Independent Tests of
Anti-Virus Software

Endpoint Prevention and Response (EPR) Product Validation Report

Palo Alto Networks Cortex XDR Pro

TEST PERIOD: OCTOBER 2021
LAST REVISION: 10TH JANUARY 2022

WWW.AV-COMPARATIVES.ORG

Content

Tested Product

Palo Alto Networks Cortex XDR Pro was tested by AV-Comparatives in October 2021. The tested version number was 7.5

Product Thumbnail

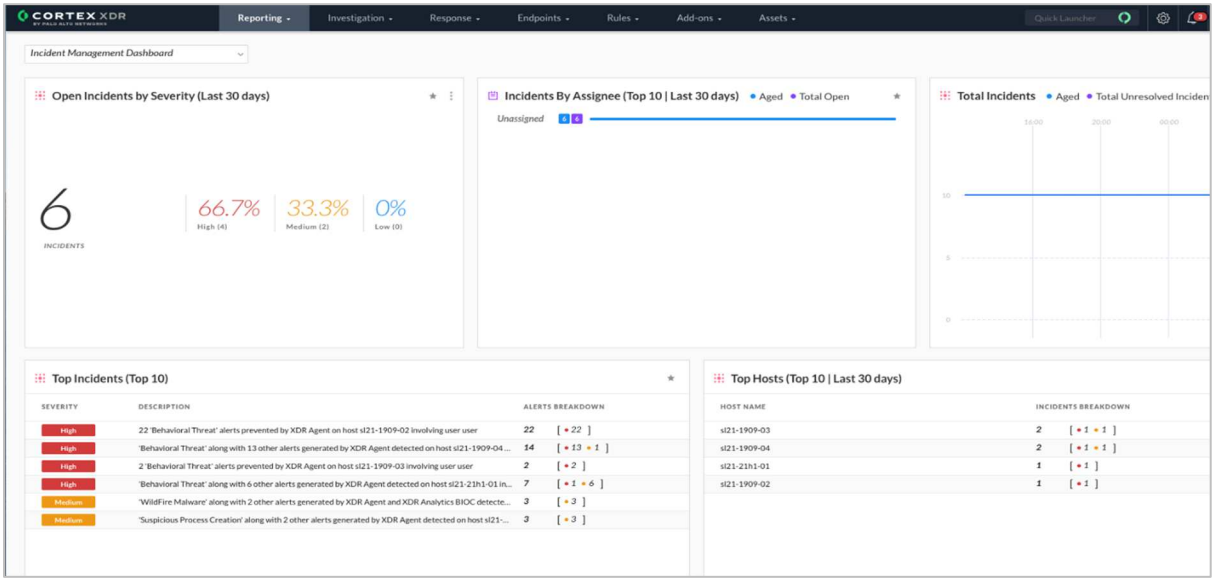

Palo Alto Networks Cortex XDR Pro

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 exceptionally well at handling threats that are targeted towards the user, in particular before the threat progresses inside the user environment. The product demonstrated several safeguards that helped in protecting the enterprise end-user against the scenarios we tested.

The product's management console was easy to use, intuitive, and provided contextual data useful for SOC analysts to ascertain which threats to prioritize. The integration with Palo Alto Networks Wildfire Sandbox offers the ability to send unknown files to the sandbox to do additional analysis and get a verdict with relative ease. Relevant threat alerts were demonstrated at the endpoint level, as well as in the cloud console, with the appropriate level of information. The product offers the ability to create different sets of behavioural rules, and excellent triaging ability for multiple users to collaborate on any given threat scenario at the same time. The endpoint agent also offers "remoting capabilities", which gives the analyst an excellent opportunity to investigate threats in real time.

The product had good mapping to MITRE's TTP, which provides low-level SOC analysts the data needed to investigate further and escalate when necessary. Alerts were prioritized and aggregated to minimize noise from all the alerts generated. The product was easy to configure and deploy in a domain or workgroup environment.

Active Response: An active response is an effective response strategy that provides detection with effective prevention and reporting capabilities.
***Note**: 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 **99.0**%.*

Passive Response: Passive response is a set of response mechanisms offered by the product with cohesive detection, correlation, reporting and actionable capabilities.
***Note**: 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**%.*

High Enterprise Savings: If most threats are detected and prevented by the EPR product at or soon after execution, and if the product provides the necessary detection information to help with an effective passive response (partially/semi-automated), it will result in high enterprise savings. The average of both active and passive response needs to be equal to or greater than 95% of the overall EPR product response rate in order to reach "High Enterprise Savings".

Description	Details
Enterprise Product Savings: **Palo Alto** prevents most attacks and offers effective passive response	**High (>95%)**
Overall **Active Response** Rate (Prevention Rate):	**99.0%**
Overall **Passive Response** Rate (Response Rate):	**100%**
Overall Operational Accuracy Result (False Positive Result):	**Pass**

Figure 1 – Executive Summary

Figure 2 depicts Palo Alto's EPR prevention & detection rates across Workflow-1 and Workflow-2, across the different phases and categories of attack. For more details on the workflows and phases, please see the appendix.

Description	Number Tested	Action Taken by the EPR
Scenarios	50	50
Phases	**Combined Prevention & Detection (T0: Time of Attack)**	**Combined Prevention & Detection (T1: 24 Hrs)**
Phase 1 (Compromise & Foothold)		
Active Response	98.0%	98.0%
Detect	100%	100%
Passive Response	100%	100%
Phase 2 (Internal Propagation)		
Active Response	100%	100%
Detect	100%	100%
Passive Response	100%	100%
Phase 3 (Asset Breach)		
Active Response	N/A[1]	N/A[1]
Detect	N/A[1]	N/A[1]
Passive Response	N/A[1]	N/A[1]
Detection Avoidance[2]	PASS	PASS
Emerging Attacks[2]	PASS	PASS
Operational Accuracy (False Positives)[2]	PASS	PASS

Figure 2 — Combined Prevention & Detection Rates of Palo Alto Networks Cortex XDR Pro

The Palo Alto Networks EPR product offered strong prevention capability, preventing 98% of the scenarios in the "Initial Access" phase of the Prevention workflow, while also offering excellent detection and reporting capabilities overall. For the 1 scenario (2%) that were able to progress to Phase 2, Palo Alto Networks prevented, detected and acted upon it in the passive response phase. Palo Alto Networks provided excellent overall active response capabilities, augmented with an effective and cohesive response strategy. Figure 3 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.

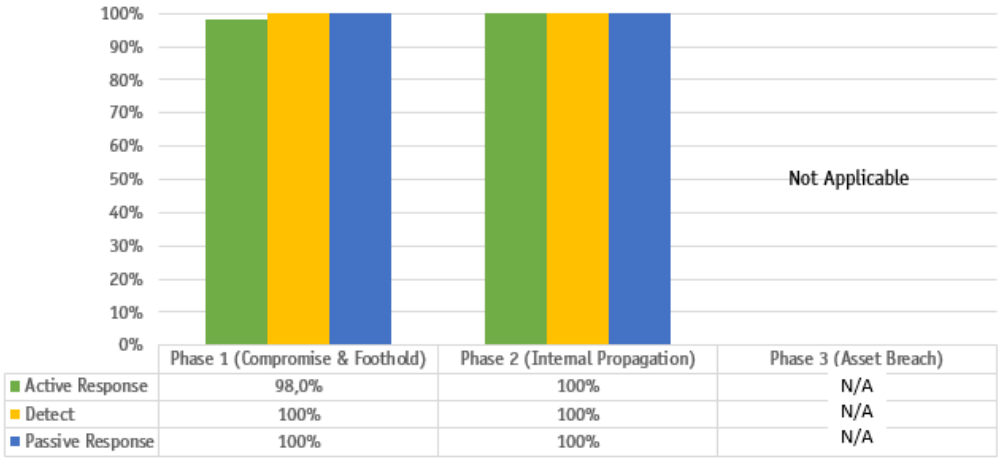

Figure 3 — Active vs Passive Response of Palo Alto Networks Cortex XDR Pro

[1] No scenario progressed to Phase 3.
[2] PASS: The EPR product had a score of 95% or better.

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 security-analyst user-based 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	98.0% (49/50)	100% (50/50)	100% (50/50)
Detect	100% (50/50)	100% (50/50)	100% (50/50)
Passive Response	100% (50/50)	100% (50/50)	100% (50/50)

Figure 4 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.

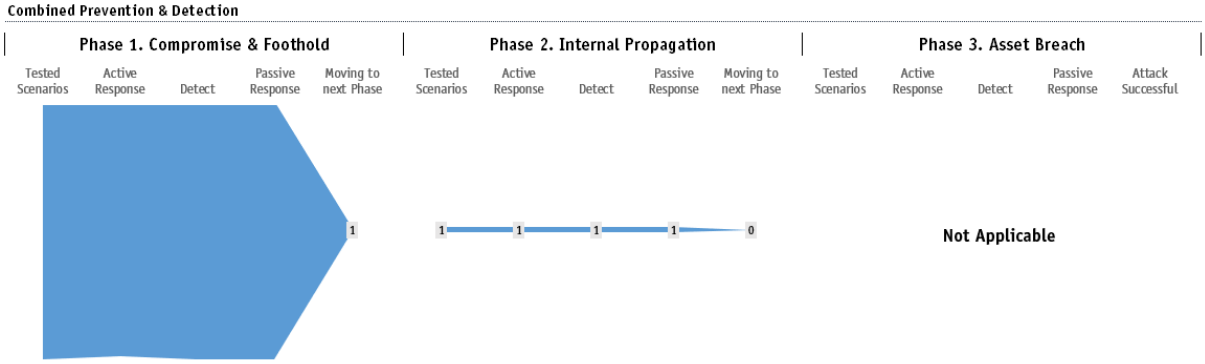

Figure 4 — EPR Efficacy per Phase of Palo Alto Networks Cortex XDR Pro

We tested a total of 50 scenarios, and only one of these was able to bypass the active response mechanism in one phase.

Phase 1:

- 49 out of 50 scenarios prevented
- 50 out of 50 scenarios detected
- 1 scenario was able to progress to Phase 2.

Phase 2:

- 1 out of 1 scenario prevented
- 1 out of 1 scenario detected
- No scenario was able to progress to Phase 3.

Phase 3:

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

EPR Test Metrics and Scoring

In our opinion, the goal of every EPR system should be to prevent threats or provide effective response capabilities as soon as possible. In other words, endpoint products that offer a high active prevention incur less costs in the event of a breach, since there is little operational overhead required to respond to and remediate the effects of a compromised system. Furthermore, EPR products that also provide a high detection rate (visibility and forensic detail) will realize additional savings because compromises do not have to be investigated manually.

Figure 5 provides an example of how the product is evaluated. For a breakdown of how the product scored, please see figures 9 through 11.

Available Ratings:

EPR Product Evaluation	Enterprise Savings
Prevents most attacks and offers effective passive response	High
Prevents most attacks, but offers weaker passive response	Medium
Weak prevention and weak passive response	Low

Figure 5 — Use-Case Scenarios Scoring

High Enterprise Savings: If most threats are detected and prevented by the EPR product at or soon after execution, and if the product provides the necessary detection information to help with an effective passive response (partially/fully automated), it will result in high enterprise savings.

Note: *The average of both active and passive response needs to be equal to or greater than 95%.*

Medium Enterprise Savings: If most threats are detected and prevented by the EPR product at or soon after execution, but with limited details surrounding the detection, it will result in a weaker passive response strategy. This is because of the operational overhead that is required to respond to and remediate the effects of a compromised system resulting in an increase in enterprise costs.

Note*: The average of both active and passive response needs to be equal or greater than 90%.*

Low Enterprise Savings: Lastly, if most threats are not prevented by the EPR product, and the product provides no details surrounding the detection, this will result in both a weaker active and a weaker passive response strategy with only low enterprise savings.

Note: *The average of both active and passive response is less than 90%.*

Reduction in TTP (Time to Prevent)

The ability of the EPR product to rapidly identify and prevent a threat, and display relevant information about it, is a very important factor. This could also be referred to as the effective reduction in active time to respond. Figure 6 provides a breakdown of Palo Alto Networks Cortex XDR Pro's overall prevention rate. This is highlighted as measured at the time of the attack (T0) and how well the product offered prevention and then at 24 hours, Time (T1) = T0 +24 Hrs.

	Time of Attack (in hours)								
Time to Prevent	0 (T0)	<1	<2	<5	<10	<15	<20	<24	24 (T1)
Phase 1	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%
Phase 2	100%	100%	100%	100%	100%	100%	100%	100%	100%
Phase 3	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Figure 6 —Time to Active Response

Immediate protection and response against new attacks is critical. Attackers use different websites to host their attacks, in order to bypass reputation engines. Therefore, products that fail to prevent or respond to an attack in a timely manner may be too late to counter a threat.

We recorded the time the threat was introduced into the test cycle and how long it took the product to prevent it. Within the 24-hour window, cumulative protection and detection rates are calculated each hour until attacks are prevented and responded to by the product.

Reduction in TTR (Time to Respond)

Time is critical when an incident that is not prevented turns into a potential breach. The timing of activities, triggering of a response, and length of a response will vary widely, depending on the capabilities of the product and the expertise of the user. Hence reduction in the passive response time becomes critical to containing any breach. The less time it takes for the EPR product to come up with the response, the better the EPR product.

	Time of Attack (in hours)								
Time to Respond	0 (T0)	<1	<2	<5	<10	<15	<20	<24	24 (T1)
Phase 1	100%	100%	100%	100%	100%	100%	100%	100%	100%
Phase 2	100%	100%	100%	100%	100%	100%	100%	100%	100%
Phase 3	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Figure 7 — Time to Respond

EPR Validation Scenario Overview

Figure 8 provides some examples of scenarios used as part of this test. We tested 50 operational enterprise scenarios comprised of several different operational workflows under normal operational environments, executed by different user personas. The intent of the EPR test was to evaluate if the tested products were able to prevent initial and ongoing attacks, without having to triage the threats, while offering active and passive response and reporting capabilities. The scenarios covered all steps of the Kill Chain and are mapped to the MITRE ATT&CK framework.

Scenario: A scenario consists of enterprise operational workflows having one or more attack samples executed using different techniques.

KillChain Phases	Delivery Exploitation Installation	Installation Command and Control	Denial of Service Action on Objectives Command and Control	MITRE Reference
Phase No	Initial Access Execution Persistence	Privilege Escalation Lateral Movement Credential Access Discovery Defense Evasion	Collection Exfiltration Impact	MITRE ID
Phase 1	Scenario 1, 2, 3 Scenario 4, 5, 6 Scenario 7, 8, 9			T1189,T1566,T1059,T1203, T1053,T1569, T1204
Phase 2		Scenario 33,34,35		T1548, T1134, T1543
Phase 3			Scenario 25,25,26,27,28,29,30,31,32	T1020, T1029

Figure 8 — Example Scenarios

The example scenarios highlighted below across the 3 phases give you an overview of how it was evaluated using a specific set of technique(s) mapped to Techniques, Tactics and Procedure (TTP). Based on good-faith vulnerability disclosure policies, we are specifically NOT disclosing all the scenarios and technique(s) used in this iteration of EPR testing.

Note: These example scenarios do not directly map to the actual tested scenario highlighted in this test report. We have highlighted only a few selected examples below so as to avoid potential product compromises. Details of the missed attacks were provided to the respective vendors after the test.

Workflow-1 Phase-1: Initial Access

Based upon EPR Prevention Workflow-1, Phase 1 (Endpoint Compromise and Foothold), we tested several scenarios using different file formats and methods, such as spear-phishing attachments and drive-by download attacks, to obtain initial access into the environment.

- Scenario 8: Initial access using a drive-by download attack. This scenario is introduced via a web browser, using a known vulnerability wherein the attacker gains access to the system of a targeted user, when the user visits a website unsuspectingly.
  MITRE reference: https://attack.mitre.org/techniques/T1189/
- Scenario 21: Initial access using spear-phishing Link. This scenario is introduced via email link using .hta files. For example, a .hta file was sent to the targeted user.
  MITRE reference: https://attack.mitre.org/techniques/T1192/

- Scenario 30: Execution through API. This scenario was emulated via a payload derived from different tools and custom-made tools. A portable executable as an email attachment was sent to the user.
  MITRE reference: https://attack.mitre.org/techniques/T1106/
- Scenario 37: Execution using PowerShell. This scenario was emulated via PowerShell files. An email with a portable executable/PowerShell file as an attachment was sent to the targeted user.
  MITRE reference: https://attack.mitre.org/techniques/T1086/
- Scenario 32: Persistence using AppCert DLLs. This scenario was emulated via different registry modifications. An email with a portable executable sent as an attachment was sent to the targeted user.
  MITRE reference: https://attack.mitre.org/techniques/T1182/
- Scenario 29: Persistence using AppInit DLLs. This scenario was emulated via different registry modifications.
  MITRE reference: https://attack.mitre.org/techniques/T1103/
- Scenario 1: Persistence using Scheduled Task. This scenario was emulated via different task scheduler task trigger mechanisms.
  MITRE reference: https://attack.mitre.org/techniques/T1053/

Workflow-1 Phase-2: Internal Propagation

If this scenario was successful, we moved into Phase 2 (Internal Propagation) and then finally Phase 3 (Asset Breach) of the prevention Workflow-1. We also tested some scenarios where an attacker is opportunistic and jumps directly from Phase 1 to Phase 3 as well.

- Scenario 33: Exploitation for Privilege Escalation. This scenario was emulated via multiple privilege-escalation vulnerabilities as well as typical methods like name-pipe impersonation.
  MITRE reference: https://attack.mitre.org/techniques/T1068/
- Scenario 1: Credential access using credential dumping.
  MITRE reference: https://attack.mitre.org/techniques/T1003/

Workflow-1 Phase-3: Asset Breach

For each of these phases we evaluated the Response Workflow-3 and Reporting Workflow-4 as stated in the methodology. **Note:** Every attempt was made to ensure that atomic test cases (ones that only look at a particular component of the ATT&CK framework) are not run as part of the workflow, wherever applicable.

- Scenario 24: End-user information collection using screen capture. This scenario was emulated by grabbing images from inside the host.
  MITRE reference: https://attack.mitre.org/techniques/T1113/
- Scenario 28: Impacting end-user using data destruction. This scenario was emulated via a payload derived from different tools and custom-made tools.
  MITRE reference: https://attack.mitre.org/techniques/T1485/

Phase-1 Metrics: Endpoint Compromise and Foothold

Phase-1 can be triggered by an attack based on the MITRE ATT&CK and other methods, and can be effectively mapped to Lockheed's Cyber Kill Chain. This workflow can be operationalized by going through the various attack phases described below.

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 to gain a foothold inside the environment. 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 phases as highlighted above and described in detail on Workflow-1 of AV-Comparatives' EPR CyberRisk Test Methodology. The resulting figures below (9, 10 and 11) showcase the product's Active Response, Detection and Passive Response capabilities against the validated attack scenarios.

Tested Scenario	Description	Active Response	Detect	Passive Response
1	MS Word Macro with CVE-2020-0668	✓	✓	✓
2	XLM Macro AutoOpen using MSBuild for compilation	✓	✓	✓
3	MS PowerPoint Macro with CVE-2020-0796	✓	✓	✓
4	MS Word macro with CVE-2020-0796	✓	✓	✓
5	MS Excel Macro with CVE-2020-0668	✓	✓	✓
6	MS PowerPoint Macro using MSBuild for compilation	✓	✓	✓
7	SYLK Macro using MSBuild for compilation	✓	✓	✓
8	Microsoft Office Word RCE Variation 1(CVE-2021-40444)	✓	✓	✓
9	Microsoft Office Word RCE Variation 2(CVE-2021-40444)	✓	✓	✓
10	MS PowerPoint Macro	✓	✓	✓
11	MS XLM Macro with In- Memory script	✓	✓	✓
12	MS Excel Macro	✓	✓	✓
13	MS Word Macro with CVE-2021-1675	✓	✓	✓
14	MS Word DotM File	✓	✓	✓
15	MS Excel with CVE-2021-1675	✓	✓	✓
16	MS PowerPoint with CVE-2021-36934	✓	✓	✓
17	MS Excel Macro	✓	✓	✓
18	MS Word DotM with CVE-2021-36934	✓	✓	✓

19	XLSM Macro with CVE-2021-1675	✓	✓	✓
20	Koadic JSE File	✓	✓	✓
21	Koadic HTA File	✓	✓	✓
22	Koadic Bat File	✓	✓	✓
23	Koadic PowerShell	✓	✓	✓
24	Caldera PowerShell	✓	✓	✓
25	Caldera Portable Executable	✓	✓	✓
26	Covenant PowerShell File	✓	✓	✓
27	Covenant Grunt Portable Executable	✓	✓	✓
28	Encoded VBE with Wiper Payload	✓	✓	✓
29	Forged Signature added to a File	✗	✓	✓
30	Keylogger Writing DLL Payload to disk	✓	✓	✓
31	Stateless MSF Writing DLL Payload to disk	✓	✓	✓
32	Keylogger via HTTP Post & Writing DLL Payload to disk	✓	✓	✓
33	CVE-2020-0683	✓	✓	✓
34	CVE-2020-0796	✓	✓	✓
35	CVE-2019-1322	✓	✓	✓
36	PowerShell ConPtyShell	✓	✓	✓
37	PowerShell Base 64 Encoded reverse shell	✓	✓	✓
38	PowerShell Simple Payload	✓	✓	✓
39	PowerShell HTA Payload	✓	✓	✓
40	PowerShell base52 stager variation 1	✓	✓	✓
41	PowerShell base52 stager variation 2	✓	✓	✓
42	PowerShell base52 stager variation 3	✓	✓	✓
43	PowerShell base52 stager variation 4	✓	✓	✓
44	PowerShell base64 stager variation 1	✓	✓	✓
45	PowerShell base64 stager variation 2	✓	✓	✓
46	PowerShell JOB Payload	✓	✓	✓
47	PowerShell New Process Payload	✓	✓	✓
48	PowerShell JOB + File Payload	✓	✓	✓
49	PowerShell JOB + File +SCT Payload	✓	✓	✓
50	In-memory File execution	✓	✓	✓

Figure 9 — Phase 1: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

✗ - Indicates the product **failed** to prevent or detect or respond to the attack in the tested scenario.

✓ - Indicates the product **successfully** prevented, detected, or responded to the attack in the tested scenario.

For an active response (preventative action) to occur, we verified whether the product made an active response during any of the three phases. Similarly, for a detection event to occur, we verified that the product saw various indicators that tied the threat to the adversary.

And finally, for the passive response to occur, we verified whether or not it was possible for the SOC analyst to respond to that threat using the product.

Palo Alto Networks performed exceptionally well at blocking the attack scenarios before the attacker was able to get a foothold inside the environment.

Phase-2 Metrics: Internal Propagation

In this phase, the EPR product should be able to prevent internal propagation. This phase is triggered when the initial identification and prevention of the threat fails. The EPR product in this phase should enable the analyst to immediately identify and correlate the internal propagation of threat in real time.

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 gets a foothold inside the environment, they try to escalate the privileges. For an active response to occur, we looked at various phases inside that method to see if there was a preventative action by the product.

For a detection event to occur, we looked at various indicators that tied the threat to the adversary. And finally, for the passive response to occur, we looked at whether or not it was possible for the SOC analyst to respond to that threat using methods provided by the product.

Tested Scenario	Description	Active Response	Detect	Passive Response
29	Forged Signature added to a File	✓	✓	✓

Figure 10 — Phase 2: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

✕ - Indicates the product **failed** to prevent or detect or respond to the attack in the tested scenario.
✓ - Indicates the product **successfully** prevented, detected, or responded to the attack in the tested scenario.

Palo Alto Networks Cortex XDR Pro was able to provide visibility and context for the threats that progressed to Phase-2. The product offered adequate options for the analyst to identify and correlate threats.

Discovery for Lateral Movement: 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 potential target of the attack. This is typically done by scanning the network.

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 ensures that they are able to access the resources they want and will not be flagged by the system's defences as an intruder. 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.

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.

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.

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.

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: Having found and extracted the target information, the attacker will try to delete or destroy all the evidence of the attack that remains within the target network. An ideal scenario for the attacker may well be one in which the victim does not even realize that the attack has taken place. Whether or not this is possible, the attacker will try to manipulate data inside the target environment to make sure that their tracks are covered as far as possible. This will ensure that the victim does not have the forensic information needed to understand the attack or trace the attacker. Data manipulation, deletion, and encryption (as used in ransomware) are the typical techniques that are used to do this.

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

Figure 11 — Phase 3: Active versus Passive Response of Palo Alto Networks Cortex XDR Pro

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

Palo Alto Product Response Mechanism

EPR products will use their response mechanisms to deal with the intrusions that have occurred inside the protected environment. As 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 to the various steps that the attacker took while attempting to breach the environment. For every step that was taken in Phase 1 and Phase 2, ESET was able to demonstrate both and active and passive responses to most of the attack techniques used, and in doing so, stop the attacker from successfully executing a full scenario.

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 if any form of intended remediation mechanism mentioned below could be completed by the analyst (Response Enablement) - based on what is supported by the product - this was evaluated and verified by AV-Comparatives as shown in the table below.

Palo Alto Networks Product Capability	End user	IT Admin[3]	SOC Analyst
System Imaging			
Patching			
System Restoration		✓	✓
Quarantine	✓	✓	✓
Network Isolation		✓	✓
Process Termination		✓	✓
Execution Prevention		✓	✓
Uninstall Services	✓	✓	✓
Shutdown or Reboot of Endpoint	✓	✓	✓
Edit Registry Keys and Values		✓	✓
Block Processes from Communication		✓	✓
Delete Files and Directories		✓	✓

Figure 12 — Response Actions (EPR Response enablement by Palo Alto Networks Cortex XDR Pro)

[3] Reported as provided by the vendor (not evaluated as part of the test).

EPR Competitive Product Differentiator (provided by Palo Alto Networks)

1. Exploit protection by techniques for "any" process you want to add (Windows, macOS, and Linux).
2. Behavioural Protection for 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 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 automated root cause analysis and 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. Consolidate endpoint security policy management and monitoring, investigation, and response across your network, endpoint, and cloud environments in one console, increasing SOC efficiency. Eradicate threats without business disruption: Shut down attacks with surgical precision while avoiding user or system downtime.
12. Protect your network against malicious insiders, policy violations, external threats, ransomware, fileless and memory-only attacks, and advanced zero-day malware.
13. Disrupt every stage of an attack by detecting indicators of compromise (IOCs), anomalous behaviour, and malicious patterns of activity.
14. 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.
15. 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.
16. Exploit protection by technique, for "any" process you want to add (Windows, macOS, and Linux)
17. Behavioural Protection for Windows, macOS, and Linux
18. Integrated sandbox analyses for all unknown samples, and displaying the full report for each of them (Windows, macOS, and Linux)
19. Integrated response actions including, Live terminal with Full CMD, PowerShell, Shell and embedded python. Also, Isolation in Bulk-Script Execution (Python) in bulk (Windows, macOS, and Linux)
20. Permission control is granular, from Super-Admin to Read-Only.

Central Management and Reporting

Management workflow is a top differentiator for any security control – if a product is difficult to manage, it will not be used. 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 supplement an existing technology or replace it completely. The management console can be deployed as cloud-based console. Communication between the Agent and Management console is done via SSL transport security. Figure 13 provides information on the capabilities of the product that was tested in this version of the EPR group test by AV-Comparatives

Reporting Features	Palo Alto Networks
Threat Visibility	
Attack Visualization	✓
Attack Timeline	✓
Attack Phases	✓
Attack Context	✓
System Visibility	
Continuous Monitoring	✓
Running applications	✓
Running processes	✓
Behaviour Monitoring (File/registry/etc..)	✓
Whitelisting capability	✓
Data Sharing	
Standards-based application programming interface (API) for access	✓
Standard output format (JSON, Syslog, CEF, etc..)	✓
Automated data export	✓
Syslog integration	✓
Splunk integration	✓
Additional reporting features	✓
Encryption of data at rest	✓
Targeted capture/e-discovery	✓
Customizable default security policies	✓
Policy and/or signature rollback	✓
Management to agent encryption	✓
Built-in-reporting capabilities for different user categories	✓
Multiple EPR analyst/user-focused workflow support	✓
Report automation	✓
Compliance reports (GDPR, PCI-DSS, etc.)	✓
Audit trail support in the management console	✓
System scanning capability	✓
Disaster Recovery	✓
Cloud marketplace support	✓
Integration with security products	✓
Enterprise recording and data storage – forensic analysis	✓

Customized reporting and management	✓
Custom reporting and filtering	✓

Figure 13 — Management: Threat Visibility, System Visibility, and Data Sharing

Palo Alto Networks 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. While providing maximum flexibility to senior analysts, the EPR should support predefined (but configurable) workflows for less-experienced personnel, who will be assigned specific tasks during an investigation.

In the following, the reporting capabilities of Palo Alto Networks Cortex XDR Pro are being listed.

IOC Integration

This is to identify the digital footprint wherein the malicious activity in 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, snort signatures or threat intelligence feeds etc. as shown in Figure 14 below.

External IOC Correlation	Product Capabilities
SIEM	✓
DNS Logs	✓
Network traffic flow logs	✓
DHCP Logs	✓
Scan results	✓[4]
YARA Signatures	✓
Multi-factor authentication logs	✓
Sandboxing logs	✓
Retrospective analysis and Logs	✓
Endpoint prevention product logs	✓[4]
Proprietary product integration (NGFW, IPS, …)	✓
Threat intelligence data assimilation	✓

Figure 14 — External data correlation supported by Palo Alto Networks Cortex XDR Pro

[4] Capability is provided also by Palo Alto Networks' endpoint product.

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 make 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 by the engineers of the vendor during setup. This configuration is typical in enterprises, which have their own teams of SOC analysts looking after their defences. The personas and the threat emulation that were run in this evaluation represent such scenarios. It is common for products of these kinds that vendor experts assist companies on the deployment and configuration best suited for the type of enterprise.

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

Palo Alto Networks: "Agent Settings", "Agent Security", "XDR Pro Endpoints", "Content Auto Update" and "Direct Server Access" were enabled. "Alert Data Dump File Size" was set to "Full". "Automatically Upload Alert Data Dump File", "Agent Upgrade" and "Network Location Configuration" were disabled. "Browser Exploits Protection", "Logical Exploits Protection", "Known Vulnerable Processes Protection" and "Operating System Exploit Protection" were set to "Block". "Exploit Protection for Additional Processes" was disabled. "Unpatched Vulnerabilities Protection" was set to "Modify Settings until the Endpoint is Patched". "Portable Executable and DLL Examination", "Office Files with Macros Examination", "Behavioral Threat Protection", "Ransomware Protection", "Malicious Child Process Protection" and "Network Packet Inspection Engine" were set to "Block". "Respond to Malicious Causality Chains", "End-User Initiated Local Scan", "Password Theft Protection" and "Monitor and Collect Forensic Data" were enabled. "Endpoint Scanning" was disabled.

Operational Accuracy (False Positives)

Operational Accuracy Tests were performed by simulating typical user activity in the enterprise environment. This included opening different file types, and browsing to different websites. Furthermore, different administrator-friendly PowerShell scripts were executed in the test environment to ensure that productivity was not affected after product installation and configuration.

The product passed all Operational Accuracy Tests.

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. The in-depth testing ran for a four-week period. 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[5] and NIST platform, so that it becomes easier to analyse the risk for a specific endpoint.

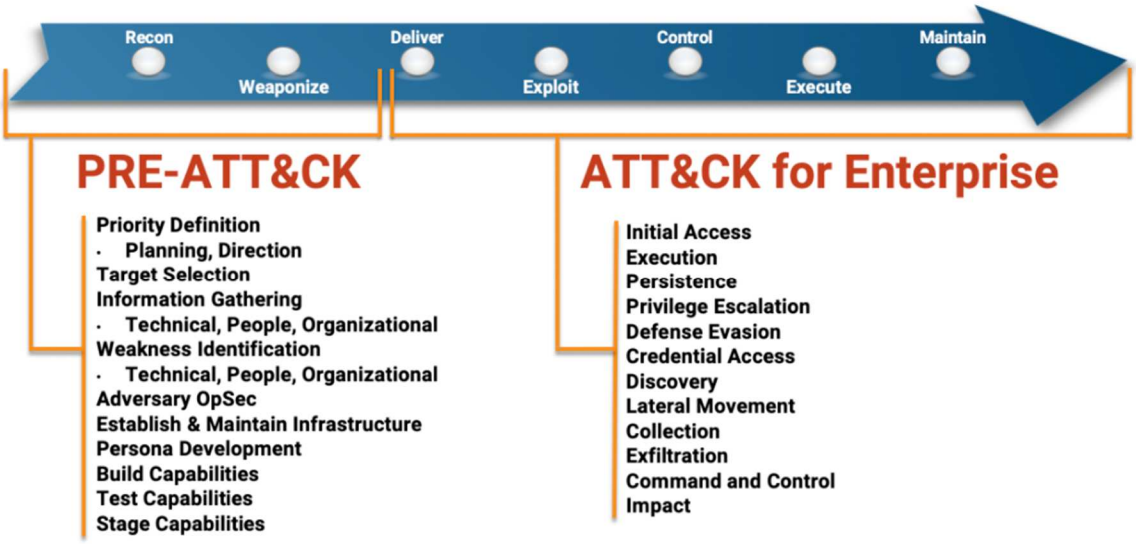

Figure 15: MITRE ATT&CK for Enterprise vs Seven Stage Cyber Attack LifeCycle[6]

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 Figure 16 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.

[5] © 2015-2021, The MITRE Corporation. MITRE ATT&CK and ATT&CK are registered trademarks of The MITRE Corporation.

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

Active Response vs Passive Response Workflow

While evaluating EPR products, the ultimate adversary is not the malware or the tools that the attacker is using, but rather the adaptive, intelligent, and motivated attacker who uses malware, threats and other tools for distraction, advancement, lateral movement, escalation and much more, all of which the EPR product is expected to prevent and respond to. Therefore, this EPR report includes security efficacy metrics around different test scenarios and product differentiating factors. This will enable enterprises to make informed decisions on the suitability of each tested product for their requirements.

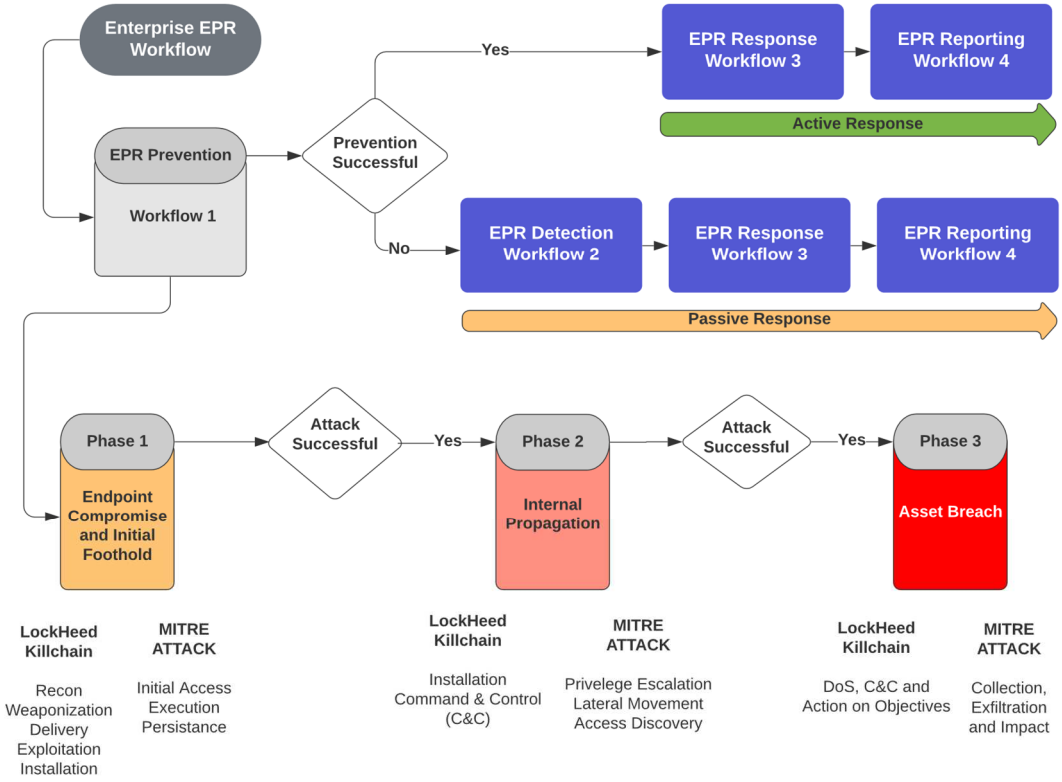

Figure 16 — Enterprise EPR Workflow Overview

Whether attacks are defined as Malicious Operations, Campaigns, Detections, Kill Chains or anything else, it is these human pathways that should be highlighted, which we are referencing as four distinct workflows in this report.

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 security analyst 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 analyst workflows and scenarios, and not on the technology used to prevent, detect or respond to it.

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 etc. 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/automated), 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 analyst. Once they have been identified, the analyst 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 following response mechanism, in order to be considered for the detection scoring metrics.

- Response based on successful identification of attack via the product's user interface (UI), which lists the attack source (http[s]/IP-based link) hosting the 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.

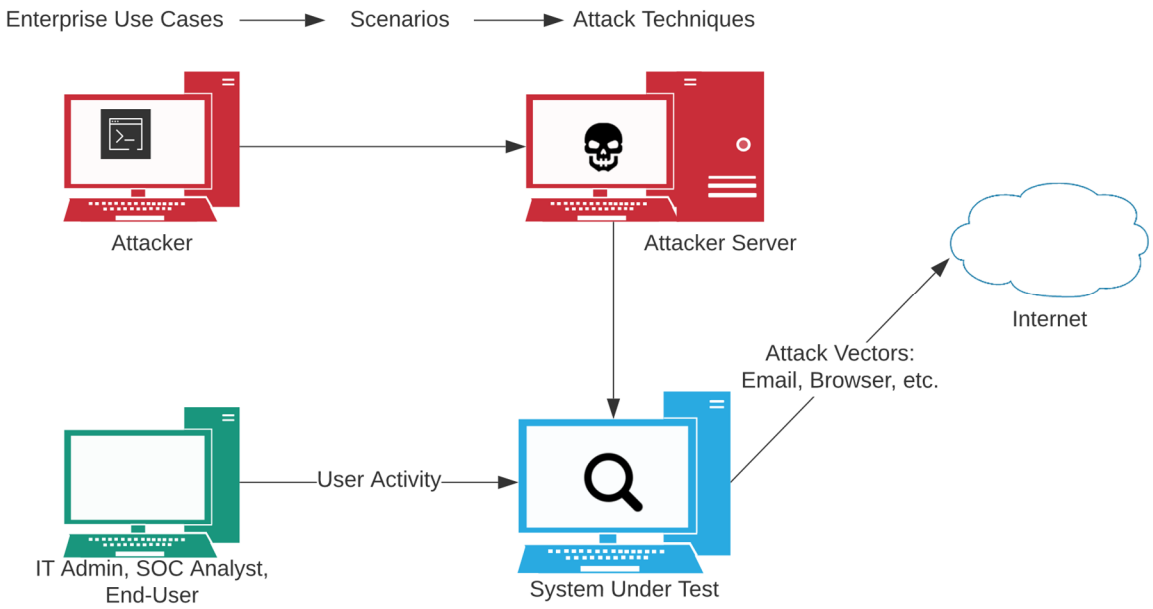

Figure 17 — EPR Test Topology Overview

All vendor 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 that are 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 in 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 LAN server. All test scenarios were executed from beginning to end, to the greatest extent possible.

Test Iteration Objective

The objective of the testing was to assess the prevention-centric workflow with specific use-cases targeted for EPR prevention Workflow-1 (referenced in the methodology) with threats that typically target enterprise users in a normal operational environment. This iteration helped us to assess the default prevention capability of the product, along with the detection mechanism. If a threat was not prevented, we evaluated if the EPR product was able to take appropriate detection and response measures in a timely manner.

The following assessment was made to validate if the EPR endpoint security product was able to prevent and detect all the attacks on the EPR Prevention Workflow-1 and Detection workflow.

- Did the prevention occur during Phase 1 (Endpoint Compromise and Foothold) of the prevention workflow?
- Did the EPR product provide us with the appropriate threat classification and threat triage, and provide an accurate threat timeline for the attack with relevant endpoint and user data?
- Did the EPR product demonstrate any negative issues in the operational accuracy test that was executed in conjunction with the attack scenarios?

Targeted Use-Cases

The user types that we considered during the test iterations were “IT Administrator”, “Regular Enterprise User”, “SOC Team Professional”, and “Analyst”. The sequence of events emulated was an enterprise-based scenario wherein 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, when executed, successfully allowed an attacker to get a foothold inside the environment and take additional steps to act upon their objectives.

During the time of testing, our analyst acted as an SOC Analyst, Administrator and an SOC Professional by logging 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 was provided to the analyst in a timely and clear way. We tested the responses available using the products in the test.

EPR Test Iteration Timeframe

The evaluation was conducted in four phases, each phase lasting a week. As weeks progressed, AV-Comparatives was able to have a detailed understanding of the product under test and attacks were crafted in such a way that they stressed the product’s true capabilities. Furthermore, Workflow-1 was conducted with an attacker-driven mindset as the attack progressed through the attack nodes to finally meet its objective. The evaluation was conducted in autumn 2021. User persona and user activities were simulated throughout the test such that they were as close to the real environment as possible.

All the attacks were crafted using open-source tools, and samples were developed using in-house expertise. Once the attacker gained initial access to the environment, they tried to be as stealthy as possible, so as not to trigger any defence mechanisms.

About this test

The 2021 Endpoint Prevention and Response (EPR) test for enterprise products performed by AV-Comparatives is currently in its second iteration this year. Participating in the main comparative report and the publication of the test is optional at the vendor's discretion.

The complex nature of the test means that automation is not possible, and so it has to be performed entirely manually, making it cost-intensive to run. Because this methodology is tailored towards the prevention, detection and response capabilities, all vendors are advised to turn on the prevention and protection capabilities (ability to block), and configure detection and response features such that they demonstrate the real-world, enterprise-class capabilities of the products deployed. The methodology supports EPR product updates and configuration changes made by cloud management console or local area network server. The intent was to execute all test scenarios from beginning to end, to the greatest extent possible. Unless absolutely warranted, vendor-recommended EPR product configurations were not updated, and vendors were contacted and findings documented, if required at all. If there were workflows mentioned in this methodology that required specific configuration changes and/or options, vendors discussed and worked with AV-Comparatives on those options during the initial setup and baselining phase.

Some vendors asked for precise details of the day and time the test would be performed, so that they could monitor the attacks in real time and interact with their products when they thought it beneficial. Because the aim of the test is to measure protection and response capabilities, we did not provide any vendors with any advance information about when the test would be performed. In real life, attackers do not tell their victims when they are going to attack, so products must provide protection all the time. We also had information requests from vendors regarding the attack methods to be used in the test.

We did however invite all the endpoint vendors who had prevention and response capabilities to be a part of the main EPR test, and invited them to provide feedback on how it might be improved. Each vendor was provided with the methodology, sample test report, and the enterprise CyberRisk Quadrant to review well in advance of the test and give their respective feedback. As a result of the feedback we received, we implemented some changes in the test methodology, where we felt that this was in the genuine interests of users and enterprise-related workflows, and where these helped to promote the general security efficacy metrics of the EPR products.

The test is very challenging, but at the same time it also reflects realistic scenarios. We have had positive feedback from many vendors' technical departments. To get an overall picture of the protection and response capabilities of any of the tested EPR products, readers should look at the results of the other tests in AV-Comparatives' Enterprise Main-Test Series[7] too.

[7] https://www.av-comparatives.org/enterprise/

Copyright and Disclaimer

AV-Comparatives
(January 2022)

Credits: Icons made by icon_king1 from freeicons.io