

WHITE PAPER

Assess Your Endpoint Security

Better Understand Capabilities Using MITRE Engenuity ATT&CK Evaluations



Executive Summary

Making a decision on endpoint security tools is complex, so access to objective vendor-neutral information is key. Organizations can use the results from the MITRE Engenuity ATT&CK evaluation to evaluate the efficacy of endpoint solutions. They also can gain a better understanding of where their current security posture stands by using MITRE tools to evaluate their existing detection coverage and gaps. Then they can compare their results against MITRE's extensive list of tactics and techniques that adversaries use in real-world attacks.

Endpoint security is all about reducing risks, so before evaluating solutions, security professionals need to make sure they have the basics in place. Best practices in security hygiene can dramatically lower threat exposure, and security professionals should have a strategy for improving their security posture and visibility. Once they have these fundamentals in place, they can use the results of third-party tests to evaluate and select solutions that best fit their needs.

Objective Product Comparisons

The constantly evolving threat landscape is characterized by substantively new versions of existing attacks as well as brand new cyber threats. It's critical that an organizations' security posture and the individual security controls on which it is built remain strong. However, cybersecurity products can be difficult to assess because they often take different approaches or use different terms to describe how they do what they do. Third-party testing that evaluates products can help organizations compare "apples to apples" so they can make more informed choices and find the right solution for their current situation and their desired security posture. However, it can be challenging to determine the value of testing as it relates to the overall security infrastructure when testing takes a "black box" approach to reporting that only shows the outcome (blocked or missed, detected or undetected) rather than the operations.

MITRE ATT&CK Evaluations

Starting in 2019, the MITRE evaluations have acted as objective and detailed tests of the capabilities of endpoint security solutions by emulating real-world cyber campaigns and their techniques and tactics. The emulation plans are sourced with public cyber-threat intelligence reporting, mapped to a subset of ATT&CK techniques, and used to replicate the behaviors that generate objective insights into how well products perform. The results not only show the technical ability of a solution to detect a specific campaign but more importantly also show the techniques and tactics that are common among many of today's cyber threats.

The MITRE Engenuity ATT&CK evaluations are powerful because they are based on the [MITRE ATT&CK Framework](#), which is a knowledge base of adversarial techniques. It provides a breakdown and classification of offensive actions taken by attackers that can be used against particular platforms, such as Windows. Unlike prior work in this area, the focus isn't on the tools and malware that adversaries use but on how they interact with systems during an operation.

To provide context, the ATT&CK Framework organizes techniques into a set of tactics. Each technique includes information that's relevant to defenders to help them understand the context surrounding events or artifacts generated by a technique in use. The relationship between tactics and techniques can be visualized in the [ATT&CK Matrix](#), which spans 11 discrete techniques. The Matrix offers a robust and granular mapping of the activity of potentially utilized cyberattacks. Each area has 10 or more tactics and spans initial access through ongoing command and control.

The 2020 round of ATT&CK Evaluations focus on emulating Carbanak, a threat group that targets banks, and FIN7, a financially motivated threat group that has primarily targeted the U.S. retail, restaurant, and hospitality sectors, often using point-of-sale (POS) malware. They are excellent test cases because they:

- Heavily utilize scripting, obfuscation, malware hiding in plain sight, and exploitation of the users behind the machine
- Use a unique spectrum of operational utilities that span both sophisticated malware as well as legitimate administration tools that are capable of interacting with various platforms



Endpoint detection and response (EDR) was the most often-cited priority when organizations were asked what their biggest endpoint security investment priorities are for the next 12-18 months.¹

In 2020, MITRE expanded the evaluations to include detection results and/or protection results. Beyond whether a technique or tactic was detected at a point it could be blocked, the key difference is that once the tactic was blocked, no further evaluation of later-stage activities was possible.

The Detection Test

For the detection tests, the evaluations cover 20 test cases. Each test case has multiple stages. The evaluations used six terms to express how the product performed for each test and also noted the data source for the detection.

- **Not Applicable.** Used if the vendor did not deploy a sensor on the test system.
- **None.** No data indicating the test behavior was detected could be seen within the product.
- **Telemetry.** The behavior could be seen but was minimally processed.
- **General.** The behavior was processed and flagged, but without detail as to why (tactic) or how (technique) the action was performed.
- **Tactic.** The behavior was processed and designated malicious, as well as enriched with the Tactic or other notation about why it was performed by the cyberattack.
- **Technique.** The behavior was processed and designated malicious, as well as enriched with the Technique or other notation about how it was performed by the cyberattack.

In cases where a detection test, technique, or tactic is marked as Not Applicable, it's a good idea to understand why. It could mean that there is a lack of operating system or similar support required for deployment. Alternatively, it could mean that the product was not designed to identify the tested behaviors. Or there could be other reasons, which may or may not be important to your organization.

In cases where the tested behavior is often associated with potentially legitimate operations, as well as malicious ones, a designation of None may not be a negative result, particularly if the risk of letting the behavior pass is also None. An example is when the cyber criminal has not yet achieved their intended outcome.

Telemetry or General are the first level of detection where behavior can be identified and recorded in the product, but with limited detail about why it was identified. This level is sufficient for organizations that lack the time or expertise to dig into the detail of exactly how the cyber criminal is attempting to conduct their mission. Detections tagged with MITRE Tactic or Technique are most valuable to seasoned security professionals who want to understand the detailed activity of the cyber criminals.

The Protection Test

The protection tests only use three terms.

- **Not applicable.** This term is the same as in the detection test. Or it can reflect the fact that the test case was blocked before the remaining techniques.
- **None.** There was no evidence that the technique was blocked or otherwise unsuccessful because of the product.
- **Blocked.** The technique was blocked and the user was informed that it was unsuccessful.

Although there are fewer terms, the interpretation is more complex because the timing of the blocking is important and can vary based on the test case. Testing for false positives is not part of the scope of the evaluation, but consider that blocking early may increase the possibilities of false positives. In contrast, blocking too late may expose the organization to a degree of risk even if the end objective is not achieved. Here is an example of each scenario.

Suppose that Test 1 was blocked at the first step 1.a.1. The result sounds great; the cyberattack was stopped at the earliest stage possible. But what if it was actually the user execution that was blocked? In that case, you would want to know the basis on which the user was prevented from accessing a file. Was there a high confidence malicious indicator or were policies set too aggressively? Alternatively, suppose blocking occurred at the end at step 2.b.5, which was exfiltration over the command-and-control channel. In that case, the product stopped the intended data breach, but it did allow step 2.b.1 remote file copy to occur, which means the attack had a malicious impact.



In this particular case, arguably the safest time to block the attack to minimize the risk of false positives given the “proof” gathered and the malicious impact given the intended action would be at step 1.a.3. This step occurs when a script attempts the first malicious file manipulation. But this information is something that can only be determined after understanding each step and sub-step of each stage. The success or failure is based on an organization’s concern about impeded legitimate user activity vs. the risk of malicious impact from a cyberattack.

An Important Reminder About the Evaluation Results

MITRE emphasizes the fact that the evaluations are not a competitive analysis and there are no scores, rankings, ratings, or “winners.” Instead, they show the detections observed and how each vendor approaches threat detection through the language and structure of ATT&CK. The evaluations can help organizations answer questions such as:

- Does this tool detect known threats to your organization?
- How does the tool present the data to your analysts?
- Can it strike the right balance between aggressive detection/protection and the potential risk of compromise?

The evaluations can tell you which vendors provide the most visibility across adversary techniques and the vendors that best address the techniques that the threats use. An evaluation can show you the insights you may receive and how often a solution is updated to cover new adversary techniques. The evaluations also can tell you if a tool uses a graphical user interface (GUI) or offers turnkey options for less experienced analysts or provides raw data that more experienced analysts may need.

However, the evaluations can’t answer questions like:

- What is the impact on systems and users?
- What is the volume of alerts and the amount of manual research and investigation that is required?
- How does the tool fit within your broader security posture? Is the tool additive or duplicative?
- Does the system erroneously block a legitimate action?
- How does the tool integrate with your other tools?
- How much does the tool cost?

Getting answers to questions like these requires additional research, testing, and considering the bigger picture in your organization.

Takeaways From the Evaluations

The greatest value of MITRE testing is that it demonstrates a product’s ability to defend against tactics and techniques that are represented by but not limited to the attack samples. Using the evaluations, you can assess your exposure to previously unknown attacks based on its tactics and techniques, as opposed to basing your exposure level on existing one-to-one or one-to-many threat intelligence or models.

CISOs can use the results of the MITRE Engenuity ATT&CK evaluations to assess gaps in their security coverage. No single solution can detect every attack or technique that may exist, but you can learn which products can detect a given type of attack. Organizations need an integrated approach so that they can:

- Detect and block threats as early in the attack as possible
- Balance confidence in an event before a block vs. early blocking
- Find out how a solution contains a threat. Does it block specific malicious actions in real time and defuse them using micro containment and process isolation, or does the system rely on network isolation to prevent lateral movement?
- Reserve “heavy-handed” containment tools such as process termination and endpoint isolation for situations where it’s appropriate
- Streamline security operations



Fortinet is participating in the MITRE 2020 round of ATT&CK evaluations, which is focused on the Carbanak+FIN7 threat actors.

Retaining Perspective

Endpoint security is more critical than ever. Fast-moving attacks like ransomware can take minutes, if not seconds, to inflict damage. The manual responses in first-generation EDR tools aren't enough anymore. The objective of a robust cybersecurity infrastructure (including endpoint security) is to reduce the overall risk exposure. Effective security policies and continuous monitoring need to be in place to discover and predict, prevent, detect and respond, and remediate attacks.

Discovery, hardening, and prevention are the foundations of security hygiene. Doing these basics correctly can dramatically reduce risk. However, every CISO will agree that while essential, prevention is never 100%.

As a result, in addition to prevention, enterprises need to have the ability to effectively detect threats early, respond and contain threats rapidly to stop a breach, and recover to return to a known good state. Ultimately, the goal is to minimize business disruption and help ensure enterprise resiliency.

MITRE ATT&CK evaluations can help organizations evaluate endpoint security solutions based on their efficacy of detection and some degree of prevention with the recent protection test. However, when selecting endpoint security to meet their needs, in addition to looking at MITRE results, organizations should:

- Practice good security hygiene with discovery and prediction to reduce the attack surface with visibility and preemptive control until patches can be applied
- Improve accuracy to reduce the impact of false positives and alert fatigue
- Focus on responses and minimizing impact with precise and automated responses for more effective risk mitigation
- Maintain availability and system stability even in the midst of an attack, particularly for operational technology (OT) and executive systems
- Use technology to its full potential; no tool will deploy, run, and maintain itself

Conclusion

MITRE ATT&CK Evaluations look beyond simple effectiveness scoring of a security product to explore its operation. The primary benefit of this approach is a better understanding of the capabilities you are deploying. In conjunction with other assessments, you can go beyond the specific sample set used for testing. MITRE ATT&CK Evaluations are an excellent resource to help select the endpoint security product that will reduce your cybersecurity risk in the areas of greatest concern while balancing the impact on people, process, and systems. The evaluations also can help you determine a solution's fit within your overall security infrastructure and posture.

However, no security solution should exist in a vacuum. Certain techniques cannot be detected in the endpoint alone. At the same time, certain behaviors can be picked up by multiple solutions. Although overlapping coverage is better than security gaps, lack of integration can cause its own challenges as well.

¹ David Gruber, "[ESG Master Survey Results: Trends in Endpoint Security](#)," ESG, March 5, 2020.

