Independent Tests of
Anti-Virus Software

EPR - Endpoint Prevention and Response CyberRisk Test Methodology 0.90

TEST PERIOD: 2020
LANGUAGE: ENGLISH
LAST REVISION: 11TH MAY 2020

WWW.AV-COMPARATIVES.ORG

Contents

Introduction

Until fairly recently, enterprise antivirus software had been oriented towards file-based threats. Such attacks involve getting malicious executable files onto the target system and running them. Endpoint protection solutions of the time were able to cope well with the threats, regardless of the source. It was not important whether the malware was introduced via the Internet, LAN or removable devices, or what type of damage was caused (data theft, system destruction or ransomware), antivirus software could detect it. However, the advent of advanced threats, such as Advanced Persistent Threats (APTs) in modern times has changed the threat landscape significantly. They employ techniques, tactics and procedures that can bypass the defences of traditional antivirus software. For example, file-less malware that uses legitimate system tools for malicious purposes can evade a scan for malicious executables on the system disk. Hence, organizations need to protect their endpoint assets against these additional threats.

The objectives of APTs typically include establishing and extending footholds within the targeted organization for the purposes of data exfiltration and/or sabotage. The attack may execute its objectives immediately, or position itself to carry them out at a future date. Most endpoint prevention and response (EPR solutions) allow IT operations and security staff to proactively work together to ensure endpoint security. They can prevent breaches, and identify, save, and analyse relevant data from the company's endpoints. Armed with this, they can go about hunting security breaches, and proactively prevent and respond, thus improving defences against future attacks. EPR solutions, and endpoint prevention systems with prevention and response capabilities, are consequently being used more and more to protect business networks. They provide an additional line of defence against these attacks.

EPR solutions (or endpoint prevention solutions with EPR capabilities) are becoming an integral part of the security-tool arsenal that IT security professionals and security operations teams are using. This is due to their highly scalable nature, faster prevention/detection/response times, along with their abilities to provide containment and unearth the stealthy activities that are typical Indicators of Compromise (IOC) or an Indicators of Attack (IOA).

What is so special about this EPR Methodology?

The Mitre ATT&CK framework is a knowledgebase of attack tactics and techniques. While the Mitre ATT&CK FRAMEWORK is an excellent tool for mapping out TTPs (tactics, techniques and procedures), the framework is threat-centric in nature and is currently only available to solutions that are oriented towards detection and response, rather than protection.

Also, it doesn't address all the use cases, EPR workflows in different scenarios, or product differentiating factors. It is also slow in adapting to emerging attacks that are evolving in the wild. It lacks a composite scoring mechanism to easily highlight strengths and weakness of tested products. Almost every methodology created by other testing firms is also threat-centric in nature and can't be used by enterprises to reduce their true operational risk.

AV-Comparatives has developed an industry changing paradigm shift towards defining EPR response according to the everyday reality of enterprise use cases and workflows. As such, AV-Comparatives is defining prevention and response as highlighted below.

1. Prevention.

The best way to respond to any threat is by preventing it. This is fundamentally the prevention capability of the EPR product. AV-Comparatives defines *prevention* as an automated, active response that kicks in 24/7, 365 days a year, without the need for human 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, rather than the technology used to prevent, detect or respond to it.

2. Passive Response

Once an Attacker is already inside the enterprise environment, traditional response mechanisms kick in, for example IOC and IOA correlation, external threat intelligence, 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 expected to prevent initial and ongoing attacks, without having to triage the threats, while offering active response and reporting capabilities. If the attack is not prevented, an EPR product should be able to assess and respond to it, thus lessening the burden on human/automated resources, thus providing better ROI in the long run. Some testing frameworks provide some useful metrics, but they fall short in correlating and evaluating the cyber-attacks. They focus on “atomic” tests, i.e. ones that only look at a particular component of the ATT&CK framework. Although atomic test-based evaluation of the ATT&CK framework is a good first step, it misses the larger picture. Real-world attack workflows interconnect at every stage from the initial execution to final data exfiltration/sabotage. Atomic test-focused execution also has a higher probability of generating false positives, thus wasting valuable time and resources by requiring further investigation.

Our testing methodology addresses those problems, so that vendors can address potential technology gaps, and enterprises can model their operational risks. Furthermore, this methodology is robust enough to address future changes as well.

EPR Inclusion Criteria, Setup and Configuration

Vendors with appropriate EPR products are invited to join the test. If you are participating with an EPR product, it will be in a standalone mode, with each vendor actively involved in the initial setup, configuration and baselining aspects. AV-Comparatives has compiled a list of popular scenarios requested by enterprises. Every vendor will be allowed to configure their own product, to the same extent that businesses are able to do when deploying it in their organizations. The details of the configurations will be included in the report.

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), along with detection and response, so that it emulates the real-world enterprise-class capabilities of these products. The methodology supports EPR product updates and configuration changes made by cloud management console or local area network server. Our intention is to go through and execute all test scenarios from beginning to end, to the greatest extent possible. Unless absolutely warranted, we will not update vendor-recommended EPR product configurations, and will ensure that we reach out to the relevant vendors, and document our findings, if such updates are required. If there are workflows mentioned in this methodology that require specific configuration changes and/or options, it is best that the respective vendor discuss these with AV-Comparatives, and work with us on these options during the initial setup and baselining phase.

Note: AV-Comparatives will give participating EPR vendors sample test cases and scenarios to trigger workflows based upon this methodology prior to setup. This will allow them to configure their products for optimal prevention, detection and response capabilities, thus reflects the operational realities of today's enterprise environment.

EPR's metrics and analyst workflow

The primary purpose of evaluating any enterprise-grade security product is to come up with quantifiable metrics that provide meaningful data. AV-Comparatives is taking a new approach towards defining the EPR metrics and workflows. It is the first testing lab in the industry to showcase individual EPR products' capabilities. AV-Comparatives maps these capabilities and workflows to enterprise use-cases. This will allow enterprises to make educated decisions on return-on-investment (ROI), risk-tolerance (based upon the EPR product's ability to cope with threats), total cost of ownership (TCO), and the overall dwell time of threats.

The analyst workflow can be further operationalized using the Cyber Kill chain and Mitre ATT&CK framework with other security controls.

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

Reduction in TTP (time to prevent)
The ability of the EPR product to rapidly identify and prevent a threat and display relevant information is a very important factor. This could also be referred to as the effective reduction in active time to respond. This could include blocking the primary vectors (email, web, SMB, etc.) and secondary vectors that are typically used for lateral movement. The faster the EPR software can prevent a threat and provide the analyst with the relevant information, the sooner the attack can be stopped, and the sooner any damage can be remediated through effective response. Reduction in TTP also means a lower time to response (TTR).

AV-Comparatives will report the TTP of each tested EPR product.

Reduction in TTR (time to respond)
Time is of the essence when an incident that has not been 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 expertise of the user and the capabilities of the product. Hence a reduction in the passive response time becomes critical to dealing with any breach.

AV-Comparatives doesn't expect the tested EPR products to come with real-time passive response, because this can bring a deluge of false positives that will severely hamper the operational environment. It does however expect EPR products to come with a reasonable response timeframe (seconds to minutes) for a potential incident. The sooner the EPR product comes up with the response, the better it is.

AV-Comparatives will report the TTR of each tested EPR product. The TTR is to be found by identifying and collecting all the event details to help analysts perform additional research and investigation to categorize the threats. The EPR product should be able to control the situation regardless of the location of the endpoint.

Threat Classification

Not all threats are of equal severity. Being able to classify attacks according to the risk each one poses is an important feature of an EPR system. While reduction in TTR and TTP gives time to prevent and respond to incident(s), threat classification gives the users the ability to understand the severity of the threat based on additional research and classification. This enables the enterprise to prioritise its remediation efforts effectively.

If there is a specific threat classification methodology or framework used by the vendor's EPR product, it is best that the vendor discuss this and work with AV-Comparatives on those details during the initial setup and baselining phase. This will allow AV-Comparatives to map the workflow outcomes to the appropriate threat classification metrics.

Threat Triage

Classifying threats according to how they will be resolved helps enterprises respond to attacks in a fast and meaningful approach. Rather than dealing with each threat individually, the admin can potentially resolve a number of similar threats all together, thus saving time and resources.

Threat Timeline

Advanced attacks typically take place over an extended time period. In order to understand the nature of a threat, it is necessary to find out which actions took place at what time. Hence, it is important to have a detailed timeline of how each attack has progressed from its initial stages to completion, along with any relevant IOCs along the way.

Endpoint and User Data

The EPR should be able to present all user and endpoint data relevant to an attack. Examples include: username of logged on user, user type (administrator or standard), hostname, IP address at time of attack, type of attack, network communications used, and other endpoints similarly affected. Such data is critically important in understanding the nature of the threat.

EPR analyst workflow

Based on feedback from enterprise clients, AV-Comparatives has developed the following workflows. They are used for enterprise environments with heterogeneous requirements. While these workflows may not encompass the needs all existing enterprise networks, they represent a minimum standard for EPR products, and will be required universally.

The EPR analyst workflow can be best summarized by the following diagram, Figure 1.

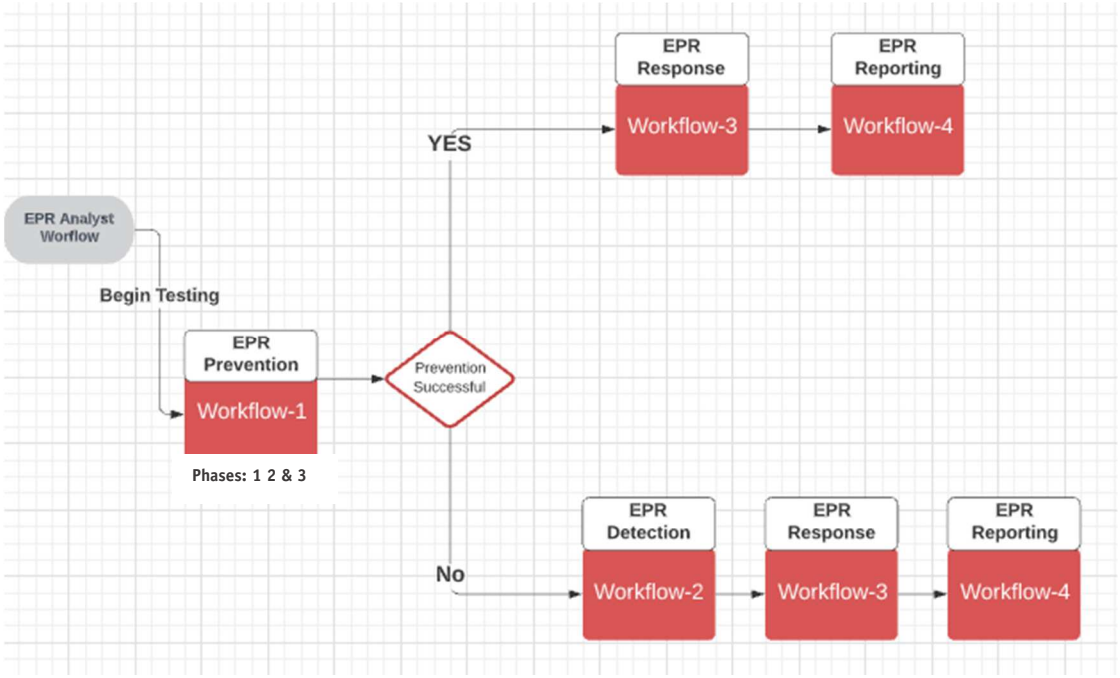

Fig 1: EPR Analyst Workflow

As shown by Figure 1, an EPR product should be able to first identify and prevent attacks using an active response, which, assuming the attack was not prevented, should be followed by a detect-and-respond passive response. It should also provide detection information and response options from a single window in its reporting capabilities.

The EPR's analyst system should be able to immediately identify the threat and correlate it with its communications to the attackers in real time. It should also detect and identify threats that were not prevented, and classify their severity and resolution techniques, along with the ability to respond to them using a specific workflow of the EPR product.

The Analyst Workflow will be triggered by an attack based on the Mitre ATT&CK framework and other methods, and can be effectively mapped to the Lockheed Cyber Kill Chain wherever applicable.

EPR Workflow 1: Prevention (Active Response) Mechanism

An EPR product should be able to first identify and prevent (active response) attacks for a specific workflow or an attack scenario.

The attack scenario in this workflow typically goes through a 3-phased approach: first, the initial endpoint compromise and gaining of a foothold by the attacker; second, internal propagation; and third, active asset breach. These 3 phases are illustrated in Figure 2 below and are documented in detail.

EPR Workflow 1: Phase 1 - Endpoint Compromise and Initial Foothold

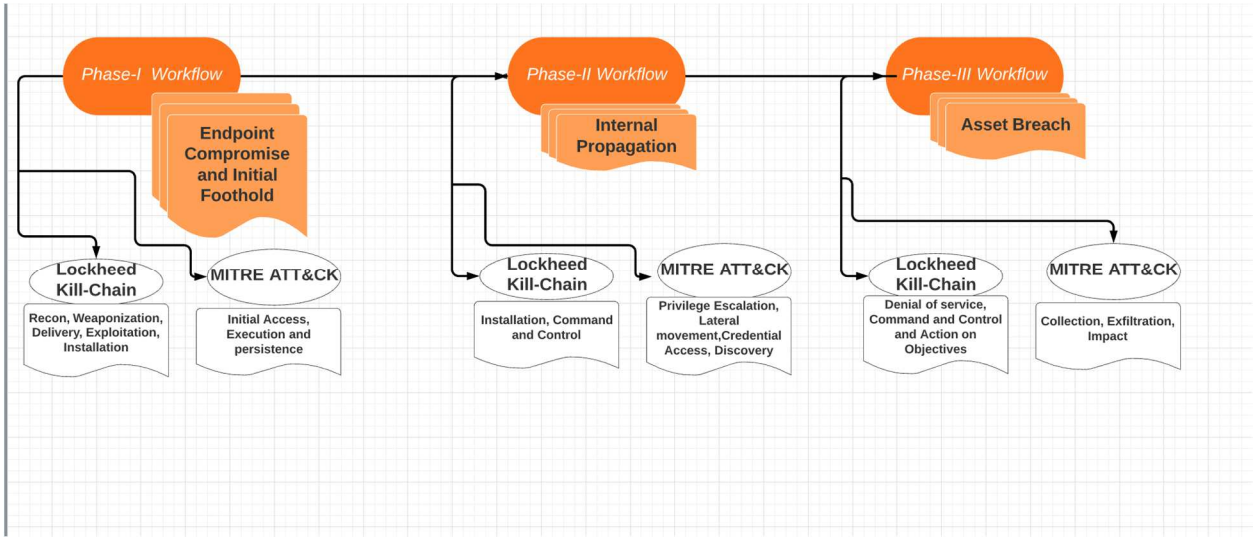

Fig 2: EPR Workflow 1: Phase 1, 2 & 3

An EPR product should be first able to identify and prevent a threat, then detect and respond to the threat on a compromised machine within minutes and provide detection information and response options from a single window. Because this is the first phase of the attack, faster the EPR prevents and enables response in this phase, more effective the EPR is. This will enable organizations to defend against the attacker before a compromise and a foothold is achieved within the enterprise infrastructure.

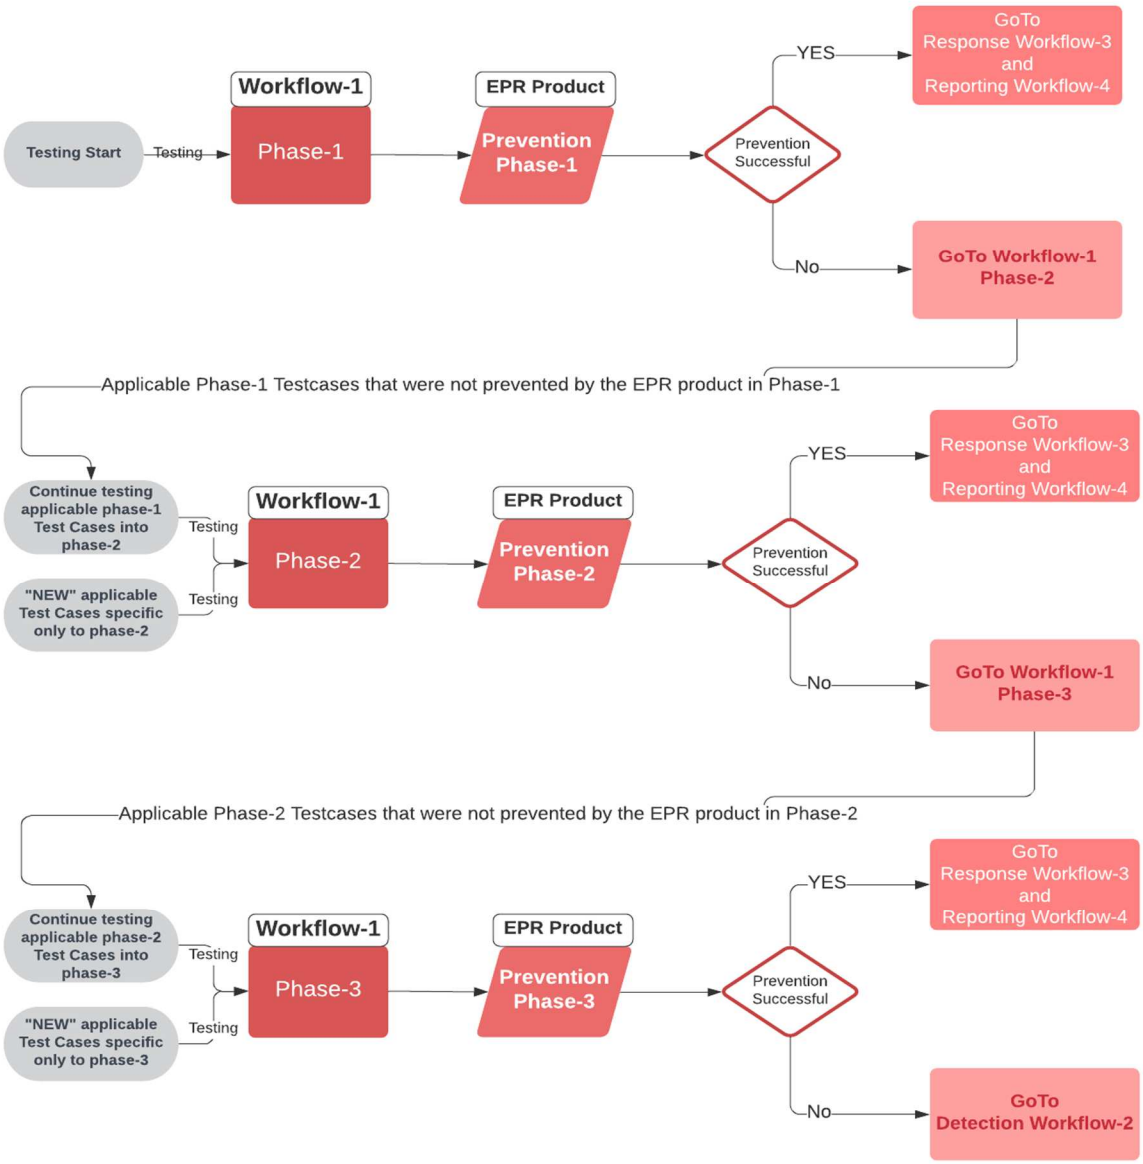

Fig 3: Detailed EPR Workflow 1, 2, 3 & 4

EPR Workflow 1 can be triggered by an attack based on the Mitre ATT&CK and other methods and can be effectively mapped to the Lockheed's Cyber Kill Chain. This workflow can be operationalised 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 typically uses operating system features to plant 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.

EPR Workflow 1: Phase 2 - 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 be able to enable the analyst to immediately identify and correlate the internal propagation of threat in real time. The analyst should be able not only to perform the necessary actions to identify, detect, classify and triage a threat, based on the data collection and analysis, but also to complete the response using the EPR product as a specific workflow tied to this phase.

This workflow can be triggered by an attack based upon Mitre ATT&CK or other methods, and can be effectively mapped to Lockheed Kill-Chain[4]. This phase of attack can be operationalized by the attacker using the steps described below.

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.

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.

EPR Workflow 1: Phase 3 - Asset Breach

The final phase of the workflow is asset breach. This is the stage where an attacker truly starts acting on its true objective. This workflow can be triggered by an attacker based upon Mitre ATT&CK or other methods and can be effectively mapped to Lockheed Kill Chain. This phase of attack can be operationalized by the attacker through the steps described below.

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 realise that the attack has taken place. Whether or not this is possible, the attacker will try to manipulate data inside the target environment to ensure 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.

EPR Workflow 2: Detection Mechanism

This workflow is triggered when the EPR products provide relevant alerts and logs. These let the analyst investigate the prevented attacks, so that he/she can discern that they are not False positives rather than genuine threats. If the EPR product is unable to prevent attacks, then the analyst can look at the detection, in order to prepare the response mechanism.

Any alerts, logs, or notifications with appropriate threat indicators will qualify as detection. Alerts or logs that do not suggest a threat will not qualify as a detection mechanism.

AV-Comparatives will report the outcomes of the detection as articulated by vendors. This will ensure that the analyst gets appropriate visibility and the ability to discern and interpret detection mechanism that best suits their environment.

In short, for every attack that is not prevented during any of the 3 phases in Workflow 1, AV-Comparatives expects the EPR product to provide detection capabilities which will help the analyst to plan and execute an effective response strategy.

EPR Workflow 3: Passive 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.

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 product to have non-automated or semi-automated passive response mechanisms.

AV-Comparatives will list the outcomes of this workflow as part of the Test Report. No single solution fits all responses that could be triggered by test cases in this methodology. It's entirely reasonable to have a variety of response mechanisms for one test case, even within a single EPR product.

A very simple example of an EPR product offering a multi-response capability is a scenario where the product both terminates a suspicious process and cuts off network communication. In this example, both the responses will be reported as part of the Test Report.

EPR Workflow 4: Reporting Capability

During an investigation, analysts may need to configure ad hoc groups of endpoints within the EPR product that could not have been anticipated in advance. Examples of such groups might be all affected endpoints, or all vulnerable endpoints. Admins should also be able to use the EPR product to review past incidents and the action taken at the time in order to decide if the same actions are applicable to the current threat. 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.

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.

Prevention Avoidance

Attackers will obviously make every effort to ensure that their attacks will not be prevented by the EPR product (or any of the victim's security measures). Product-specific EPR evasion measures are currently outside the scope of this methodology. However, AV-Comparatives will make use of common content-masquerading techniques and employ trusted apps and processes in order to mimic detection avoidance. This will determine whether the tested EPR products can prevent attacks that deliberately use obfuscation techniques. This will give a better understanding of the products' prevention capabilities in realistic enterprise scenarios. This section will be independent of the workflow listed above and will focus on testing resiliency of the products against more advanced attacks.

Emerging Attacks

Targeted attack techniques are not static in nature but evolve continually. As EPR products improve their defensive capabilities, attackers change to different methods to try to circumvent them. In acknowledgement of this, AV-Comparatives will update the attack methods used for the tests, as they develop. Fileless attacks and machine-learning poisoning are examples of currently evolving attack techniques. This section will be independent of the workflow listed above and will focus against more novel attacks.

Operational Accuracy Test

A security product that reports 100% of malicious attacks, but also reports legitimate (non-malicious) actions, can be hugely disruptive/noisy. AV-Comparatives will use appropriate tools and techniques to ensure that the tested EPR products do not raise significant numbers of alerts with legitimate applications and processes. This section of the methodology will be performed in conjunction with Workflow and other independent section as much as possible. This will ensure that EPR products aren't heavily biased towards prevention and response by sacrificing operation accuracy in enterprise environment.

EPR Scoring Criteria

The different EPR Workflows mentioned above, along with prevention and detection avoidance, emerging attacks and the false alarm test, will be used as the scoring mechanism to report the overall prevention and response metrics of the EPR product. We will also be measuring and reporting on the TTP (time to prevent) and TTR (time to respond) metrics of each of the participating EPR products.

All this will be included as a part of each product's individual test report, and the comparative test report as well. Only qualifying products will be evaluated and represented in the comparative report. As a part of the test, we will also be reporting on some of the key EPR capabilities of each product from a functional standpoint, so that enterprises are well-informed.

Glossary of Terms

Cyber Kill chain

"The Cyber Kill Chain® framework is part of the Intelligence Driven Defense® model for the identification and prevention of cyber intrusions activity. The model identifies what the adversaries must complete in order to achieve their objective."

MITRE ATT&CK®

MITRE ATT&CK® is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations.

Appendix

Sample Selection, Curation and Sourcing:
Sample selection and sourcing will be based on the realities of the threat landscapes during the testing window. Publicly and privately available threat intelligence report, sources, techniques will be collected and assessed to see the viability of inclusion for the test. Enterprise feedback will also be solicited for Sample selection and curation purpose.

Endpoint and User Data
An EPR product should at minimum list the following data about the endpoint:

- Endpoint Host Name
- Endpoint IP Address and FQDN as applicable
- Logged on users and user types
- Attack relevant information: Process Name, File Name, Hash, Process Network Communication, Parent-Child Relationship, and File and Process operations, wherever applicable.
- Registry Information and Registry Operations wherever applicable.
- Threat Classification

Copyright and Disclaimer

For more information about AV-Comparatives and the testing methodologies, please visit our website.

AV-Comparatives
(July 2020)