exabeam

White Paper

A New Approach to Turbocharge Your Threat Detection and Response

Behavioral analytics combined with MITRE ATT&CK techniques drive SOC productivity and effectiveness

Introduction

As first responders, security analysts are painfully aware that the complexity of cyberattacks is on a steep rise. Exploits are getting more automated as attackers leverage tools to simultaneously assail related vulnerabilities in a vast range of targets. Security operations center (SOC) teams are struggling to keep up—furiously switching between tools as they attempt to investigate, contain, and respond to security alerts—all while hoping nothing slips through the cracks.

Sound familiar?

It's an unfortunate reality that breaches continue to prevail. As a result, it's become prudent to consider augmenting legacy approaches to threat detection. SOCs not only need the appropriate tools; they also need a standard way to communicate and collaborate about the attacks they are detecting, investigating, and responding to. This white paper describes how the MITRE ATT&CK framework enables this objective. It provides a common taxonomy for understanding the various tactics, techniques, and procedures (TTPs) adversaries employ and how to use them for more effective threat detection efforts. The paper also describes enhanced results when adding behavioral analytics to threat detection with MITRE ATT&CK by using the capabilities of a modern SIEM.

What are TTPs, why they are indispensable for modern detection?

TTPs provide a description of activities used by an adversary. They describe the "what and how" of an attack. Using TTPs enables security analysts to look for attack patterns instead of the artifacts left after as a result of an attack. Attack artifacts are often referred to as "indicators of compromise" (IOCs); they are merely pieces of evidence observed on a network or on operating systems that indicate some level of intrusion has occurred.

The figure below shows the varying levels of effort needed to detect different types of threat indicators. In the diagram, all levels below "Tools" represent IoCs. While they are the easiest to spot, using IoCs for threat detection is simply not enough. For example, IoCs are inherently reactive, so they are usually valid just for a short period as hackers change their attack infrastructure to avoid detection. IoCs also lack context about what a hacker was trying to achieve. Reasons like these make IoCs prone to high rates of false positives when solely relied on for threat detection.

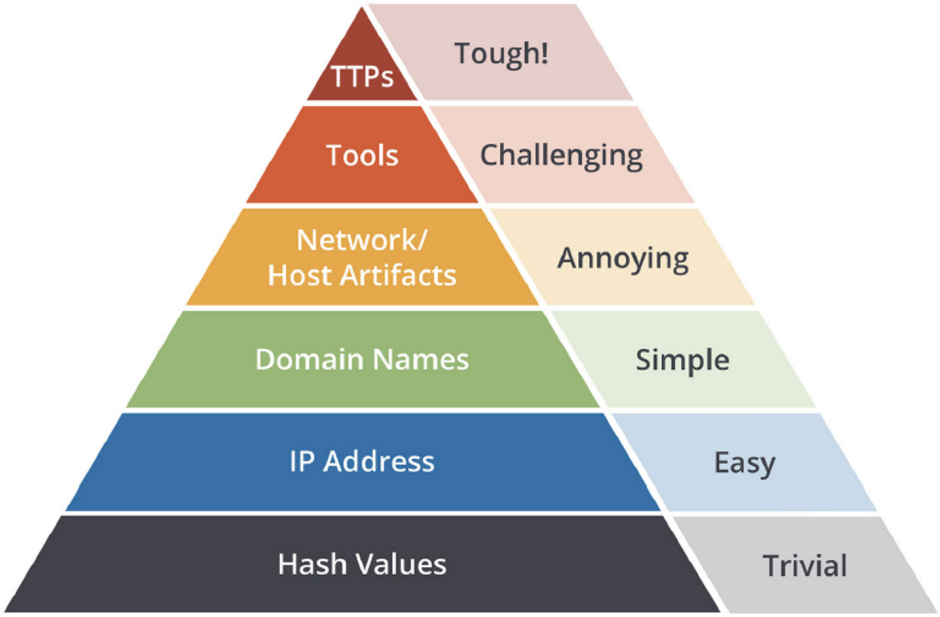

Figure 1 – Pyramid of Pain, showing the level of effort required to detect different types of threat indicators
Source: David J Bianco's **blog**

IoCs are also ineffective for threat hunting because there are so many of them. An attack's forensics typically show hundreds or thousands of IoCs, and sometimes many more. As you go up the pyramid, the threat indicators become more valuable, but also more difficult to detect. This paper will describe how to use TTPs instead of IoCs to greatly improve detection and threat hunting efficacy, and how to leverage behavioral analytics to further compound the effectiveness of this approach. It revolves around hunting threats and attack patterns with behavioral analysis guided by the MITRE ATT&CK Framework.

SOC teams require a common framework that aligns TTPs with their security tools and provides a standard language to use when hunting for threats and discussing attack patterns.

A security analyst's best friend: Introducing MITRE ATT&CK

MITRE ATT&CK provides the knowledge captured from millions of attacks on enterprise networks and systems and maps these tactics, techniques, and procedures to a common framework. It provides a common taxonomy and knowledge base that the security community can use in communication, as well as in their efforts for detection, investigation, and response. This functional junction also helps security vendors design threat hunting tools and detection methods capable of identifying specific tactics and techniques within the framework.

MITRE ATT&CK organizes TTPs into a simple matrix. Tactics are listed across the top, with individual techniques that achieve that tactic listed below in each corresponding column. Tactics are presented from left to right in the general order of an attack sequence.

Tactics

Techniques

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
Drive-by Compromise	AppleScript	.bash_profile and .bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation	Account Discovery	AppleScript	Audio Capture	Commonly Used Port	Automated Exfiltration	Data Destruction
Exploit Public-Facing Application	CMSTP	Accessibility Features	Accessibility Features	Binary Padding	Bash History	Application Window Discovery	Application Deployment Software	Automated Collection	Communication Through Removable Media	Data Compressed	Data Encrypted for Impact
External Remote Services	Command-Line Interface	Account Manipulation	AppCert DLLs	BITS Jobs	Brute Force	Browser Bookmark Discovery	Distributed Component Object Model	Clipboard Data	Connection Proxy	Data Encrypted	Defacement
Hardware Additions	Compiled HTML File	AppCert DLLs	AppInit DLLs	Bypass User Account Control	Credential Dumping	Domain Trust Discovery	Exploitation of Remote Services	Data from Information Repositories	Custom Command and Control Protocol	Data Transfer Size Limits	Disk Content Wipe
Replication Through Removable Media	Control Panel Items	AppInit DLLs	Application Shimming	Clear Command History	Credentials in Files	File and Directory Discovery	Logon Scripts	Data from Local System	Custom Cryptographic Protocol	Exfiltration Over Alternative Protocol	Disk Structure Wipe
Spearphishing Attachment	Dynamic Data Exchange	Application Shimming	Bypass User Account Control	CMSTP	Credentials in Registry	Network Service Scanning	Pass the Hash	Data from Network Shared Drive	Data Encoding	Exfiltration Over Command and Control Channel	Endpoint Denial of Service
Spearphishing Link	Execution through API	Authentication Package	DLL Search Order Hijacking	Code Signing	Exploitation for Credential Access	Network Share Discovery	Pass the Ticket	Data from Removable Media	Data Obfuscation	Exfiltration Over Other Network Medium	Firmware Corruption
Spearphishing via Service	Execution through Module Load	BITS Jobs	Dylib Hijacking	Compile After Delivery	Forced Authentication	Network Sniffing	Remote Desktop Protocol	Data Staged	Domain Fronting	Exfiltration Over Physical Medium	Inhibit System Recovery

Figure 2 – The MITRE ATT&CK framework organizes hacking tactics and techniques into a matrix. Tactics span the top row, and the techniques that can achieve that tactic are listed below in each column.

Into action: How to use the MITRE ATT&CK framework?

Security analysts often wish to associate security events and abnormal activities with the relevant threat actors, intrusion sets and campaigns. This could be for threat hunting, incident investigation or general knowledge building. To that end, MITRE ATT&CK is the definitive knowledge base that maps different tactics and techniques used in attacks to the **threat groups**[1] and tools associated with them. Analysts can also find the tactics used by these groups should they want to search for potential activity by specific threat actors.

For example, by looking up APT3[2] in the ATT&CK database, one learns that the attack group is associated with the Chinese military and primarily targets U.S. government organizations and political organizations in Hong Kong. The database also notes this group has used **LaZagne**, **PlugX**, **SHOTPUT**, and **RemoteCMD** software in their past attacks.

In addition to being an effective learning tool and a common framework for analysts to communicate about attacks, MITRE ATT&CK is also useful for guiding the detection, investigation, and threat hunting efforts of analysts.

Partnering for success: Detecting abnormal TTPs with behavioral analytics

And now for the tricky part: while TTPs are a good thing for analysts because they illustrate how an attack happens, the digital evidence alone revealing TTPs cannot tell you if that activity is related to specific malicious action–or should be attributed to normal workflow performed by enterprise users.

For example, analysts may be familiar with how attackers could maliciously leverage processes for account creation, screensaver activity, remote desktop access, etc., but these processes are also part of the normal everyday activities in enterprise IT.

To distinguish the bad from the good, MITRE-related tools used by SOC analysts must be smart enough to detect and alert only when the behavior is malicious or has bad intent. The inability to make this distinction means analysts will end up with a lot of false positives!

Behavioral analytics monitors all user and asset behavior and applies machine learning to baseline what behavior is normal. This application of machine learning is what enables it to identify deviations from normal activity, and thus accurate detection of malicious TTPs. User and entity behavior analytics (UEBA) can leverage TTPs defined in the MITRE ATT&CK framework to tag anomalous events to make it easier for security analysts to hunt for threats.

For example, consider an attacker logging onto a service designed to accept remote connections, such as telnet, SSH, or VNC. An adversary typically uses this technique to access the network and then move laterally within it to reach high-value assets. This approach is a TTP, defined as **Remote Services**[3] in the MITRE ATT&CK framework. With a legacy SOC tool, TTP detection would be created using a static correlation rule. Once configured all occurrences of remote connections would be flagged by this rule because static correlation rules have no understanding of the normal operating circumstances that may involve remote connections. As a result this rule would create a large number of false positives, and consequently likely cause analysts to ignore alerts generated by the rule. However, combining MITRE ATT&CK detection with user and behavioral analytics can help analysts hone in on TTPs which occur in their environment that are genuinely abnormal, and thus more likely to represent real threats.

[1] **https://attack.mitre.org/groups**
[2] **https://attack.mitre.org/groups/G0022/**
[3] **https://attack.mitre.org/techniques/T1053/**

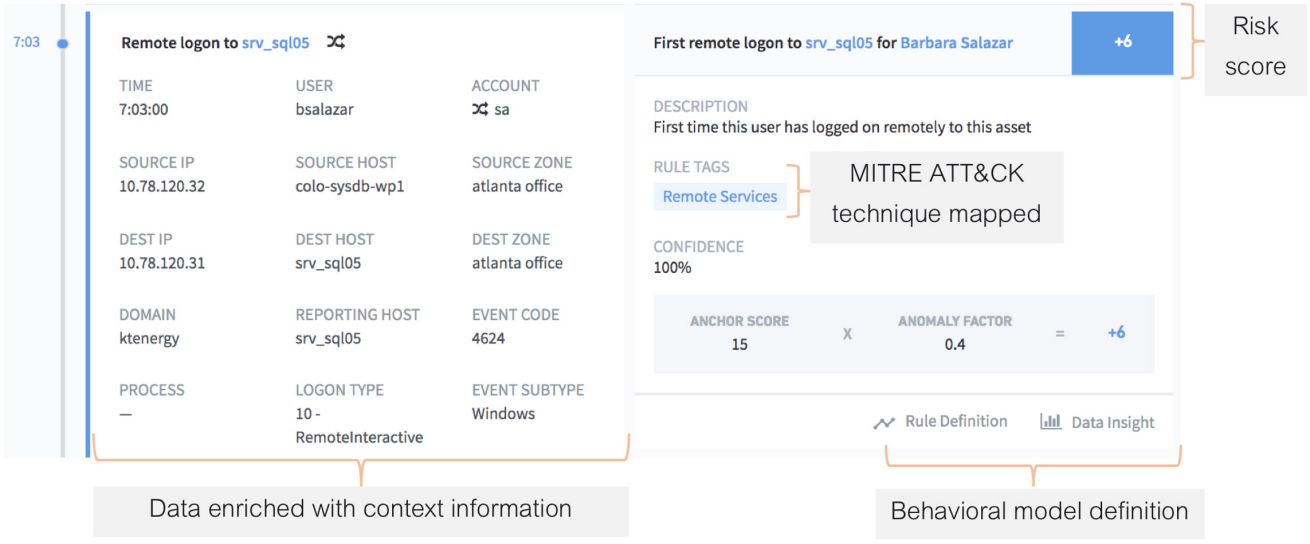

Figure 3 – This screenshot shows the first-time logon to a server by the user as anomalous. This is part of the user timeline along with the MITRE technique tagged "Remote Services", the risk score, and contextual information surrounding the event.

The key to this approach is that behavioral analysis learns the behavior of users and assets. The ML-based models establish a baseline of an organization's normal behavior, enabling algorithms to easily detect deviation from the baseline. For example, assume the model detects that it's the first time a user has remotely logged onto a particular SQL server. Since this has never occurred before, it inherently carries more risk than an activity that happens regularly. Figure 3 shows an abnormal remote login within a UEBA tool. It displays the event risk score, MITRE technique used and the context dynamically built around the user and assets in question. The SOC analyst is provided with complete, contextually enriched information about this behavior with the risk reasons and evidence for this alert. No more guessing the meaning of obscure alerts and trace data. Imagine the potential time savings and the efficiency gains!

Behavioral analysis is used to track deviation from normal baseline behavior and to detect threats in real time.

UEBA does not flood SOC dashboards with individual alerts, since the occurrence of a TTP by itself does not provide enough proof that it's a threat. This is where user timelines come into play. The next section describes how machine-built timelines within UEBA solutions can automatically aggregate all relevant events and alerts along with contextual information to give a complete picture of the attack.

An alert of a TTP by itself does not provide proof of an attack. Analysts need to understand the context in which it occurred—ideally in an incident timeline—to get a complete picture of the threat.

A better lens: Investigating an attack with behavioral analytics and the MITRE framework

The incident timeline provided by a modern behavioral analytics tool is the operational point of integration with the MITRE framework. Instead of becoming distracted by potentially irrelevant TTPs, the timeline provides useful one-click access to all of the context surrounding a potentially damaging incident. To understand how a timeline helps with threat investigation, let's consider its functional attributes.

The machine-built incident timeline automatically stitches together all behavior by users and assets and contextually presents the data with highlighted risk reasons and risk scores. UEBA enriches the data with context from various sources such as AD, LDAP, host-to-IP mapping and dynamic peer grouping.

Visual presentation of these data makes it easier to see the full attack chain. Legacy tools tend to swamp SOC analysts with an alert for every technique used. This creates an impractical and ineffective situation where analysts have to manually assemble these disparate alerts to "hopefully" make sense of a situation – whether or not malicious activity actually exists. For analysts in a large organization getting thousands of alerts daily, be assured that hope plays a major role in successful threat investigation! A UEBA's machine-built timeline exposes to investigators all the evidence they need with pinpoint accuracy. It reduces both false positives and mean time to detect/respond to real threats.

To enable collaboration, many UEBA tools have labeled TTPs that are identified in the MITRE ATT&CK framework. This helps to show where specific events map to the overall framework (i.e. tactics, and the kill chain). MITRE labels include a description of the attack and a link to the framework. Having direct hyperlinks to the MITRE ATT&CK knowledge base for any abnormal TTPs discovered helps analysts understand the implications of the techniques they detect and provides them with a resource for additional learning.

Threat investigation compared

Legacy indicators of compromise vs. analytics & TTPs

	IoC-based threat investigation	Advanced analytics & TTP-based threat investigation
Detect	Alert-based. Prone to false positive rates due to the high number of IoCs and their short-lived window of effectiveness.	Detects based on abnormal behaviors. Zeroes in on abnormal occurrences of TTPs.
Search	Query language-based. Analysts usually search for known IoCs. Results are raw logs.	Intuitive, point-and-click UI interface. Analysts can search for TTPs or IoCs. Results are machine-built timelines.
Pivot	Query language-based. Analysts usually search for known IoCs. Results are raw logs.	No need to understand the underlying attack change search parameters. Can search by MITRE ATT&CK tags to quickly zero in on abnormal occurrences of specific techniques or tactics.
Prioritize	Analysts must manually determine which alerts are worthy of further investigation. High numbers of alerts and low contextual information often result in wasting investigation cycles on false positives.	Automatically identifies abnormal TTPs and sorts them by risk score to prioritize the highest risk items for analyst review.
Investigate	Investigation of discovered threats requires manual assembly of evidence, which takes hours, days or weeks and may be prone to error.	Automatically stitches together all relevant user and entity activity into incident timelines that pinpoint anomalous behavior and follow lateral movement.

More precise threat hunting using behavioral analytics with TTPs

Implementing UEBA solutions for threat hunting alongside the MITRE ATT&CK framework can yield powerful results. Behavioral analytics allows analysts to zero in on abnormal TTPs, as opposed to all of the TTPs occurring in an environment. Threat hunting using these search results provides the alerts, events, and incidents the investigator is looking for. In addition, it also provides the complete machine-built timeline complete with all risk reasons, risk scores where all events are tied together and rule tags mapped to MITRE techniques. This is a vastly more efficient way to threat hunt when compared to the traditional IoC-based approach taken by legacy tools and traditional SIEMs (see figure 4 below).

The integration dramatically reduces mean-time-to-response (MTTR) as analysts are presented with key evidence–machine-built timelines for every user and device in your enterprise.

Greatly reduce mean-time-to-response (MTTR) with behavioral analytics as the foundation layer for threat detection, machine-built timelines for rapid investigation, and natively-integrated threat hunting to easily hunt for abnormal TTPs.

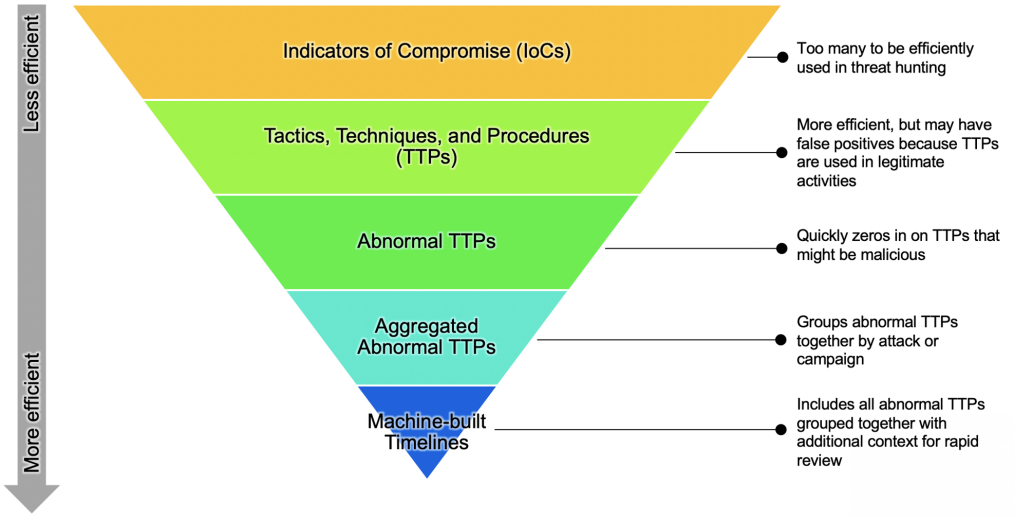

Figure 4 – Shows the efficiency of various threat hunting approaches, ranked from least to most efficient.

Conclusion

When detecting and resolving security threats, time is your most precious resource. Add to this, the rising sophistication of attack sequences is ill-met by the legacy tools used in today's SOC. It has become clear that the old way of manually addressing a deluge of alerts and manually attempting to stitch together an event timeline is slow and impractical ... if it works at all. Legacy tools and IoC-based approaches make it difficult for an analyst to quickly and fully understand the scope of an incident when there are multiple users, processes, devices, and network connections involved. Resolving the attack sequences requires SOC analysts to see the complete picture. **By using behavioral analytics to identify anomalous activity, and automatically mapping it to the techniques identified in the MITRE ATT&CK framework, responders can quickly detect, trace and respond to an attacker before they cause significant damage.**

About Exabeam

Exabeam is a global cybersecurity leader that adds intelligence to every IT and security stack. We are reinventing the way security teams use analytics and automation to solve threat detection, investigation, and response (TDIR), from common security threats to the most critical that are difficult to identify. The Exabeam Security Operations Platform is a comprehensive cloud-delivered solution that leverages machine learning and automation using a prescriptive, outcomes-based approach to TDIR. It is designed and built to help security teams detect external threats, compromised users and malicious adversaries, minimize false positives, and make security success the norm. For more information, visit www.exabeam.com.

exabeam

To learn more about how Exabeam can help you visit exabeam.com today.