Independent Tests of
Anti-Virus Software

Endpoint Prevention and Response (EPR) Product Validation Report

Palo Alto Networks Cortex XDR Pro

TEST PERIOD: JUNE - SEPTEMBER 2023
LAST REVISION: 16TH OCTOBER 2023

WWW.AV-COMPARATIVES.ORG

Contents

Tested Product

Palo Alto Networks Cortex XDR Pro was tested as part of AV-Comparatives' Endpoint Prevention and Response (EPR) Test in summer 2023. The product version number was 8.0.2

Product Thumbnail

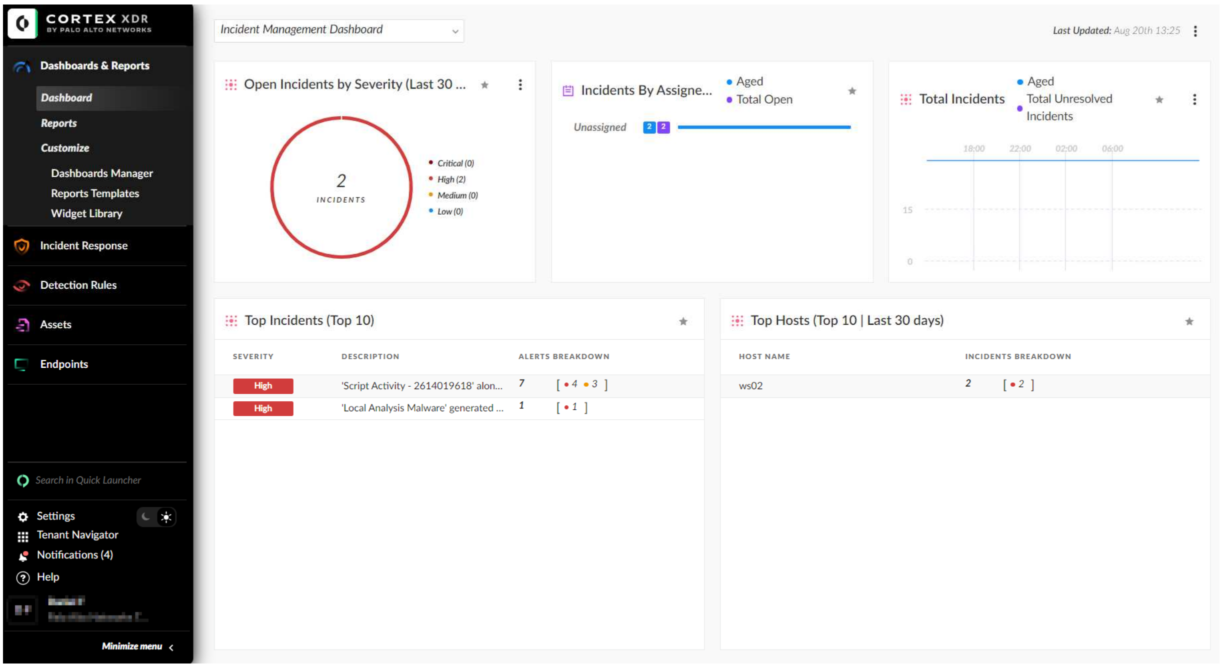

Top Incidents (Top 10)

SEVERITY	DESCRIPTION	ALERTS BREAKDOWN	
High	'Script Activity - 2614019618' alon...	7	[● 4 ● 3]
High	'Local Analysis Malware' generated ...	1	[● 1]

Top Hosts (Top 10 | Last 30 days)

HOST NAME	INCIDENTS BREAKDOWN	
ws02	2	[● 2]

Palo Alto Networks Cortex XDR Pro management console

Palo Alto Networks EPR Product: Executive Summary

Palo Alto Networks Cortex XDR Pro was tested by AV-Comparatives to validate if the product could provide effective enterprise prevention and response capabilities.

Palo Alto Networks Cortex XDR Pro did well at handling threats that are targeted towards enterprise users, in particular before the threats could progress inside and infiltrate the organisation's network. The product demonstrated several safeguards that helped in protecting the enterprise systems and network against the scenarios we tested. It should be noted that the product has very good correlation and post-detection capabilities that can terminate malicious processes in the event that they were not stopped by some other protection mechanism in an earlier phase.

The integration with Palo Alto Networks Wildfire Sandbox offers the ability to send unknown files to the sandbox, where additional analysis can be performed, and a verdict reached, with relative ease. Relevant threat alerts were demonstrated at the endpoint level, as well as in the cloud console, with an appropriate level of information. The product offers the ability to create different sets of behavioural rules, and good triaging ability for multiple users to collaborate on any given threat scenario at the same time. The endpoint agent also offers "remoting capabilities", which allows the analyst to investigate threats in real time.

The product had good mapping to MITRE's TTPs, thus providing low-level SOC analysts with the data needed to investigate further and escalate when necessary. Alerts were prioritized and aggregated, so as to minimize noise from all the alerts generated. The product can be easily configured and deployed in a domain or workgroup environment.

Active Response (Prevention): This occurs when the product stops the attack automatically, and reports it. Palo Alto Networks had an Active Response to **50/50** scenarios across all the phases tested. This resulted in a cumulative Active Response rate of **100**%.

Passive Response (Detection): This occurs when the product does not stop the specific attack phase, but reports suspicious activity. Palo Alto Networks had a Passive Response to **50/50** scenarios across all the phases tested. This resulted in a cumulative Passive Response rate of **100**%.

Operational Accuracy Costs: These occur when legitimate programs/actions are blocked/detected. Palo Alto Networks had **few costs** arising from imperfect Operational Accuracy.

Workflow Delay Costs: These arise e.g. when the user has to wait while a file is being analysed by the product. Palo Alto Networks had **no costs** relating to workflow delays.

Description	Details
EPR Certification Level Reached:	**Strategic Leader**
Overall **Active Response** Rate (Prevention Rate):	**98.7%**
Overall **Passive Response** Rate (Response Rate):	**98.7%**
Operational Accuracy Costs:	**Low**
Workflow Delay Costs:	**None**

Executive Summary

The table below depicts Palo Alto Networks' EPR prevention & detection rates across the different phases and categories of attack. For more details on the workflows and phases, please see the appendix.

Description	Number Tested
Scenarios	50
Phases	**Combined Prevention & Detection**
Phase 1 (Compromise & Foothold)	
Active Response (Prevention)	96.0%
Passive Response (Detection)	96.0%
Phase 2 (Internal Propagation)	
Active Response (Prevention)	100%
Passive Response (Detection)	100%
Phase 3 (Asset Breach)	
Active Response (Prevention)	N/A
Passive Response (Detection)	N/A
Operational Accuracy Costs	Few
Workflow Delay Costs	None

Combined Prevention & Detection Rates

Palo Alto Networks prevented 96% of the scenarios in Phase 1 (Compromise and Foothold). For the 2 scenarios (4%) that were able to progress to Phase 2 (Internal Propagation), Palo Alto Networks detected and acted upon all of them in this phase. Hence, none of the scenarios progressed to Phase 3.

The graphic below breaks down Palo Alto Networks' Active versus Passive Response capabilities for the duration of the test.

"**Not Applicable**" indicates that no test scenario was able to progress to Phase 3.

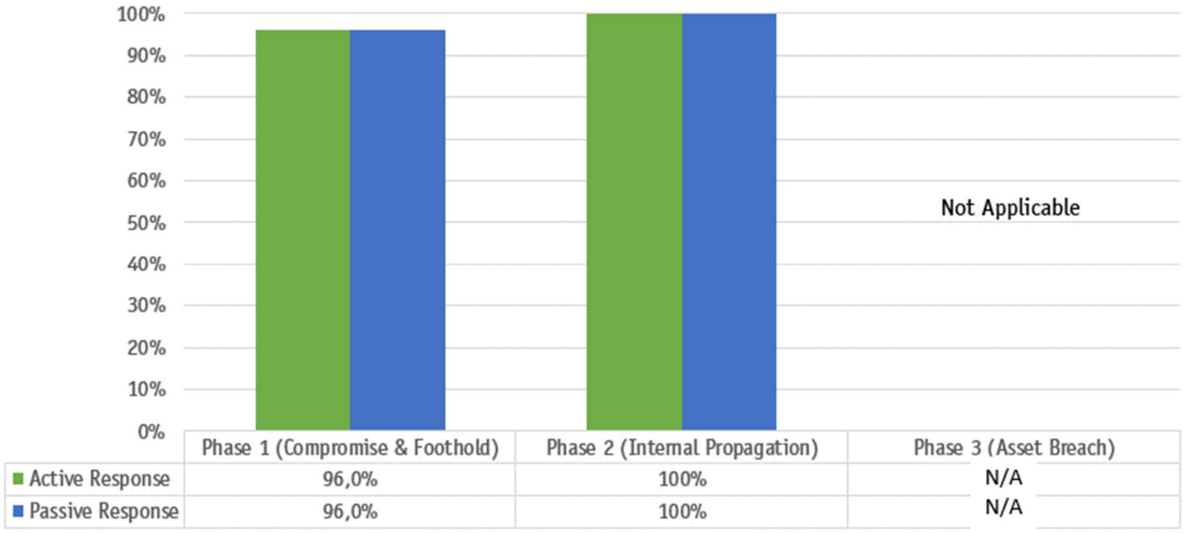

Active vs Passive Response of Palo Alto Networks Cortex XDR Pro

Modern threats usually come with layers of techniques to evade prevention and response, such as encryption, obfuscation, anti-analysis, packing, file-less malware, exploit, and privilege escalation.

AV-Comparatives' Enterprise EPR methodology covers some of the most prevalent enterprise scenarios and system-administrator EPR workflows, specifically requested by enterprises based on inquiries and primary research.

Cumulative Prevention and Response by phases

Response Type	Phase 1 Only	Phase 1 & 2	Overall (Phase 1, 2 & 3)
Active Response	96.0% (48/50)	100% (50/50)	100% (50/50)
Passive Response	96.0% (48/50)	100% (50/50)	100% (50/50)

Cumulative Prevention and Response by Phase

The graphic below depicts Palo Alto Networks' Active and Passive Response capabilities in the three attack phases tested.

"**Not Applicable**" indicates that no test scenario was able to progress to Phase 3.

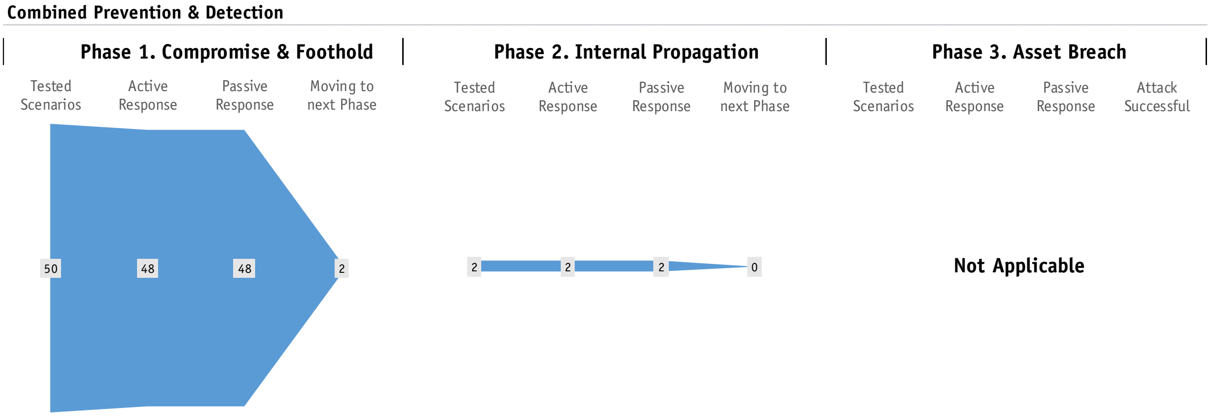

EPR Efficacy per Phase of Palo Alto Networks Cortex XDR Pro

Phase 1:

- 48 out of 50 scenarios prevented.
- 48 out of 50 scenarios detected.
- 2 scenarios were able to progress to Phase 2.

Phase 2:

- 2 out of 2 scenarios prevented.
- 2 out of 2 scenarios detected.
- No scenario was able to progress to Phase 3.

Phase 3:

- Not applicable, because no scenario was able to progress to Phase 3.

MITRE ATT&CK Matrix for Enterprise

The diagram below[1] shows the entire MITRE ATT&CK Matrix for Enterprise[2]. The column headings represent the ATT&CK Tactics[3] (aims), while the boxes below them represent the ATT&CK Techniques[4] used to achieve those goals. Our EPR test covers the entire attack chain shown here, using the most realistic possible scenarios. Across the 50 attack scenarios used in this EPR test, we tried to employ all of the Techniques shown in the green boxes below.

The Tactics relate to our 3 attack Phases as follows:
Phase 1 = Initial Access, Execution, Persistence
Phase 2 = Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement
Phase 3 = Collection, Command and Control, Exfiltration, Impact

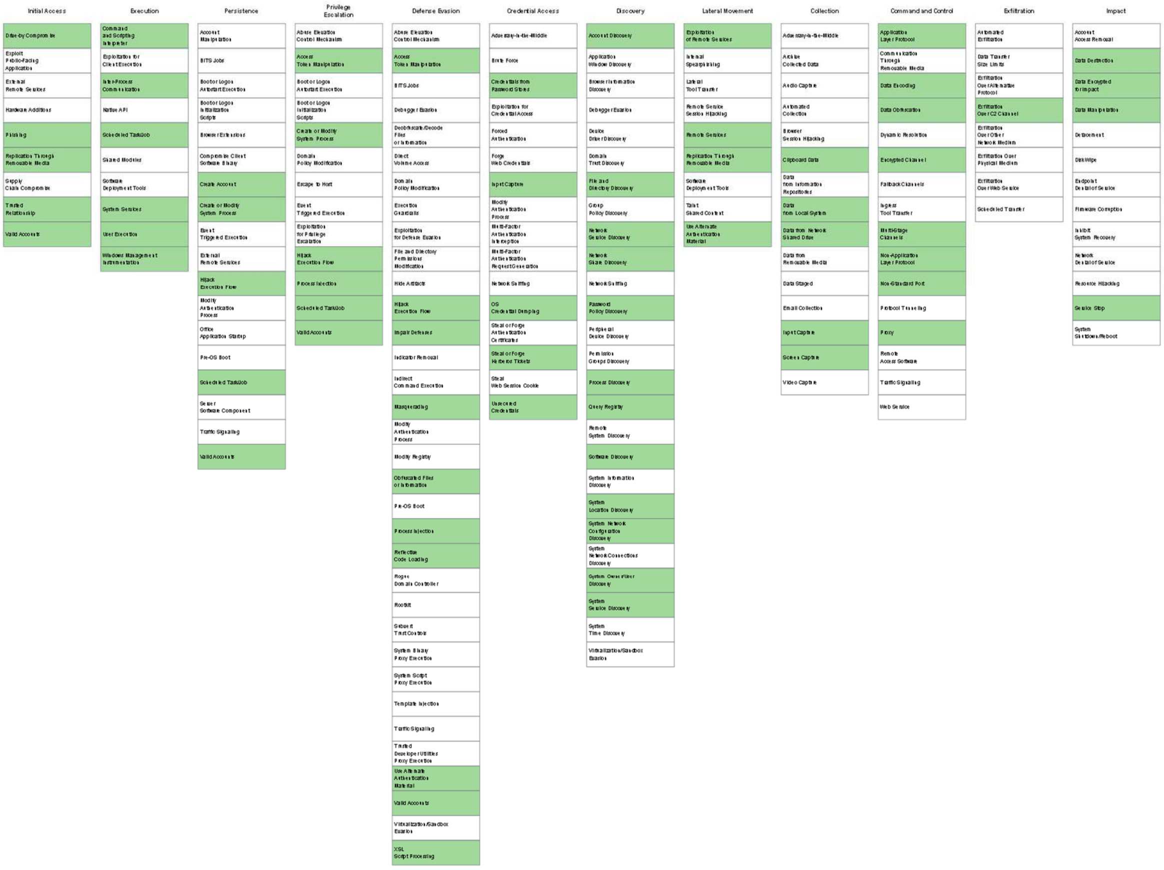

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
🟩 Drive-by Compromise	🟩 Command and Scripting Interpreter	Account Manipulation	Abuse Elevation Control Mechanism	Abuse Elevation Control Mechanism	Adversary-in-the-Middle	🟩 Account Discovery	🟩 Exploitation of Remote Services	Adversary-in-the-Middle	🟩 Application Layer Protocol	Automated Exfiltration	Account Access Removal
Exploit Public-Facing Application	Exploitation for Client Execution	BITS Jobs	🟩 Access Token Manipulation	🟩 Access Token Manipulation	Brute Force	Application Window Discovery	Internal Spearphishing	Archive Collected Data	Communication Through Removable Media	Data Transfer Size Limits	🟩 Data Destruction
External Remote Services	🟩 Inter-Process Communication	Boot or Logon Autostart Execution	Boot or Logon Autostart Execution	BITS Jobs	🟩 Credentials from Password Stores	Browser Information Discovery	Lateral Tool Transfer	Audio Capture	🟩 Data Encoding	Exfiltration Over Alternative Protocol	🟩 Data Encrypted for Impact
Hardware Additions	Native API	Boot or Logon Initialization Scripts	Boot or Logon Initialization Scripts	Debugger Evasion	Exploitation for Credential Access	Debugger Evasion	Remote Service Session Hijacking	Automated Collection	🟩 Data Obfuscation	🟩 Exfiltration Over C2 Channel	🟩 Data Manipulation
🟩 Phishing	🟩 Scheduled Task/Job	Browser Extensions	🟩 Create or Modify System Process	Deobfuscate/Decode Files or Information	Forced Authentication	Device Driver Discovery	🟩 Remote Services	Browser Session Hijacking	Dynamic Resolution	Exfiltration Over Other Network Medium	Defacement
🟩 Replication Through Removable Media	Shared Modules	Compromise Client Software Binary	Domain Policy Modification	Direct Volume Access	Forge Web Credentials	Domain Trust Discovery	🟩 Replication Through Removable Media	🟩 Clipboard Data	🟩 Encrypted Channel	Exfiltration Over Physical Medium	Disk Wipe
Supply Chain Compromise	Software Deployment Tools	🟩 Create Account	Escape to Host	Domain Policy Modification	🟩 Input Capture	🟩 File and Directory Discovery	Software Deployment Tools	Data from Information Repositories	Fallback Channels	Exfiltration Over Web Service	Endpoint Denial of Service
🟩 Trusted Relationship	🟩 System Services	🟩 Create or Modify System Process	Event Triggered Execution	Execution Guardrails	Modify Authentication Process	Group Policy Discovery	Taint Shared Content	🟩 Data from Local System	Ingress Tool Transfer	Scheduled Transfer	Firmware Corruption
🟩 Valid Accounts	🟩 User Execution	Event Triggered Execution	Exploitation for Privilege Escalation	Exploitation for Defense Evasion	Multi-Factor Authentication Interception	🟩 Network Service Discovery	🟩 Use Alternate Authentication Material	🟩 Data from Network Shared Drive	🟩 Multi-Stage Channels		Inhibit System Recovery
	🟩 Windows Management Instrumentation	External Remote Services	🟩 Hijack Execution Flow	File and Directory Permissions Modification	Multi-Factor Authentication Request Generation	🟩 Network Share Discovery		Data from Removable Media	🟩 Non-Application Layer Protocol		Network Denial of Service
		🟩 Hijack Execution Flow	🟩 Process Injection	Hide Artifacts	Network Sniffing	Network Sniffing		Data Staged	🟩 Non-Standard Port		Resource Hijacking
		Modify Authentication Process	🟩 Scheduled Task/Job	🟩 Hijack Execution Flow	🟩 OS Credential Dumping	🟩 Password Policy Discovery		Email Collection	Protocol Tunneling		🟩 Service Stop
		Office Application Startup	🟩 Valid Accounts	🟩 Impair Defenses	Steal or Forge Authentication Certificates	Peripheral Device Discovery		🟩 Input Capture	🟩 Proxy		System Shutdown/Reboot
		Pre-OS Boot		Indicator Removal	🟩 Steal or Forge Kerberos Tickets	Permission Groups Discovery		🟩 Screen Capture	Remote Access Software		
		🟩 Scheduled Task/Job		Indirect Command Execution	Steal Web Session Cookie	🟩 Process Discovery		Video Capture	Traffic Signaling		
		Server Software Component		🟩 Masquerading	🟩 Unsecured Credentials	🟩 Query Registry			Web Service		
		Traffic Signaling		Modify Authentication Process		Remote System Discovery					
		🟩 Valid Accounts		Modify Registry		🟩 Software Discovery					
				🟩 Obfuscated Files or Information		System Information Discovery					
				Pre-OS Boot		🟩 System Location Discovery					
				🟩 Process Injection		🟩 System Network Configuration Discovery					
				🟩 Reflective Code Loading		System Network Connections Discovery					
				Rogue Domain Controller		🟩 System Owner/User Discovery					
				Rootkit		🟩 System Service Discovery					
				Subvert Trust Controls		System Time Discovery					
				System Binary Proxy Execution		Virtualization/Sandbox Evasion					
				System Script Proxy Execution							
				Template Injection							
				Traffic Signaling							
				Trusted Developer Utilities Proxy Execution							
				🟩 Use Alternate Authentication Material							
				🟩 Valid Accounts							
				Virtualization/Sandbox Evasion							
				🟩 XSL Script Processing							

MITRE ATT&CK Tactics and Techniques covered by this EPR Test

For a magnified view of the above table, please click here: https://www.av-comparatives.org/wp-content/uploads/2023/09/EPR2023.svg

An example scenario might look like this: phishing mail with script payload is sent to user on Workstation A – internal discovery is performed – access to C$ share on Workstation B is found – lateral movement to Workstation B – network admin session on Workstation B is found – LSASS dumped to obtain admin credentials – lateral movement to Server 1 – defence evasion used to bypass security product on Server 1 – credit-card data found – data is extracted via open C2 channel.

[1] Generated with https://mitre-attack.github.io/attack-navigator/
[2] https://attack.mitre.org/matrices/enterprise/
[3] https://attack.mitre.org/tactics/enterprise/
[4] https://attack.mitre.org/techniques/enterprise/

Phase 1 Metrics: Endpoint Compromise and Foothold

The Phase 1 content of the executed attacks can be described by means of MITRE ATT&CK and other frameworks. The following Tactics are part of this phase.

Initial Access: Initial access is the method used by the attacker to get a foothold inside the environment that is being targeted. Attackers may use a single method, or a combination of different techniques. Threats may come from compromised websites, email attachments or removable media. Methods of infection can include exploits, drive-by downloads, spear phishing, macros, trusted relationships, valid accounts, and supply-chain compromises.

Execution: The next goal of the attacker is to execute their own code inside the target environment. Depending upon the circumstances, this could be done locally or via remote code execution. Some of the methods used include client-side execution, third-party software, operating-system features like PowerShell, MSHTA, and the command line.

Persistence: Once the attacker gets inside the target environment, they will try to gain a persistent presence there. Depending upon the target operating system, an attacker may use operating-system tools and features. These include registry manipulation, specifying dynamic-link-library values in the registry, shell scripts that can contain shell commands, application shimming, and account manipulation.

Palo Alto Networks Cortex XDR Pro was subjected to the various attack steps as highlighted above and described in detail in AV-Comparatives' EPR CyberRisk Test Methodology. The resulting table below showcases the product's Active Response and Passive Response capabilities for the attack scenarios in Phase 1.

Tested Scenario	Description	Active Response	Passive Response
1	Metasploit Framework - Binary Direct SysCalls	✓	✓
2	Metasploit Framework - Binary Asynchronous Procedure Call Injection	✓	✓
3	Metasploit Framework - Binary Indirect SysCalls	✓	✓
4	Metasploit Framework - Visual Basic Script	✓	✓
5	Metasploit Framework - Staged MSIexec	✓	✓
6	Metasploit Framework - JavaScript DLL Sideload	✓	✓
7	Metasploit Framework - Staged DLL via Rundll32	✓	✓
8	Metasploit Framework - PowerShell Script with AMSI and ETW Patch	✓	✓
9	Metasploit Framework - Staged HTA	✓	✓
10	Metasploit Framework - Visual Basic Script and AMSI Patch	✓	✓
11	PowerShell Empire - Masqueraded Binary Indirect SysCalls	✓	✓
12	PowerShell Empire - Binary UUID Exec	✓	✓
13	PowerShell Empire - Visual Basic Script with obfuscated strings	✓	✓
14	PowerShell Empire - Stageless MSIexec	✓	✓
15	PowerShell Empire - Stageless Visual Basic Script	✓	✓
16	PowerShell Empire - Excel Shellcode Injection via VBS	✓	✓
17	PowerShell Empire - Stageless DLL via Rundll32	✓	✓
18	PowerShell Empire - PowerShell Script with AMSI Patch	✓	✓
19	PowerShell Empire - Stageless HTA	✓	✓

20	PowerShell Empire - Visual Basic Script	✓	✓
21	Commercial Framework - Masqueraded Binary Indirect SysCalls Shellcode	✓	✓
22	Commercial Framework - Masqueraded Binary NTAPI and ETW Bypass	✓	✓
23	Commercial Framework - Process Injection into Excel via PPT Macro	✓	✓
24	Metasploit Framework - Binary with Invalid Code Signature and UUID Exec	✓	✓
25	Metasploit Framework - Masqueraded Binary and ETW-Patch	✓	✓
26	Metasploit Framework - Obfuscated JavaScript DLL Sideloading	✓	✓
27	Metasploit Framework - Obfuscated Visual Basic Script non-standard port	✓	✓
28	Metasploit Framework - Packed MSIexec non-standard port	✓	✓
29	Metasploit Framework - Binary Process Hollowing and ETW-Patch	✓	✓
30	Metasploit Framework - Encrypted DLL via Rundll32	✓	✓
31	Metasploit Framework - Stageless obfuscated PowerShell Script	✓	✓
32	Metasploit Framework - Obfuscated HTA	✓	✓
33	Metasploit Framework - Obfuscated Visual Basic Script shellcode fetch	✓	✓
34	Metasploit Framework - Binary NTAPI	✓	✓
35	Metasploit Framework - JavaScript DLL Sideload NTAPIs	✓	✓
36	PowerShell Empire – Obfuscated .PIF file and ETW-Patch	✓	✓
37	PowerShell Empire - Masqueraded obfuscated .SCR file SysCalls	✗	✗
38	PowerShell Empire - HTML file (.chm) process injection into Office process	✓	✓
39	PowerShell Empire - Visual Basic Script shellcode fetch	✓	✓
40	PowerShell Empire - Packed MSI	✓	✓
41	PowerShell Empire - Binary DLL Sideloading (Process Hollowing)	✓	✓
42	PowerShell Empire - DLL shellcode fetch via rundll32	✓	✓
43	PowerShell Empire - Heavily Obfuscated PowerShell Script	✓	✓
44	PowerShell Empire - Stageless obfuscated HTA	✓	✓
45	PowerShell Empire - Visual Basic Script Win32 APIs	✓	✓
46	PowerShell Empire - Packed MSI	✓	✓
47	PowerShell Empire - JavaScript DLL Sideload via MSIexec	✓	✓
48	Commercial Framework - Encrypted JavaScript DLL Sideload	✓	✓
49	Commercial Framework - Masqueraded Binary with obfuscated shellcode	✗	✗
50	Commercial Framework - Encrypted Control Panel Applet Application	✓	✓

Phase 1: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

✗ - Indicates the product **failed** to prevent/detect the attack in the tested scenario during this phase.
✓ - Indicates the product **successfully** prevented/detected the attack in the tested scenario during this phase.

In 48 out of 50 test scenarios in Phase 1, Palo Alto Networks provided both a Passive Response (detection) and an Active Response (prevention). For two test cases, there was neither an Active nor a Passive Response.

Phase 2 Metrics: Internal Propagation

In this phase, the EPR product should be able to prevent internal propagation. This phase is triggered if the attack is not stopped in Phase 1. The EPR product in this phase should enable the system administrator to immediately identify and track the internal propagation of the threat in real time. We have explained below the relevant Tactics from the MITRE ATT&CK Framework.

Privilege Escalation: In enterprise networks, it is standard practice for users (including system admins on their own personal computers) to use standard user accounts without administrator privileges. If an enterprise endpoint is attacked, the logged-on account will not have the permissions the attacker requires to launch the next phase of the attack. In these cases, privilege escalation must be obtained, using techniques such as user-access token manipulation, exploitation, application shimming, hooking, or permission weakness. Once the adversary has got a foothold inside the environment, they will try to escalate the privileges. For an active response to be credited, we looked at various phases inside each method to see if there was a preventative action by the product.

Defense Evasion: The attacker's aim is to carry out their objectives without being detected or blocked. Defense Evasion consists of measures used to ensure that the attack remains undiscovered. This could include tampering with security software, obfuscating processes, and abusing e.g. system tools so as to hide the attack.

Credential Access: This is a method used by the attacker to ensure their further activities are carried out using a legitimate network user account. This means that they can access the resources they want, and will not be flagged as an intruder by the system's defences. Different credential-access methods can be used, depending on the nature of the targeted network. Credentials can be obtained on-site, using a method such as input capture (e.g., keyloggers). Alternatively, it might be done using the offline method, where the attacker copies the entire password database off-site, and can then use any method to crack it without fear of discovery.

Discovery: Once the attacker has gained access to the target network, they will explore the environment, with the aim of finding those assets that are the ultimate target of the attack. This is typically done by scanning the network.

Lateral Movement: The attacker will move laterally within the environment, so as to access those assets that are of interest. Techniques used include pass the hash, pass the ticket, and exploitation of remote services and protocols like RDP.

Tested Scenario	Description	Active Response	Passive Response
37	PowerShell Empire - Masqueraded obfuscated .SCR file SysCalls	✓	✓
49	Commercial Framework - Masqueraded Binary with Obfuscated Shellcode	✓	✓

Phase 2: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

✗ - Indicates the product **failed** to prevent/detect the attack in the tested scenario during this phase.

✓ - Indicates the product **successfully** prevented/detected the attack in the tested scenario during this phase.

In both test scenarios in Phase 2, Palo Alto Networks provided both a Passive Response (detection) and an Active Response (prevention).

Phase 3 Metrics: Asset Breach

The final phase of the workflow is asset breach. This is the stage where an attacker starts carrying out their ultimate objective. We have explained below the relevant Tactics from the MITRE ATT&CK Framework.

Collection: This involves gathering the target information – assuming of course that information theft, rather than sabotage, is the object of the exercise. The data concerned could be in the form of documents, emails or databases.

Command and Control: A Command-and-Control mechanism allows communication between the attacker's system and the targeted network. This means that the attacker can send commands to, or receive data from, the compromised system. Typically, the attacker will try to mask such communications by disguising them as normal network traffic.

Exfiltration: Once the attacker has reached the objective of collecting the target information, they will want to copy it covertly from the targeted network to their own server. In almost all cases, exfiltration involves the use of a command-and-control infrastructure.

Impact: This can be defined as the direct damage done to the targeted organisation's network. It includes the manipulation, disruption or destruction of operational systems and/or data. This might be an end in itself (sabotage), or a means of covering up data theft, by making it more difficult to investigate the breach.

Tested Scenario	Description	Active Response	Passive Response
N/A	N/A	N/A	N/A

Phase 3: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

As previously mentioned, Phase 3 scenarios were **N/A** (**not applicable**) for Palo Alto Networks, as the threats had already been prevented in a previous phase.

Operational-Accuracy and Workflow-Delay Costs

Costs arising from imperfect operational accuracy and workflow delays are calculated as follows.

Costs arising from imperfect operational accuracy

Operational accuracy testing was performed by simulating a typical user activity in the enterprise environment. This included opening clean files of different types (such as executables, scripts, documents with macros) and browsing to different clean websites. Furthermore, different administrator-friendly tools and scripts were also executed in the test environment to ensure that productivity was not affected by the respective product configuration used for the test.

To assess operational accuracy, each product is tested with a battery of clean scenarios. Over-blocking or over-reporting of such scenarios means that a product reaches high prevention and detection rates, but also causes increased costs. Where legitimate programs/actions are blocked, the system administrator will have to investigate, restore/reactivate any blocked programs etc, and take steps to prevent it happening again. The principle of "The boy who cried wolf" may also apply; the greater the number of false alerts, the more difficult it becomes to recognise a genuine alert.

Products are then assigned to one of five Groups (None, Low, Moderate, High, and Very High, whereby lower is better), according to the number of affected scenarios. These are shown in the table below.

		Operational Accuracy	
Group	Number of affected scenarios	*Active Response Multiplying Factor*	*Passive Response Multiplying Factor*
None	0	x0	x0
Low	1	x1	x0.75
Moderate	2-3	x5	x3.75
High	4-5	x10	x7.5
Very High	6+	x20	x15

Multiplying factors for Operational Accuracy costs

The costs arising from imperfect Operational Accuracy are worked out using Cost Units of USD 1.72 million. The number of Cost Units a product is deemed to have caused is calculated using a Multiplying Factor. This varies according to the Group, and also whether the scenario was affected by an Active Response (action blocked), or by a Passive Response (action not blocked, but detection alert shown in the console). The Multiplying Factor for an erroneous Passive Response is always three-quarters of that of an erroneous Active Response, because less time and effort is required to resolve the problem.

How this works in practice is best explained by looking at the table above. Products in the "None" Group have a Multiplying Factor of 0 for both Active and Passive Responses, therefore Operational Accuracy costs are zero. Products in the "Low" Group (1 affected scenario) have a Multiplying Factor of 1 for erroneous Active Responses, but only 0.75 for an erroneous Passive Response. Hence, a product with one erroneous Active Response incurs one Cost Unit, while a product with one erroneous Passive Responses only incurs 0.75 Cost Units. If a product had 2 affected scenarios, one being an Active Response, the other a Passive Response, it would incur 8.75 Cost Units (5 for the Active Response, and 3.75 for the Passive Response).

Costs arising from workflow delays

Some EPR products will cause delays in the user's workflow because they e.g. stop the execution of a previously unknown file and send it to the vendor's online sandbox for further analysis. Due to this behaviour, execution is stalled, and the user is not able to proceed till the analysis comes back from the sandbox. We noted the delay caused by such analysis, for both scenarios (clean and malicious). Where a product caused significant delays when analysing a scenario, this was penalised. The analysis time for each product was calculated as follows. For *clean* scenarios, we took the longest observed delay for any one scenario. So, for example, a product with two delays - of 2 minutes and 10 minutes respectively - for *clean* scenarios would have a recorded time of 10 minutes. For *malicious* scenarios, we took the average of all the delays. So, a product with two delays - of 2 minutes and 10 minutes respectively - for *malicious* scenarios, would have a recorded time of 6 minutes. Products are then assigned to one of five Workflow Delay Groups (None, Low, Moderate, High and Very High), depending on how long the respective delay is. These are shown in the table below.

Group	Delay Caused (in minutes)	Workflow Delay Multiplying Factor
None	under 2	x0
Low	2-5	x0.5
Moderate	6-10	x2.5
High	11-20	x5
Very High	over 20	x10

Multiplying factors for Workflow Delay costs

The costs of these delays are calculated using the same Cost Units as for operational accuracy. Again, there is a multiplying factor, which varies according to the Workflow Delay Group. Products in the Low Workflow Delay Group have a Multiplying Factor of 0.5, hence incurring costs of 1 Cost Unit; products in the Very High Workflow Delay Group have a Multiplying Factor of 10, thus incurring costs of 10 Cost Units. Products in the latter category would be disqualified from certification, due to the excessive costs incurred.

Results

The costs arising from imperfect Operational Accuracy and Workflow Delays are shown below:

	Operational Accuracy		Workflow Delays
	Active Response	*Passive Response*	
Palo Alto Networks	Low	None	None

Combined results table for Operational Accuracy and Workflow Delays

Palo Alto Networks incurred some (minor) Operational Accuracy costs for Active Responses.

EPR Competitive Product Differentiator (provided by Palo Alto Networks)

1. Behavioural Protection for Windows, macOS, and Linux.
2. Exploit protection by techniques for "any" process you want to add (Windows, macOS, and Linux).
3. Integrated sandbox analyses for all unknown samples and displaying the full report for each of them (Windows, macOS, and Linux).
4. Live terminal with Full CMD, PowerShell, Shell and embedded python (Windows, macOS, and Linux).
5. Isolation in Bulk-Script Execution (Python) in bulk (Windows, macOS, and Linux).
6. Leverage AI-based local analysis and Behavioural Threat Protection, including custom behavioural detection and prevention rules to stop the most malware, exploits, and fileless attacks in the industry.
7. Collect and correlate data from Palo Alto Networks and third-party tools to detect, triage, investigate, hunt, and respond to threats.
8. Use always-on AI-based analytics and custom rules to detect advanced persistent threats and other covert attacks.
9. Simplify investigations with stitching across various alert sources providing a unified incident engine, resulting in a 98% reduction in alerts and lowering the skill required to triage alerts.
10. Integrated Advanced Query Language, which allows complex queries against data stored in Cortex XDR.
11. Simplify response with recommended next steps for remediation. You can rapidly recover from an attack by removing malicious files and registry keys, as well as restoring damaged files and registry keys.
12. Enable behavioural analytics on logs collected from third-party firewalls while integrating third-party alerts into a unified incident view and root cause analysis for faster, more effective investigations.
13. Permission control is granular, from Super-Admin to Read-Only. Customizable Role and Scope based access is available.
14. Automatically analyses every executable on Windows/Mac/Linux and generates detailed behavioural accessible reports as part of the malware prevention process.
15. Incident management with an automatic score system based on context to improve analyst workflow.
16. Simple Automation capabilities that bridge the gap between analyst workflow and full SOAR use.
17. Identity Threat Module with multiple Identity provider support

Product features

In this section, we provide an overview of the products' features and the associated services provided by their respective vendors. Please note that in each case, these refer only to the specific product, tier and configuration used in our test. A different product/tier from the same vendor may have a different feature set. On the following pages we describe the General features, Product Response, Management and Reporting, IOC Integration features, Support features, Support features and then provide a feature list showing which products support these features.

General features

This section looks at general features such as phishing protection, web access control, device control, interface languages, and supported operating systems.

Product Response Mechanism

EPR products will use their response mechanisms to deal with the intrusions that have occurred inside the protected environment. At a minimum, an EPR product is expected to allow the correlation of endpoints, processes and network communications, as well as the correlation of external IOCs with the internal environment. EDR capabilities were tested and examined by using the detection and response capabilities of the product. We were able to examine the events that correlated with the various steps that attacker took while attempting to breach the environment.

The EPR product should enable complete visibility of the malicious artifacts/operations that make up the attack chain, making any response-based activities easy to complete. This means that where any form of intended remediation mechanism is available in the product (Response Enablement), this mechanism is shown below. Please note that the capabilities shown below only apply to the specific product/version used in this test. A vendor might offer additional features as an add-on or in another product.

Central Management and Reporting

Management workflow is a top differentiator for enterprise security products. If a product is difficult to manage, it will not be used efficiently. The intuitiveness of a product's management interface is a good determiner of how useful the product will be. Minutes saved per activity can translate into days and even weeks over the course of a year.

Management: Threat Visibility, System Visibility, and Data Sharing

The ability to provide threat context is a key component of an EPR product. This visibility can be critical when organizations are deciding whether to either supplement an existing technology or replace it. The management console can be deployed as physical appliance, virtual appliance, or cloud-based appliance. A full trail of audit logs is available in the management console. Communication between the agent and management console is done via SSL. The following tables provide information on the applicable capabilities of each of the tested products.

EPR Product Reporting Capabilities

An EPR platform should have the ability to unify data, that is to say, bring together information from disparate sources, and present it all within its own UI as a coherent picture of the situation. Technical integration with the operating system and third-party applications (Syslog, Splunk, SIEM or via API) is an important part of this. An EPR system should be able to offer response options appropriate to the organization.

IOC Integration

This is to identify the digital footprint by means of which the malicious activity on an endpoint/network can be identified. We will examine this use case by looking at the EPR product's ability to use external IOCs including Yara signatures or threat intelligence feeds etc. as shown in the table below.

Support features

Free, basic human support for deployment: this means real-time communication with a member of the support staff, who will talk you through the deployment process and can provide immediate answers to any basic questions you have. Of course, many vendors will provide user manuals, videos and premium (paid-for) deployment support services instead/in addition.

Professionally assisted training: this includes any form of interactive training with an instructor. A few vendors include professional training as part of the license fee paid for 5,000 clients, while others charge additionally for it. Some other vendors might only offer videos and other online material for self-training.

Feature List

Below you can find the list of features. Please note that this only applies to the test product and version (8.0.2).

Feature List	
Product Name	Palo Alto Networks Cortex XDR Pro
Supported languages - endpoint client	English, German, Japanese, Spanish, French, Chinese
Supported languages - management console	English
List price for 5000 clients / 5 years (without any discount)	$ 1 750 000
Product Features for 5,000 endpoints	
Do you also offer a managed version (MDR) of the tested product in your portfolio?	✓
General Features	
Third-party scan engine used (in addition to its own)	proprietary
Phishing protection for web browsers	□
Web access control	✓
External device control	✓
Sandbox feature	✓
2-factor authentication	optional
Right-click on-demand scan	✓
Lock settings	✓
Lock uninstalling	✓
Supported Operating Systems	
Microsoft Windows	✓
□ Windows 7	□
□ Windows 8	□
□ Windows 10	✓
□ Windows 11	✓
Virtual environments (such as VMware, HyperV)	✓
Apple macOS	✓
Linux	✓
Google Android	✓
Apple iOS	✓
Response Actions	
Quarantine	✓
Delete Files and Directories	✓
Process Termination	✓
Shutdown or Reboot of Endpoint	✓
Edit Registry Keys and Values	✓
Network Isolation	✓
User Isolation	□
Execution Prevention	✓
Block Processes from Communication	✓
Uninstall Services	✓
System Restoration	✓
System Imaging	✓
Patching	□
Guided Response Available	✓

Reporting Features	
Attack Visualization	✓
Attack Timeline	✓
Attack Context	✓
Continuous Monitoring	✓
Running applications & process	✓
Behaviour Monitoring (File/registry/etc..)	✓
Whitelisting capability	✓
Data Sharing Features	
Customizable default security policies	✓
Customized reporting and management	✓
Custom reporting and filtering	✓
Report automation	✓
Standard output format (JSON, Syslog, CEF, etc..)	✓
Splunk & Syslog integration	✓
Automated data export	✓
Policy and/or signature rollback	✓
System scanning capability	✓
Integration with security products	✓
Standards-based application programming interface (API) for access	✓
Disaster Recovery	✓
Audit trail support in the management console	✓
Management to agent encryption	✓
Encryption of data at rest	✓
Multiple EPR system-administrator/user-focused workflow support	✓
Enterprise recording and data storage –forensic analysis	✓
Built-in-reporting capabilities for different user categories	✓
Cloud marketplace support	✓
Compliance reports (GDPR, PCI-DSS, etc.)	✓
External Data Correlation	
Threat Intelligence data assimilation	✓
SIEM	✓
Proprietary product integration (NGFW, IPS, ...)	✓
YARA Signatures	✓
Support of IoC upload	✓
Sandboxing logs	✓
Scan results	✓
Retrospective analysis and logs	✓
Endpoint prevention product logs	✓
Multi-factor authentication logs	✓
Network traffic flow logs	✓
DNS Logs	✓
DHCP Logs	✓
Support	
Is free, basic, human support for the deployment process included in the licence for 5,000 endpoints?	✓
Assisted training for the IT staff in portfolio	✓
Supported languages of support	English

EDR Telemetry

For IT security professionals, especially those on the blue team, understanding the telemetry[5] capabilities of antivirus (AV) and endpoint detection and response (EDR) solutions[6] is paramount. Telemetry offers a comprehensive view of endpoint activity, enabling a deeper grasp of security alerts. This knowledge is crucial for swift threat response and invaluable for forensic investigations, allowing teams to trace and analyse attack evolution. Telemetry also serves a proactive role, helping identify new attack vectors and the tactics, techniques, and procedures used by adversaries.

However, it goes beyond defence. Telemetry comprehension allows teams to refine configurations, reduce false positives, and optimize operations. In an era prioritizing data privacy, it's essential to ensure telemetry remains compliant with stringent regulations. Detecting potential security gaps becomes easier with telemetry insights, aiding in pinpointing areas requiring additional protection or tools. Additionally, assessing data collection's impact on system performance ensures a seamless user experience.

Armed with this data, integrating AV and EDR insights into security information and event management (SIEM) solutions becomes more seamless. Furthermore, this foundational knowledge fosters enhanced collaboration, enabling blue teams to work cohesively with other departments, such as red teams or IT operations, to bolster the organization's security posture.

This data should be readily accessible and investigated by customers when using the respective products. Some vendors transparently provide this information in their documentation[7], empowering users to maximize the data/product for their defence strategies. Please note that this data pertains solely to the product/tier assessed in this report; the vendor may offer other products/tiers with additional telemetry features and support. The listed data was verified and provided by the vendors.

LEGEND	
✓	Implemented
✗	Not Implemented
~	Partially Implemented
Logs	Via Windows EventLogs (EDR is inspecting Windows event logs to collect the telemetry)
Telemetry	Via EnablingTelemetry (Additional telemetry that can be enabled easily as part of the EDR product but is not ON by default.)

[5] https://kostas-ts.medium.com/edr-telemetry-project-a-comprehensive-comparison-d5ed1745384b

[6] https://docs.google.com/spreadsheets/d/1ZMFrD6F6tvPtf_8McC-kWrNBBec_6Si3NW6AoWf3Kbg/htmlview

[7] https://github.com/tsale/EDR-Telemetry/wiki#product-documentation-references

Telemetry Feature Category	Sub-Category	Implementation
Process Activity	Process Creation	✓
	Process Termination	✓
	Process Access	✓
	Image/Library Loaded	✓
	Remote Thread Creation	✓
	Process Tampering Activity	✓
File Manipulation	File Creation	✓
	File Opened	✓
	File Deletion	✓
	File Modification	✓
	File Renaming	✓
User Account Activity	Local Account Creation	✓
	Local Account Modification	✓
	Local Account Deletion	✓
	Account Login	✓
	Account Logoff	✓
Network Activity	TCP Connection	✓
	UDP Connection	✓
	URL	✓
	DNS Query	✓
	File Downloaded	✓
Hash Algorithms	MD5	✓
	SHA256	✓
	IMPHASH	✗
Registry Activity	Key/Value Creation	✓
	Key/Value Modification	✓
	Key/Value Deletion	✓
Schedule Task Activity	Scheduled Task Creation	✓
	Scheduled Task Modification	✓
	Scheduled Task Deletion	✓
Service Activity	Service Creation	✓
	Service Modification	✓
	Service Deletion	✓
Driver/Module Activity	Driver Loaded	✓
	Driver Modification	✓
	Driver Unloaded	✗
Device Operations	Virtual Disk Mount	✓
	USB Device Unmount	✓
	USB Device Mount	✓
Other Relevant Events	Group Policy Modification	✓
Named Pipe Activity	Pipe Creation	✓
	Pipe Connection	✓
EDR SysOps	Agent Start	✓
	Agent Stop	✓
	Agent Install	✓
	Agent Uninstall	✓
	Agent Keep-Alive	✓
	Agent Errors	✓
WMI Activity	WmiEventConsumerToFilter	✓
	WmiEventConsumer	✓
	WmiEventFilter	✓
BIT JOBS Activity	BIT JOBS Activity	✗
PowerShell Activity	Script-Block Activity	✓

Overview of EDR technologies

In the dynamic field of cybersecurity, IT security professionals need a deep understanding of antivirus (AV/EPP) and endpoint detection and response (EDR) systems, which are crucial for comprehensive defence strategies. One key aspect is understanding how different AV and EDR systems implement essential technologies[8]. The following information offers a high-level overview of these technologies, highlighting their importance in the ever-changing cybersecurity landscape. These technologies encompass the Antimalware Scan Interface (AMSI), User-Mode Hooking, Callbacks, and Kernel Drivers.

1. **Antimalware Scan Interface (AMSI):** AMSI in Windows is an API set designed for enhanced malware detection. Integrated into components such as PowerShell, Windows Script Host, and .NET, it intercepts scripts post-deobfuscation at runtime. AMSI communicates directly with the system's antimalware solution, forwarding content for analysis. As an interface, it's agnostic to the specific antimalware vendor. Its integration ensures real-time threat detection, even for dynamically executed content.

2. **User-Mode Hooking:** User-mode hooking intercepts function calls in application-level processes in Windows. By overwriting a function's start, calls are redirected to a custom function. For instance, an EDR might hook `CreateFileW` in kernel32.dll, redirecting it to its own DLL. When an application uses `CreateFileW`, it's first processed by the EDR's function, allowing real-time monitoring or restrictions before proceeding with the original call.

3. **Callbacks:** EPP/EDR solutions leverage kernel callback routines for deep system monitoring. These routines notify registered callbacks when specific OS events occur. By tapping into these events, EPPs/EDRs observe real-time system behaviour. For instance, an EPP/EDR might monitor process creation events. When a new process starts, the callback inspects its details and origin. This allows the EPP/EDR to quickly detect, assess, and respond to potential threats.

4. **Kernel Drivers:** EPP/EDR solutions employ kernel drivers to deeply integrate with the operating system for advanced threat mitigation. Minifilter drivers, part of the Windows Filter Manager, allow EPP/EDR tools to monitor, modify, or block operations on files and data streams. This is crucial for real-time scanning and access restrictions. ELAM (Early Launch Anti-Malware) drivers, on the other hand, start early during the boot process, ensuring that only legitimate, signed drivers are loaded, thereby preventing rootkits or bootkits from compromising the system. Collectively, these drivers ensure comprehensive protection from boot-up to system operation.

This information equips IT security professionals with valuable insights for making informed decisions about cybersecurity solutions. Whether you need a comprehensive understanding or a quick reference, these insights empower you to navigate the complex world of IT security effectively.

It's important to note that these are just some of the technologies employed in modern cybersecurity, and others may also be included in the arsenal of IT security professionals. The absence or presence of a certain technology does not necessarily mean that a product is worse or better. The effectiveness of a cybersecurity strategy depends on its holistic approach and adaptability to evolving threats. The listed data was verified and provided by the vendors.

[8] https://kwcsec.gitbook.io/the-red-team-handbook/techniques/defense-evasion/basics/iocs/high-level-overview-of-edr-technologies

EDR Technology	Description	Palo Alto
Antimalware Scan Interface (AMSI)	This is a standard interface that allows applications and services to integrate with any antimalware product present on a machine.	✓
Event Tracing for Windows (ETW)	This is a mechanism for tracing and logging events that are raised by both user-mode applications and kernel-mode drivers.	✓
Microsoft Threat Intelligence (EtwTi)	This is a mechanism for tracing and logging events using Microsoft Threat Intelligence.	✓
User Space API-Hooking	This is a technique used to intercept API function calls in user space. This can be used by EPP/EDR solutions to monitor and potentially block suspicious behaviour.	✓
Kernel Space API-Hooking	Similar to user space API hooking, but this intercepts API function calls in the kernel space.	✓
Kernel Callback Routines	These are functions that the kernel calls when certain events or conditions occur. EPP/EDR solutions can use these to monitor system events.	✓
Filter Driver	This is a type of driver used to monitor and potentially modify the behaviour of device drivers. EPP/EDR solutions may use this to monitor for suspicious device behaviour.	✓
Minifilter Driver	This is a specific type of filter driver that can be used to monitor and potentially modify the behaviour of file system operations.	✓
Early Launch Antimalware (ELAM) Driver	This is a driver that starts early in the boot process to scan drivers for malware before they're loaded.	✓

Palo Alto Networks Product Configurations and Settings

In business environments, and with business products in general, it is usual for products to be configured by the system administrator, in accordance with vendor's guidelines. Therefore, we asked vendors to request us to implement any changes they wanted to the default configuration of their respective products. Results presented in this test were only accomplished by applying the respective product configurations as described here.

The configurations were applied together with the engineers of the respective vendors during setup. This configuration is typical in enterprises, which have their own teams of security staff looking after their defences. It is common for products of this kind that vendor experts assist companies on the deployment and configuration best suited for the type of enterprise.

Below we have listed relevant non-default settings (i.e. settings used by the vendor for this test).

Palo Alto Networks: Under "Agent settings", in "XDR Pro Endpoints", "XDR Pro Endpoint Capabilities" were enabled. Under "Malware Profile", "Portable Executable and DLL examination", "Behavioural Threat Protection" and "Ransomware Protection" were set to "Quarantine". "Treat Grayware as Malware" was enabled.

Appendix

Endpoint Prevention Response vs MITRE ATT&CK Framework

This EPR product report is a comprehensive validation of features, product efficacy and other relevant metrics to guide your risk assessment. A total of 50 scenarios were executed against real-world enterprise use-cases. These scenarios comprised several prevention and detection workflows operating under normal operational environments by different user personas. The results for the validation can be efficiently and effectively mapped to the MITRE ATT&CK® Platform[9] and NIST platform, so that it becomes easier to operationalize the risk regarding a specific endpoint.

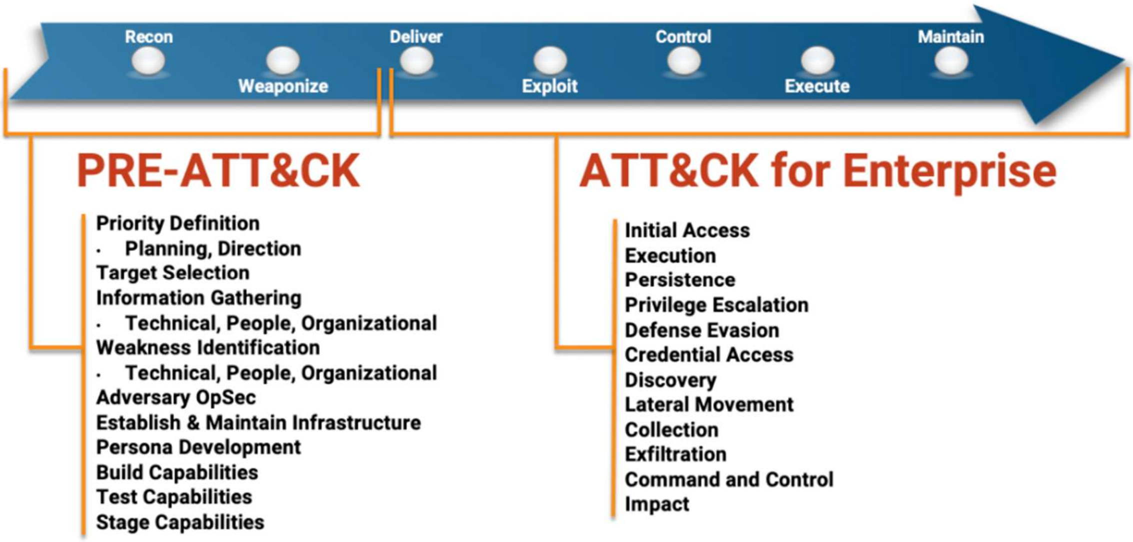

MITRE ATT&CK for Enterprise vs Seven Stage Cyber Attack LifeCycle[10]

AV-Comparatives has developed an industry-changing paradigm shift by defining a real-world EPR methodology reflecting the everyday reality of enterprise use cases and workflows to be used for mapping the kill-chain visibility to the MITRE ATT&CK framework.

As illustrated in the graphic on the next page, we moved away from "atomic" testing, i.e. tests that only look at a particular component of the ATT&CK framework, and instead evaluated the EPR products from the context of the entire attack kill-chain, with workflows interconnecting at every stage from the initial execution to final data exfiltration/sabotage.

[9] © 2015-2023, The MITRE Corporation. MITRE ATT&CK and ATT&CK are registered trademarks of The MITRE Corporation.

[10] Source: https://attack.mitre.org/resources/enterprise-introduction/

EPR Testing Workflow

The graphic below provides a simplified overview of the test procedure used:

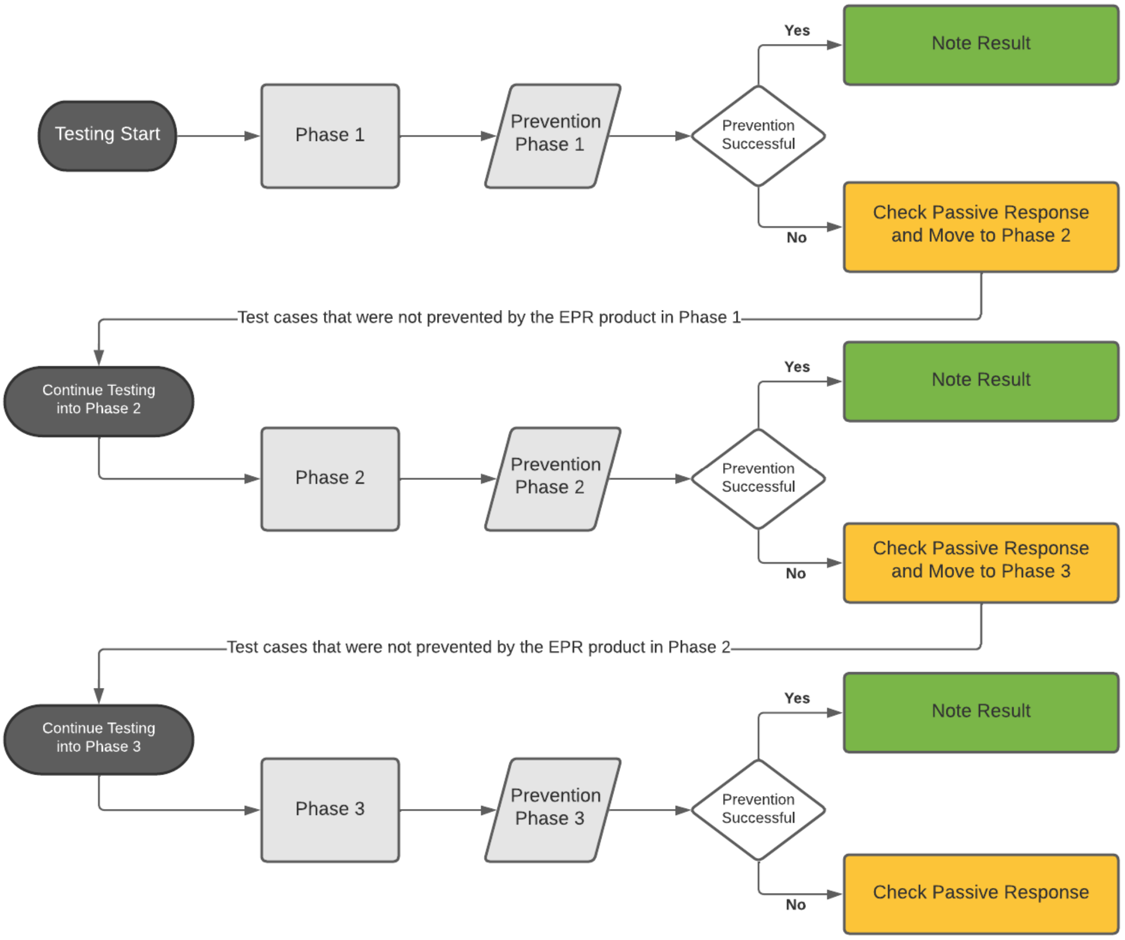
Testing Start
Phase 1
Prevention Phase 1
Prevention Successful
Yes
Note Result
No
Check Passive Response and Move to Phase 2
Test cases that were not prevented by the EPR product in Phase 1
Continue Testing into Phase 2
Phase 2
Prevention Phase 2
Prevention Successful
Yes
Note Result
No
Check Passive Response and Move to Phase 3
Test cases that were not prevented by the EPR product in Phase 2
Continue Testing into Phase 3
Phase 3
Prevention Phase 3
Prevention Successful
Yes
Note Result
No
Check Passive Response

Enterprise EPR Workflow Overview

Prevention (Active Response)

The best way to respond to any threat is by preventing and effectively reporting on it as soon as possible. AV-Comparatives defines prevention as an automated, active response that kicks in 24/7, 365 days a year, without the need for human intervention, but with quantifiable metrics and reporting data points that can be leveraged for effective analysis.

An EPR product should be able to initially identify and prevent a threat on a compromised machine. The incident should be detected, identified, correlated, and remediated from a single pane of glass (centralized management system) through an effective passive response strategy (partially/fully automated) ideally in real time. Furthermore, the system administrator should be able classify and triage a threat based on the data collection and analysis, and be able to close out a response using the EPR product with a specific workflow.

An active response, as defined in this test, is an effective response strategy that provides detection with effective prevention and reporting capabilities. This should all be done in an automated way with no manual intervention. This can be done through a multitude of technologies and mechanisms, for example: signature-based models, behaviour-based models, ML-based models, transaction rollbacks, isolation-based mechanisms, and so forth. This definition is technology-agnostic because it focuses on the outcomes of the various system-administrator workflows and scenarios, and not on the technology used to prevent, detect or respond to it.

Detection (Passive Response)

Passive response, as defined in this test, is a set of response mechanisms offered by the product with cohesive detection, correlation, reporting and actionable capabilities. Once an attacker is already inside the enterprise environment, traditional response mechanisms kick in, for example IOC and IOA correlation, external threat intel and hunting. AV-Comparatives defines these response mechanisms as Passive Response. The precondition for passive response is the detection of a potential threat by EPR products.

EPR products are typically expected to prevent initial and ongoing attacks without having to triage, while offering active response and reporting capabilities. If the attack is missed or not prevented, EPR products should then be able to assess and respond to attacks, thus providing lesser burden on resources (human/automation) and providing better ROI in the long run.

The range of available response capabilities of an EPR product is extremely important for organizations that need to review threats/compromises in multiple machines across multiple locations. An EPR product should be able to query for specific threats using the intelligence data provided to the system administrator. Once they have been identified, the system administrator should be able to use the EPR product to initiate responses based on the type of infection. AV-Comparatives expects EPR products to have non-automated or semi-automated passive response mechanisms.

Correlation of Process, Endpoint and Network

The EPR product should be able to identify and respond to threats in one or more of the following ways:

- Response based on successful identification of attack via the product’s user interface (UI) that lists attack source (http[s]/IP-based link) that hosts compromised website/IP).
- Exploit identification (based upon the CVE or generic detection of threat)
- Downloaded malware file
- Malware process spawning
- Command and control activity as part of the single chain of attacks

EPR Validation Overview

AV-Comparatives have come up with the following topology and metrics to accurately assess the capabilities of endpoint prevention and response (EPR) products.

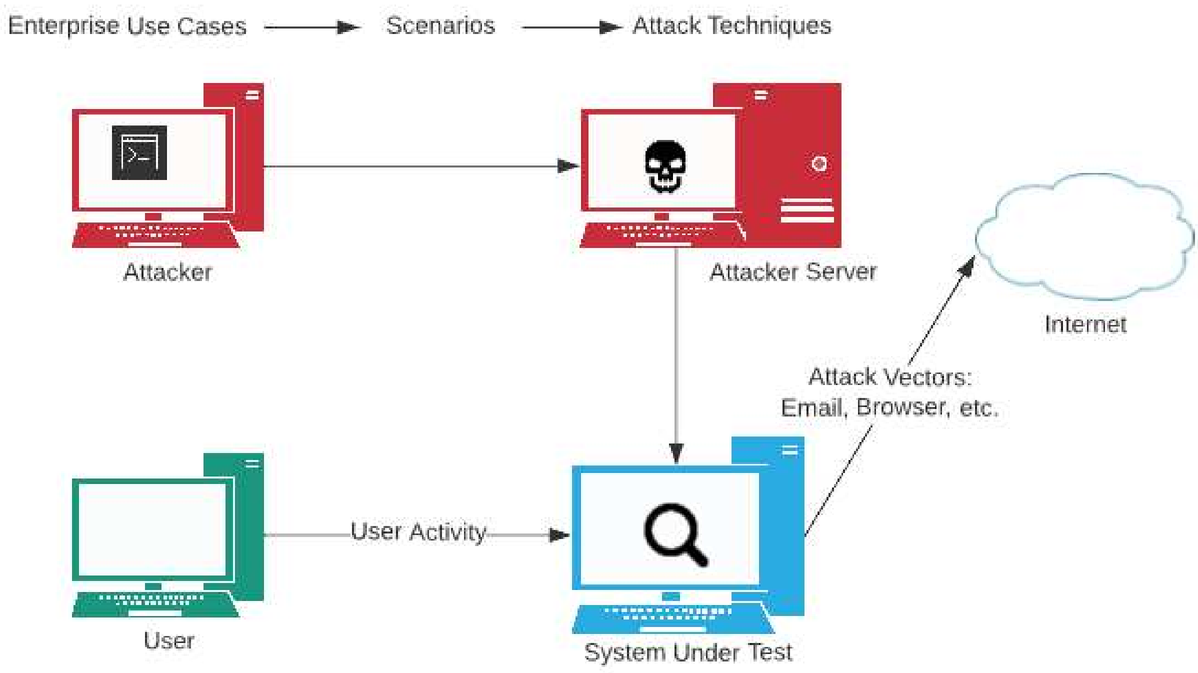

EPR Test Topology Overview

All the tested vendors' EPR products were deployed and evaluated in a standalone mode, with each vendor actively involved in the initial setup, configuration, and baselining aspects. AV-Comparatives evaluated a list of 50 scenarios, as often requested by analysts and enterprises, highlighting several enterprise-centric use cases. Every vendor was allowed to configure their own product, to the same extent that organizations are able to do when deploying it in their infrastructure. The details of the configurations are included at the beginning of this report.

Because this methodology is tailored towards the prevention, detection and response capabilities, all vendors activated their prevention and protection capabilities (ability to block), along with detection and response, so that they emulate the real-world enterprise-class capabilities of these products.

The testing supported EPR product updates and configuration changes made by cloud management console or local area network server. We went through and executed all test scenarios from beginning to end, to the greatest extent possible.

Test Objective

The following assessment was made to validate if the EPR endpoint security product was able to react appropriately to each scenario.

- In which attack phase did the prevention/detection occur? Phase 1 (Endpoint Compromise and Foothold), Phase 2 (Internal Propagation) or Phase 3 (Asset Breach)?
- Did the EPR product provide us with the appropriate threat classification and threat triage, and demonstrate an accurate threat timeline of the attacks with relevant endpoint and user data?
- Did the EPR product incur any additional costs due to imperfect Operational Accuracy or workflow delays?

Targeted Use-Cases

The sequence of events emulated was an enterprise-based scenario where in the system-level user received a file in an email attachment and executed it. In some cases, the emails were benign, while in others they were not. The malicious email attachments, if successfully executed, allowed an attacker to get a foothold inside the environment and take additional steps to act upon their objectives.

During testing, we logged into the EPR product management and the individual test system consoles, to observe, analyse and document what kind of activity is recorded by the product. For instance, if there is an attack, are there any alerts or events, and are these true positives or true negatives?

For true positive alerts, we further investigated whether the subsequent response in terms of event correlation, triages, threat classification and threat timeline were provided to the system administrator in a timely and clear way. We tested the responses as available by products under the test.

The test was conducted in summer 2023, and used an attacker-driven mindset as the attack progressed through the attack nodes to finally meet its objective. User activities were simulated throughout the test such that they were as close to a real-life environment as possible.

All the attacks were crafted using open-source and commercial tools[11]/frameworks, and were developed using in-house expertise. The reason why we included commercial C2 frameworks[12] is that these are frequently misused by attackers[13] in real-life APTs; not using them would cause a „blind spot“ and lead to a false sense of security. Due to license agreement restrictions, we took measures to prevent samples created by commercial C2 frameworks from being distributed to the EPR vendors. These restrictions are made to prevent vendors from focussing on the tools instead of the techniques.

To illustrate the test procedure, we provide below an example of how a typical targeted attack might work. The attacker sends a script payload (containing some defence evasion techniques such as DLL sideloading) via a phishing mail to Network User A on Workstation A. After getting a foothold in the targeted network with the User Account A, internal discovery is performed. This involves enumerating user privileges, user groups, installed security products etc. Through this process it can be seen that the compromised User Account A has access to the C$ share on Workstation B, meaning that the account has local admin privileges on this workstation. With the knowledge gained from internal discovery, the attacker moves laterally from Workstation A to Workstation B. They then continue with internal discovery on Workstation B. This enables them to find a network administrator's open user session on Workstation B. To take advantage of this, the attacker dumps the LSASS process, and is thus able to steal the administrator's credentials. After doing this, they discover that the compromised administrator account has access to Server 1. The attacker then uses this compromised admin account to move laterally from Workstation B to Server 1, and then compromise this server. Here they perform further internal discovery, and also use some defence evasion techniques to bypass the installed security product (e.g. by patching AMSI and ETW). At the end of this procedure, they are able to identify credit-card data on Server 1, which they extract via an open C2 channel.

[11] https://attack.mitre.org/software/
[12] https://redcanary.com/threat-detection-report/trends/c2-frameworks/
[13] https://www.trendmicro.com/en_us/research/22/j/black-basta-infiltrates-networks-via-qakbot-brute-ratel-and-coba.html

Copyright and Disclaimer

This publication is Copyright © 2023 by AV-Comparatives®. Any use of the results, etc. in whole or in part, is ONLY permitted after the explicit written agreement of the management board of AV-Comparatives prior to any publication. AV-Comparatives and its testers cannot be held liable for any damage or loss, which might occur as result of, or in connection with, the use of the information provided in this paper. We take every possible care to ensure the correctness of the basic data, but a liability for the correctness of the test results cannot be taken by any representative of AV-Comparatives. We do not give any guarantee of the correctness, completeness, or suitability for a specific purpose of any of the information/content provided at any given time. No one else involved in creating, producing or delivering test results shall be liable for any indirect, special or consequential damage, or loss of profits, arising out of, or related to, the use or inability to use, the services provided by the website, test documents or any related data. For more information about AV-Comparatives and the testing methodologies, please visit our website.

AV-Comparatives
(October 2023)

Credits: Icons made by icon_king1 from freeicons.io