

FIVE STAGES TO CREATE A CLOSED-LOOP SECURITY PROCESS WITH MITRE ATT&CK

Before the [MITRE ATT&CK](#) framework was introduced, there were multiple security frameworks for strategic thought: [ISO-17799](#), it's successor [ISO-27000](#), [Cobit](#), [NIST](#), and others. However, these frameworks showed very little comprehensive work on the actual tactical needs of security operations teams. Frameworks like NIST touched on it, but as far as an actual security framework that allowed for realistic testing and provided a basis for improvement of real-time security operations process and technology, it truly did not exist until the MITRE ATT&CK framework was introduced.

According to MITRE,

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

While this is true, the framework is much more than that.

MITRE ATT&CK has two primary components: techniques, tactics, and procedures (TTPs) and adversary emulation planning (AEPs) guidelines. TTPs are the knowledge base of MITRE ATT&CK techniques that are not only front and center in the framework, but also most commonly used. They are incredibly valuable to establish a common vocabulary around which security analysts and vendors can discuss attacks and remediation techniques. Adversary emulation planning guidelines are suggested processes designed to identify and shore up weaknesses in your security posture and technologies. Though they are especially important, they are often overlooked by security practitioners versed in the "trench warfare" of day-to-day security operations. MITRE ATT&CK TTPs are useful in their own right, but they are even more effective when coupled with adversary emulation planning guidelines.

Adversary emulation planning guidelines are the most important feature when constructing a valuable, closed-loop tactical security effort tailored for your industry. **TTPs change, but the general process of bringing TTPs together into a real-world attack simulation whose efficacy can be measured is the real gem.** AEPs allow you to construct a real-world attack simulation in conjunction with TTPs that you can execute against your enterprise security infrastructure in a red team simulation. They enable you to note where and when attacks are seen and not seen, when alerts fire or are silent, and when objectives are achieved or stymied. These efforts can be measured quantitatively so that any gaps in your defenses can be quickly analyzed and filled. This brings you added visibility into your environment and can help reduce security teams load by filling gaps in your defense before attackers reach them. This document is a primer on the most rudimentary elements and processes necessary to build a simple, closed-loop improvement cycle around MITRE ATT&CK that allows for tangible, real world improvement in detection capabilities.

KEY TAKEAWAYS

- 1 » The Cybereason team has identified five essential stages to implement an efficient, iterative defense with MITRE ATT&CK around adversaries, defense posture, and security operations.
- 2 » It is crucial to understand the connection between techniques, tactics, and procedures, adversary emulation planning guidelines, and adversary groups to achieve an integrated, productive security strategy targeted to your industry.
- 3 » Following these steps brings visibility into your environment and reduces your security team load by filling gaps in your defense.

SECURITY RECOMMENDATIONS

- 1 » Create an adversary emulation plan for your red team and threat hunters to work from. Structuring a plan around the MITRE ATT&CK framework allows for a common language and repeatable process.
- 2 » It is crucial to understand the connection between techniques, tactics, and procedures, adversary emulation planning guidelines, and adversary groups to achieve an integrated, productive security strategy targeted to your industry.
- 3 » Following these steps brings visibility into your environment and reduces your security team load by filling gaps in your defense.

DANIELLE WOOD

SENIOR DIRECTOR, SECURITY SERVICES

CYBEREASON

DIVING DEEPER INTO MITRE ATT&CK

The MITRE ATT&CK website provides [tactics](#), [groups](#), and [planning tools](#) for constructing repeatable processes around security improvement. These steps are available on the MITRE ATT&CK website, though they are not organized in a step-by-step framework like NIST.

CHOOSING WHAT TO DEFEND AGAINST

The [groups](#) section of MITRE ATT&CK provides intelligence on almost 80 identified attacker groups coupled with known techniques used by each group and commonly targeted vertical markets and organizations. You should select specific groups you want to target with your emulation plans, starting with the groups that pose the most immediate threat to your organization. For example, a healthcare organization would likely start with a group like [Deep Panda](#) (MITRE ATT&CK ID G0009), as they are well-known for their intrusion into healthcare company [Anthem](#) (Table 1). When you select a specific group, the MITRE ATT&CK site lets you drill down into a list of the groups common techniques by identifier and a list of their commonly used software and malware.

Organization Type	Example Adversary Group
FINANCE	APT19
HEALTHCARE	DeepPanda
MANUFACTURING	menuPass
LEGAL	APT19
OIL AND GAS	OilRig
HIGHER EDUCATION	Turla
GOVERNMENT	BRONZEBUTLER
CRITICAL INFRASTRUCTURE	Dragonfly2.0

Table 1: Example adversary groups for particular vertical markets and organization types.

We believe it is crucial to take a hard look at your enterprise and identify key adversary groups for you to focus your defense. One could argue that if you can detect all of the items in the ATT&CK framework, you can defend against attacks by any of the adversary groups identified by MITRE ATT&CK in the framework. While this is technically true, and visibility is incredibly important, many of the individual items in the ATT&CK TTPs are not malicious and prone to a significant number of false positives, which can lead to noise fatigue and a loss of SecOps efficiency and efficacy. In maintaining a good defense, it's crucial to reduce manual effort while maintaining visibility. For example, the [Account Discovery techniques](#) (MITRE ATT&CK ID T1087) list alone is 33 items long and includes benign actions like running "net user /domain". Alerting every time this command is run in a domain has the potential to create a large number of alerts, which could decrease your security operations team efficiency. In addition, without any context around the command such as who ran it, what its parent process was, whether remote access was involved, etc., your security operations team will likely have a hard if not

impossible time identifying whether the command was run maliciously or not. In the end, your team will tune the alert out rather than rely on it, which defeats the purpose of having the alert in the first place.

A better approach is to test existing controls against a fully developed attack simulation that takes into account various techniques and validates against the existing infrastructure. This lets you see where context should be added, where logs may not be recorded and need to be, and where you should be adding new policies and technologies. Low fidelity alerts like “alert me when net user / domain is run” are only useful in context as higher fidelity alerts such as “alert me when net user / domain is run by a non-shell process or by a domain user under a shell whose parent tree doesn’t contain explorer when that user is not a member of domain admins”.

USING GROUPS TO DEVISE AEPS

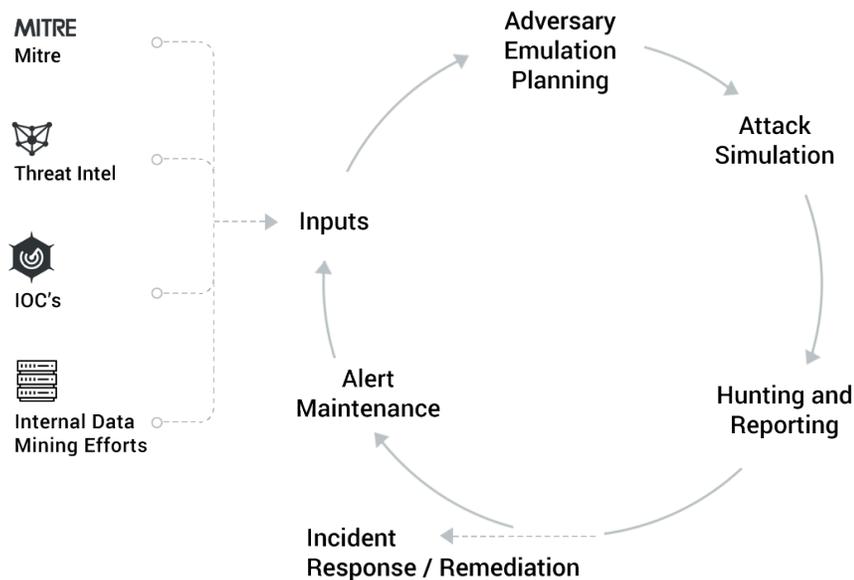
You can use this group information and the [APT 3 adversary emulation plans](#) as a guide to construct your adversary emulation plan. From this, you can easily create an attack scenario for your red team. Groups can be organized into a simple table to show the priority and status of any AEP.

Groups can be organized into a simple table to show the priority and status of any AEP, as well as the progress of the attack simulation construction.

Group	Priority	Threat	AEP Status	Attack Scenario Construction	Due	Owner
DEEP PANDA	High	High	Completed	In Progress	1/12/19	John Smith
APT3	High	High	Completed	Completed	12/28/18	John Smith
MENUPASS	High	High	Not Started	Not Started	None	None
ORANGEWORM	High	High	Not Started	Not Started	None	None

Table 2: An example set of relevant attack groups and their associated priorities and status

HOW TO MOST EASILY & EFFECTIVELY USE MITRE ATT&CK



Using the MITRE ATT&CK framework to create a closed-loop security process around adversaries, defense posture, and security operations doesn't have to be difficult. We have deconstructed the process into five stages:

01. INPUTS

02. ADVERSARY EMULATION PLANNING

03. ATTACK SIMULATION

04. HUNTING AND REPORTING

05. ALERT MAINTENANCE

These five distinct stages to the process flow should cover the majority of your security efforts.

STAGE 1: INPUTS

When it comes to creating an effective improvement cycle for hunting, alerting, and response, you should always have more inputs than solely MITRE ATT&CK. Additional, more traditional feeds can inform the cycle with data for more effective decisions on alerting and defenses.

THREAT INTEL

Outside [threat intelligence](#) is useful for two key reasons: new attack TTPs and attack validation and identification. Threat intel can be used for creating one-off attack simulations based on recent attacks like campaigns executed by [APT 39](#), or even more established attacks like [NotPetya](#) or [WannaCry](#). Alternatively, it can be used to validate information from the MITRE ATT&CK group list or pinpoint when specific malicious groups are executing previously known or new campaigns.

IOC'S

Indicators of compromise (IOCs) are probably the least useful input when [constructing overall defenses](#), but are definitely helpful when identifying intrusions by various groups. IOCs such as domain names and file hashes can be added to AEPs to identify malicious groups on the fly and beef up security from a static signature perspective. For example, you can add flags for the unique hashes associated with a specific adversary group tool to add context to static alerts.

DATA MINING

Data mining is an incredibly helpful tool for hunters and defenders when identifying new attack patterns. Unfortunately, most shops are unable to take advantage of these capabilities due to infrastructure constraints. Data mining is an incredibly complex, specialized task that takes a lot of expertise and resources. A lack of infrastructure like data lakes, indexing utilities, parallel processing facilities, etc., and inability to retain the expertise to construct it, are a large challenge for most organizations. However, if the option is available, using tools like [Splunk](#), [Elasticsearch](#), [Hadoop](#), and others make this type of deep data mining very productive and can yield dividends in both hunting and threat identification efforts.

STAGE 2: ADVERSARY EMULATION PLANNING

It's crucial to construct AEPs for each of the adversary groups most likely to attack your enterprise. While it would be optimal to create AEPs for all of the groups identified in the MITRE ATT&CK framework, the best use of your resources is to focus on the groups that target your organization or data. If your organization has the resources to manage AEPs for all the adversary groups, that's ideal, but not easily achieved. You should refresh AEPs at least annually. A guide for constructing an adversary emulation plan and an example adversary emulation plan for ATP3 [are available from MITRE ATT&CK](#).

Track the status of your AEPs in a simple table. Each individual TTP your team is addressing should have an associated planning status based on where it is in your team's execution.

PLANNING STATUS

- 01. DOCUMENTED:** The TTP has been properly documented for the adversary group.
- 02. CODED:** The TTP has been coded into an actual exploit for use by the red team.
- 03. EXECUTED:** The coded TTP was successfully executed.
- 04. SUCCESSFUL/NON-SUCCESSFUL:** The TTP execution did or did not complete its goal.
- 05. DETECTED/NOT DETECTED:** The TTP execution was successful, and it was either detected or not detected.

On the next page, is a simple and basic example of an adversary emulation plan. This table should provide you with an easy-to-track layout for keeping up to date. However, as your plan evolves, we recommend adding more context and information to this layout. As mentioned before, an emulation plan should be a series of steps, not individual techniques. Therefore, we recommend adding a timeline, a hierarchy, TTP type, notes, and more details as needed.

DEEP PANDA EXAMPLE ADVERSARY EMULATION PLAN

ID	Name	Planning Status	Time Rationale
T1015	Accessibility Features	Deep Panda has used the sticky-keys technique to bypass the RDP login screen on remote systems during intrusions. For example, use the above technique on your internet-facing servers.	Documented
T1066	Indicator Removal from Tools	Deep Panda has updated and modified its malware, resulting in different hash values that evade detection.	Documented, Coded
T1086	PowerShell	Deep Panda has used PowerShell scripts to download and execute programs in memory, without writing to disk.	Documented, Coded
T1057	Process Discovery	Deep Panda uses the Microsoft Tasklist utility to list processes running on systems.	Documented, Coded
T1117	Regsvr32	Deep Panda has used regsvr32.exe to execute a server variant of Derusbi in victim networks.	Documented
T1018	Remote System Discovery	Deep Panda has used ping to identify other machines of interest.	Documented
T1064	Scripting	Deep Panda has used PowerShell scripts to download and execute programs in memory, without writing to disk.	Documented
T1100	Web Shell	Deep Panda uses Web shells on publicly accessible Web servers to access victim networks.	Documented
T1077	Windows Admin Shares	Deep Panda uses net.exe to connect to network shares using net use commands with compromised credentials.	Documented, Coded
T1047	Windows Management Instrumentation	The Deep Panda group is known to utilize WMI for lateral movement.	Documented, Coded

Table 3: Adversary emulation plan examples and their status

STAGE 3: ATTACK SIMULATION

Utilize either internal or external red team resources to construct attack simulations that follow the emulation plans in both technology and process as closely as possible. It is vital that your red team exercises simulate actual attack resources such as external command and control systems, proper infiltration and exploitation, and data exfiltration. Failure to execute the steps of the adversary emulation plan can result in missed steps in the event of an actual attack, which can have extreme consequences.

Automated adversary emulation tools, like MITRE's [CALDERA](#), are useful to consider when building your attack simulation. These tools can emulate post-compromise adversarial behavior for your red or blue teams. This can give your red team the freedom to focus on the tasks they deem most important, or automate parts of the test in the absence of the necessary manpower.

STAGE 4: HUNTING AND REPORTING

HUNTING

It is important that you document all resources used by your red team and that you maintain constant communication. You want to ensure that the real attack executions are not lost in the noise generated by red team activity. Any successful detections should be noted and documented for evaluation at the end of the attack simulations.

In a hunt framework, the best use of AEPs and attack simulations is two-fold. First and foremost, they inform your hunt operations so your team can look for techniques in the real world on a day-to-day basis. Second, the AEPs provide a roadmap for automating the identification of attacks with a high degree of fidelity. **This is all only possible if your organization has the capability to detect the right TTPs.** If you are unable to detect the TTPs, this is an opportunity to look into new tooling or data collection methods.

At a minimum, red teams should use adversary emulation plans and TTPs for execution and should actively report on the success of their activities. They can use automated adversary emulation tools to aid their efforts. If the red team is not detected at any point, your security operations team should evaluate immediately to determine the cause. There are many reasons this could happen, from too much alerting noise to a lack of data, or simply human error. Your evaluation **should inevitably** result in tooling or process improvements. It's important to emphasize that reporting on red team aspects should be **faultless** and **rankless**. Following these principles will lead to better results from reporting and a more collaborative spirit in process improvement.

REPORTING

Each simulation needs to be quantitatively scored, even by something as simple as a scheme based on the number of TTPs used against the number of TTPs detected, including hunting results. All TTPs, detected or not, need to be categorized according to the ATT&CK category and the general detection method best suited based on your internal architecture. In addition to providing a measurable metric, you may optionally gamify the scoring to provide incentives to your team. Some individual TTP scores can be nullified at your discretion if reliable, high-fidelity detection is not possible or other detections mitigate the issue.

Reports on your red team activities need to include descriptions of the executed attack plans, the results of the attacks, and the remediation steps you should take to close any gaps. Each attack should be fully documented in the report with the specific technique used, where those activities were logged, detected, and prevented, and any methods that should be used to improve detection. Any remediation recommendations should be coupled with a priority level derived from a combination of the likelihood of exploitation and the potential damage of the exploit. In many cases, you may have existing technology with telemetry that shows the effects of an attack, but you may not have alerts that utilize that telemetry. A simple example of the minimum reporting for each technique is in the table below, which can serve as a good starting point for reporting.

ATTACK EMULATION PLAN : AEP 20190107 (DEEP PANDA)

Technique	Activity	Exploitation Results	Detections (Detected/Telemetry/Missed)	Remediation	Priority
TT1015	Sticky Keys	Replaced SetHC.exe with cmd.exe	File Write seen by SIEM; Detected as malop when replacement was executed via powershell.	None Needed	None
TT1066	Unique Binary Malware	Execution via Powershell	File Write in Telemetry from SIEM; No Detection	Add unsigned binary execution from temp to malop ruleset	Medium

Table 4: The minimum reporting your red team should provide on an attack

Your reports need to include technical details with any recommendations from your AEP executions. It is important to note that identifying individual TTPs may not be possible atomically, and as such you may need to gather additional contextual information from the execution. For example, running an unsigned binary in a large enterprise, as with the example shown in TT1066 above, would likely result in a lot of false positives if alerted on individually. However, when coupled and correlated with other details such as the execution chain, network activity, etc., the alert may become high fidelity.

STAGE 5: ALERT MAINTENANCE

Construct a process and technology improvement plan based on the results and the TTPs identified during your red team activities. Process improvement plans should be flexible enough to incorporate the results of several simulations, as changes per simulation can significantly influence technology decisions.

Improving your alerts stems from the quality of your reporting. When identifying remediation steps in your final red team report, you should include the means of detection and methods of prevention. This can lead to much more impactful alert improvements.

Some TTPs can easily be misconstrued as common actions, like a user admin creating a new user account from the command line. This can be difficult to identify, or can cause you to drown in alerts. In order to identify them properly, you should track the effects of **the actions that follow the events**. This will give you more context to understand what is going wrong and why.

In order to track alert management related to the attack execution, add remediated tracking measurements to the AEP reporting table (Table 3). This can include information on the system to be modified, the status of modifications, and their owner. It's important to note that, when looking to add more tools to cover security defense gaps, a much more thorough evaluation may be necessary. This might include additional red team testing, budgeting, PoCs, and more.

USING MITRE ATT&CK FOR SUBSTANTIAL SECURITY ADVANCEMENT

Multiple security frameworks were available before MITRE ATT&CK, but they were all missing a key component: complete explanations on the tactical needs of security operations teams. This makes MITRE ATT&CK a multi-faceted tool for teams and a crucial addition to the security space. Many organizations use MITRE ATT&CK effectively as a testing tool for their products using TTPs, but this is a limited representation of the frameworks powerful capabilities. With AEPs, MITRE ATT&CK lets you create real-world attack simulations melded with TTPs. You can execute against your infrastructure in a red team simulation to identify what attacks are identified, what alerts are sent, and what objectives are achieved. This gives you valuable visibility into your system so you are able to build a closed-loop improvement cycle for your security operations.

This white paper is meant to be a primer on implementing a closed-loop security process. However, actually implementing it can get a lot more complicated. If you're looking to use this method in your environment and have questions, [get in touch with our team](#).