ENDPOINT PROTECTION

PALO ALTO NETWORKS TRAPS/CORTEX XDR V6.2

Q1 2020 PRODUCT RATING

AA

Overview & Outlook

During Q1, 2020, NSS Labs performed an independent test of the Palo Alto Networks Traps / Cortex XDR v6.2.

Management extensive and flexible. Overall protection impressive; low false positive rate; excellent resistance to evasion. Excellent protection against malware, drive-by exploits, and handcrafted (targeted) attacks. Some weakness defending against social exploits embedded in documents.

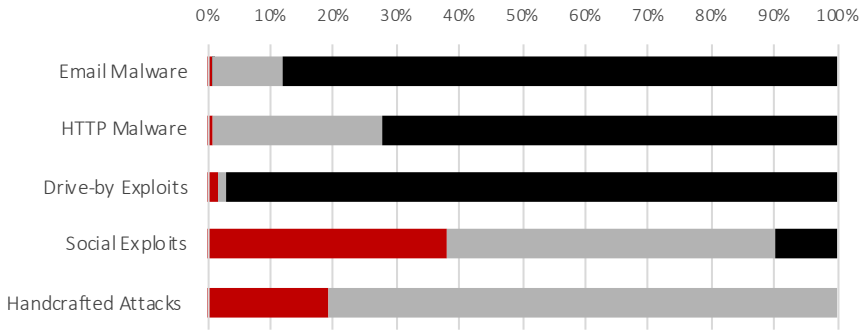

Summary of Results

MANAGEMENT AA

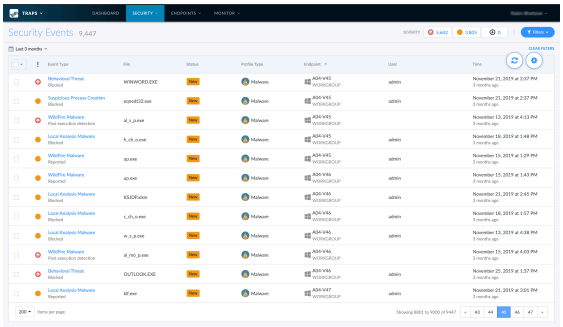

Deployment of the Palo Alto Network Traps agent was uneventful and ongoing operational tasks were easy to carry out. The management console supports role-based access control (RBAC) and third-party authentication via Active Directory with support for multi-factor authentication. We found it to be straightforward to define and save multiple security profiles, which were then combined to create a security policy. Policies may be applied to a single endpoint, groups of endpoints, or globally via Active Directory organizational units or static groups that are created and managed within the XDR console. The policy mechanism is diverse and supports most use cases, enabling customization such as whitelisting and blacklisting. Inheritance (nested rules) is supported. Logging and alert handling is robust, facilitating fast and accurate consumption of data. XDR captures events and changes made during an attack and automatically correlates separate activities into a single alert. The system provides built-in reports as well as the ability to generate custom reports in PDF.

FALSE POSITIVE RATE 3/645 (0.5%) AA

With a false positive rate of 0.5%, the Palo Alto Network Traps/Cortex XDR is unlikely to introduce much additional work for administrators.

RESISTANCE TO EVASION 49/49 (100%) AAA

The endpoint protection was capable of detecting and blocking malware and exploits when subjected to numerous evasion techniques.

BLOCK RATE 2,237/2,282 (98.03%) AA

ATTACKS	RATING	BLOCKED ON DOWNLOAD	BLOCKED ON EXECUTION	TOTAL BLOCKED	DETECTED	UNBLOCKED & UNDETECTED
Email Malware	AA	1,347	169	1,516	2	13
HTTP Malware	AA	306	115	421	-	3
Drive-by Exploits	AA	248	4	252	-	4
Social Exploits	BB	5	26	31	-	19
Handcrafted Attacks	A	-	17	17	-	4
TOTAL	AA	1,906	331	2,237	2	43
				98.03%	0.09%	1.88%

Results indicate that the product is highly capable. We found the protection against the vast majority of classic malware attacks to be excellent, as well as the protection against drive-by exploits. Protection against handcrafted (targeted) attacks was strong. However, the product underperformed when asked to protect against social exploits that were embedded in documents, blocking 31 of 50 attacks.

TOTAL COST OF OWNERSHIP $195,000

Expected Costs (2,500 Agents)	
Initial Purchase Price	$65,000
Annual Cost of Support/Maintenance	$0
Other Annual Cost (AV, IPS, Cloud, etc.)	$0
3-Year Total Cost of Ownership	**$195,000**
Total Cost Year 1	$65,000
Total Cost Year 2	$65,000
Total Cost Year 3	$65,000

Testing was based on the Advanced Endpoint Protection Test Methodology v4.0 (available at www.nsslabs.com)

Table of Contents

Table of Figures

Security

The threat landscape is evolving constantly; attackers are refining their strategies and increasing both the volume and complexity of their attacks. Enterprises now are having to defend against everyday cybercriminal attacks as well as targeted attacks and even the rare advanced persistent threats (APTs). For this reason, we tested using multiple commercial, open-source, and proprietary tools to employ attack methods that are currently being used by cybercriminals and other threat actors. We increased the levels of difficulty as we tested, beginning with common attacks, escalating to targeted attacks, and then applying obfuscation techniques to see if we could evade defenses. We then recorded whether the endpoint protection blocked and logged threats accurately and how frequently it triggered false positives.

Tuning and False Positives

Blocked	Detected	Rating
3/645 (0.5%)	-	AA

This test includes a varied sample of legitimate application traffic that may be falsely identified as malicious (also known as false positives). As part of the initial setup, we tuned the endpoint protection as it would be by a customer. Every effort was made to eliminate false positives while achieving optimal security effectiveness and performance, as would be the aim of a typical customer deploying the device in a live network environment. To ensure that the vendor did not deploy unrealistic (overly aggressive) security policies that blocked access to legitimate software and websites, we tested the endpoint protection against 645 false positive samples, including but not limited to the following file formats: .exe, .jar, .xls, .xlsm, .accdb, .css, .pdf, .doc, .docx, .zip, .DLL, .js, xls, .chm, .rar, .lnk, .cur, .xrc., .slk, .ppt, pptx, .iqy, .htm.

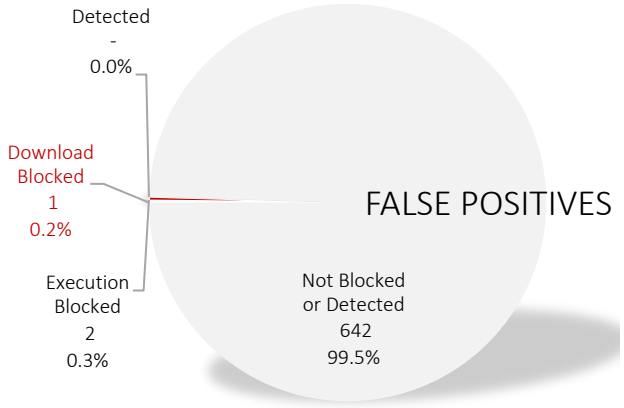

Figure 1 – False Positives

Resistance to Evasions

Blocked	Detected	Missed	Rating
49/49 (100.0%)	-	-	AAA

Threat actors apply evasion techniques to disguise and modify attacks at the point of delivery in order to avoid detection by security products. Therefore, it is imperative that endpoint protection correctly handles evasions. If an endpoint protection platform fails to detect a single form of evasion, an attack can bypass protection.

Our engineers verified that the endpoint protection was capable of detecting and blocking malware when subjected to numerous evasion techniques. To develop a baseline, we took several attacks that had previously been detected and blocked. We then applied evasion techniques to those baseline samples and tested. This ensured that any misses were due to the evasions and not the underlying (baseline) attacks.

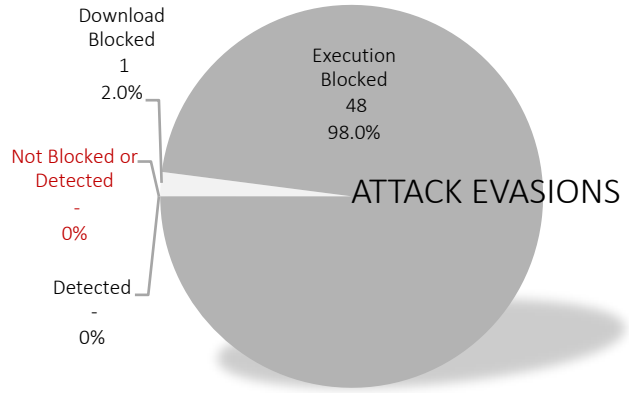

Figure 2 – Resistance to Evasions

For example, we applied an evasion technique called *process injection* where the original file is extracted from the binary and code is injected into a legitimate/trusted target process (i.e., Google Chrome). The malicious execution then occurs under the context of the target process (Chrome). Once these process injections techniques ran, we tried to further elude the detection by introducing anti-sandbox/anti-discovery evasions that employed techniques to determine whether or not the malware was on a user's machine; whether or not a security product was present; whether or not debugging or sandboxing was occurring; etc.

Malware Delivered over Email

Blocked	Detected	Missed	Rating
1,516/1,531 (99.0%)	2/1,531 (0.1%)	13/1,531 (0.8%)	AA

One of the most common ways in which users are compromised is through malware delivered over email. For several years, the use of social engineering has accounted for the bulk of cyberattacks against consumers and enterprises. Socially engineered malware attacks often use a dynamic combination of social media, hijacked email accounts, false notification of computer problems, and other deceptions to encourage users to download malware. One well-known social engineering attack method is spear phishing. Cybercriminals use hijacked email accounts to take advantage of the implicit trust between contacts and deceive victims into believing that the sender is trustworthy. The victim is tricked into opening the email attachment, which then launches the malicious malware program.

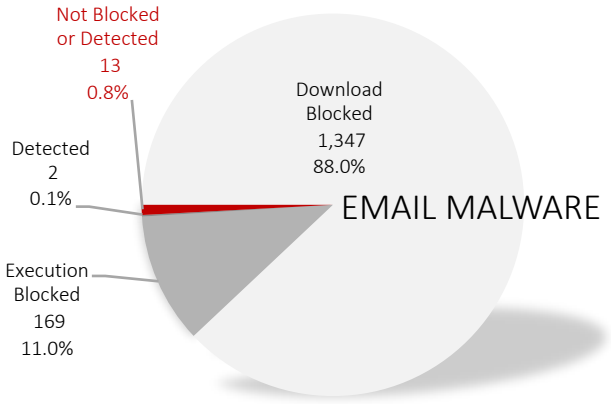

Figure 3 - Malware Delivered over Email

To test how well the endpoint protection is able to protect against this type of attack, malware was emailed to the user. The desktop client then retrieved the email and opened/executed the malware. If the malware was blocked, the corresponding time was recorded. We deployed a CentOS 7.7.1908 Linux mail store with kernel 3.10.0-957.5.1.el7.x86_64 running Dovecot v2.2.36 for IMAP as the mail server. Victim machines consisted of a combination of 32-bit and 64-bit Windows 7 endpoints and 64-bit Windows 10 endpoints.

Malware Delivered over HTTP

Blocked	Detected	Missed	Rating
421/424 (99.3%)	-	3/424 (0.7%)	AA

One of the more widespread threats to the enterprise involves attackers using websites to deliver malware. In these web-based attacks, the user is deceived into downloading and executing malware. For example, an employee may be tricked into downloading and installing a malicious application that claims it will “speed up your PC.”

In cases where an attacker is aiming for a large number of victims, the attacker may hijack widely used reputable websites to distribute the malware. However, in cases where an attacker plans to target specific individuals, the attacker typically would use an industry-specific “watering hole” plus one or more social engineering techniques to deceive a user into unknowingly installing malware.

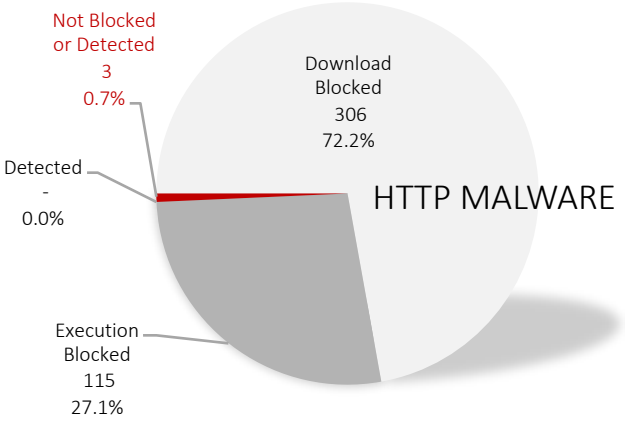

Figure 4 – Malware Delivered over HTTP

We tested the capability of the endpoint protection to protect against malware that was downloaded over HTTP and then executed (if the download was not blocked) using 424 malware samples against live victim machines running a combination of 32-bit and 64-bit Windows 7 endpoints and 64-bit Windows 10 endpoints, with various versions of Google Chrome, Mozilla Firefox, Microsoft Internet Explorer, and Microsoft Edge. Browser reputation systems were disabled so that the endpoint protection was not inadvertently credited for protection offered by a web browser.

Drive-by Exploits

Blocked	Detected	Missed	Rating
252/256 (98.4%)	-	4/256 (1.6%)	AA

While there are millions (or hundreds of millions) of malware samples in circulation at any given point in time, they are frequently delivered by exploits that target consumer desktops known as drive-by exploits.

In a drive-by exploit, an employee visits a website containing malicious code that exploits the user's computer and installs malware without the knowledge or permission of the user. An example of this would be where an employee visits WSJ.com (Wall Street Journal), which is inadvertently hosting an advertisement that contains an exploit. Another example (that we frequently observe in the wild) is where a user navigates to a URL and then is re-directed without interaction to a web page serving malicious content. Using this technique, a single exploit can silently deliver and install millions of malware samples to unsuspecting victims' computers.

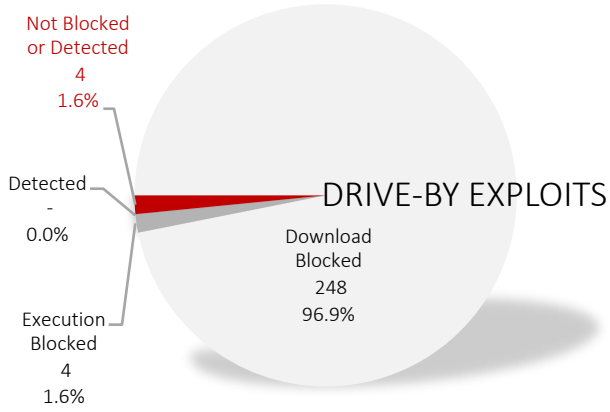

Figure 5 – Drive-by Exploits

To test how well the solution was able to protect against drive-by exploits, victim machines were deployed running 32-bit Windows 7 (version 6.1 (Build 7601: SP1) and 64-bit Microsoft Windows 10 (version 1709 (Build: 16299.15) with Microsoft Office (Office 16.0.7341.2032) and various versions of Google Chrome, Mozilla Firefox, Microsoft Internet Explorer, and Microsoft Edge. Depending on the victim machine, one or more of the following applications was installed: Java 8 Update 231, Microsoft Silverlight 5.1.20125, Adobe Flash Player 18.0.0.160, Adobe Reader DC 2017.012.20093, Adobe Reader 9.40, Java 6 Update 27, Adobe Flash Player 32.0.0.238, Java 8 Update 221, Microsoft Silverlight 5.1.50918, Adobe Flash Player 32.0.0.223, Java 8 Update 211, Adobe Flash Player 32.0.0.207, Internet Explorer 11, Internet Explorer 10, and Internet Explorer 9. Browser reputation systems were disabled so that the endpoint protection was not inadvertently credited for protection offered by a web browser.

While vulnerabilities are patched and defenses against exploits incorporated into new versions of operating systems (i.e., Windows), many organizations cannot easily upgrade due to financial, technical, or other constraints. As of January 2020, NetMarketShare[1] reports OS market share for Windows 7 (released 11 years ago in 2009) at 25.56% and for Windows 10 (released in 2015) at 57.08%.

Research has shown that oftentimes the most valuable assets have the most stringent change control to avoid business interruption. This creates a challenging dynamic whereby the most valuable assets tend to be the most difficult to defend (e.g., older OS, unpatched, etc.). Therefore, as vulnerabilities are patched and defenses against exploits are incorporated into new versions of operating systems (i.e., Windows)—which makes exploitation of computers more difficult—the value of endpoint protection is often associated with its ability to protect older, unpatched, and generally more vulnerable systems.

[1] https://netmarketshare.com

Social Exploits

Blocked	Detected	Missed	Rating
31/50 (62.0%)	-	19/50 (38.0%)	BB

Social exploits combine social engineering (manipulating people into doing what you want them to do) and exploitation (malicious code designed to take advantage of existing deficiencies in hardware or software systems, such as vulnerabilities or bugs). An example of this would be an email with "Your Bonus" as a subject line and containing a malicious spreadsheet labeled "bonus.xlsx" (which the employee opens).

As with drive-by exploits, these attacks are limited to specific operating systems and/or applications. However, the exploits contained within Excel spreadsheets or Word documents may target kernel functions or common functions such as object handling, which provides attackers with a wide attack surface. As such, sending social exploits through mass email (phishing), could yield profit as the number of victims would be large, albeit smaller than in the case of malware since exploits would have technical dependencies.

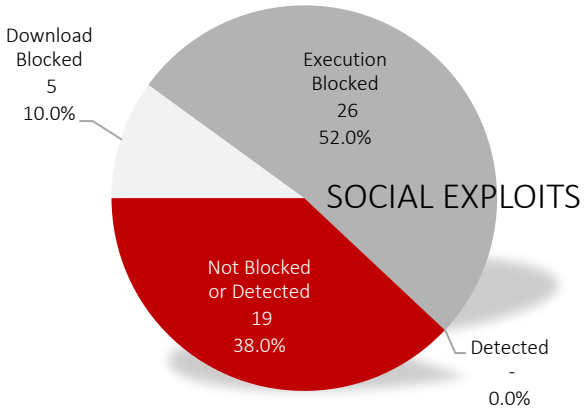

Figure 6 – Social Exploits

To test how well the product was able to protect against social exploits, we deployed 19 victim machines. All of the machines were running Windows 10 version 1709 (OS Build 16299.15). Machines were configured with Internet Explorer 11 (version 11.15.16299.0 – Update Version 11.0.47) and Microsoft Office 2016 (version 16.0.7431.2032).

Handcrafted (Targeted) Attacks

Blocked	Detected	Missed	Rating
17/21 (81.0%)	-	4/21 (19.0%)	A

The aim of this test was to see which endpoint products were able to protect customers while under adverse conditions dictated by the attacker. In this case, we wanted to find out which products could block new handcrafted (unknown) malware while being prevented from accessing cloud services.

What happens, for example, if an employee goes on a business trip to China where Internet traffic is tightly controlled? In such a scenario, access to the corporate VPN is likely blocked and the security software on the employee's laptop may not be able to receive updates or communicate in general. What happens if the employee's laptop is attacked with targeted malware?

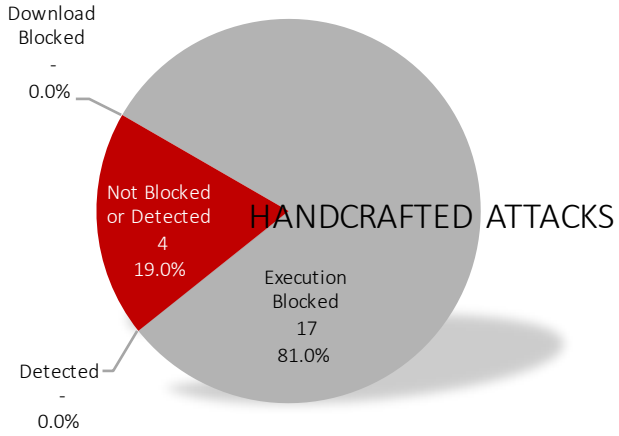

Figure 7 – Handcrafted (Targeted) Attacks

For the purposes of this test, handcrafted (targeted) malware was created by modifying the source code of keyloggers, ransomware, and destructoware, and then recompiling the binary so that it was new to the products being tested. We then attempted to infect a host (e.g., a laptop) with the malware and recorded whether or not the endpoint protection blocked the attack.

Because creating samples in this manner is a painstaking and time-consuming exercise, we tested only a handful of targeted samples; results should be viewed with this in mind.

Central management is available is offered as a cloud service and is accessed from a variety of supported web browsers.

Management & Reporting Capabilities

Rating

AA

Authentication

The management console supports role-based access control (RBAC) with numerous roles for users, such as “Hub Roles,” “Viewer,” “Security Admin,” “Privileged Security Admin,” “IT Admin,” “Privileged IT Admin,” “Deployment Admin,” and “No Role(default).” Authentication and RBAC are integrated with Active Directory with support for multi-factor authentication (MFA).

Policy

The management system provides several pre-defined policies as well as the ability to create and save multiple security policies, with no pre-defined limit. Administrators create profiles that can be cloned and modified from the default template. Whitelist and blacklist rules can be created and incorporated into policies. Policies may be applied to a single endpoint, groups of endpoints, or globally via Active Directory organizational units or static groups that are created and managed within the XDR console. When malicious activity is detected, the rule or module that triggered the event is presented in the security alert.

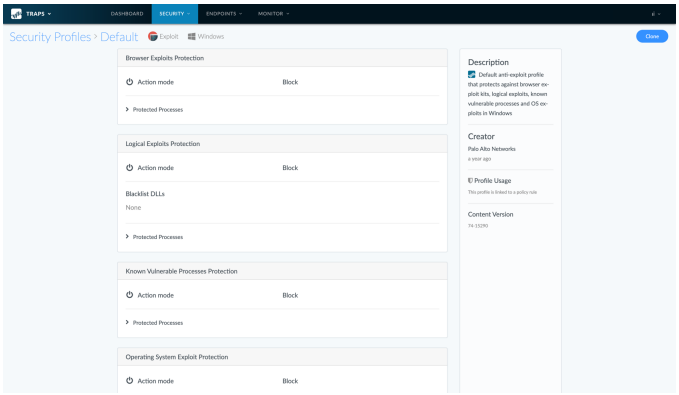

Inheritance (nested rules) is fully supported, including creation of groups and sub-groups such that sub-groups can inherit certain aspects of configuration and policy definition from parent groups. Versioning enables administrators to view the policy version of each endpoint. Bulk operations are supported, albeit via a slightly awkward mechanism.

Anti-tampering is handled by “service protection” that controls the registry, file system, and other parts of the local system.

Logging

Logging occurs centrally to the Cortex Data Lake; log time is normalized using a common time zone. When an incident occurs, the file, registry, process and network, data is captured and stored in the cortex data lake which can then be accessed from the XDR management interface.

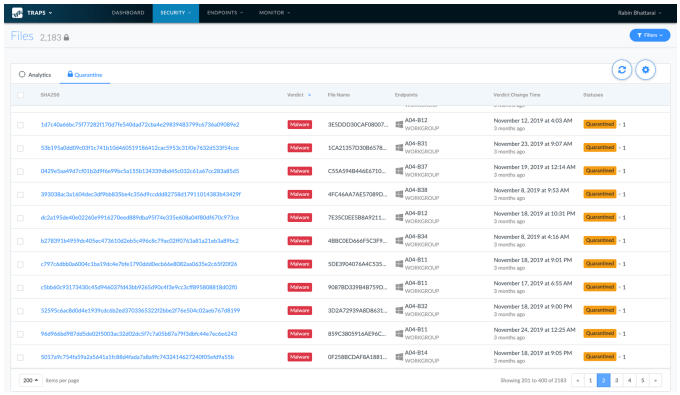

Policy changes and who made them, as well as any administrative actions such as session login/logout, user login/logout, successful authentication, and unsuccessful authentication attempts are all logged alongside disk quota and status of services. In the event of a service outage or hardware failure, the UI provides real-time information on all cloud hosted services. When an endpoint is power cycled, it is captured through the window event logs in the Cortex Data Lake. However, the system does not alert administrators when an endpoint is in a bad state, down, not reporting, or has detected issues; the current workaround is for administrators to create a custom report and schedule it to run as desired.

Change Control

The management system provides a change control/audit log that contains the username, date and time of the changes, as well as details of the change. Neither revision history nor policy rollback is supported.

Alert Handling

Palo Alto Networks captures all events and changes made during an attack and automatically correlates separate activities into a single alert that includes the computers involved, IP addresses, users, malicious samples, a common parent process in the attach chain, time, etc. In addition, analysts can manually add any event to an incident.

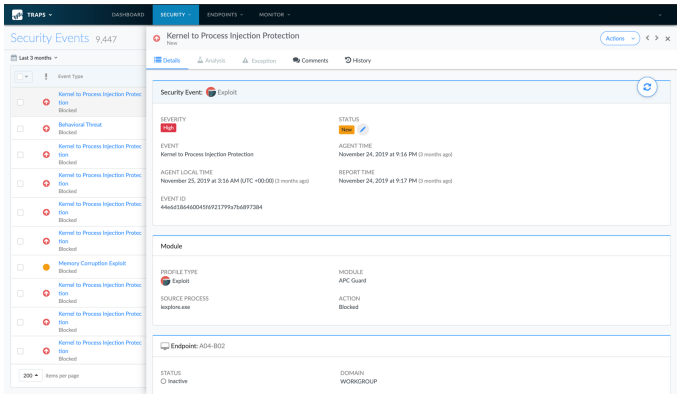

A variety of filters can be applied, including searching by "Endpoint ID," "Endpoint Name," "Event ID," "Event Type," "Platform," "Process/File Name," "SHA256," "Status," or "User," as well as searching by high, medium, or low severity levels. An analyst can select an alert and view all activity (including other alerts) associated with it, such as the endpoint, the user, the time, or a common parent process in the process chain.

An analyst may also select a particular piece of data and summarize on that data field (i.e., file, network, registry, process, or event log activity). These alerts are grouped into incidents in XDR.

Incidents can be annotated and tracked to resolution; analysts can assign an incident, change its status, and add notes.

Reporting

XDR provides pre-defined dashboards and reports as well as the ability to create and customize additional dashboards and reports that can be exported as PDFs. On-demand reporting allows administrators to query a customized range of time in order to audit the environment. Built-in reports cover typical requirements such as a list of top attacks, top source/destination IP addresses, top targets, etc. Reports can be scheduled to run daily, weekly, or monthly and can be delivered via email distribution.

Data may be forwarded (i.e., to SIEMs) in IETF syslog message format.

Total Cost of Ownership (TCO)

Implementation of security solutions can be complex, with several factors affecting the overall cost of deployment, maintenance, and upkeep. All of the following should be considered over the course of the useful life of a product:

- Initial Purchase – The cost of acquisition
- Maintenance/Subscription – Fees paid to the vendor for ongoing use of software and access to updates
- Technical Support – Fees paid to the vendor for 24/7 support

3-Year Total Cost of Ownership

Calculations are based on vendor-provided pricing information. Where possible, the 24/7 maintenance and support option with 24-hour replacement is used, since this is the option typically selected by enterprise customers. Prices include the purchase and maintenance costs for 2,500 software agents

- Year 1 Cost is calculated by adding purchase price + first-year maintenance/support fees.
- Year 2 Cost consists only of maintenance/support fees.
- Year 3 Cost consists only of maintenance/support fees.

Expected Costs for Palo Alto Networks Traps / Cortex XDR – 2,500 Agents	
Initial Purchase Price	$65,000
Annual Cost of Support/Maintenance	$0
Other Annual Cost (AV, IPS, Cloud etc.)	$0
3-Year Total Cost of Ownership	$195,000
Total Cost Year 1	$65,000
Total Cost Year 2	$65,000
Total Cost Year 3	$65,000

Figure 8 – 3-Year TCO (US$)

Test Environment

- BaitNET™ (NSS Labs Proprietary)
- 32-bit Microsoft Windows 7 (Version 6.1 (Build 7601: SP1)
- 64-bit Microsoft Windows 7 (Version 6.1 (Build 7601: SP1)
- 64-bit Microsoft Windows 10 (version 1607 (Build: 14393.0)
- 64-bit Microsoft Windows 10 (version 1709 (Build: 16299.15)
- Adobe Acrobat Reader 19.021.20061
- Adobe Flash Player 18.0.0.160
- Adobe Flash Player 32.0.0.207
- Adobe Flash Player 32.0.0.223
- Adobe Flash Player 32.0.0.238
- Adobe Reader 9.40
- Adobe Reader DC 2017.012.20093
- Google Chrome 78.0.3904.70
- Kali (Kernel release 4.19.0-kali1-amd64)
- Microsoft Internet Explorer 9.0.8112.16421
- Microsoft Internet Explorer 10.0.9200.16438
- Microsoft Internet Explorer 11.0.14393.0
- Microsoft Office Professional 2013 version 15.0.5119.1000 (Microsoft Word, Excel, PowerPoint, Access, etc.)
- Microsoft Office Professional 2016 version 16.0.7341.2032 (Microsoft Word, Excel, PowerPoint, Access, etc.)
- Microsoft Silverlight 5.1.20125
- Microsoft Silverlight 5.1.50918
- Oracle Java 6 Update 27
- Oracle Java 8 Update 181
- Oracle Java 8 Update 211
- Oracle Java 8 Update 221
- Oracle Java 8 Update 231
- Rapid7 Metasploit (v5.0.46-dev)
- VMware vCenter (Version 6.7u2 Build 6.7.0.30000)
- VMware vSphere (Version 6.7.0.30000)
- VMware ESXi (Version 6.7u3 Build 14320388)
- Wireshark version 3.0.3

Appendix

NSS LABS RATINGS	
RATING	DEFINITION
AAA	A product rated 'AAA' has the highest rating assigned by NSS Labs. The product's capacity to meet its commitments to consumers is extremely strong.
AA	A product rated 'AA' differs from the highest-rated products only to a small degree. The product's capacity to meet its commitments to consumers is very strong.
A	A product rated 'A' is somewhat more susceptible to sophisticated attacks than higher-rated categories. However, the product's capacity to meet its commitments to consumers is still strong.
BBB	A product rated 'BBB' exhibits adequate protection parameters. However, sophisticated or previously unseen attacks are more likely to negatively impact the product's capacity to meet its commitments to consumers.
	A product rated 'BB,' 'B,' 'CCC,' 'CC,' and 'C' is regarded as having significant risk characteristics. 'BB' indicates the least degree of risk and 'C' the highest. While such products will likely have some specialized capability and protective characteristics, these may be outweighed by large uncertainties or major exposure to adverse conditions.
BB	A product rated 'BB' is less susceptible to allowing a compromise than products that have received higher-risk ratings. However, the product faces major technical limitations, which could be exposed by threats that would lead to its inability to meet its commitments to consumers.
B	A product rated 'B' is more susceptible to allowing a compromise than products rated 'BB'; however, it currently has the capacity to meet its commitments to consumers. Adverse threat conditions will likely expose the product's technical limitations and expose its inability to meet its commitments to consumers.
CCC	A product rated 'CCC' is currently susceptible to allowing a compromise and is dependent upon favorable threat conditions for it to meet its commitments to consumers. In the event of adverse threat conditions, the product is not likely to have the capacity to meet its commitments to consumers.
CC	A product rated 'CC' is currently highly susceptible to allowing a compromise. The 'CC' rating is used when a failure has not yet occurred but NSS Labs considers a breach a virtual certainty, regardless of the anticipated time to breach.
C	A product rated 'C' is currently highly susceptible to allowing a compromise. The product is expected to fail to prevent a breach and to not have useful forensic information compared with products that are rated higher.
D	A product rated 'D' is actively being breached by known threats and is unable to protect consumers. For non-specialized products, the 'D' rating category is used when protecting a consumer is unattainable without a major technical overhaul. Unless NSS Labs believes that such technical fixes will be made within a stated grace period (often 30-90 calendar days), the 'D' rating also is an indicator that it is a virtual certainty that existing customers using the product have already experienced a breach—whether they know it or not—and should take immediate action.

Authors

Rabin Bhattarai, Thomas Skybakmoen, Vikram Phatak

Test Methodology

NSS Labs Advanced Endpoint Protection (AEP) Test Methodology v4.0 is available at www.nsslabs.com.

Contact Information

NSS Labs, Inc.

3711 South Mopac Expressway
Building 1, Suite 400
Austin, TX 78746

info@nsslabs.com

www.nsslabs.com