Implementing Intrusion Kill Chain Strategies

by Creating Defensive Campaign Adversary Playbooks

Rick Howard | Ryan Olson

ABSTRACT

This paper extends the work of the Lockheed Martin research team on intrusion kill chains (the identification and prevention of cyber intrusions) in 2010. The theory has languished in the network defender community not because it is not the right idea, but because most InfoSec teams do not have the resources to implement it. What has prevented the success of the intrusion kill chain strategy is a standard framework to collect the intelligence associated with specific adversaries, to share and consume that standardized intelligence with trusted partners, and then to automatically process that intelligence and distribute new prevention controls to the network defender's security stack. The adversary playbook is that framework.

***Keywords and Concepts**: Adversary Playbooks, Adversary Campaigns, Adversary Playbook Visualizations, Automatic Intelligence Sharing, Cyber Strategic Defense Initiative (Automatic Deployment of Prevention Controls), Data Islands, Defensive Campaigns, DevSecOps Automation Layers, How Big is the Adversary Problem?, Intrusion Kill Chain Analysis, Sharing Technical and Tactical Attack Details vs Sharing Strategic Defensive Campaigns*

SETTING THE STAGE

Sometime in the early 1990s, the Internet became useful to commercial enterprises, academic institutions, and government operations. Soon after, criminals, spies, warriors, and troublemakers of all sorts discovered that it might be a useful avenue through which to pursue their activities. That was about the time when it became necessary to have network defenders within all organizations dedicated to protecting the enterprise. From the beginning, security practitioners installed their own systems designed to detect and prevent the use of malicious tools by cyber adversaries. Looking back, that was shortsighted. By focusing on individual attack tools and the indicators of compromise left in their wake, with no understanding the adversary's broader goals, the network defender community was left with no way to know if their defensive plans were working. We could tell if we stopped a specific malicious tool with our defensive systems but had no idea if we prevented the success of the cyber adversary's ultimate goal.

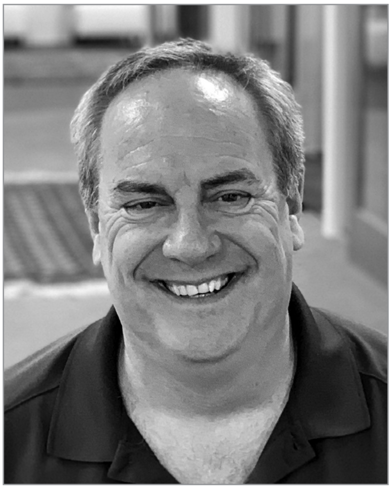

Rick Howard is the Chief Analyst, Chief Security Officer, and Senior Fellow at The CyberWire, a cybersecurity media network. His previous jobs include the Palo Alto Networks CSO, the TASC CISO, the iDefense GM, the Counterpane Global SOC Director, and the Commander of the U.S. Army's Computer Emergency Response Team. He was one of the founding players that created the Cyber Threat Alliance, and he also created the Cybersecurity Canon, a Rock & Roll Hall of Fame for cybersecurity books. Rick holds a Master of Computer Science degree from the Naval Postgraduate School and an engineering degree from the U.S. Military Academy. He also taught computer science at the Academy from 1993 to 1999. He has published many academic papers on technology, security, and risk, and has contributed as an executive editor for two books: *Cyber Fraud: Tactics, Techniques and Procedures* and *Cyber Security Essentials*.

When the research team at Lockheed Martin published their now-famous 2010 white paper on the Cyber (Intrusion) Kill Chain®,[1] the network defender community registered a new method to defeating the cyber adversary. Instead of installing one prevention control designed to defeat a single malicious tool, we could install prevention controls designed to defeat specific adversaries at each step of their attack sequence. Today, we know that hackers and hacker groups must string a series of actions across the intrusion kill chain in a campaign to achieve their purpose. Our aim should not be to stop the use of one technical tool with no context about what the adversary is trying to accomplish. It should be to stop the overall success of the attacker's entire campaign.

Unfortunately, the intrusion kill chain theory languished. Most network defenders understood the importance of the concept but could not muster the resources to deploy the tactics required to implement it. We needed to extend the theory and create a framework so that network defenders could build infrastructure to support it. The adversary playbook is one of those frameworks.

ADVERSARY PLAYBOOK DESCRIPTION

An adversary playbook collates all known intelligence on the hacker group's attack sequence: tactics, techniques, indicators of compromise, attack time frame, and context about motivation as well as attribution. It provides a standard framework designed to collect cyber adversary actions across the intrusion kill chain and eases the burden of sharing that collection with other network defenders. It further facilitates the automatic consumption of that intelligence on the other end, allows the receiver to write code to absorb it systematically, and provides the means to automatically deploy new and updated security controls to their already deployed defensive posture within their DevSecOps infrastructure.

Ryan Olson is the Vice President of Threat Intelligence for Palo Alto Networks. He leads Unit 42, a team responsible for the collection, analysis, and production of intelligence on adversaries targeting organizations around the world. His area of expertise is detecting and identifying actors and groups conducting cyber-crime and cyber-espionage operations. Ryan is a contributing author to *Cyber Fraud: Tactics, Techniques and Procedure*, and the primary author of *Cyber Security Essentials*. Before joining Palo Alto Networks, Ryan served as Senior Manager in Verisign's iDefense Threat Intelligence service. Ryan is a named inventor on two patents related to malware analysis and threat intelligence collection. Ryan holds a Bachelor of Science degree in Management Information Systems from Iowa State University and a Master of Science degree in Security Informatics from The Johns Hopkins University.

The five characteristics of an adversary playbook include the following:

1. Description of a hacker or hacker group's goals.
2. Timestamp of a hacker or hacker group's campaigns.
3. Collection of tactics and techniques they employed across the intrusion kill chain using the MITRE ATT&CK® framework.[2]
4. Aggregated indicators of compromise left behind as they execute their attack sequence.
5. Intelligence data set stored in a STIX™[3],[4] object designed to facilitate automatic intelligence consumption and deployment of security controls.

PLAYBOOKS VS. CAMPAIGNS

One adversary playbook might consist of several campaigns spread out over time. Network defenders describe campaigns in three ways: campaigns attempted in the past, campaigns currently running, and campaigns running in parallel. These descriptors are important because they create the opportunity to compare and contrast adversary behavior over time. When adversaries devise an attack sequence—a campaign—and run it against a victim, they may decide to change parts of the sequence for various reasons: efficiency, prevention control avoidance, new tools, etc. When they make those changes, however, they do not change the entire sequence. They only change the bits that need adjustment. The implication then is that the bulk of prevention controls that a network defender deploys against a specific campaign will likely apply to other campaigns run by the same adversary group. Even if the adversary leverages some new zero-day vulnerability somewhere in the attack sequence, with a vulnerability that nobody has ever heard about before, network defenders will have a good chance of preventing the adversary from being successful because of the other prevention controls already deployed against this playbook will still work.

Collecting all campaigns into an adversary playbook also facilitates the assessment of any new attack sequences. If the InfoSec team already knows which prevention controls it has in place for campaign one, when campaign two emerges, the task of evaluating whether the organization is vulnerable to the new campaign becomes easier. The team already knows what it has in place and can make decisions regarding how fast to respond to any new tactics. If the change in campaign two bypasses the already deployed defensive controls from campaign one, that is a higher priority than if the bulk of prevention controls are still valid.

How Many Active Playbook Campaigns Are Hackers Running on the Internet?

Since adversary playbooks contain every tactic and technique for specific attack sequences in various campaigns, network defenders can answer some important Critical Information Requirements (CIRs).[5] For example, one useful CIR asks how many tactics and techniques of all known adversaries are there? Another is how many adversary campaigns are hackers running on any given day? The InfoSec community already has a good answer to the former—and a decent estimate for the latter.

MITRE researchers have been collecting and documenting attacker tactics and techniques across the intrusion kill chain since 2013.[6] As of this writing, they are currently tracking 12 tactics and 330 techniques.[7] Of course, these numbers change over time as the researchers refine their collection mechanisms and develop insight into the problem space. The striking fact is how low the number is. Because of the volume of cyberattacks that are public knowledge these days, it seems like threat actors utilize millions of techniques to break into systems. In reality, hackers reuse a handful of tried and true techniques because network defenders have failed to deploy prevention controls against them. Malicious actors, therefore, do not need to create millions of new techniques. The old ones work just fine.

The answer to how many adversary campaigns hackers are running on the Internet on any given day is an estimate, and like the number of tactics and techniques out there, the number is likely smaller than expected. The Cyber Threat Alliance is a group of ~28 cybersecurity vendors who share adversary playbook information.[8] Their Algorithms and Intelligence Committee is staffed by some of the brightest intelligence minds in the commercial sector. For the past four years, their estimate of the volume of live adversary campaigns on the Internet on any given day has been under 250.[9] Unit 42 is Palo Alto Networks' Threat Intelligence Team, and for the last two years, it has been publishing adversary playbooks for public consumption. As of this writing, it has published ~22 adversary playbooks, which include ~50 campaigns. The observations by the Cyber Threat Alliance and Unit 42 estimate with 95% confidence that the number of active campaigns attackers are running on any given day is between 50 and 250.[10]

The InfoSec community has been treating the problem with the opposite assumption: that the volume of live attack sequences is so large, we cannot possibly keep up with it. If adversaries are running fewer than 250 campaigns every day that uses the same 330 techniques, then the conventional wisdom is completely wrong. It is possible for the community to keep up with active attack campaigns. It is possible to deploy prevention controls more rapidly than the adversary can develop new tactics. The obstacles that prevent us from doing so are not about scale but about a willingness to share known adversary's attack sequences with our peers, along with the difficulty of automating the response once we have that intelligence. We designed adversary playbooks to facilitate the latter.

For the former, there are two schools of thought in the network defender community that mostly align with the policies of government cyber intelligence groups and everybody else. For government intelligence groups, their mission is more significant in that they are trying to help government leaders influence the international political and security environment. For everybody else, we are just trying to prevent material impact on our organizations. The differences between the two are stark. For the government side, some of the intelligence they collect comes from espionage operations. As such, they have a vested interest in protecting their sources and methods. For everybody else, most of the intelligence collected is from one's network and sharing partners, and it makes sense to share with trusted partners as efficiently as possible. For the government, it makes sense to support that sharing so that they do not have to give up their sources and methods.

One argument against sharing is that if adversaries discover what the network defenders know about them, then they will change their attack sequence, but that is the point of efficiently sharing threat intelligence. Instead of the network defender community scrambling to react to every newly discovered technical technique, we want to cause the adversaries to expend additional resources attempting to find new attack techniques that work. The key is agility in sharing new intelligence quickly and deploying new security controls to our infrastructure with speed and efficiency. The adversary playbook model supports that concept.

PLAYBOOK DATA ELEMENTS

Playbooks consist of two types of data: observables and context. Observables are digital objects or clues left behind by the adversary that give network defenders notice that there might be an intruder. We find them on all the data islands where our employees operate: on laptops and mobile devices inside the traditional perimeter and out in public, on servers within data centers, on SaaS (software as a service) supporting infrastructure, and on various public cloud infrastructures that provide PaaS (platform as a service) and IaaS (infrastructure as a service). Finding these observables on these data islands means that an attacker either executed an attack sequence in the past or is busy executing one currently. Context is intelligence derived from the observable. In other words, what do analysts know—or what can they assume—when they find an observable?

Consider the information included in Table 1. It lists the observables and derived context that one team of network defenders witnessed during an unsuccessful attack campaign by a hacker group we call DragonOK. By derived context, we mean that InfoSec analysts observed a malicious email arriving in an employee's inbox with the subject, "Your Purchase Order," and assumed that the attackers used spear-phishing as their delivery mechanism. They found the malicious Word document with its unique hash, "020f5692b998...," and derived that the attackers leveraged a known vulnerability, "CVE-2015-1641," for their exploit code. They observed the portable executable file, "12d88fbd4960...," and derived its name, "Nflog," and its function, a remote access tool (RAT). Finally, the analysts recorded the command and control domain name, "www.dppline[.]org," and derived that the attackers used the standard HTTP communications protocol for command and control purposes.

Table 1. Adversary Playbook Data from a Single Attack by the DragonOK Threat

Intrusion Kill Chain Phase	Data
Delivery	Observable: "Email Subject: Your Purchase Order"
	Context: TTP: Spear Phishing
Exploitation	Observable: Sample – Word Document: 020f5692b998...
	Context: Exploited Vulnerability: CVE-2015-1641
Installation	Observable: Sample – Portable Executable: 12d88fbd4960...
	Context: Malware Name: Nflog
	Context: Malware Type: Remote Access Trojan
Command and Control	Observable: Domain Name: www.dppline[.]org
	TTP: Standard Application Layer Protocol – HTTP

Intrusion Kill Chain Analysis to Support a Defensive Campaign

A domain name that malware uses to support its command and control function is an observable. This kind of intelligence is valuable for blocking a specific attack technique and for "connecting the dots" between two separate attack sequences when adversaries reuse tools and infrastructure. Unfortunately, the time-to-live period of this observable is often short. Once the network defender community becomes aware of it, an attacker will stop using it. Alternatively, the higher-level "context" data elements within an adversary playbook are much longer lived, but they may not be as valuable to network defenders in defeating the attack or creating a defensive campaign.

Analyzing the data from the table above, network defenders might decide to block traffic destined to the associated command and control domain name, www.dppline[.]org, preventing malware already inside the network from communicating with the attacker. This is certainly a worthwhile action to take, but it will likely be a temporary solution. Once the attackers behind DragonOK notice that no traffic is coming into their server, they will probably change their command and control server to a different domain. Advanced adversaries change their command and control domains on a regular and automated cadence anyway to prevent this specific defensive measure. A longer-term action would be to deploy the Microsoft patch for

"CVE-2015-1641." This would prevent future attacks by DragonOK and other adversary groups who exploit the same vulnerability. Still, it would not prevent DragonOK from further actions along the intrusion kill chain spectrum if they were already inside. Neither of these defensive tactics offers a robust defensive campaign against DragonOK. This is the reason for intrusion kill chain analysis. The act allows network defenders to find gaps in their defensive posture against specific adversaries.

Let us examine the same data in another way. Figure 1 shows us the DragonOK attack techniques and their corresponding intrusion kill chain phases.

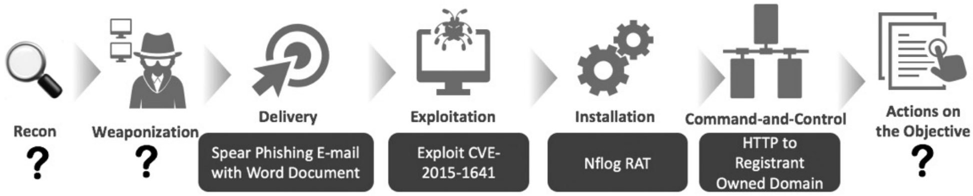

Figure 1. Intrusion kill chain view of a DragonOK attack

This view makes it more apparent that we are missing some elements of the attack. Based on our observations of a single attack, we only have information about four of the attack sequence phases. Of course, the goal of building an adversary playbook is not to look at a single attack, but at all the attacks attributed to the same adversary. The adversary playbook identifies past tactics and techniques and those likely to be used in the future. If other organizations that have observed attacks from DragonOK share additional data with us in the same format, we can build a complete picture (Figure 2).

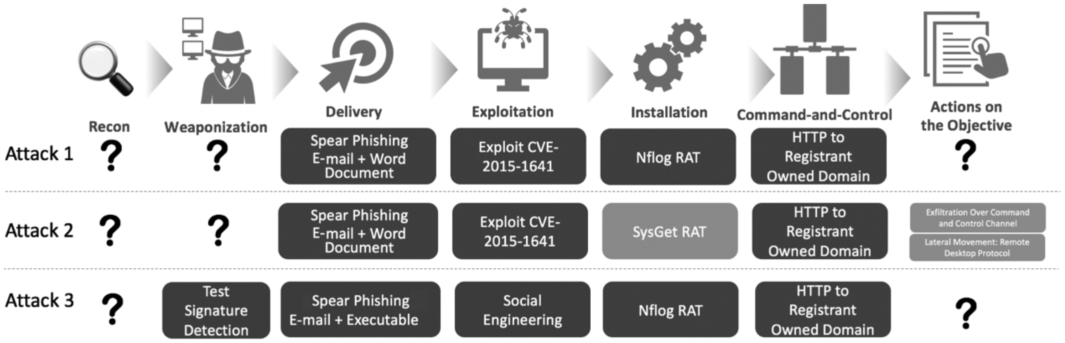

	Recon	Weaponization	Delivery	Exploitation	Installation	Command-and-Control	Actions on the Objective
Attack 1	?	?	Spear Phishing E-mail + Word Document	Exploit CVE-2015-1641	Nflog RAT	HTTP to Registrant Owned Domain	?
Attack 2	?	?	Spear Phishing E-mail + Word Document	Exploit CVE-2015-1641	SysGet RAT	HTTP to Registrant Owned Domain	Exfiltration Over Command and Control Channel Lateral Movement: Remote Desktop Protocol
Attack 3	?	Test Signature Detection	Spear Phishing E-mail + Executable	Social Engineering	Nflog RAT	HTTP to Registrant Owned Domain	?

Figure 2. Attack sequence view of three DragonOK playbooks

This picture in Figure 2 remains incomplete, but now we know more about the DragonOK adversary playbook. Attack 2 indicates this attack sequence uses a different remote administration tool (RAT), called SysGet, during the installation phase, compared to Nflog in Attack 1 and tells us more information about what the attackers do once they breach a network. Attack 2 indicates, in the "Actions on the Objective" column, that the attackers exfiltrate data over a command and control channel and move laterally within the victim's network using the

Remote Desktop Protocol (RDP). Attack 3 shows us more ways the threat actor delivers its malware and how it might evade antivirus protection. In the "Delivery" column of Attack 3, the attackers use spear-phishing to deliver malicious code. Then, as shown in the "Exploitation" column, they use social engineering to trick the victim into running that code. Visualizations of other adversary playbooks can be found at the Unit 42 Playbook Viewer site.[11]

If a single group of network defenders, operating alone, observed Attack 1, its options for preventing the success of DragonOK in its networks would be limited and likely would not work. By combining and sharing the intelligence gathered by other network defender groups for other DragonOK campaigns, however, the entire InfoSec community could build a more robust defensive campaign specifically designed to thwart the DragonOK playbook.

We designed these visualizations for two purposes: we wanted to help analysts understand the value of grouping adversary intelligence into playbooks, but more importantly, we designed the playbooks to be readable by a machine to facilitate the network community's automatic sharing of this intelligence.

ADVERSARY PLAYBOOK DESIGN: THINGS TO CONSIDER

Table 2 shows a summary of the DragonOK attack information in a tabular form. This version of playbook information, boiled down to the essentials for automatic consumption, is not long or particularly verbose. Human analysts who try to read this information will likely find it wanting. That is why it is essential to include reference material, which gives more detail on

Table 2. Tabular Form of DragonOK Playbook

Adversary: DragonOK	
Recon	UNKNOWN
Weaponization	UNKNOWN
Delivery	• Spear Phishing with Word Attachment • Spear Phishing with EXE Attachment
Exploitation	• Exploit Known Vulnerability – CVE-2015-1641 • Social Engineering
Installation	• Tool: Nflog • Tool Type: Remote Administration Tool (RAT) • Tool: SysGet • Tool Type: Remote Administration Tool (RAT) • Tool: IsSpace • Tool Type: Remote Administration Tool (RAT) • Tool: TidePool • Tool Type: Remote Administration Tool (RAT)
Command and Control	• Standard Application Layer Protocol
Actions on Objectives	UNKNOWN

named elements. For instance, intelligence analysts might like to share the discovered DragonOK remote administration tools: NFlog and SysGet. Providing reference links to this more detailed information is not essential to automatic intelligence sharing, but it is useful for developing a more robust picture of adversary behavior.

One of the significant barriers that has inhibited intelligence sharing from the beginning[12] is that the network defender community could not agree on a standard language or format to transfer the information. Common sense dictates that to facilitate information exchange, network defenders must agree on what to call things. If one person uses the term “Keylogging” to describe capturing keys pressed on a keyboard, but another uses the broader term “Input Capture,” the entire network defender community could be talking about the same attack technique, but nobody would know.

This is where MITRE’s Adversarial Tactics, Techniques, and Common Knowledge model and framework come in.[13],[14] MITRE ATT&CK includes hundreds of techniques in a Wiki-like format (Figure 3) to provide names, descriptions, and links to examples of adversaries using specific tactics inside an organization’s networks.

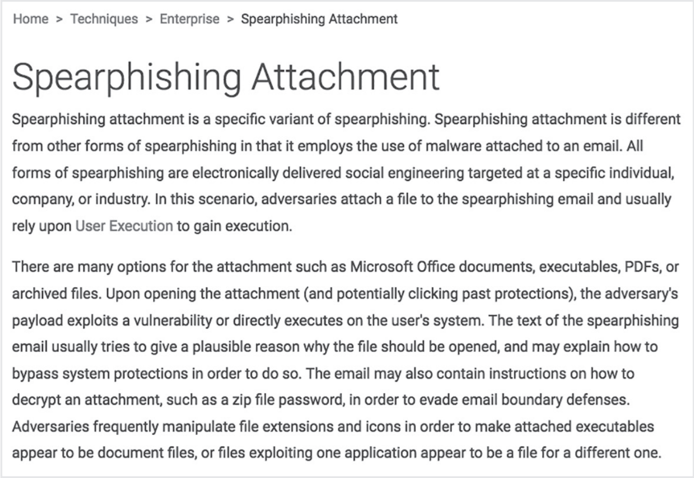

Home > Techniques > Enterprise > Spearphishing Attachment

Spearphishing Attachment

Spearphishing attachment is a specific variant of spearphishing. Spearphishing attachment is different from other forms of spearphishing in that it employs the use of malware attached to an email. All forms of spearphishing are electronically delivered social engineering targeted at a specific individual, company, or industry. In this scenario, adversaries attach a file to the spearphishing email and usually rely upon User Execution to gain execution.

There are many options for the attachment such as Microsoft Office documents, executables, PDFs, or archived files. Upon opening the attachment (and potentially clicking past protections), the adversary's payload exploits a vulnerability or directly executes on the user's system. The text of the spearphishing email usually tries to give a plausible reason why the file should be opened, and may explain how to bypass system protections in order to do so. The email may also contain instructions on how to decrypt an attachment, such as a zip file password, in order to evade email boundary defenses. Adversaries frequently manipulate file extensions and icons in order to make attached executables appear to be document files, or files exploiting one application appear to be a file for a different one.

Figure 3. ATT&CK description of the spear-phishing attachment technique

The tabular format of the playbook in Table 2 is closer to something a machine can read as compared to the intrusion kill chain diagram shown in Figures 1 and 2, but what we need to be able to exchange this information is a machine-readable format.

PLAYBOOKS IN STIX

There have been many efforts to build a common language to facilitate information sharing from both the open-source and commercial communities. In recent years, though, the network defender community seems to have embraced STIX™ (Structured Threat Information eXpression) to be, at least, the common language to which all others must talk. This is evident by the fact that the most famous and well-respected information sharing organizations—like the Financial Sector Information Sharing and Analysis Center (FS-ISAC), the Cyber Threat Alliance, the Defense Industrial Base Information Sharing and Analysis Organization, IBM, and Palo Alto Networks, to name a few—have all adopted it.[15]

STIX allows for the exchange of many forms of threat intelligence, from a simple list of IP addresses to descriptions of assets involved in an incident. With an adversary playbook, our goal is associating adversaries with the tactics and techniques they employ at specific phases of the intrusion kill chain. Three core elements in STIX are necessary for encoding information for an adversary playbook.

The "Threat Actor" element is the characterization of a specific adversary. It does not need to include identifying information about individual actors, but it does need to include a consistent code name or identifier that one can associate with this adversary. The Threat Actor element is what lets the recipient know with which adversary the remaining elements should be associated.

"TTPs" (tactics, techniques, and procedures) are representations of what an adversary does when it conducts its attack. Does it scan the Internet looking for hosts that are vulnerable to an SSH, or does it send targeted spear-phishing email messages to your CFO? STIX allows broad descriptions of TTPs, but to be incorporated into a playbook, we suggest a predefined set of descriptions like those in MITRE ATT&CK be used.

STIX 1.2 does not have a mechanism to specifically reference MITRE ATT&CK TTPs, but they can be included by adding custom fields or by overloading the included Common Attack Pattern Enumeration and Classification (CAPEC) reference to point to MITRE ATT&CK TTP identifiers instead. MITRE has already created MITRE ATT&CK definitions for TTPs STIX 2.0. (see STIX 1.x vs 2.x box).

Indicators convey specific observable patterns in STIX. They tell us what to look for in our networks and on our endpoints when we are trying to identify an attack. STIX 1.x uses the CybOX (Cyber Observable eXpression) standard for defining specific types of observables, but STIX 2.x has incorporated these observables directly into the standard.

Whether STIX 1.x or 2.x is chosen to encode playbook data, the elements described above are the minimum you need to include when building a package for exchange. Details about the impact of an intrusion or the types of organizations targeted are valuable, but the Threat Actor, TTP, and Indicator data are critical.

Why Do We Need Adversary Playbooks?

We designed the adversary playbook to make it easier to share threat intelligence with trusted partners in a meaningful and efficient way. We also designed it to reduce the impediments of automatically processing that intelligence on the receiving end, allowing network defenders to make decisions faster than the hacker. By adopting the adversary playbook construct, cyber intelligence practitioners can leverage actionable intelligence in a machine-readable format designed for the activities that follow.

Intelligence Collection and Capture. Generally, all intelligence teams are unique, regardless if they work in similar industries or government sectors. Team size, financial resources, organizational mission, and the boss's CIR (Commander Information Requirements)[16] all contribute to team uniqueness. This is one of the main reasons it has taken so long to develop a universal standard format for storing cyber intelligence. For cyber intelligence teams, the adversary playbook provides an industry-accepted format to store raw information on adversary behavior across the intrusion kill chain in a manner that is easily shared with other cyber intelligence teams.

Intelligence Distribution. To see a mostly complete view of the elephant (i.e., a comprehensive view of adversary activity), it is incumbent upon intelligence teams to swap information on adversary attack sequences in real time with trusted partners. Combining the intelligence with that of two or more trusted partners fills in the gaps of what one intelligence team knows. Distributing that intelligence to them in a machine-readable format allows those partners to process it automatically for their use without having to dedicate humans to the endeavor.

Intelligence Consumption. Intelligence teams consume threat intelligence products from trusted sharing partners in a format and language that facilitate automatic processing. The value of information sharing is thus realized because InfoSec teams can concentrate on more strategic tasks, like designing defensive campaigns or updating defensive campaigns for all known cyber adversaries, instead of manually crunching through written reports in documents, slide decks, spreadsheets, and emails.

DevSecOps Security Control Deployment. Network defenders understand the value of the DevSecOps infrastructure-as-code philosophy. They know it is imperative that whatever processes and procedures their DevOps teams pursue, they should go right along with them in a "shift left" kind of way. Security is one of the operational silos that the DevOps movement is designed to strike down. However, after years of advocating for a DevOps or DevSecOps vision, Gene Kim, author of several DevOps books, says:

> [I]ncredible problems still remain. In other words, someone could embrace fully all the principles and patterns espoused in *The Phoenix Project* (a book about the DevOps philosophy listed within the Cybersecurity Canon Hall of Fame[17]) … but I think one of the problems is that there is still all these … invisible structures required to make developers productive.[18]

For network defenders, one set of invisible structures prohibiting automatic response is unformatted intelligence products. They cannot very well automate their response to incoming intelligence if a human is required for each piece. Once intelligence products come into the organization in an understood and agreed-upon framework, it becomes possible to automatically deploy prevention controls to the organization's deployed security infrastructure. This goal has been out of reach in the InfoSec community, but with the adoption of adversary playbooks as a best practice, the community can start to move toward achieving it. DevSecOps security control deployment becomes achievable now.

Defensive Campaign Design and Deployment. As intelligence teams share and consume more information on adversary campaigns over time, the operational picture of how the adversary operates on the Internet becomes more apparent. It is possible to design a comprehensive defensive campaign tailored to a specific adversary playbook within the network defender's DevSecOps infrastructure. InfoSec teams design these defensive campaigns to defeat the adversary's ultimate objective. In terms of material impact, there is a sizable difference between an adversary group compromising a single laptop on the victim's network as a key step in its attack sequence and that same group succeeding in exfiltrating customer data that might eventually materially impact the victim's organization. It is not enough to only try to stop the former. It is desired but insufficient. InfoSec teams must be successful at preventing the latter, and the design of all defensive campaigns must reflect that. The technology needed for network defenders to accomplish these goals is not yet ready. The first step is for all of us in the network defender community to adopt the adversary playbook concept as a common language to communicate what we know about the adversary's purpose.

Figure 4 shows a potential future model of cyber conflict represented by three color tones: light - security infrastructure and protected data, dark - network defender actions, and medium - adversary actions. The labeled arrows show in which direction information and action flow. The key on the right provides additional details.

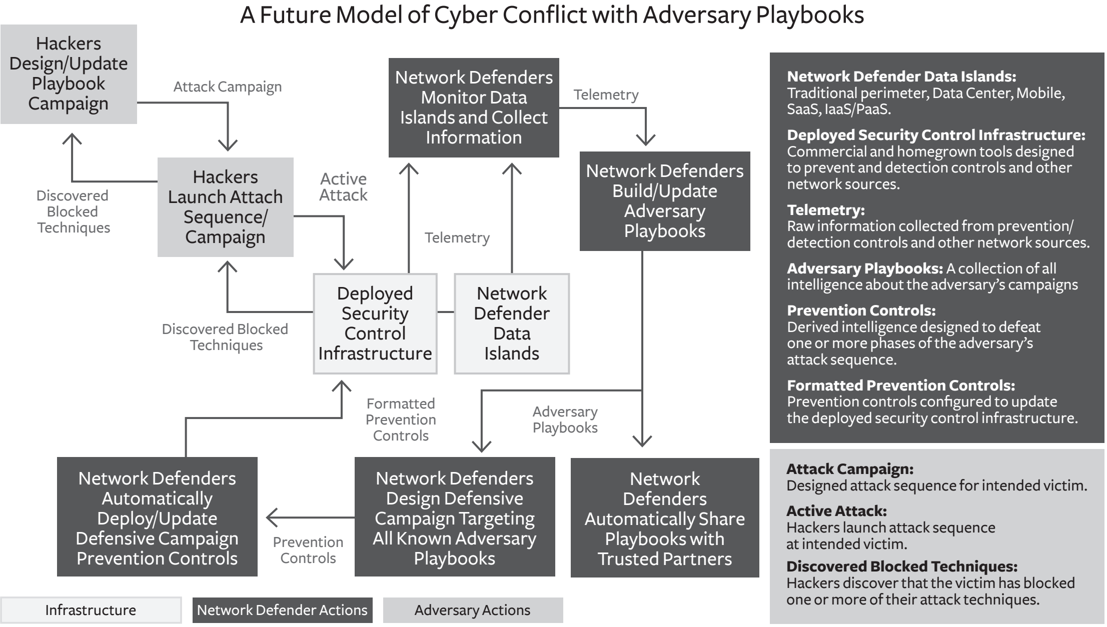
A Future Model of Cyber Conflict with Adversary Playbooks
Hackers Design/Update Playbook Campaign
Attack Campaign
Hackers Launch Attach Sequence/ Campaign
Discovered Blocked Techniques
Active Attack
Network Defenders Monitor Data Islands and Collect Information
Telemetry
Telemetry
Network Defenders Build/Update Adversary Playbooks
Deployed Security Control Infrastructure
Network Defender Data Islands
Discovered Blocked Techniques
Formatted Prevention Controls
Adversary Playbooks
Network Defenders Automatically Deploy/Update Defensive Campaign Prevention Controls
Prevention Controls
Network Defenders Design Defensive Campaign Targeting All Known Adversary Playbooks
Network Defenders Automatically Share Playbooks with Trusted Partners
Infrastructure
Network Defender Actions
Adversary Actions

Network Defender Data Islands: Traditional perimeter, Data Center, Mobile, SaaS, IaaS/PaaS.

Deployed Security Control Infrastructure: Commercial and homegrown tools designed to prevent and detection controls and other network sources.

Telemetry: Raw information collected from prevention/ detection controls and other network sources.

Adversary Playbooks: A collection of all intelligence about the adversary's campaigns

Prevention Controls: Derived intelligence designed to defeat one or more phases of the adversary's attack sequence.

Formatted Prevention Controls: Prevention controls configured to update the deployed security control infrastructure.

Attack Campaign: Designed attack sequence for intended victim.

Active Attack: Hackers launch attack sequence at intended victim.

Discovered Blocked Techniques: Hackers discover that the victim has blocked one or more of their attack techniques.

Figure 4. Model utilizing adversary playbooks in cyber conflict

The model shows that whichever side has the most agility will win. If hackers can deploy their attack campaigns more rapidly than network defenders can deploy their prevention controls, they will likely succeed. Conversely, if network defenders can collect telemetry, organize it into adversary playbooks, share those playbooks with their trusted partners, design defensive campaigns to thwart them, and deploy those defensive campaigns on their existing infrastructure faster than the hackers can act, then the network defender will likely succeed in preventing material impact to their organization due to cyberattacks. The network defender's only hope of being more agile than their cyber adversaries is to automate the deployment of prevention controls to the already-deployed security control infrastructure. To be specific, network defenders need four automation layers in their DevSecOps infrastructure:

1. **Adversary Playbook Consumption**—the ability to automatically consume adversary playbook intelligence products from their trusted sharing partners.
2. **Adversary Playbook Sharing**—the ability to share internally derived adversary playbook intelligence products automatically with their trusted sharing partners.
3. **Defensive Campaign Staging**—the decisions of the InfoSec team about how to thwart the adversary playbook efficiently at each phase of the intrusion kill chain and staging that information in a way that facilitates automatic deployment.
4. **Defensive Campaign Deployment**—leveraging the defensive campaign staging area by automating the deployment of security controls to the network defender's already-deployed security control infrastructure.

Building defensive campaigns and supporting automation layers has the added benefit of helping network defenders identify the gaps and redundancies in their prevention control toolset. If the intelligence team discovers that, after it completes its intrusion kill chain analysis, there is no way to stop the successful completion of the adversary's ultimate mission, this might indicate that the organization needs another prevention tool. Likewise, after the InfoSec team has deployed and maintained several defensive campaigns, it may discover some security tools within their DevSecOps arsenal that are not often used or are redundant controls for a specific phase of the attack sequence. That might be an indicator that the organization has too many tools deployed.

The industry-standard MITRE ATT&CK framework has shown us that the number of techniques used by hackers is under 400.[19] By collecting the techniques of all known hacker groups, intelligence teams can see which techniques are used most often. If the bulk of hacker groups mostly use the same handful of techniques repeatedly, the InfoSec teams could prioritize their defensive campaigns on those techniques first. For instance, the four adversary playbooks in Figure 5 identify the same hacker technique in the Exploitation phase of the attack sequence. Building defensive campaigns that prevent this exploit from working protects the organization from four different adversary groups at once.

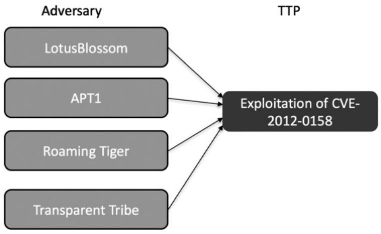

Figure 5. Multiple adversaries use the same TTP

Product managers behind many commercial security tools designed them to be successful against various adversary tactics and techniques. For example, security vendors created commercial off-the-shelf (COTS) spam tools to thwart adversaries from using email as a delivery tool. Others created anti-exploitation tools to prevent adversaries from using exploitation techniques on the endpoint. Deploying these commercial tools and updating them with the latest response based on new intelligence serves as the basis for all network defender prevention programs. Analyzing the aggregate hacker playbooks will provide network defenders insight into what kinds of tools they will need. Figure 7 demonstrates that all network defenders need some anti-exploitation tool.

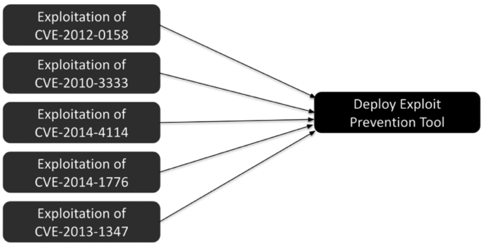
Exploitation of CVE-2012-0158
Exploitation of CVE-2010-3333
Exploitation of CVE-2014-4114
Exploitation of CVE-2014-1776
Exploitation of CVE-2013-1347
Deploy Exploit Prevention Tool

Figure 6. One defense may be effective against multiple TTPs

Additionally, by building playbooks for your top 10 (or more) adversaries and evaluating their tactics and techniques against your possible defenses, you can identify which technologies, processes, or policies will have the most impact on defending your organization from the significant threats you face. Figure 7 demonstrates that it might be possible for the InfoSec team to reduce the myriad of adversary tactics and techniques to a handful of generic defenses as an added layer to defensive campaign strategies.

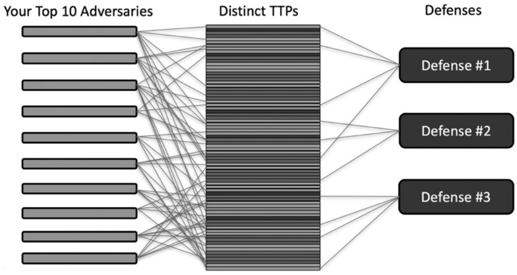

Figure 7. Identifying overlap between your top adversaries, their TTPs and your defenses

CURRENT STATE

Unit 42[20] did the initial work on adversary playbook development some five years ago. They brought that work to the Cyber Threat Alliance[21] when the security vendor intelligence-sharing group was just forming. The adversary playbook concept is baked into the Cyber Threat Alliance's DNA. Members share adversary playbook intelligence products so their common customers do not have to do it themselves. They have become a collection of trusted sharing partners. Because they are security vendors, when they receive the daily intelligence from the other vendors, they develop prevention controls for their own product sets and deliver them to their customer base. Aside from this handful of security vendors, no one else in the network defender community has adopted the adversary playbook concept as a best practice yet, and no one has come close to building defensive campaigns for all known adversary attack sequences. There is still much work to be done.

Figure 5 shows a potential future model of cyber conflict. To carry out this vision, the network defender community must transform its approach from manually responding to cyberattacks to embracing the philosophy of the DevSecOps model. The community has to get comfortable with automated responses to cyberattacks. It also must let go of the notion that InfoSec teams should respond to technical threats observed on their networks without consideration for the cyber adversaries' objectives.

NEXT STEPS

To achieve the vision of the DevSecOps model, the network defender community should pursue the following short-term activities:

- **Join the Cyber Threat Alliance.** Each of us in the network defender community already has a set of commercial security vendors we use to defend our data islands. The Cyber Threat Alliance is nonprofit organization working to improve the cybersecurity of the global digital ecosystems by enabling high-quality cyber threat sharing among companies and organizations. We must educate the network community regarding the benefits of the Cyber Threat Alliance. Even if our organization does not now have the resources to work toward this vision internally, using security vendors that do will spread the adversary playbook as a best practice within the community. The Cyber Threat Alliance has the added benefit of putting the burden on each security vendor to deploy prevention controls designed to defeat all known adversary attack sequences. This is one way we can promote and encourage the standard.
- **Encourage Government Organizations and Standards Bodies to Adopt the Adversary Playbook Model.** Whenever possible, urge government entities in charge of national cyber policy and government InfoSec teams to adopt adversary playbooks as a best practice.

- **Build and Share Adversary Playbooks with Trusted Partners.** If your organization is not sharing cyber intelligence with a trusted partner, find one. Make it your business to determine how your organization can make the adversary playbook model a reality in your organization. Find ways to share your internally developed adversary playbooks with your security vendors, especially if they are members of the Cyber Threat Alliance.
- **Encourage the Information Sharing and Analysis Centers (ISACs) to Adopt the Standard.** If you already belong to an information-sharing group, like the ISAC for your business sector, encourage the group's leadership to adopt the adversary playbook standard too. Find a way for your ISAC membership not only to share adversary playbooks with themselves but also to share their adversary playbook intelligence products with the Cyber Threat Alliance. In this way, the ISAC helps its members enhance their DevSecOps projects and helps vendors provide prevention controls to the products that their members already use.
- **Support and Adopt the MITRE ATT&CK Framework Standard.** For your intelligence efforts, use the MITRE ATT&CK framework to develop a universal standard for the community.[22],[23]
- **Support the Oasis Standards Group for STIX.** The Organization for the Advancement of Structured Information Standards (OASIS) is a nonprofit, international consortium that manages the open-source standards for STIX.[24] We believe the OASIS STIX standard is the way forward for future DevSecOps work.
- **DevSecOps Automation Layers.** Start building your own DevSecOps infrastructure to support these layers: Adversary Playbook Consumption, Adversary Playbook Sharing, Defensive Campaign Staging, and Defensive Campaign Deployment.

CONCLUSION

The network defender community began their work in the 1990s by trying to prevent, or at least, detect, the tools that cyber adversaries were using to penetrate their networks. That was short-sighted. Instead of trying to stop individual tools used with no context about what the adversary was trying to accomplish, we should have been trying to stop the success of the adversary's campaign. The famous 2010 Lockheed Martin white paper on the Cyber (Intrusion) Kill Chain® gave us the means. It advocated for the defeat of the entire adversary's campaign by deploying prevention and detection controls at every stage of the attack sequence. Currently, the commercial security vendor community believes there are fewer than 250 active campaigns at any one time, which is not a large problem space. What has prevented the success of the intrusion kill chain strategy is a standard framework to collect the intelligence associated with specific adversaries, to share and consume that standardized intelligence with trusted partners, and then to automatically process that intelligence and distribute new prevention controls to the network defender's security stack. The adversary playbook is that framework.

NOTES

1. Eric Hutchins, Michael Cloppert, and Rohan Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," Lockheed Martin Corporation, 2010, accessed May 19, 2020, https://lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf.
2. "MITRE ATT&CK® Navigator," MITRE, accessed November 5, 2019, https://mitre-attack.github.io/attack-navigator/enterprise/.
3. John Wunder, "STIX 2.0 Finish Line," MITRE Blog, April 12, 2017, https://www.mitre.org/capabilities/cybersecurity/overview/cybersecurity-blog/stix-20-finish-line.
4. "Introduction to STIX," Oasis, accessed May 19, 2020, https://oasis-open.github.io/cti-documentation/stix/intro.
5. Major Michael Barefield, "Commander's Critical Information Requirements (CCIR): Reality Versus Perception," School of Advanced Military Studies, United States Army Command and General Staff College, Fort Leavenworth, Kansas, 1992-1993, https://apps.dtic.mil/dtic/tr/fulltext/u2/a264509.pdf.
6. Chris Brook, "What is the MITRE ATT&CK Framework?" Data Insider, DigitalGuardian, October 24, 2019, https://digitalguardian.com/blog/what-mitre-attck-framework.
7. "MITRE ATT&CK® Navigator," MITRE.
8. "Our Sharing Model," Cyber Threat Alliance, accessed October 7, 2020, https://www.cyberthreatalliance.org/our-sharing-model/.
9. Conversations with Cyber Threat Alliance (CTA) leadership from 2014 to 2019.
10. Ibid.
11. Playbook Viewer, Unit 42, Palo Alto Networks, accessed November 5, 2019, https://pan-unit42.github.io/playbook_viewer/.
12. "Sharing Timely, Relevant and Actionable Intelligence Since 1999," FS-ISAC, accessed November 5, 2019, https://www.fsisac.com/who-we-are.
13. "MITRE ATT&CK® Navigator," MITRE.
14. Brook, "What is the MITRE ATT&CK Framework?"
15. "Products," OASIS, accessed November 5, 2019, https://wiki.oasis-open.org/cti/Products.
16. Barefield, "Commander's Critical Information Requirements (CCIR).
17. "Cybersecurity Canon: A Rock and Roll Hall of Fame for Cybersecurity Books," sponsored by Palo Alto Networks, accessed December 20, 2019, https://cybercanon.paloaltonetworks.com/.
18. Mark Miller, "The Unicorn Project with Gene Kim—a transcription of the DevSecOps Podcast, recorded October 16, 2019," The DevSecOps Podcast Series, October 16, 2019, accessed December 20, 2019, https://cdn2.hubspot.net/hubfs/4132678/DSO%20Days%20-%20Transcriptions/The%20Unicorn%20Project%20with%20Gene%20Kim%20-%20DevSecOps%20Podcast%20Series%20Transcription.pdf.
19. "MITRE ATT&CK® Navigator," MITRE.
20. Conversations with CTA leadership, 2014–2019.
21. "Our Sharing Model," CTA.
22. "MITRE ATT&CK® Navigator," MITRE.
23. "MITRE ATT&CK®," MITRE, accessed December 20, 2019, https://attack.mitre.org/.
24. "OASIS Cyber Threat Intelligence (CTI) TC," by OASIS, accessed December 20, 2019, https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=cti.