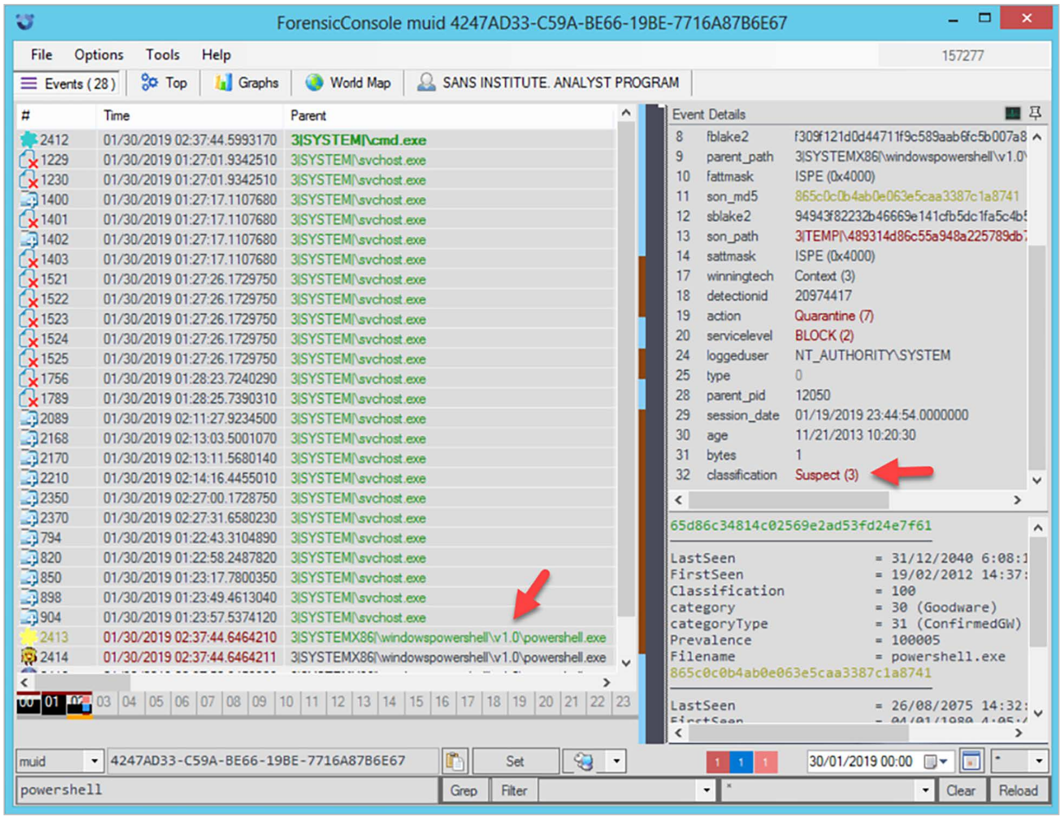

#	Time	Parent
2412	01/30/2019 02:37:44.5993170	3|SYSTEM|\cmd.exe
1229	01/30/2019 01:27:01.9342510	3|SYSTEM|\svchost.exe
1230	01/30/2019 01:27:01.9342510	3|SYSTEM|\svchost.exe
1400	01/30/2019 01:27:17.1107680	3|SYSTEM|\svchost.exe
1401	01/30/2019 01:27:17.1107680	3|SYSTEM|\svchost.exe
1402	01/30/2019 01:27:17.1107680	3|SYSTEM|\svchost.exe
1403	01/30/2019 01:27:17.1107680	3|SYSTEM|\svchost.exe
1521	01/30/2019 01:27:26.1729750	3|SYSTEM|\svchost.exe
1522	01/30/2019 01:27:26.1729750	3|SYSTEM|\svchost.exe
1523	01/30/2019 01:27:26.1729750	3|SYSTEM|\svchost.exe
1524	01/30/2019 01:27:26.1729750	3|SYSTEM|\svchost.exe
1525	01/30/2019 01:27:26.1729750	3|SYSTEM|\svchost.exe
1756	01/30/2019 01:28:23.7240290	3|SYSTEM|\svchost.exe
1789	01/30/2019 01:28:25.7390310	3|SYSTEM|\svchost.exe
2089	01/30/2019 02:11:27.9234500	3|SYSTEM|\svchost.exe
2168	01/30/2019 02:13:03.5001070	3|SYSTEM|\svchost.exe
2170	01/30/2019 02:13:11.5680140	3|SYSTEM|\svchost.exe
2210	01/30/2019 02:14:16.4455010	3|SYSTEM|\svchost.exe
2350	01/30/2019 02:27:00.1728750	3|SYSTEM|\svchost.exe
2370	01/30/2019 02:27:31.6580230	3|SYSTEM|\svchost.exe
794	01/30/2019 01:22:43.3104890	3|SYSTEM|\svchost.exe
820	01/30/2019 01:22:58.2487820	3|SYSTEM|\svchost.exe
850	01/30/2019 01:23:17.7800350	3|SYSTEM|\svchost.exe
898	01/30/2019 01:23:49.4613040	3|SYSTEM|\svchost.exe
904	01/30/2019 01:23:57.5374120	3|SYSTEM|\svchost.exe
2413	01/30/2019 02:37:44.6464210	3|SYSTEMX86|\windowspowershell\v1.0\powershell.exe
2414	01/30/2019 02:37:44.6464211	3|SYSTEMX86|\windowspowershell\v1.0\powershell.exe

8	fblake2	f309f121d0d44711f9c589aab6fc5b007a8
9	parent_path	3|SYSTEMX86|\windowspowershell\v1.0\
10	fattmask	ISPE (0x4000)
11	son_md5	865c0c0b4ab0e063e5caa3387c1a8741
12	sblake2	94943f82232b46669e141cfb5dc1fa5c4b5
13	son_path	3|TEMP|\489314d86c55a948a225789db7
14	sattmask	ISPE (0x4000)
17	winningtech	Context (3)
18	detectionid	20974417
19	action	Quarantine (7)
20	servicelevel	BLOCK (2)
24	loggeduser	NT_AUTHORITY\SYSTEM
25	type	0
28	parent_pid	12050
29	session_date	01/19/2019 23:44:54.0000000
30	age	11/21/2013 10:20:30
31	bytes	1
32	classification	Suspect (3)

```
65d86c34814c02569e2ad53fd24e7f61
________________________________
LastSeen                 = 31/12/2040 6:08:1
FirstSeen                = 19/02/2012 14:37:
Classification           = 100
category                 = 30 (Goodware)
categoryType             = 31 (ConfirmedGW)
Prevalence               = 100005
Filename                 = powershell.exe
865c0c0b4ab0e063e5caa3387c1a8741
________________________________
LastSeen                 = 26/08/2075 14:32:
```

muid 4247AD33-C59A-BE66-19BE-7716A87B6E67

powershell

30/01/2019 00:00

Figure 14. Threat Hunting Platform

Response and Quarantine

In our scenario, Panda Adaptive Defense 360 prevented attacks. But what happens in a case where the attack is not initially prevented? In this case it is important to identify the compromised system and then to contain the attack. To that end, we investigated the capability of Panda Adaptive Defense 360 to contain a machine. Panda Security could not have made the process easier. Isolating a computer is a three-step process:

1. Click on the *Computers* menu.
2. Search for and select the computer you want to isolate.
3. Click on *Isolate*, as shown in Figure 15.

If an organization utilizes network forensics or incident handling tools, the processes behind those tools can be allowed under the "Advanced options" link shown in Figure 15. The result is a computer that is locked down so that infections or unauthorized activity do not remain rampant within an organization. The process of putting a machine in isolation took slightly less than 10 seconds during our testing.

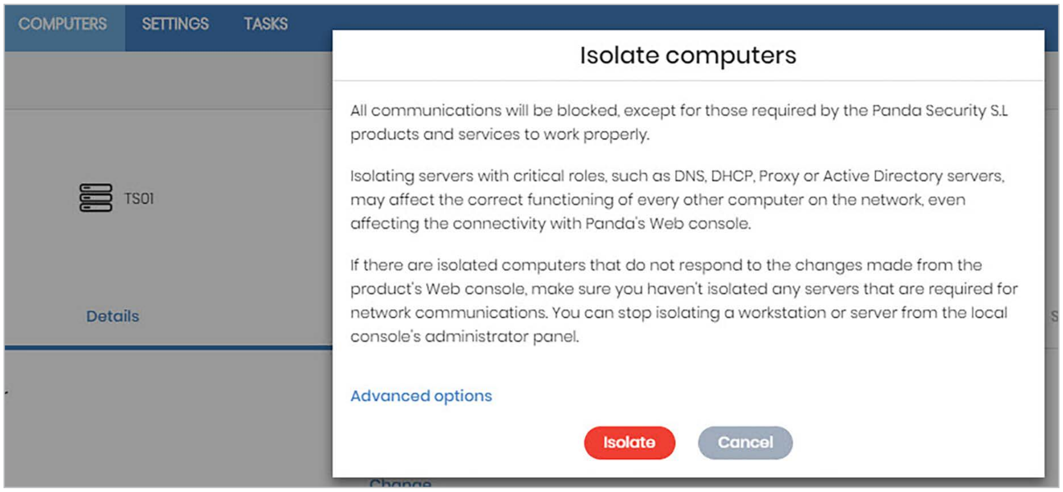

Isolate computers

All communications will be blocked, except for those required by the Panda Security S.L products and services to work properly.

Isolating servers with critical roles, such as DNS, DHCP, Proxy or Active Directory servers, may affect the correct functioning of every other computer on the network, even affecting the connectivity with Panda's Web console.

If there are isolated computers that do not respond to the changes made from the product's Web console, make sure you haven't isolated any servers that are required for network communications. You can stop isolating a workstation or server from the local console's administrator panel.

Figure 15. Client Isolation

Conclusion

Endpoints are the new perimeter, and securing them is far from easy. As discussed throughout this paper, endpoint security requires a solution that scales, is easy to maintain and provides a comprehensive integration into the endpoint itself. The deeper the integration, the more organizations need a solution that shields them from complexity and tuning requirements. While some organizations have dedicated staff that can fine-tune an endpoint solution, many organizations struggle to do so.

Knowing this, organizations should perform a self-assessment of their approach to endpoint security. A few questions worth asking are:

- Is maintaining multiple endpoint security solutions draining internal resources or causing performance issues?
- Is the organization concerned about advanced malware threats?
- Is there a lack of visibility into what the organization's endpoints are doing?
- Does the organization need to optimize and accelerate the tasks performed by security teams or SOCs, thus avoiding the need to spend time validating security alerts with new security practitioners?
- Does the organization need to increase its maturity in terms of security and risk management, allowing the security team to focus on more strategic decisions?

If the answer is "yes" to any of these questions, an alternative approach may be warranted. In our opinion, the Panda Adaptive Defense 360 platform addresses each of these concerns and more. We appreciate the opportunity and time we had to assess Panda Security's endpoint solution. Between the 100% Attestation Service and the ability to quickly pivot or hunt for internal threats, as well as having the ability to respond, patch and automatically manage agents, we found Panda Adaptive Security 360 to be an appealing solution for organizations of varying sizes and skill levels.

Panda Adaptive Defense 360 helps organizations focus on results rather than the in-between steps. Why have full-time staff deploy products when they should be wielding them as cyber defense weapons and shields? From Day One, Panda Adaptive Defense 360 equips organizations to see results.

About the Author

Justin Henderson is a certified SANS instructor who authored the SEC555 (SIEM with Tactical Analytics) course and co-authored SEC455 (SIEM Design and Implementation) and SEC530 (Defensible Security Architecture). He is a member of the SANS Cyber Guardian Blue Team who is passionate about making defense fun and engaging. Justin specializes in threat hunting via SIEM, network security monitoring and ad hoc scripting.

Sponsor

SANS would like to thank this paper's sponsor:

panda security

Last Updated: March 26th, 2019

Upcoming SANS Training

Click here to view a list of all SANS Courses

SANS Cyber Security Middle East Summit	**Abu Dhabi, AE**	**Apr 04, 2019 - Apr 11, 2019**	**Live Event**
SANS London April 2019	**London, GB**	**Apr 08, 2019 - Apr 13, 2019**	**Live Event**
Blue Team Summit & Training 2019	**Louisville, KYUS**	**Apr 11, 2019 - Apr 18, 2019**	**Live Event**
SANS Riyadh April 2019	**Riyadh, SA**	**Apr 13, 2019 - Apr 18, 2019**	**Live Event**
SANS Boston Spring 2019	**Boston, MAUS**	**Apr 14, 2019 - Apr 19, 2019**	**Live Event**
SANS Seattle Spring 2019	**Seattle, WAUS**	**Apr 14, 2019 - Apr 19, 2019**	**Live Event**
FOR498 Battlefield Forensics Beta 1	**Arlington, VAUS**	**Apr 15, 2019 - Apr 20, 2019**	**Live Event**
SANS FOR585 Madrid April 2019 (in Spanish)	**Madrid, ES**	**Apr 22, 2019 - Apr 27, 2019**	**Live Event**
SANS Northern Virginia- Alexandria 2019	**Alexandria, VAUS**	**Apr 23, 2019 - Apr 28, 2019**	**Live Event**
SANS Muscat April 2019	**Muscat, OM**	**Apr 27, 2019 - May 02, 2019**	**Live Event**
SANS Pen Test Austin 2019	**Austin, TXUS**	**Apr 29, 2019 - May 04, 2019**	**Live Event**
Cloud Security Summit & Training 2019	**San Jose, CAUS**	**Apr 29, 2019 - May 06, 2019**	**Live Event**
SANS Bucharest May 2019	**Bucharest, RO**	**May 06, 2019 - May 11, 2019**	**Live Event**
SANS Security West 2019	**San Diego, CAUS**	**May 09, 2019 - May 16, 2019**	**Live Event**
SANS Perth 2019	**Perth, AU**	**May 13, 2019 - May 18, 2019**	**Live Event**
SANS Milan May 2019	**Milan, IT**	**May 13, 2019 - May 18, 2019**	**Live Event**
SANS Dublin May 2019	**Dublin, IE**	**May 13, 2019 - May 18, 2019**	**Live Event**
SANS Stockholm May 2019	**Stockholm, SE**	**May 13, 2019 - May 18, 2019**	**Live Event**
SANS New Orleans 2019	**New Orleans, LAUS**	**May 19, 2019 - May 24, 2019**	**Live Event**
SANS Northern VA Spring- Reston 2019	**Reston, VAUS**	**May 19, 2019 - May 24, 2019**	**Live Event**
SANS Autumn Sydney 2019	**Sydney, AU**	**May 20, 2019 - May 25, 2019**	**Live Event**
SANS MGT516 Beta Two 2019	**San Francisco, CAUS**	**May 20, 2019 - May 24, 2019**	**Live Event**
SANS Amsterdam May 2019	**Amsterdam, NL**	**May 20, 2019 - May 25, 2019**	**Live Event**
SANS Hong Kong 2019	**Hong Kong, HK**	**May 20, 2019 - May 25, 2019**	**Live Event**
SANS Krakow May 2019	**Krakow, PL**	**May 27, 2019 - Jun 01, 2019**	**Live Event**
SANS Atlanta 2019	**Atlanta, GAUS**	**May 28, 2019 - Jun 02, 2019**	**Live Event**
SANS San Antonio 2019	**San Antonio, TXUS**	**May 28, 2019 - Jun 02, 2019**	**Live Event**
Security Writing NYC: SEC402 Beta 2	**New York, NYUS**	**Jun 01, 2019 - Jun 02, 2019**	**Live Event**
SANS London June 2019	**London, GB**	**Jun 03, 2019 - Jun 08, 2019**	**Live Event**
SANS Zurich June 2019	**Zurich, CH**	**Jun 03, 2019 - Jun 08, 2019**	**Live Event**
Enterprise Defense Summit & Training 2019	**Redondo Beach, CAUS**	**Jun 03, 2019 - Jun 10, 2019**	**Live Event**
SANS Kansas City 2019	**Kansas City, MOUS**	**Jun 10, 2019 - Jun 15, 2019**	**Live Event**
SANS 2019	**OnlineFLUS**	**Apr 01, 2019 - Apr 08, 2019**	**Live Event**
SANS OnDemand	**Books & MP3s OnlyUS**	**Anytime**	**Self Paced**