

A SANS Whitepaper

Measuring and Improving Cyber Defense Using the MITRE ATT&CK Framework

Written by **John Hubbard**

July 2020

Sponsored by:

ThreatQuotient

Introduction: What Is MITRE ATT&CK and Why Is It Important?

The Adversarial Tactics, Techniques & Common Knowledge (ATT&CK)[1] project by MITRE is an initiative started in 2015 with the goal of providing a "globally-accessible knowledge base of adversary tactics and techniques based on real-world observations."[2] Since its inception, ATT&CK has taken the information security industry by storm. Many vendors and information security teams the world over have moved to adopt it with blinding speed—and for good reason: It is one of the most exciting, useful and needed efforts within InfoSec in recent memory. ATT&CK provides a key capability that many organizations have struggled with in the past: a way to develop, organize and use a threat-informed defensive strategy that can be communicated in a *standardized* way across partner organizations, industries, vendors and products.

Throughout this paper, we describe what ATT&CK is, where it is going, ways of using the information, and how to leverage the now-extensive collection of information and ecosystem of tools surrounding ATT&CK to develop, bolster and assess your own defenses.

[1] Both MITRE ATT&CK® and ATT&CK® are registered trademarks of The MITRE Corporation.

[2] https://attack.mitre.org

©2020 SANS™ Institute

Before ATT&CK existed, assessing an organization's security posture could be a bit of an exercise in frustration. Sure, you could utilize threat intelligence and verify that you could detect specific attack methods, but there was always the lingering question, "What if I missed something?" Teams that made an attempt to verify attack methods could easily fall victim to a false sense of security and overconfidence in their ability to defend. After all, it is hard to know what you *don't* know.

Fortunately, the ATT&CK project came along with the lofty goal of eliminating this problem. After an incredible amount of work by the MITRE team, the ATT&CK knowledge base was created, and it alleviated much of this assessment anxiety. ATT&CK accomplished this by creating a categorized list that will eventually include *all* known attack methods and marrying it with threat intelligence on groups that use them, the software that implements them, and the mitigations and detection methods that control their use. The goal of the ATT&CK project is to be a living dataset that is continuously updated with all new information as soon as it can be verified by the industry—one that security teams can trust to be complete.

Given this goal and assuming that MITRE continues to fulfill its end of the promise, information security teams can now assess themselves against the body of knowledge that ATT&CK provides, with greater confidence that they are covering all the necessary ground and not missing any important "unknown unknowns." Therefore, through this complete and ever-updating enumeration of industry-verified attack methods and related information, the ATT&CK knowledge base delivers an incredibly useful and much more confidence-inspiring tool for assessing gaps in cyber defense.

The MITRE ATT&CK Knowledge Base

Before jumping into ATT&CK for security operations, in this section we describe the key concepts necessary to understanding the MITRE ATT&CK framework, its layout and its variants.

Matrix Layout and Contents

The most important aspect to understand about the ATT&CK knowledge base is the way it organizes the data it encompasses. The key pieces of the project are the "tactics" and "techniques" enumerated by the various matrices that group these techniques into platforms and environments where they are relevant. When exploring a matrix's techniques, keep in mind that each technique is a container of information that holds additional properties and links to multiple categories of useful information pertaining to how and where to use that technique.

Tactics and Techniques

Techniques are the key data type that the ATT&CK project centers around. A *technique* is a unique method that MITRE or the information security community has identified as being used by attackers to achieve some specific higher-level intrusion goal, or *tactic* (shown at the top of each column in Figure 1). Some examples of tactics include *Persistence, Command and Control* and *Defense Evasion*. All tactics include an organized list of *techniques*, which are specific ways of achieving that higher-level objective. Examples of techniques enumerated with the Persistence tactic include *Bootkit*, *Logon Scripts* and *New Service*.

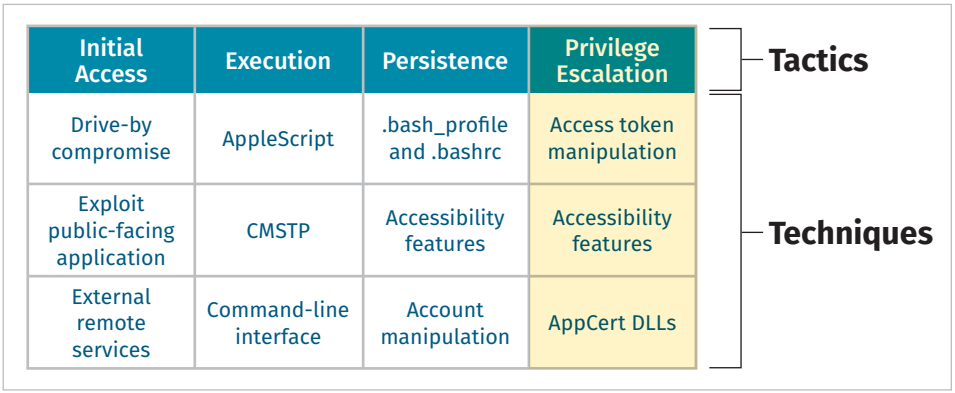


Figure 1. Tactics and Techniques Layout in ATT&CK[3]

Each technique has a set of structured data that is associated with it and includes:

- A unique four-digit **identifier** in the form of *T####*, such as T1037 for Logon Scripts
- A **tactic** or tactics with which the technique is associated. (A technique can be listed under more than one tactic.)
- **Platforms** the technique is applicable to, such as Windows or Linux
- **System** or **permission requirements** for attackers to use that technique
- **Defense strategies bypassed**, such as whitelisting
- **Data sources** that can identify the use of the technique
- **Mitigations** and **detection** methods for preventing or identifying the technique an attacker is using

Figure 2 shows an example of the metadata associated with the Scheduled Task technique.

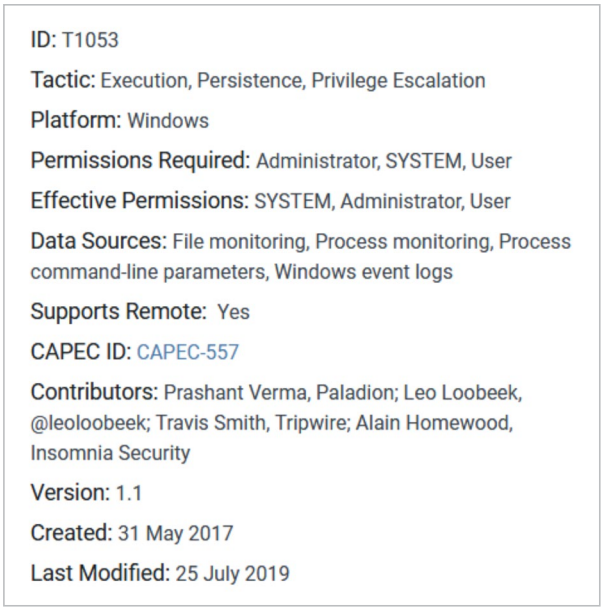


Figure 2. Metadata Associated with the Scheduled Task Technique[4]

Sub-Techniques: An Important Shift on the Horizon

All users of the ATT&CK knowledge base, both consumers and vendors, need to be aware of an important fundamental shift in the organization of ATT&CK that has recently occurred—the release and reorganization of the matrices for "sub-techniques." *Sub-techniques* take the previous structure of the matrix and introduce potential additional child techniques under each technique, where appropriate.

For example, in the sub-techniques release, the Persistence tactic has one technique—listed as Scheduled Task/Job—which now is a parent-level technique. Under this parent-level technique are sub-techniques

[3] https://attack.mitre.org

[4] https://attack.mitre.org/techniques/T1053/

with the names *At (Windows)*, *At (Linux)*, *Scheduled Task*, *Launchd* and *Cron*. This hierarchy is intuitive because they are all the same method of persistence in a sense (scheduling a task or job on a system, for example), but there are multiple unique ways of accomplishing that goal. Because each sub-technique would require a separate detection analytic implemented in a unique way and given the fact that ATT&CK exists to help teams assess their detection capability, it makes sense to split these items and list them separately under the higher-level Scheduled Task/Job parent technique. Figure 3 shows the layout of the ATT&CK matrix with the introduction of sub-techniques.

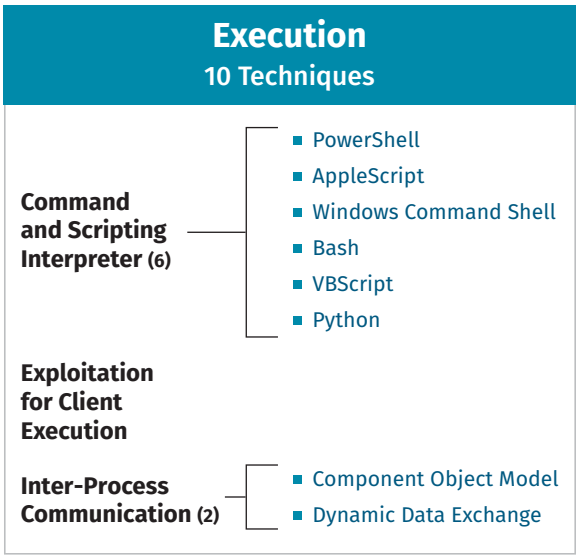


Figure 3. Execution Tactic with Multiple Techniques and Sub-Techniques

Splitting techniques this way better enumerates each possibility and enables more granular tracking within vendor tools, use cases and detection analytics. In addition to this refactoring of techniques, each tactic and technique in existence since the beginning of the project was revisited and reconsidered to question whether it still should exist in light of this change, meaning some items that previously existed may have shifted form or no longer be present. This was an additional important step for the team to take because some of the original techniques no longer matched with the spirit of what ATT&CK had evolved into but remained to preserve backward compatibility. In this new release of ATT&CK, this is no longer true.

The shift to sub-techniques is likely to be one of the most important and large-scale shifts the ATT&CK framework has undergone since its inception. As a result of this change, you will need to revisit any tools or data tracking you or your vendors provide with ATT&CK tactic and technique mappings to make them compatible with the new format. Despite the short-term pain this change may cause, in the long run it will undoubtedly be a huge boon for the ATT&CK knowledge base and its organization.

Additional Categories

Each technique in the ATT&CK knowledge base has additional key information tied to it, including the following useful items:

- **Mitigations—***Mitigations* are a list of methods that can interrupt attacker attempts to perform that specific technique. Mitigations have a numeric identification scheme in the form of *M#####*, similar to techniques.
- **Groups—***Group pages* list all techniques each known group has been reported to use in past attack campaigns, as well as software used by or attributed to them. *Groups* are sets of related attack campaigns attributed to a named actor within the threat intelligence community. Examples include the infamous APT1, Darkhotel and Turla. Each group also has a list of associated groups and potential aliases that vendors may use for the same threat actor, as well as groups that may have partial overlap with the named group based on open source intelligence reporting.

- **Software—**The software category enumerates the tools and open source software that attackers have used to conduct the behavior listed in the matrices. Like Groups, software also has an associated software property because certain tools may have partial overlap with others, leading to potential confusion depending on when and where that software's use was noted. Software also has a numeric identification scheme in the form of *S####*, similar to techniques and mitigations.
- **Data sources—***Data sources* list potential sources of information for detecting the usage of a given technique. At the time of writing, data sources are purely a list of standardized named sources shown under each technique, but the 2020 ATT&CK roadmap notes that a more formal organization of data sources is coming. This additional organization will undoubtedly improve the ability to perform certain types of gap analysis, as described later in this paper.

Flavors of MITRE ATT&CK

The ATT&CK knowledge base is not just one matrix. It is comprised of multiple matrices including ATT&CK Enterprise, ATT&CK Mobile, ATT&CK ICS and PRE-ATT&CK. While ATT&CK Enterprise gets the most attention, you should be aware of all versions of ATT&CK. Depending on your organization, one or more of these matrices may be relevant to your cyber defense strategy:

- **ATT&CK Enterprise—**ATT&CK Enterprise is the most commonly referenced matrix. It mostly contains techniques that attackers use for the *post-exploitation* stage portion of an intrusion. The information is broken into the following platforms:
  - **Operating systems—**Microsoft Windows, macOS and Linux
  - **Cloud platforms—**Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP)
  - **Cloud services—**Microsoft Office 365, Microsoft Azure Active Directory and generic SaaS platforms
- **ATT&CK Mobile—**The Mobile matrix covers techniques involving access and network-based effects that adversaries can use *without* device access. It encompasses techniques for Android and iOS.
- **ATT&CK ICS—**ATT&CK for ICS is the knowledge base specific to the tactics and techniques that attackers may use while operating within an ICS network.
- **PRE-ATT&CK—**While the other ATT&CK matrices aim to enumerate tactics and techniques used as part of the post-exploitation attack stages (except for the Initial Access tactic), PRE-ATT&CK enumerates the *pre-exploitation phase*. These are tactics, such as Technical Weakness Identification, Target Selection and Technical Information Gathering, that adversaries may perform as they prepare for and hone their attack methods. Note that the 2020 roadmap indicates that MITRE is targeting an eventual merging of the matrices into a single ATT&CK model.

Using MITRE ATT&CK to Improve Operations

These next two sections cover the most common methods of utilizing information from ATT&CK to improve organizations' security operations capabilities.

Cyber Threat Intel

One primary use of the ATT&CK knowledge base is a way to know your enemy—a way to organize and display threat intelligence related to attack group tactics, techniques and procedures (TTPs). Assuming we can predict future attackers' actions based on observations of previous TTPs, having these TTPs listed in a structured and usable way with supporting details becomes an especially useful tool for cyber defenders and threat intel teams alike. Because one of the main goals of ATT&CK is to enable threat-informed defense, threat intelligence mapping is one of the main activities in which organizations utilizing the framework should be participating.

Defenders have two primary ways to utilize ATT&CK for threat intelligence: as a consumer of the data and as a producer. Being a consumer of the data, which *every* organization should be doing, is utilizing the data that has already been created to improve defensive decision making. Taking this information and building on it as a producer of additional intelligence is the second method, and teams with the capability and capacity should look to engage in this way as well.

Being a consumer of ATT&CK information starts with narrowing the threat landscape to specific groups an organization can reasonably presume have interest in its data, assets or resources. To reduce the threat landscape, research past attacks on similar organizations and industry peers, and identify the groups attributed to those attacks. Once threat groups of interest are identified, you can leverage the Groups dataset to look at the TTPs for those groups. While some techniques may not overlap, it's highly likely that others will. By looking at TTPs common across groups that you presume will attack your organization, you can begin to form a prioritized list of detection and prevention capabilities that your security operations team must have. (Later in this paper, we discuss the ATT&CK Navigator tool, which makes this activity quick and easy.) This is a basic utilization of the data already created by the MITRE team and is highly recommended, even for the smallest of teams.

The second recommended activity is going beyond what is already known about these groups and producing your *own* threat intelligence information to add to the dataset. This activity requires that organizations give analysts the time and training to parse through available incident reporting (both closed and open source, internal and external) to extract data and map it against the ATT&CK matrices. In practice, this means reading these reports line by line; highlighting tools, techniques, tactics and group names; and extracting the information to further feed the information your team has about the organization's presumed adversaries. MITRE's new TRAM tool (discussed in the "Tools and Resources" section), which is a still-in-development effort, helps analysts to partially automate this process. With the additional information, your decision making should improve, because analysis of the attacker TTPs has been put through your organization's "filter" of context.

Data Source Gap Identification

While using the ATT&CK matrix for mapping cyber threat intelligence looks outward at the threat environment, the next common use for the matrix is inward-looking. Because each technique is listed with information about how teams identify, detect and mitigate the technique, extracting this information is an outstanding way for teams to understand their own ability to defend and prioritize plans for improvement.

The first step in this process is programmatically extracting data source information for the techniques of interest or for the matrix as a whole. There are multiple ways to accomplish this using the APIs that MITRE provides or other open source tools on GitHub. Once complete, comparing the data sources to which your team has access and the groups of users and systems that have access to those data sources can highlight important collection and visibility gaps for key attack techniques. If, for example, your collected threat intel points to the Scheduled Task technique (see Figure 4) as a primary technique used by groups attacking your organization, you will want to know whether or not you can detect it. The data sources listed in the technique—*File monitoring*, *Process monitoring*, *Process command-line parameters* and *Windows event logs*—give you that answer. If your team has none of these data sources available, or if they are available only on a subset of the systems in the environment, your next logical move should be to prioritize the correction of this problem. Whether you collect these new information sources through built-in OS logging or through augmentation with new security tools (network monitoring, network detection and response [NDR], host-based IDS/IPS, endpoint detection and response [EDR], and so on) is a separate issue to solve. But you have, at least, completed the most important step: identifying the most important missing data. Having this information in a clearly communicable fashion can help justify the additional effort and potential costs related to implementing the new data collection.

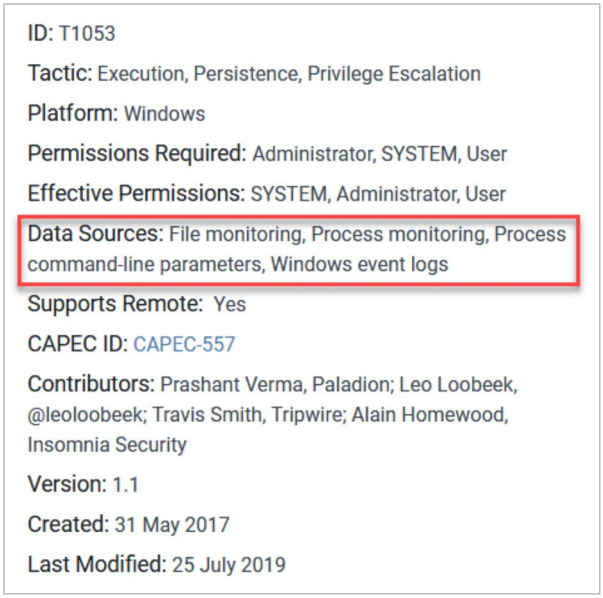


Figure 4. Data Sources for Scheduled Task Technique[5]

While collecting the data sources required is a great achievement, it's still just the first step in the process. After obtaining the data and sending it to a centralized collection system, such as a SIEM, the next step is to find an appropriate analytic tool that can be applied to highlight when an attacker is using that technique. MITRE makes this step easy for many techniques with its prewritten Cyber Analytics Repository (CAR) and even offers open source analytics options such as the BZAR project, which contains a set of Zeek/Bro scripts for detections of some ATT&CK techniques.[6,7] While not all techniques have an entry in CAR, it is a great place for teams to look for guidance when they're starting to implement new detection capabilities because many of the examples that do exist have analytic logic written in pseudo-code (such as in the PowerShell analytic example in Figure 5), as well as in EQL, Sysmon, Splunk and other product-specific languages.

[5] https://attack.mitre.org/techniques/T1053/

[6] https://car.mitre.org

[7] https://github.com/mitre-attack/bzar

Alternatively, these analytics may be built into vendor tools. Purchasing a solution with pre-created analytics that will highlight attack techniques during use provided the right data is a simple way to jump-start this second part of the process. If your team cannot find a reference analytic, the next level of activity would be to develop a new one, verify its operation, and, ideally, share that information with the community.

```
• process = search Process:Create
• powershell = filter process where
  (exe == "powershell.exe" AND parent_exe != "explorer.exe" )
• output powershell
```

Figure 5. CAR-2014-04-003—Pseudo-code to Catch PowerShell Processes Not Launched Interactively from **`explorer.exe`**[8]

Analytic Testing

After performing an analysis of the external threat environment as well as an inward-looking assessment of data collection capabilities and their apparent coverage of ATT&CK techniques, it's time to put the operation to the test. Teams can and *should* test at multiple levels of abstraction. Techniques (and soon-to-be sub-techniques) need to be individually and atomically verified. This alone, however, is not enough to get the full picture of the organization's defensive capabilities. Knowing that a single analytic is not possible or present is important, but it is the *chain* of missing items that an adversary can potentially put together that will lead to the full compromise. Therefore, we recommend testing at a higher level of abstraction based on the intelligence stored in the ATT&CK knowledge base through red and purple team exercises, as well as using ATT&CK to guide adversary emulation.

Atomic Analytics Assessment

What is worse than not having analytics to detect an attack technique? Having analytics that you think will work but that do not function properly in reality. The first and most granular step in alert testing is the atomic evaluation of your analytics. As most SOC analysts know, the constant state of flux in the operational environment paired with the consistent tuning to reduce false positives means that rules that worked for years can suddenly cease to function, so atomic testing is used as the solution to this problem.

Atomic testing of analytics often takes the form of running a singular command-line command or singular action that will trigger an alert in the SIEM, IDS or EDR, and both commercial and open source tools facilitate these tests. Regardless of the implementation type, the most important factor is the sustainable and dependable testing and *continuous retesting* of each analytic in your arsenal that corresponds to a specific technique or sub-technique, ensuring that it still functions as expected. In an ideal system, any time a change is made that could affect the functionality of an analytic, an atomic test would be initiated to verify that no unexpected and negative consequences have occurred.

[8] https://car.mitre.org/analytics/CAR-2014-04-003/

Red and Purple Teaming and Adversary Emulation

Building upon the atomic testing of analytics is stringing these tests together into something more representative of an actual attack. While a full description of red and purple team tactics could fill another paper, for the purposes of this discussion, there are two general ways in which emulating attack testing is most commonly done.

The initial testing method most teams should pursue is the purple team assessment, which typically happens in a cooperative, interactive and iterative fashion. Purple team assessments often involve penetration testers or red teamers stepping through each attack technique in their arsenal using methods that potentially vary from the atomic tests. The goal is to assess whether analytics trigger across a broadly varying number of methods, such as sending phishing emails with 10 different types of malicious attachments. While atomic testing may cover some of these methods, letting penetration testers throw the newest and sneakiest methods at the analytics may highlight unknown weaknesses and ensure that the analytics are as robust as desired.

Through design of purple team campaigns combining atomic tests from each tactic in the ATT&CK matrix, purple team assessments can highlight chains of attack techniques that, when used together, could lead to a complete compromise. These assessments, testing both network- and host-based data sources, can show how an adversary might successfully accomplish initial delivery and exploitation stages to post-exploitation command and control and exfiltration techniques. Purple team testing should be performed against both sets of analytics that the SOC expects to work, as well as additional undetectable techniques with an eye toward highlighting important gaps. The goal is to give the security operations team an initial base level of assurance that its analytics will perform against a realistic adversary. After the team performs reasonably well in this type of purple team assessment, it is time to move on to the second type—red teaming.

Red team testing is often a threat-model–driven assessment of attacker tactics and techniques, and the goal is to highlight whether an adversary could potentially reach the most important data, assets or users of the environment undetected. In contrast to purple teaming, red teaming will likely not walk through every variation of a technique, just techniques that the attacking team expects to work in the environment. Using the blue team's anticipated ATT&CK coverage and detection capabilities, red teams should select items that were previously tested in atomic or purple team form, and that the SOC has confidence will work in a real-world scenario. This is a key point for extracting the most value from a red team assessment, and it should be used as the culmination of the more purposeful and atomic testing that was previously run as additional validation of detection in an unexpected and unannounced attack scenario.

Because red team tests are often unannounced, they are the next rung up on the realism ladder, after atomic and purple team testing. These tests simulate an unanticipated attack from an actual adversary by using the items from ATT&CK that the blue team has now verified to some extent. Red team exercises are unique from the previous test types in that they test both the blue team's set of analytics *and* their ability to evaluate and respond to (what should appear as) real alerts in a timely fashion.

The top rung of the realism ladder in testing is adversary emulation. These tests aim to, as faithfully as possible, simulate an attack from a *specific* threat group, one that the organization has previously identified as a threat. Again, maintaining your own threat intelligence that aligns with ATT&CK is an enormous help when it comes to planning the methods that will be used for adversary emulation tests. Red teams can reference this information when tailoring their attacks to look like a specific threat group. MITRE CALDERA is a tool that can help to plan, facilitate and even automate portions of these types of tests.[9]

In many ways these types of tests are similar to penetration testing or red teaming, but they're scoped now to the attack techniques that your highest risk attackers are known to use. The peak of assurance as a blue team is being able to quickly and confidently respond to an adversary emulation test, because it shows your team has optimized detection and mitigation resources, as well as worked out the processes necessary to react to potential intrusions from the most dangerous groups. Figure 6 summarizes key aspects of the testing types.

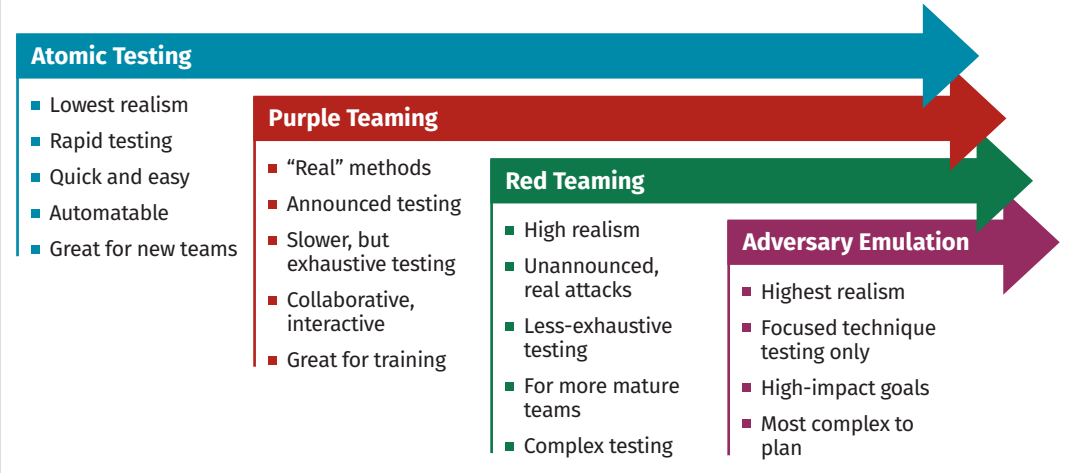


Figure 6. Spectrum of Testing Types Based on ATT&CK

For scheduling these types of tests, we recommend scheduling assessments for every quarter or every six months, depending on the team size and other responsibilities. For atomic testing, any change to the threat environment, techniques, tools or data sources creates an opportunity for adversaries to slip by undetected. Therefore, complex tests should occur frequently enough to assure the team of its capabilities while not pulling team members from responding to real events. Any tools that lower the time required for planning and executing these tests can easily pay for themselves in this regard. For atomic and more automatable testing, tying the release of potentially disruptive events to a trigger for retesting of specific analytics is a great way to ensure teams stay ahead of their adversaries. SOAR, breach and attack simulation solutions, and other security-focused automation tools can facilitate this type of continuous testing for organizations looking to minimize testing overhead. By combining both atomic, focused testing and unannounced adversary simulation, a successful blue team can create a solid sense of real-world capability to detect and respond to an attacker.

[9] https://github.com/mitre/caldera

Keys to Success

When implementing the ATT&CK knowledge base into security operations for the first time, you can shorten the time it takes to realize a return on your hard work. In this section, we cover some of the best practices that bring value as quickly as possible.

Leveraging Your Security Tools to Get Started

Because many security teams use vendor-supplied security tools and appliances as part of daily security operations, working with the built-in capabilities of these systems is a natural place to begin integrating ATT&CK models into the environment. Vendor products such as EDR suites, intrusion detection systems, SIEMs, and many more security tools now come with signature sets already classified into categories that label alerts with the corresponding ATT&CK tactics and techniques they represent. This classification makes it easy to immediately start creating metrics and labeling the activity the SOC is alerted to with ATT&CK techniques.

As alerts are sent from the organization's various security appliances, these classifications related to ATT&CK should travel with the alerts until their resolution as either a false positive or true incident. At the end of the week, month or other metrics period, with this labeling still attached, an organization can poll information from all resolved incidents and look at the ATT&CK techniques that were observed as used against it in its own environment. This is the *best* form of threat intelligence—information sourced from actual attacks within the organization that have already occurred—and having that intelligence easily and automatically created through vendor tools empowers the security team to act on it quickly and decisively. With these metrics available, security teams can display observed activity on tools such as MITRE ATT&CK Navigator and make a visual map of the most prevalent techniques in the environment based on closed incidents. This information can be fed back to the threat intelligence function and used to source additional information about who might be behind the attacks, as well as provide the all-important justification for the budget needed to bolster defenses in an area that may be poorly covered.

Utilizing Network- and Host-Based Data Sources

The ATT&CK knowledge base lists a multitude of data sources, some based primarily on network data while others are best identified through host-derived data. When attempting to provide coverage across all items, be sure to consider both options because each has its own strength and weakness.

Network-based data sources (such as NetFlow, open source Zeek, transaction data recorded from security appliances, NDR tool info and full packet capture) pulled "off the wire" from a tap or switch mirror port have the advantage of telling the truth about what is happening on the network because these data extraction points are highly unlikely to be affected by attacker activity because they are out-of-path—it's unlikely attackers even

know about their existence. If there is lateral movement between endpoints, for example, as long as the tap for the data is in a position to see it, that information *will* be reported to the SOC. The downside is that, increasingly, network data is becoming more difficult to use because of the prevalence of encryption. While some protocols may be decrypted on the fly and recorded in plaintext, many organizations cannot or do not implement these capabilities, especially with traffic from one internal source to another. This blinds defenders to some of what is happening, but those organizations that choose to decrypt traffic for analysis can take great confidence in knowing that all activity is included in the dataset. As a partial workaround to this issue, some vendors are developing tools that can infer the presence of malicious traffic without decryption, relying instead on traffic metadata, flow patterns and other fingerprints of malicious activity that are discernable through observation.

On the flip side, host-based data (such as process creation logs, antivirus, EDR tool info and host intrusion prevention suites) can provide incredible detail on what is occurring on each endpoint, assuming the reported data can be trusted. These tools record crucial security-relevant information, such as tying network traffic to the processes that created it, those processes' hashes, signature information, reputation, activity, and more. This level of detail is an incredible boon to security teams because most organizations, up until the availability of technologies like EDR, were unlikely to have this level of visibility. The downside is the incredible volume of information that can be produced and the endpoints' susceptibility to control by attackers that have compromised it. Whether or not this will have an effect on a team's ability to detect attack endpoint-based ATT&CK techniques depends on the endpoint tool's resiliency to tampering and the team's associated ability to sift through the high volume of endpoint data and highlight the most important suspicious actions across the network to SOC analysts. In addition, EDR is less likely to be a viable solution in some environments and asset types, such as unmanaged BYOD endpoints, IoT, ICS and other infrastructure devices.

Therefore, both endpoint and network-based data are complementary, providing two different views on attacker activity, and security teams typically use both types of data for the most comprehensive technique visibility. Additionally, newer efforts such as the development of the open Community ID[10] standard, already supported by many security tools, help defenders pivot to different views of the same transaction across multiple data sources and will continue to make utilization of both types of data more convenient. One further note is that some tactics are more easily detected on endpoints, such as Persistence, Privilege Escalation and Execution, while others are more network-centric, such as Command and Control, Exfiltration, and Lateral Movement. Organizations finding that their weaknesses lie chiefly in one of these tactics may be able to more efficiently improve all techniques within the tactic by focusing on the relevant (network or host) data type best aligned to that set of techniques.

[10] https://suricon.net/wp-content/uploads/2019/01/SuriCon2018_Kreibich.pdf

Using MITRE ATT&CK for Historical Measurements and Trending of Capabilities

A final key to success centers around the tracking of the security operations team's progress over time. While measuring the team's abilities at any given moment is of key importance, do not forget that the story can be told by tracking progress over time!

While a SOC is an expensive operation to run at both the human and technological levels, it should ultimately pay for itself in the additional protection and prevention it provides from disruption. Unfortunately, communicating this value is not always easy. Fortunately, doing MITRE ATT&CK-based assessments provides a way not only to objectively measure the team's abilities (e.g., "We can detect or prevent *n*% of techniques in the attack matrix known to be used by our adversaries."), but also to demonstrate and communicate the *improvement in these numbers over time*. There are multiple ways security teams can show improvement that align with the assessment techniques that were previously mentioned. Here are some of the options:

- An increase in the *number* of techniques that atomic and automated analytic testing can detect
- An increase in the *percentage* of techniques you can detect known to be used by your adversaries
- The results of purple team testing over time—How many techniques were missed vs. detected vs. prevented?
- The results of unannounced red team and adversary emulation testing—Was the team caught? How quickly? How quickly did the SOC respond?

A blue team that can demonstrate improvement in these metrics over time can easily demonstrate the value it brings to the business. This value communication inherently leads to increased funding, better tools and improved skills for those on the team, feeding a virtuous cycle in which team members, management and business all benefit.

Common Challenges

When using ATT&CK in a new program or continuing to develop capabilities, teams should be aware of a few common pitfalls. Following the advice in this section will keep teams on a smooth path to success.

Do Not Try to Do Everything at Once

One of the initial problems that teams run into when first adopting the ATT&CK framework is the overwhelming number of options of techniques to focus on. To overcome this, the recommendation is to not try to "boil the ocean," but instead focus on the most important techniques identified during threat intelligence gathering. Pick a number of techniques (perhaps the top 10) and set a short sprint goal for the team to perfect before moving on to the next set of options. Focusing on a small set of items in a short time frame and iterating through them by priority is much more likely to be successful than dumping every technique on the whole team and saying, "Have this done in a year."

Find a Balance in Assessment Detail

A common issue many teams find themselves tackling is that coverage of a given technique will be partial—partial in that the team either sees it from one subnet or asset type only (desktops vs. servers), or that data may be recorded about a technique but that data is not centralized for reporting. This problem leaves teams debating how to objectively assess their detection and prevention coverage of that technique and annotate this partial coverage.

A common solution is to develop levels of coverage, perhaps assigning numbers to various stages of data collection and centralization provided or tagging each technique with the network, user, security tool or asset populations for which it is available. This approach is rational and provides additional fidelity beyond a binary yes or no, but taken too far, it can lead the team into frustration. Advice for those who would like to more granularly track coverage is to come up with a usable system that provides meaningful data without complicating tracking so much that it becomes self-prohibitive. A solution that can highlight gaps without burdening those using the metrics with unnecessary detail strikes the right balance between pointing to the right area for improvement without drowning analysts and managers in unnecessary classifications and processes. A simple solution could rank capabilities on four levels—None, Partial, Most and Complete. Remember, don't let perfect be the enemy of the good—a 90% solution is still 90% better than a 0% solution, and any significant improvement is at least partially useful.

Consistent and Automatic Updates

An additional concern, especially with knowledge bases, such as ATT&CK, that are by definition a moving target, is staying up to date with the most recently released techniques and associated data. Teams that move to implement assessment and tracking based on ATT&CK should develop a process to be notified immediately—and ideally automatically—of additions to the matrix and ensure this data becomes available to act on as quickly as possible. In addition, teams should ask their vendors how quickly they plan to update tools and signature sets upon the release of new ATT&CK tactics, techniques and other changes to the knowledge base.

MITRE provides the ATT&CK dataset in the structured STIX 2.0 form through multiple avenues. Teams are encouraged to leverage the MITRE ATT&CK TAXII server to poll for data changes and updates.[11]

[11] https://attack.mitre.org/resources/working-with-attack/

Tools and Resources

As ATT&CK continues to grow, a wealth of tools provided by MITRE, as well as vendor-oriented and open source tools, are available to help organizations quickly jump into a threat-informed defense inspired by ATT&CK. Two of the key tools from MITRE are the ATT&CK Navigator web app and the still-beta-phase but highly promising Threat Report ATT&CK Mapping (TRAM) tool.

ATT&CK Navigator[12]

As an adopter of the ATT&CK model, one tool you absolutely must be familiar with is ATT&CK Navigator. ATT&CK Navigator is a web-based representation of the matrix that enables you to visualize techniques of interest using colors and numbers. Each technique can be assigned a color and/or score, which is applied to it in a layer. Layers may be created for techniques used by different threat groups, detection capabilities or anything else. You may think that this would be a tedious task, and it might be if MITRE had not built the information on Software and Groups directly into Navigator itself! Because the data is included, making a layer for a group such as APT is as simple as clicking on a checkbox, which fills in the information automatically.

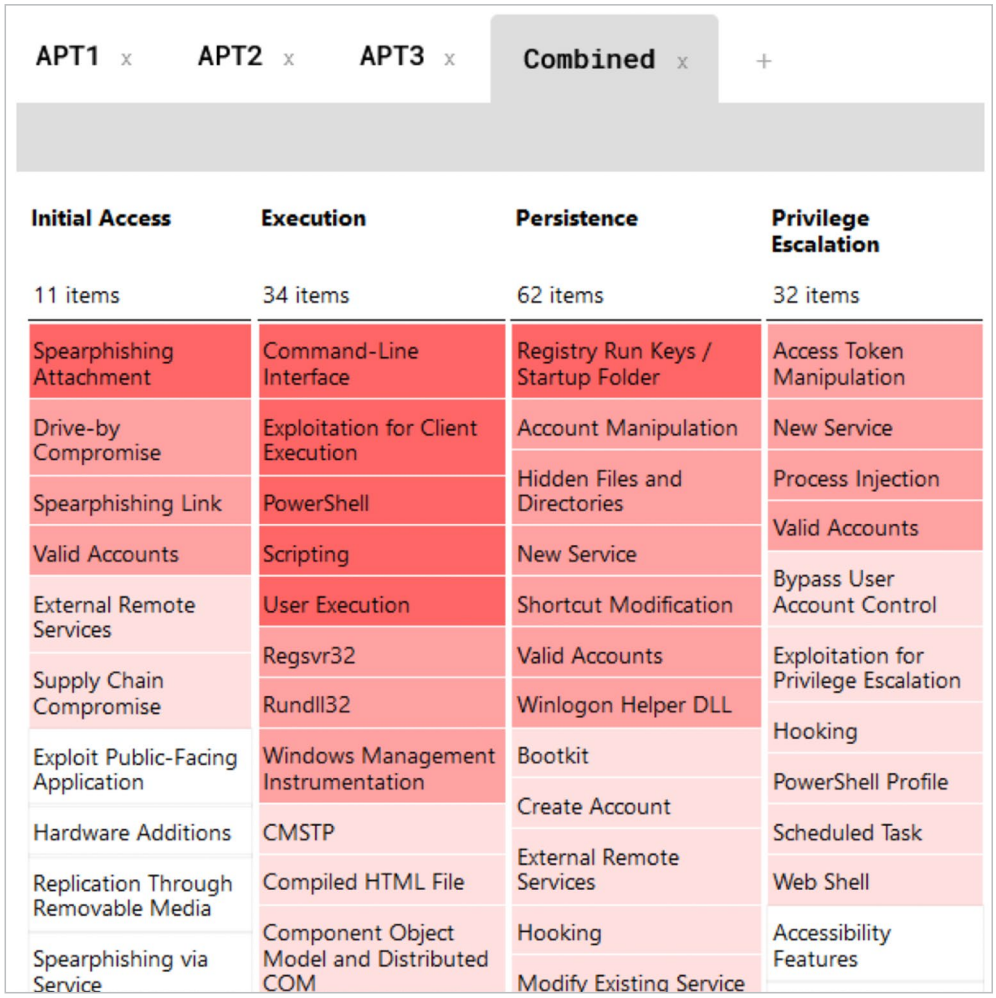


Figure 7. ATT&CK Navigator with Three Scored Layers Combined[13]

After you have built separate layers containing the data from each group or capability you would like to chart, layers can be combined using mathematical formulas to see the resulting set of items. In Figure 7, the example APT1, APT2 and APT3 threat groups were created, each having its own layer (shown in the upper left of the image). A fourth layer was then created by adding the scores per technique of each of the three layers, creating a new layer that shows the overlap of technique usage of these three groups. If your organization knows that these three groups are its biggest threat, Navigator now gives you a way to create a visual map of exactly which techniques to prioritize.

[12] https://mitre-attack.github.io/attack-navigator/

[13] https://mitre-attack.github.io/attack-navigator/enterprise/

TRAM[14]

One of the ongoing operational activities teams and threat intelligence vendors will need to engage in to keep ATT&CK Software, Groups and techniques up to date is reading through intelligence reports and extracting key data. MITRE realizes that manually parsing a PDF for all mentions of techniques, group names and software can be a tedious task, and thus it set out to solve the problem with automation and natural language processing (NLP). From this goal, TRAM was born.

The TRAM software has a web-based interface that intel analysts can use to submit a URL to be analyzed. TRAM then parses through the text, applies the NLP matching ruleset and automatically highlights sections of the report that appear to be references to MITRE techniques. Instead of having to read through an article manually to identify any of the hundreds of ATT&CK techniques, the TRAM interface gives analysts a list of potential matches from the article and lets them approve or reject them. This approach considerably speeds up the process of extracting threat intel from reports, improves analyst accuracy and consistency, and reduces fatigue. Although TRAM is not yet available for general release, keep an eye on this software—it has enormous potential to speed up extraction of intelligence from threat reports in the future.

Conclusion

Throughout this paper, we discussed the most popular and recommended methods of leveraging the ATT&CK knowledge base to improve security operations and threat intelligence capabilities. In summary, your team can utilize the pregathered information to find and set priorities for attack groups and techniques that may be used against the organization, as well as develop and plot its own internal data to supplement what is provided. This arms your defensive team with the best possible knowledge of what techniques and tactics attackers have and will likely use against the organization. After assessing the threat environment, your team can then use the built-in data source information to paint a picture of potential defensive capabilities: What attacks, in theory, should your team be able to detect and prevent? Do they line up with the attacks you expect to see from the first part of this activity? Where key information is absent, you must make a concerted effort to collect the data and implement analytics for these techniques. Tools such as ATT&CK Navigator can help make visualization of needs easy, while open source and other vendor-supplied security appliances and software can help fast-track the alignment of data needed and data you actually collect and run against analytics.

Once your team has started to create its threat-informed defense—the primary goal of ATT&CK—the final step is to test it, and test it *continuously*. Through repeated (and ideally automated or semi-automated) atomic testing, red and purple teaming, and adversary

[14] https://github.com/mitre-attack/tram

emulation, your team can rest easy knowing that its data, tools and processes have been verified to perform as expected under both controlled and unexpected attack conditions, giving your team the best possible chance at real-time attack detection and mitigation.

Finally, remember that the ATT&CK knowledge base is not only an enormous dataset, but also an ever-growing one. While it can seem overwhelming at first to teams new to using it, items can easily be prioritized and simple systems can be designed to help make rapid improvements. Even if systems are not perfect, that does not mean they can't be incredibly useful and create value for the SOC—again, don't let perfect be the enemy of the good.

Using these key ideas and strategies, we have shown how using ATT&CK can take your team in a positive, objective direction, one that is informed by threat intelligence. This allows you to not only better defend but also quantify the improvement, demonstrate those improvements with evidence and set your security operations team on the path to success for the long term.

About the Author

John Hubbard is a certified SANS instructor and the author of three courses: SEC450: Blue Team Fundamentals: Security Operations and Analysis, MGT551: Building and Leading Security Operations Centers, and SEC455: SIEM Design & Implementation. As a security operations center (SOC) consultant and speaker, John specializes in security operations, threat hunting, network security monitoring, SIEM design and defensive process optimization. His mission to improve blue teams led him to partner with SANS to help develop the next generation of defensive talent around the world.

Sponsor

SANS would like to thank this paper's sponsor:

