A HOUSE DIVIDED:

RE-MAPPING MITRE ATT&CK TO NETWORK DETECTION AND RESPONSE

This paper underscores the significance of how machine learning is helping to address the challenge of threat detection in a world of encrypted east-west traffic in networks and datacenters.

Summary

This white paper demystifies the rise of network detection and response solutions from legacy network IDS and maps the endpoint-centric ATT&CK model to NDR.

Author Information

Alissa Valentina Knight
Partner
Knight Ink
1980 Festival Plaza Drive
Suite 300
Las Vegas, NV 89135
ak@knightinkmedia.com

Publication Information

This white paper is sponsored by Lastline, Inc.

Initial Date of Publication:
December 2020
Revision: 0.1

A HOUSE DIVIDED

RE-MAPPING MITRE ATT&CK TO
NETWORK DETECTION AND RESPONSE

TABLE OF CONTENTS

TABLE OF CONTENTS

HISTORY

ON MITRE ATT&CK

A HISTORY ON MITRE ATT&CK

Before MITRE ATT&CK, organizations sought out a framework to build their cybersecurity programs around according to best practice. Best practices defined by standards organizations, such as ISO 27001 published by the International Standards Organization and the NIST CSF framework developed by the National Institute of Standards and Technology were the most prevalent.

Frameworks, such as NIST CSF and ISO 27001 help guide decision makers in what gaps exist in their technical and administrative cybersecurity controls when making investment decisions in their cybersecurity programs.

However, no other standard or security control framework in history has seen such widespread adoption by organizations than the MITRE ATT&CK framework. Buyers use ATT&CK to determine if there are any gaps in their security controls while vendors align their products to being able to detect specific ATT&CK techniques making it easier for buyers to determine which products fill their gaps.

THE MITRE CORPORATION

The MITRE Corporation is a federally funded organization dating back to its roots in 1958 in defense and intelligence based in both Bedford, Massachusetts and McLean, Virginia. MITRE's mission is to make a more secure world by solving today's contemporary challenges across a wide array of mission areas of systems engineering, advanced technologies, acquisition effectiveness, and cybersecurity.

MITRE performs these functions through federally funded research development centers (FFRDCs), individual non-profits funded by individual government sponsors that are created to solve specific national, global, military, and civilian complex challenges.

ATT&CK HISTORY

Before I can decompose ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) into its ancillary parts, I need to first explain the Fort Meade Experiment and introduce you to Blake Strom. Strom was an incident responder with the Department of Defense (DoD) out of college. Strom gained notoriety in investigating high profile incidents involving classified and unclassified networks across different DoD components.

About the same time Strom went to work for MITRE, MITRE Corporation had begun operating an internal project called The Fort Meade eXperiment (FMX). FMX was built on the internal corporate network of MITRE itself but instrumented with network probes and sensors deployed across the network and system endpoints where actual MITRE users were performing their job functions. FMX invited red teams (penetration testers) to attempt to reach specific objectives within this production enclave of MITRE in a sort-of capture the flag (CTF) experiment but with the ability to record the tactics, techniques, and procedures (TTPs) of those red teamers with significant fidelity. The purpose of this exercise was to gamify adversary emulation in order to more quickly detect advanced persistent threats (APTs) using real-world threat scenarios, tools, and TTPs.

Unlike historical approaches based on theory, the behaviors observed in the FMX environment were based on real-world TTPs categorized and used by both the red teamers and blue teamers as they timed their ability to more quickly achieve their actions on objectives. In what would eventually become known as ATT&CK, these TTPs were cataloged into the first model published in September of 2013 based on the TTPs affecting Microsoft Windows.

Today, the ATT&CK Enterprise model now contains a beta version of sub-techniques and as of October 2019, now comprises 314 techniques as of this writing across 12 tactics in Windows, Linux, Mac, and Cloud workloads. ATT&CK has since been expanded and now comprises three new models with the introduction of PRE-ATT&CK, Mobile, and ICS (industrial control system).

DEMYSTIFYING TACTICS, TECHNIQUES, AND PROCEDURES

There is quite a lot of confusion between the concepts of tactics, techniques, and procedures and their idiosyncratic differences. Simply put, tactics (goals) are a series of strategies or tasks employed to achieve a specific end while technique is the unique way or methodology in which those specific tactics are performed. Perhaps more difficult to distinguish is the introduction of the concept of procedure, which is the specific order or manner in which those actions are performed.

The ATT&CK matrices (by technology domain) are architected to categorize these specific TTPs horizontally and vertically in a tabular framework with tactics being ordered horizontally across columns and techniques ordered vertically in rows as illustrated in Figure 1. This can be represented simply as a set of actions (techniques) used to achieve a specific goal (tactics).

THE RISE OF PRE-ATT&CK

Realizing that the Enterprise matrix left out the tactics and techniques used by adversaries before a foothold is achieved on the network, MITRE developed the PRE-ATT&CK matrix to address these steps prior to the Initial Access tactic in the Enterprise matrix. PRE-ATT&CK effectively covers the reconnaissance and weaponization stages of Lockheed's Cyber Kill Chain Model,

In order to understand the distinction between PRE-ATT&CK, MITRE's adaptation of the Lockheed KCM, and the KCM itself, the below diagram illustrates the relationship between all four models.

Figure 1: ATT&CK Enterprise Model

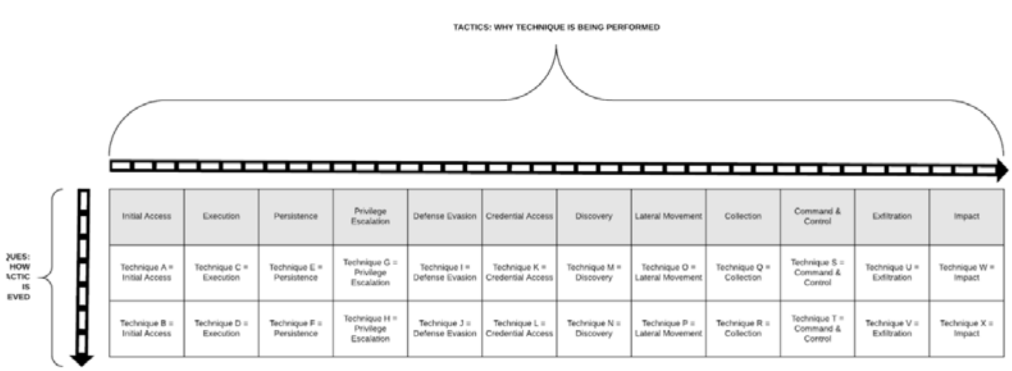

Source: Knight Ink

NETWORK THREAT DETECTION'S

IDENTITY CRISIS

NETWORK THREAT DETECTION’S IDENTITY CRISIS

There was once a world in which north-south and east-west traffic was rarely encrypted and passed over clear-text protocols, such as Telnet, File Transfer Protocol (FTP), and Hypertext Transport Protocol (HTTP). Use of these clear text protocols would later become taboo over more secure protocols that employed encryption. These protocols such as Secure Shell (SSH), secure FTP (SFTP)/secure copy (SCP), and Hypertext Transport Protocol Secure (HTTPS) employ encryption, such as transport layer security (TLS). The use of clear text protocols before they were replaced by those that employed encryption made analyzing the traffic at the network layer trivial as the the headers and the packet payloads could be searched for specific keywords or patterns.

In order to identify indicators of compromise (IoCs) on the network layer, tools would be developed which would eventually be released as open source and freely available for download, such as Shadow, Snort IDS, Snort-Inline, and Suricata. Companies also commercialized and productized some of these open-source tools, such as Sourcefire and SecurityOnion. Eventually, companies like Internet Security Systems (ISS), Top Layer, and Intruvert brought their network IDSs to market, which were later supplanted altogether by Unified Threat Management (UTM) systems at the edge.

The early ancestors of today’s NDR solutions evolved to machine learning models from pattern matching/signature-based detection systems, such as Shadow and Snort over the last two decades, which we’ll quickly digress to.

THE RISE OF NETWORK IDS

Shadow, or Secondary Heuristic Analysis for Defensive Online Warfare (SHADOW) developed at the Naval Surface Warfare Center (NSWC) was a combination of Perl scripts that the administrator would process PCAP dump files through generated by tcpdump on remote sensor stations. These dump files would be transferred to the analysis station where the Perl scripts that shipped with Shadow were waiting to process the files looking for specific indicators of compromise.

PCAP files are generated by a packet sniffer (tcpdump wireshark, etc) that listens on a promiscuous mode network interface card (NIC), passively capturing data packets off a wired or wireless network and storing them to a file on disk for later analysis.

Shadow would later become unmaintained abandonware giving rise to a new project and global community of developers and signature creators for what would later be called Snort, developed by Martin Roesche.

Snort's success was not just in its ability to perform pattern matching against packet payloads and headers, but also using what were called Snort Preprocessors -- extensible plugins capable of performing analysis across multiple fragmented packets in a TCP stream.

An example Snort rule in the Emerging Threat ruleset is illustrated in Figure 3 on the next page that would fire on likely bot activity when the keyword /NICK followed by USA is found in a string of a packet payload. This command is commonly found on Internet Relay Chat (IRC) for setting the nickname of a user that's typically indicative of bot activity.

THE RISE OF NETWORK IDS

Figure 3: Example Snort/Suricata rule looking for likely bot activity

```
drop tcp $HOME_NET any -> $EXTERNAL_NET any (msg:"ET
TROJAN Likely Bot
Nick in IRC (USA +..)"; flow:established,to_server;
flowbits:isset,is_proto_irc; content:"NICK "; pcre:"/NICK
.*USA.*[0-9]{3,}/i"; classtype:trojan-activity;
reference:url,doc.emergingthreats.net/2008124;
reference:url,www.emergingthreats.net/cgi-
bin/cvsweb.cgi/sigs/VIRUS/TROJAN_IRC_Bots;
sid:2008124; rev:2;)
```

Source: Open Infosec Foundation

However, as with everything in life, things change and evolve over time, adapting to changes in the environment. This change in network IDS technology was largely propelled by increased adoption of encrypted protocols in both east-west and north-south traffic of an enterprise network rendering pattern-matching detection systems largely ineffective since they couldn't apply those rules against encrypted traffic. According to a report by Gigamon, 81% of enterprise web traffic is encrypted and according to Gartner, more than 50% of new malware campaigns use various forms of encryption and obfuscation.

An example clear text packet where encryption is not being used is illustrated in Figure 4. In Figure 5 I've provided a sample HTTPS packet where TLS is being used for encryption to compare and contrast the differences between the two packets at both the header and payload layers of the datagrams. As you can see from both figures, applying signatures against a packet in Figure 5 on the next page would not produce any corresponding alerts due to the fact that the payload is encrypted.

The rise in encrypted east-west and north-south traffic causing network blind spots for cybersecurity teams wasn't the only thing pushing legacy network IDS solutions out of the internal network.

The significant number of false positives these solutions were generating as a result of the signatures and lack of context awareness were the biggest motivators for the market to seek out an alternative approach to network threat detection. False positives created systemic event fatigue causing security analysts to ignore real events as false positives and created longer mean time to detection (MTD) and mean time to response (MTR) -- rendering network IDS solutions ineffective.

Figure 4: Example HTTP packet in Wireshark

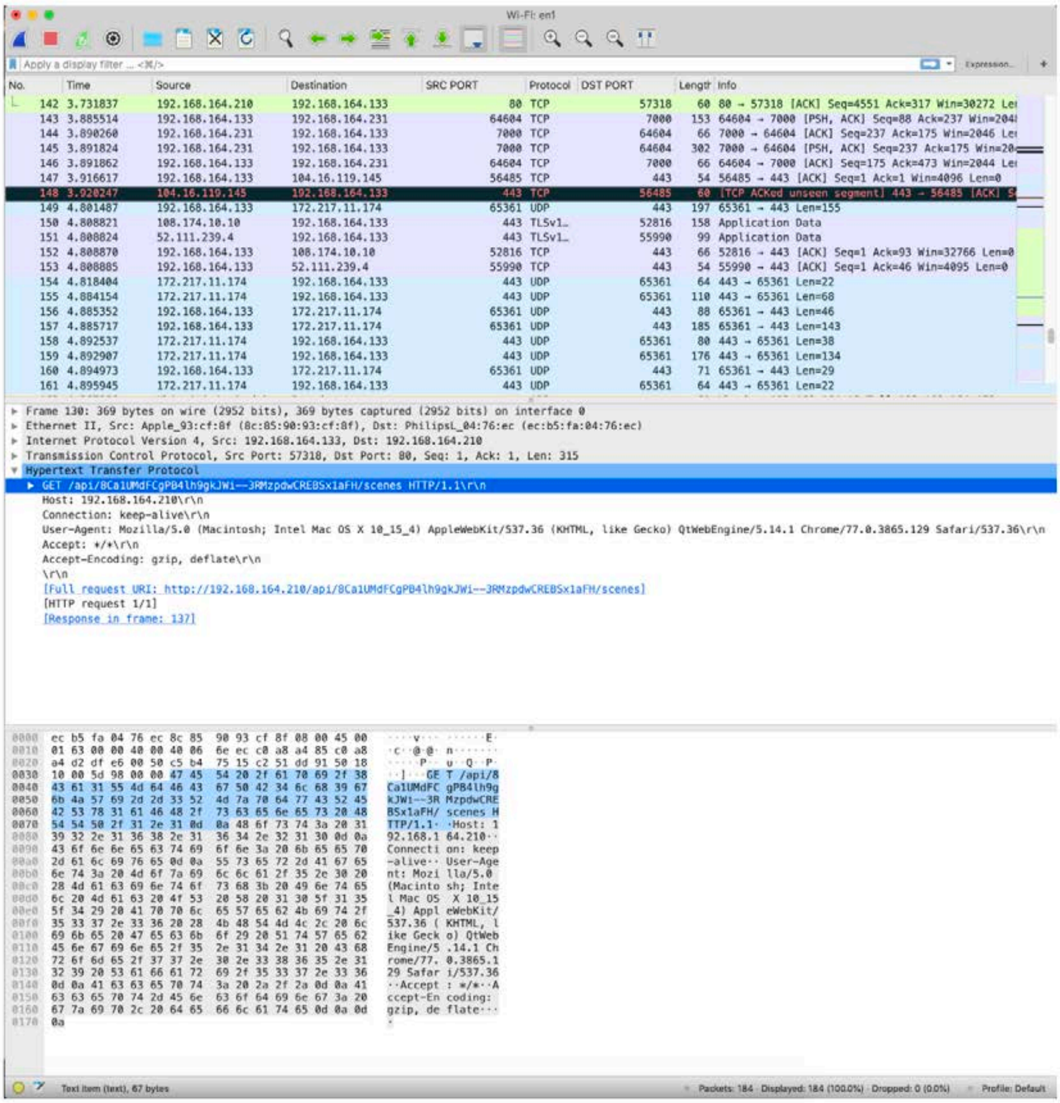

Source: Knight Ink

Figure 5: Example TLS encrypted packet

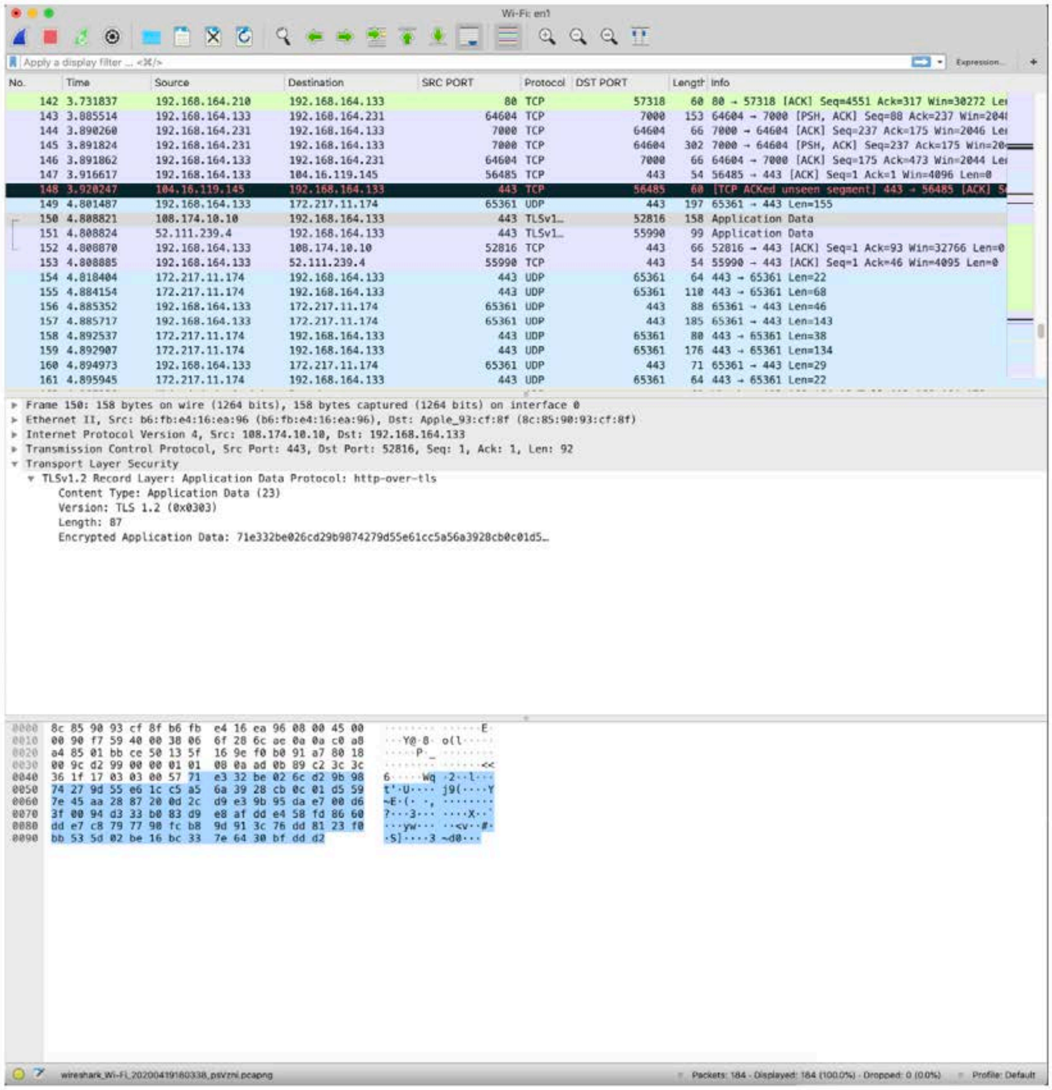

Source: Knight Ink

SANDBOXING

The fact of the matter is, the days of signature-based detection alerting on “known knowns” as the only layer to threat detection on the network have been all but forgotten. Automated analysis of malware in order to detect zero-day threats that have never been seen is what should be the defining factor in evaluating any NDR solution for buyers.

Sandboxing for automated malware analysis has already proven its value in previous large-scale infections of zero day malware, such as the more recent spread of ransomware like Wannacry.

Sandboxes have come a long way since their first inception, impelled mainly by the advancements made in fooling evasion techniques used by malware to detect if its running in a sandbox. The ability for a sandbox to execute malware in a controlled, instrumented environment without the malware detecting the sandbox is what defines its efficacy.

The ability for a sandbox to go undetected by malware comes down to its ability to perform emulation or virtualization. With virtualization, malware can easily look for virtual device drivers and other indicators in order to detect it has been executed in a potential sandbox.

With emulation, it’s much more difficult for the malware to detect the sandbox as even the OS and system calls can be emulated. When emulating, the sandbox is able to provide a response to the syscalls made by the malware to make it think they were successful.

Only one sandbox technology implements emulation as an alternative to virtualization, and that’s Lastline, which is why so many of its own competitors white label their sandbox for their own use. It simply works and works well, proving its ability to detect zero-day malware undetected when the Wannacry ransomware outbreak happened.

BECOMING NETWORK TRAFFIC ANALYSIS

In February of 2019, industry analyst firm Gartner published its first look into a new product category it termed Network Traffic Analysis (NTA). This new market of solutions as defined by Gartner are vendors who've entered the market beginning in 2016 that apply behavioral analysis to network traffic to detect suspicious behavior using a combination of machine learning, advanced analytics, and rule-based detection to detect suspicious activity.

While the market guide did not specifically call out the approach to leverage ML to detect IoCs in encrypted traffic, it did highlight this functionality for vendors in the vendor profiles section.

Specifically, companies that could perform encrypted traffic analysis on encrypted data without needing to decrypt the traffic was an obvious reference to a non-signature-based approach to detection.

The concept of NTA is simple -- analyze raw traffic and/or flow records to build models of an expected behavior for individual nodes and users and alert on deviations from it. An example of this would be clustering algorithms where like objects are clustered together based on identified features and anything outside of that (an outlier) would cause the system to alert as illustrated in Figure 5.

Figure 5: Illustration of clustering in machine learning

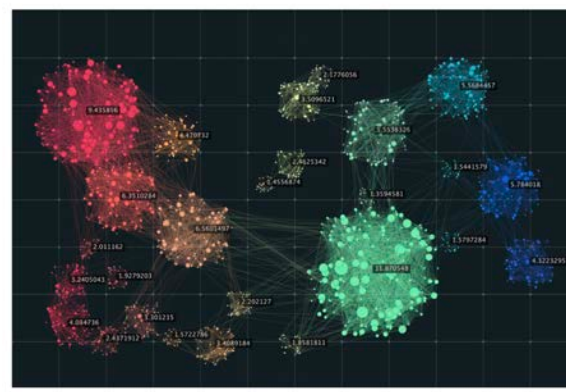

Source: Great Learning

There are different types of learning algorithms in ML with the purposes of this paper to focus specifically on supervised and unsupervised learning.

Unsupervised learning as demonstrated in the previous figure is an area of ML that deals with unlabeled data -- broadly speaking, datasets on some shared attributes and detecting anomalies in the dataset that don't match.

An example of this would be a node 192.168.0.1 that always accepts traffic inbound on port 80 from 9-5pm but all of a sudden at 2am initiates an outbound connection on port 6667 (IRCD).

This would be an example of an outlier in a clustering algorithm. Common unsupervised models include k-means clustering, principal component analysis, and autoencoders.

BECOMING NETWORK TRAFFIC ANALYSIS

In supervised learning, you're specifying the features the models should look for.

Specifically, supervised learning is performed using ground truth, or prior knowledge of what the values should be. This makes the data that vendors train their solutions on when using supervised learning models paramount to the efficacy of the solution. Meaning, the data they are training with should be relevant, contemporary, and rich in features making the concept "garbage in-garbage out" very relevant here.

Common algorithms used in supervised learning include regression, naive bayes, support vector machines, neural networks, and random forests.

But do these algorithms and whether or not one is better than the other even mean anything to buyers? I propose that they don't and will explain why below.

Vendors who've adopted supervised over unsupervised learning is really inconsequential, just as the type of algorithms they chose to use is just as irrelevant to buyers. It's my opinion that asking these questions in the pre-sales process is truly irrelevant and is something so arcane to anyone who isn't a data scientist that the person asking probably won't understand the answer anyway.

It's with great emphasis that I suggest the most important questions in ML for a vendor is asking:

1. What features the vendor trains their models on; and
2. How rich, relevant, and contemporary is the training data they are using, how often is it updated, and where are they getting it?

ENTER

NETWORK DETECTION AND RESPONSE

ENTER NETWORK DETECTION AND RESPONSE

In August of 2015, Anton Chuvakin at Gartner blogged about the concept of a SOC nuclear triad. The tools in his triad covered a holistic view of the entire enterprise using a SIEM, visibility into threats on the network with network forensics tools (NFT), and threats on the endpoint using EDR. The triad was later renamed the SOC visibility triad.

In September of 2018, Chuvakin then blogged about the lack of use of the original term NFT, instead referring to a new acronym picked up by several vendors in their marketing material as network detection and response (NDR). The SOC visibility triad was then updated to contain NDR as the third tool, replacing NFT. Then, in February 2019, Gartner released the first market guide for network traffic analysis. NTA and NDR have now become synonymous with one another, referring to the same market of products and has now begun being referenced in new Gartner publications.

Figure 6: Gartner's SOC Visibility Triad

Source: Gartner

Vendors with network threat detection solutions have now begun abandoning the use of the term NTA except for a select few while many have pivoted to the NDR branding completely, which is the product category name I've adopted moving forward and as a reference to these class of solutions in this paper.

NDR products attempt to provide visibility into the network, real time detection of threats, and investigative tools for analysts to take action against threats or even automate responses on behalf of the analyst. These solutions take network traffic off of the network via virtual traffic mirroring in a cloud service provider (CSP), SPAN port or network TAP, and learn from typical traffic to/from nodes and users and alert to deviations.

Many NDR solutions take different approaches to analyzing encrypted traffic, including the challenges introduced by TLS 1.3 with perfect forward secrecy creating blind spots for some NDR solutions in the market.

The debate over the use of the term NTA or NDR may still not be over as the impetus for adopting the term NDR was to give appropriate attention to NTA solutions capable of performing automated response to threats. However, I've yet to find a solution today that refers to itself as a NTA solution incapable of providing response actions to threats. I digress, but suffice to say, with Gartner recently updating its SOC visibility triad to include NDR as the product category name instead of NTA, we just may see the 2019 NTA Market Guide take a new name in 2020.

REMAPPING

MITRE ATT&CK TO NDR

REMAPPING MITRE ATT&CK TO NDR

Why a remapping of the MITRE ATT&CK to NDR solutions anyway? Because the very roots of the ATT&CK Enterprise matrix as discussed was originally born out of a need to categorize tactics and techniques detected at the endpoint.

Before mapping the NDR solution space to the areas of the ATT&CK Enterprise matrix, it's important I first explain what each tactic addresses from the initial access to pivoting.

Initial Access
Initial access addresses techniques used by adversaries to gain their initial "beach head" on the victim's network. These techniques range from an individual being coerced into providing access to the adversary either through providing credentials over the phone or in person (social engineering known as voice phishing, or vishing), inserting a USB stick into their computer they found in a parking lot that contains a backdoor, a spear phish by clicking on a malicious link or file attachment in an email, or smishing where the victim clicks on a malicious link sent to them via a SMS text message.

These forms of attacks address the soft target of humans, while other techniques include the exploitation of a vulnerable service. Examples include an exploitable vulnerability in a web server or DNS server, a third-party supplier with a VPN connection to the target, or simply a malware infection introduced by self-propagation.

Execution
These set of techniques is the actual execution of a payload or adversary-controlled code that provides an interactive command interpreter with the target system or network. Examples of this would be a command prompt on a target host with superuser privileges or even the execution of commands via an application programmable interface (API) on a remote API endpoint.

Persistence
In an advanced persistent threat (APT) attack, adversaries will above all make sure they have continued access to the target system or network across system reboots, changes to cybersecurity controls, or changed passwords. These consist of backdoors via command and control networks effectively enabling persistent access to the target.

Privilege Escalation
These techniques, when successful, elevate an adversary's restricted user access to be able to execute commands requiring superuser privileges or access systems they aren't allowed to access. Privilege escalation can be the capturing and reuse of Kerberos tickets for an enterprise admin/domain admin in a Microsoft Windows domain, the cracking or guessing of a root user's password, access to the private key and password for a superuser account, or exploiting a vulnerability in an executable file or service that runs as a superuser account. The ultimate objective of privilege escalation is to give unlimited access to an adversary that started out with restricted access to the system or network.

REMAPPING MITRE ATT&CK TO NDR

Defense Evasion
The adversary leverages these techniques in order to evade detection by detective and preventative technical security controls such as leveraging a payload that disables or shuts down memory-resident antimalware agents or encryption of command-and-control traffic. These techniques enable the adversary to go undetected for long periods of time to continue to pivot within the network, exfiltrate data, and compromise more accounts and systems.

Effective defense evasion is what helped adversaries stay undetected in some of the longest recorded APT investigations, such as the Target and Equifax breaches where dwell times were in the magnitude of months, not days or weeks. According to the FireEye M-Trends report for 2020, From October 1, 2018 to September 30, 2019, the global median dwell time was 56 days.

Credential Access
These techniques enable an adversary to steal usernames and passwords for valid credentials to access a target system or network. The techniques used in this category include the capturing of credentials via keystroke loggers or dumping credentials from SAM hives on Windows hosts.

Discovery
This tactic employs techniques where the adversary is learning their target environment, identifying systems, mapping network infrastructure, and understanding the local network and the remote networks connected to it. Techniques include network sniffing to identify what hosts talk to one another, users in the environment, and what services and protocols the hosts use to communicate.

Lateral Movement
Once an adversary establishes a foothold on the network, the effort to discover target devices, hosts, and users results in pivoting within the network (east-west traffic direction) also called lateral movement.

Examples of lateral movement include using remote access applications, such as RDP (remote desktop protocol) for graphical interactive sessions between hosts or SSH (secure shell) to remote Linux/Unix hosts.

Collection
This tactic includes techniques employed by adversaries to harvest data they expect to exfiltrate out of the network, typically to sell for profit on the dark web. Data is typically copied to staging servers on the local network, encrypted and compressed, and then exfiltrated using different protocols.

REMAPPING MITRE ATT&CK TO NDR

Command and Control

Command and control, also referred to as C2, are techniques used by adversaries to establish and maintain remote control of a system, typically using a tool such as a remote access tool (RAT). C2 traffic is typically north-south in directionality, from the internal hosts under the adversary's control to C2 servers under their control on the internet. C2 traffic is typically encrypted to prevent network detection and response solutions from detecting it.

Techniques in this category often blend into other tactics such as persistence, collection, credential access, and execution.

Exfiltration

Exfiltration techniques enable the adversary to actually pilfer data out of the target network from their staging servers. This is the removal of data (data loss or data spill) from the target network where it is then either used by the adversary to support other objectives or simply to profit from it in the sale of that data on the dark web.

Examples of exfiltration can include the exfiltration of usernames and passwords, personally identifiable information (PII), personal healthcare information (PHI), or even payment card information.

Impact

Impact techniques include actions taken by the adversary to affect the confidentiality or integrity of data in the target network. This can include not just exfiltration of the data, but also destruction of the data once it has been exfiltrated. An even worse case scenario would be an unauthorized modification to data, such as unknown changes to drug recipes for a pharmaceuticals company.

CON CLUSION

AUTHOR'S
FINAL THOUGHTS

CONCLUSION

In this first of a multi-part series, we explored the history of MITRE, the ATT&CK matrices, and demystified network detection and response solutions. Furthermore, we discussed the different tactics found in the enterprise ATT&CK matrix, defining initial access all the way to impact, and the techniques adversaries employ within each category.

In the next part of this white paper series, we'll align the Lastline solution to the MITRE ATT&CK and through real-world live-fire exercises, demonstrate how it detects and responds to attacks detected by it in each of the ATT&CK categories so you, the reader, can determine its efficacy to fill the NDR gap in your cybersecurity program.

ABOUT THE AUTHOR

Alissa Knight is a partner at Knight Ink and blends influencer marketing, content creation in writing and video production, go-to market strategies, and strategic planning for telling brand stories at scale in cybersecurity.

She achieves this through ideation to execution of content strategy, storytelling, and execution of influencer marketing strategies that take cybersecurity buyers through a brand's custom curated journey to attract and retain them as long-term partners.

Alissa is a published author, having published the first book on hacking connected cars and am working on a new series of books into hacking and securing APIs and microservices.

ABOUT KNIGHT INK

Firm Overview
Knight Ink is a content strategy, creation, and influencer marketing agency founded for category leaders and challenger brands in cybersecurity to fill current gaps in content and community management. We help vendors create and distribute their stories to the market in the form of written and visual storytelling drawn from 20+ years of experience working with global brands in cybersecurity. Knight Ink balances pragmatism with thought leadership and community management that amplifies a brand's reach, breeds customer delight and loyalty, and delivers creative experiences in written and visual content in cybersecurity.

Amid a sea of monotony, we help cybersecurity vendors unfurl, ascertain, and unfetter truly distinct positioning that drives accretive growth through amplified reach and customer loyalty using written and visual experiences.

Knight Ink delivers written and visual content through a blue ocean strategy tailored to specific brands. Whether it's a firewall, network threat analytics solutions, endpoint detection and response, or any other technology, every brand must swim out of a red sea of competition clawing at each other for market share using commoditized features. We help our clients navigate to blue ocean where the lowest price or most features don't matter.

We work with our customers to create a content strategy built around their blue ocean then perform the tactical steps necessary to execute on that strategy through the creation of written and visual content assets unique to the company and its story for the individual customer personas created in the strategy setting.

Contact Us
Web: www.knightinkmedia.com
Phone: (702) 637-8297
Address: 1980 Festival Plaza Drive, Suite 300, Las Vegas, NV 89135