CHAPTER

1

What Is an Active Defender?

Chinese general Sun Tzu is perhaps best known for his fifth century BC seminal treatise, *The Art of War*, in which he presents the basic principles of warfare and explains to military leaders how to fight. In particular, the text emphasizes ways to outsmart one's opponent rather than resorting to physical combat, and as a means of ensuring victory if a battle ensues. As Sun Tzu said, "To know your enemy, you must become your enemy." In other words, you must do more than have a superficial understanding of your enemies' behaviors and tactics: you must deeply learn their motivations, thought processes, problem-solving approaches, and patterns as well. Likewise, in the *Republic*, Plato said that "the one who is most able to guard against disease is also most able to produce it unnoticed." Conversely, one who is able to embrace and embody the hacker mindset enough to take appropriate action to contain and excise attackers, as well as better prevent intrusion in the first place, can more easily grasp the nature of the risks and threats in their environments. Acquiring this deep understanding is what transforms a traditional defender into an Active Defender.

The Hacker Mindset

The Active Defender is someone who is aware of and actively engaged with the offensive security community in order to cultivate the hacker mindset.

But what is the hacker mindset? First, let me address the term *hacker* and then we can address the mindset they cultivate.

When I use the word *hacker* I am referring to the original use of the term coined at the Tech Model Railroad Club of MIT in the 1950s: someone who applies ingenuity to create a clever result, called a *hack*.[i] While the media has taken to using hacker in an exclusively pejorative way, tying it to someone who breaks into computer networks to steal or vandalize, nothing about this original definition implies anything illegal or immoral. I will not spend time rehashing the debate over this word, but simply say that this is what I mean when I use it here. In contrast, the word *attacker* is what will refer to someone who breaks into computer networks for the purposes of theft or to cause damage.

With the term *hacker* defined, we can now move to a discussion of the hacker mindset. Generally speaking, this mentality is one of curiosity, creativity, patience, persistence, agility, and nonlinear thinking. Hackers by their nature are curious individuals and believe strongly in freely sharing information. At a young age, they're often the ones who take apart their toys to see how they work. They love to experiment and make something do things it was never intended to do, just because they can. Most people ask why? A hacker is more likely to ask why not? When problem solving, they frequently consider the unintended or overlooked properties of a thing or situation, leading to a new and unique solution only a hacker could produce. As avid security researcher Casey Smith noted in a recent presentation, when paraphrasing a quote by Jon Erickson, "Hackers get their edge from knowing how things really work."[ii]

For example, some of the earliest technological hackers, known as phreakers because they were focused on phone networks, spent hours reading obscure telephone company technical documentation and listening to the patterns of tones used in these networks to figure out how calls were routed. This research led to a full understanding of how telephone networks operated, and as a result, these hackers determined how to make free long-distance calls. After some experimentation, they realized that by producing a 2,600 hertz tone, they could manipulate the phone line into thinking it was idle, thereby allowing a free long-distance call to be placed. This tone could be created by someone with perfect pitch whistling or by a small plastic toy whistle that happened to be a prize found in boxes of Cap'n Crunch cereal at the time. Later, tone-generating devices called blue boxes, which could be readily assembled from some parts at Radio Shack, made this process even easier.

Hackers approach problems in unique and creative ways, often thinking outside the box and challenging assumptions about how things work. When tackling a new problem, they may not focus on a particular goal but rather just start pulling at a thread to see where it goes. Even when there is a goal, the journey, as they say, may be far more interesting to a hacker than the destination. They are not afraid to get lost in tangents and are known to follow the rabbit

hole to wherever it leads regardless of any failures they might experience. For every success a hacker has, they likely have a multitude of failures from either continuously trying new techniques or trying new combinations of existing ones. They often have an enhanced way of looking at the world and can make connections others cannot, often by asking one simple question: what if? To view the world in this way requires literally questioning everything you know or think you know in order to defeat confirmation bias. In other words, to prevent interpreting new evidence as just a confirmation of one's existing beliefs or theories requires this level of skeptical analysis.

Perhaps the most important element of the hacker ethos is the fact that hackers investigate and explore simply because of their enthusiasm for learning and the love of the investigative journey. As a result, hackers love to share what they're learning and are willing to spend copious amounts of time doing so. In other security communities, whether in defense or offense, this joy, discovery, and inquisitiveness often becomes subservient to the duties and necessity of the job. Unfortunately, that not only leads these practitioners to burn out and struggle harder to get through each day but also be less inclined to share information or teach others what they know.

Traditional Defender Mindset

One critical problem with traditional defenders is that they don't typically think in this way. As John Lambert, Distinguished Engineer, Microsoft Threat Intelligence Center says, "Defenders think in lists. Attackers think in graphs. As long as this is true, attackers win."[iii]

While hackers and attackers are definitely not the same thing, the hacker mindset is certainly more in line with how attackers think. The difference, however, is one of intent. The hacker, as already described, is curious and wants to understand the possibilities, while the attacker wants to use these possibilities to cause harm.

Regardless, whether they realize it or not, defenders often focus on lists as part of what are considered best practices. For example, the CIS Critical Security Controls, NIST standards, and PCI controls are just some of the lists they often are required to follow, whether for compliance reasons or simply because they've been taught these controls are expected as industry standards.

The problem with thinking in lists is that, broadly speaking, it is linear thinking. Consider an example where a defender is tasked with hardening a web server. The defender can follow all the best practices, such as documenting an inventory of what hardware and software is installed, scanning the machine for vulnerabilities and mitigating anything discovered, running the latest fully patched firmware and operating system, locking the system down to only communicate

over port 443, and even restricting what and who has access to the machine both locally and over the network. At the end of the day, the defender believes they have done everything they can do to protect this system and considers it as secure as possible. They move on. The attacker, reflecting on this same system, begins to wonder what they could do with that open port that perhaps it was not meant to do. Someone with the hacker mindset is more likely to recognize the possibilities than the traditional defender.

Even the SANS incident response process, where you have separate, discrete phases of preparation, investigation, containment, eradication, remediation, and lessons learned, is an example of this limited type of thinking. The investigative approach they provide is a linear process in that it's more of a flowchart walk-through rather than having the investigator's actions informed by a deep understanding of the activity observed and its implications as well as an understanding of attacker mindset, goals, and anticipated moves.

In contrast, thinking in graphs is holographic in nature. In other words, it draws connections among disparate data points and uses those connections to inform the attackers actions in wise fashion at a holistic and instinctual level. Each data point informs and enriches all the others such that what you wind up with is essentially a fancy mesh network of information rather than a traditional flat Ethernet network, thus making it more efficient and resilient.

Of course, that's not to say that the process SANS provides is ultimately not useful—it absolutely is. However, I'd argue that it's only meant to be a starting point, providing a sense of the components that need to be considered. An investigation typically begins as the result of some form of initial activity that causes an alert to fire or a defender to become suspicious. With this preliminary evidence as a starting point, you should be looking for context, which in turn will typically result in additional evidence. As a result, there are many paths an investigation could take. However, only someone who cultivates the hacker mindset and has the right experience will be able to efficiently uncover the graph an attacker left behind. Otherwise, they remain stuck, unable to connect the dots.

Another way of explaining this idea is that as a defender trying to piece together the narrative of what an attacker did on a system, you typically have no insight into what the process looked like for the attacker. You look for artifacts that point to activities, and you generally assume those activities were done with intention and were successful unless you see reason to believe otherwise. However, someone who spends enough time within the offensive security space to begin to cultivate the hacker mindset understands that, despite their best efforts, attempts made by attackers often fail in unexpected ways.

Getting from Here to There

As a traditional defender becomes more engaged with the offensive security community, they gain emergent insight into how hackers, and consequently

attackers, operate. It opens the door to a world that they didn't know exists, and, if they remain involved, can lead them to becoming Active Defenders. However, unlocking that knowledge requires making a choice.

Long before Neo was asked to choose between the blue and the red pill in the Matrix and saw just how far down the rabbit hole goes, the Greek philosopher Plato had his famous cave. This emblem of a theory of knowledge highlights the blind spots we all possess and parallels the choice we are forced to make. The "Allegory of the Cave" tells the story of a group of prisoners chained together inside a cave facing a blank wall their entire lives. They are unable to turn their heads. Behind the prisoners is a fire. Between the fire and the prisoners is a raised walkway with a low wall. As a result, anyone or anything that moves along the walkway appears as shadows on the wall the prisoners are facing. The prisoners perceive these shadows to be all the things that exist in the world. It is, of course, the only world they know. The story further considers what would happen if one of these prisoners were to escape their bonds, discover the real world that lies beyond those shadows, and then try to return to liberate their former cave-mates. It posits that the rest of the prisoners would reject the knowledge of this "reality," choosing instead to stay in the safety and darkness of the world as they know it.

As defenders, we cannot afford to stay in that darkness if we want to be increasingly effective. We are only as good as what we know—and until we can shine a light on our own blind spots, we cannot begin taking the first steps in breaking the shackles of passivity and transition toward becoming an Active Defender. By continuing to spend time immersed in the offensive security community, we can begin to better recognize the connections into and within our networks that make a variety of attacks possible. Consequently, this understanding facilitates attacker behavior discovery, which is a key part of what makes the Active Defender special. It is key because it is what peels away the blinders and exposes the previously unknown knowledge gaps.

Being an Active Defender requires more than the ability to break through the darkness. They also must understand the difference between the ability to "see" and the ability to "observe" the various pieces of information that they discover, whether in log files, security tools, or incoming intelligence. Simply "seeing" something is a passive function, whereas "observing" something is an active function. This ability is part of the hacker mindset and is crucial to making the connections often overlooked by traditional defenders.

No less illustrious a detective than Sherlock Holmes himself, Arthur Conan Doyle's master of deduction, hints at this difference in "A Scandal in Bohemia." John Watson expresses amazement with how the results of Holmes's investigation seem plain and simple when explained but seem to elude him, despite the fact that they both see through perfectly capable eyes. Holmes gently admonishes him, saying of his friend, "You see, but you do not observe." Both men have trudged up the front steps of 221B Baker Street countless times, and have

seen the same steps, but only Holmes knows that the steps are 17 in number, because he has "both seen and observed."

To be clear, though, what Holmes has in mind does not merely equate the idea of observation with memorization of what has been seen. Instead, it involves consciously ingesting the data presented and recognizing the relevant pieces as actual information that can provide context. It is important to understand that data by itself is unstructured with no real meaning. It's not until that data is coupled with context that it becomes information. In other words, what were once simply points of data, with the appropriate context, become organized, analyzed, and interpreted to be useful information. Thus, Active Defenders must not merely uncover the correct artifacts and see them; they must also draw upon their understanding of the artifacts themselves, the context in which the artifacts were found, their knowledge of attacker motivation and behavior, and countless other pieces of information in order to not merely see the data but to observe its significance.

Active Defenders take the joyful and creative exploration of the hacker mindset and apply it to their defensive roles. By their nature, they have a proactive security ethos. In other words, it is in their character to actively seek out ways for their organizations to become more secure, whether by traditional means such as network segmentation or more innovative ways such as taking extra steps to understand how specific attack vectors work.

They view the notion of "security" as a verb, something that they actively work at doing or being rather than a destination or an ultimate goal to achieve. They understand that an organization can never truly be 100 percent secure but that being situationally aware and engaged with the offensive community can help improve their security posture. Furthermore, they recognize that excellence in security is not a matter of achieving a particular and static level of mastery but rather requires regularly practicing this proactive behavior.

What I mean by excellence here comes from the Greek philosopher Aristotle. He believed that excellence consists of being the best possible version of whatever you are, which means you are good at doing whatever it is you are supposed to do. An excellent knife is good at cutting, and an excellent fish fork is good at removing fish meat from small bones. For humans, excellence is more complicated than performing one simple function, but it involves being the best person you can be. This requires both the moral character to see the wisest course of action in any given situation and the technical knowledge to translate that vision into effective action. Similarly, in security, being excellent requires both the wisdom to discern the most appropriate path to take and the technical knowledge to get us there.

Active Defender Activities

First and foremost, as I've said, the Active Defender engages with the offensive security community. They do so in a variety of ways, both in person and online, such as by attending local or national conferences, going to meetups, or participating in forums and on social media. Chapter 2, "Immersion into the Hacker Mindset," will explore some of these options and how to get involved with them. However, they also participate in a number of proactive activities that make them stronger defenders, such as threat modeling, threat hunting, and attack simulations.

Threat Modeling

The Active Defender regularly participates in threat modeling in terms of both the business they are protecting and the technology it employs. Threat modeling, at a very basic level, is a proactive, risk-based evaluation of what attacks are likely or possible for a given person, environment, or situation with an eye toward evaluating particular scenarios. It requires considering what you or your organization has that is of value to a potential attacker. Where these elements are cyber assets most critical to the accomplishment of the organization's mission, they're often referred to as the "crown jewels."[iv]

Threats include anything that has the potential to cause harm (whether reputational or technical), do damage, or in the case of a company, interrupt business. Risk involves the likelihood that a threat could be successful.

Threat modeling also involves considering what vulnerabilities, or weaknesses, a threat can use against you or your organization. Attackers never act without motive, and it is impossible to separate the attack from what it is the attacker is hoping to gain. Therefore, understanding what a threat actor might hope to gain can help you identify where the greatest need for security is in your organization. It can also help determine, to some extent, the path that an attacker might take after landing on a system.

Typically, formal threat modeling occurs at the beginning of a project, whether building new software or deploying a new system to prevent introducing new risks into an organization. Several formal methodologies are available to use for performing an evaluation of this nature. For example, Microsoft offers a model called STRIDE, which is an acronym that stands for Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege. STRIDE specifically focuses on finding potential threats against software.[v]

CVSS, the Common Vulnerability Scoring System, is another well-known model created by NIST (National Institute of Standards and Technology) that provides a numerical score for the severity of vulnerabilities. Furthermore, many

books and articles have been written about the topic of threat modeling, mostly intended for an audience of people who create software, such as developers, as well as those responsible for building operational systems, such as analysts or architects. Most recently, a working group of security and privacy professionals came together to develop a methodology-neutral guide called the Threat Modeling Manifesto.[vi] They wanted a way to share their collective threat modeling knowledge to "inform, educate, and inspire other practitioners to adopt threat modeling as well as improve security and privacy during development."[vii]

Although these formal approaches exist, the reality is that each of us threat models regularly, just without referring to it with that name. Anytime we do something involving potential threats, such as traveling or even crossing the street, we consider the threats we might face. For example, when planning a shopping trip for the Friday after Thanksgiving, which is commonly known as Black Friday here in the United States, we must consider several threats such as the possibility of being trampled at large box stores, the potential for vehicle vandalism, potential loss from theft, and the amount of time spent waiting in lines versus opportunities for theft. Contemplating each of these concerns in advance may cause us to take different actions than we would have otherwise. Therefore, I believe the idea of threat modeling can be extended to be useful beyond developers and architects to also include all practitioners of defensive security as we've already defined it, including system administrators, network administrators, and other security professionals. In Chapter 5, "Implementing Defense While Thinking Like a Hacker," we'll examine this topic further and discuss what threat modeling as an Active Defender with the hacker mindset engaged looks like.

Threat Hunting

The Active Defender also may engage in threat hunting regularly. Threat hunting is a human-centric activity by defenders who proactively search through data to find threats that evade traditional security measures.[viii] It forces them to look beyond the alerts generated by SIEMs and other tools, which, as we've already discussed, can lead to tunnel vision. As a result, it is possible to find previously undetected threats, thereby reducing dwell time and limit attacker persistence.

Dwell time is the period between when a system gets compromised and when the compromise is detected. Threat hunting can also help uncover previously overlooked misconfigurations, detect poor user/administrator practices, reduce the overall attack surface, and identify gaps in existing defenses. Ideally, with this information documented, these issues can be resolved. Furthermore, by identifying these threats, the defender may even be able to adjust existing detections or create new ones for their environment. We'll discuss the topic of building effective detections in Chapter 7.

The process used for hunting should be one that is readily repeatable and not random. Furthermore, despite what some may claim, threat hunting does not require the defender to have a significant amount of experience or for the business to have a high level of security maturity. Even a defender who is entry level can threat hunt and do so regardless of the sophistication of the organization's security program. Almost all organizations will have some logs or other places to hunt, even if they're not the ideal central logging repositories of advanced security programs.

Threat hunting does require having at least some knowledge about the environment. However, one of the hidden advantages to engaging in threat hunting is that this activity ultimately helps increase the defender's knowledge of the network and the datasets available. As a result, their responses are typically more accurate and response times are reduced. We'll explore what threat hunting looks like as an Active Defender in Chapter 5. Keep in mind, though, that what gets uncovered during a threat hunt can lead to the need for incident response. The defender should be prepared to take the necessary steps to transition to the appropriate person or team if the need arises.

Attack Simulations

Attack simulations are another area that the Active Defender pursues whenever feasible. As the name implies, attack simulations attempt to create actions and behaviors that mimic what an attacker might be expected to do against the environment being defended. While not a real attack, it allows the defender to get a feel for what a particular type of attack might look like against their environment and whether any detections for it alert as expected.

Utilizing attack simulations allows the defender to continuously test these detections, develop new ones, tune configurations, and compare security products to get a feel for what is more effective in a particular environment against different kinds of attacks. There are a number of different options for employing attack simulations, including both free software such as Atomic Red Team, offered by Red Canary, and paid software such as the SCYTHE adversary emulation platform.

Active Defense

Now that you're starting to get a feeling for what an Active Defender is and what they do, I want to explore the distinction between the type of active defense they practice compared to that proposed by other entities.

"Active Defense" for the Active Defender

The origin of the phrase *active defense*, which is briefly mentioned in the introduction in conjunction with Robert M. Lee's white paper, comes from a paper about the 1973 Arab/Israel war written by US Army General William E. DePuy. It was used to describe the necessity for ground troops to be able to be mobile rather than remain static while defending their position. Specifically, he stated that the purpose of active defense is to "wear down the attacker by confronting him successively and continuously with strong combined arms teams and task forces fighting from mutually supported battle positions in depth throughout the battle area."[ix]

Unfortunately, despite being added to an Army Field Manual in 1976, debate continued about what the term really meant. In 2015 the Department of Defense adopted *active defense* to mean "the employment of limited offensive action and counterattacks to deny a contested area or position to the enemy" in the context of traditional warfare, which further confused matters. As a result, use of the term *counterattacks* in this definition was taken to literally mean "hack-back" with regard to the cybersecurity space.

However, even in the traditional warfare notion of the term "active defense," the word *counterattack* was only meant to occur inside the defended area. For example, the military provides what is known as force protection to safeguard DoD personnel within their defensive zone, including their families, resources, and critical information.[x] These preventative mitigations do not extend beyond the border of protection. In other words, defenders can take proactive steps for protection within their space but not step across this border to go after an attacker.

A similar example from popular culture, for those unfamiliar with military concepts, can be found in the movie *Home Alone*. While the main character sets up a variety of booby traps inside his house to deter a pair of burglars from robbing it, he does not pursue the burglars beyond the entrance to the home itself. Instead, as Rob Lee states, "counterattack in cybersecurity would be more properly reflected in the concept of incident response where personnel 'counterattack' by containing and remediating a threat."[xi]

Thus, rather than "hacking back," Lee's description of active defense for cybersecurity involves the original intent of the term *active defense*, with a focus on the flexibility, intelligence acquired, and application of this intelligence within their networks and, as explained earlier, is what I take to be at the heart of the Active Defender's actions.

Another Take on Active Defense

In contrast, John Strand defines active defense as "the employment of limited offensive action and counterattacks to deny a contested area or position to the enemy."[xii] He describes the actions that make up active defense as being proactive, anticipatory, and reactionary against attackers. In his book *Offensive*

Countermeasures: The Art of Active Defense, Strand outlines three active defense categories: annoyance, attribution, and attack.[xiii]

Annoyance

Annoyance, which is also known as "cyber deception" in certain organizations, involves finding ways to make an attack more time consuming, such as making a website appear like something else, providing nonexistent DNS records, or including a random URL generator on a website. In addition, he recommends the use of honeypots. A honeypot is a decoy that consists of data or a service that looks legitimate to an attacker but is meant to detect and deflect their efforts, thereby increasing the efforts of an attacker by attacking systems that are not real.

Attribution

Attribution entails attempting to figure out who is attacking and, ideally, how to track any intellectual property that has been stolen by an attacker.

Attack

The category of attack he proposes begins similarly to the traditional warfare concept of counterattack that happens within the defender's boundaries. It is done by drawing the attacker into the systems being defended and encouraging them to steal intellectual property and take it back to their own systems. However, this last step is where what Strand has in mind diverges significantly from that meaning of counterattack. What the attacker does not know is that by stealing this data, they've taken more than what they intended. The files contain malware and, as a result, allow defenders to gain access to the attacker's system.

He does caution that while these techniques can be extremely effective, they do require significant caution, preparation, and potential coordination with legal counsel and law enforcement. I agree that there can be some merit to his concept of counterattacks in very specific circumstances. However, the form of active defense practiced by the Active Defender I have in mind would predominantly stick to the strategies he offers within the categories of deception and attribution. The Advanced Active Defender, who will be discussed in a subsequent chapter, might embrace the attack strategies he describes, but only (as Strand himself recommends) with blessings from management and legal and even coordination with law enforcement where appropriate.

Active Defense According to Security Vendors

Several well-known security vendors also offer devices or software that they claim use some form of "active defense," but they often provide limited insight as to what is meant by that term. Fortinet, a cybersecurity company that sells

physical firewalls, antivirus software, intrusion prevention systems, and endpoint security components, defines active defense as "the use of offensive tactics to outsmart or slow down a hacker and make cyberattacks more difficult to carry out."[xiv] They claim that it involves using deception technology that misdirects attackers, causing them to waste time and processing power and, although it can sometimes involve striking back at the attacker, they assert that this action should be reserved for law enforcement agencies that have the authority to do so.

Similarly, Illusive offers a product called its Active Defense Suite. This product claims to provide visibility into lateral movement by "creating a hostile environment and accelerating the time to detection for an attacker that has established a beachhead."[xv]

Acalvio Technologies claims to be "a pioneer in Active Defense strategies" and embraces something it calls "deception-based detection" that falls slightly outside of the description.[xvi] Founded in 2015, it describes its product Shadowplex as an alternative to passive defense. Fundamentally, it claims the product detects attacks; impedes attacks through obfuscation and deception; gathers intelligence through an attacker's tactics, techniques, and procedures; and remediates assets.

The problem with all of these technologies is that they are each automated products, not taking the context of the environment they are protecting into account. While it is possible these tools could be helpful in the defense process, at the end of the day true active defense requires human intervention to provide and apply this context.

Active > Passive

Active defense has several significant advantages over passive defense. Being passive, as we saw earlier, implies a distinct lack of involvement or direct interaction on the part of the defender. It can cause tunnel vision from strictly focusing on the information from certain systems like a SIEM and alert fatigue from the constant barrage of alerts and blinking lights as well as burnout from insufficient staffing.

Being active, on the other hand, is an interactive practice that provides insight into the effectiveness of existing defenses, the potential for new defense capabilities, and an ongoing cycle of responding to threats internal to the network, learning from them, and then applying the knowledge gained across the environment, providing better detection and prevention where possible. Being active avoids many of these problems because the emphasis is on proactively seeking out threats and looking at the environment holistically rather than focusing on constant alerts.

Active Defense by the Numbers

One set of statistics that is particularly illustrative in showing how active forms of defense outshine their passive counterparts involves threat hunting. According to the SANS 2021 Threat Hunting Survey, organizations saw a 10–20 percent improvement in their overall security posture as a direct result of their active threat hunting programs.[xvii] Furthermore, 37 percent of these same organizations saw a marked reduction in their overall attack surface. Forty-three percent of them also saw a significant improvement in the creation of more accurate detections and fewer false positives. Nearly 27 percent saw a significant decrease in breakout time, which is the time between initial compromise and lateral movement, as well as a 25 percent improvement in the ability to detect data exfiltration. They also saw a 29 percent decrease in the resources needed for remediation, such as staff hours and expenses.

While these statistics are specific to the survey from 2021, a similar pattern can be seen for many of the previous surveys. With these kinds of statistics, it's easy to see that the benefits of an active approach to security are undeniable.

Active Defense and Staffing

While it cannot directly address insufficient staffing issues, an active defensive environment is one in which people are more likely to want to work for several reasons. An organization that encourages proactive security practices typically has leadership that will fight for the appropriate resources and staffing necessary to do the job properly, not to mention allowing flextime for research and encouraging real-time information sharing. As a result, it is not uncommon for security practices in these spaces to be more efficient, thereby potentially needing fewer people. For example, we've already discussed that threat hunting can reduce the time to discovery of a breach and reduce dwell time as well as increase speed and accuracy during incident response.

Active Defender > Passive Defender

There are a number of different reasons why the Active Defender is better at defending their organizations than their traditional defender counterparts. Let's examine a few of them now.

Relevant Intel Recognition

Because the Active Defender is engaged at this level, they have a much greater ability to recognize reliable, relevant intel when it is discovered directly or provided to them by an outside entity. Defenders regularly receive intelligence

from a variety of sources, including but not limited to social media, federal agencies (FBI, Secret Service, CISA), vendors, ISACs (information sharing and analysis centers), and scanning services like Shadowserver. This intel might be in the form of a data feed, an email, or a report.

The problem is that not all threat intelligence is created equally. The glut of information being provided is often out of date, non-actionable, or not relevant to a particular environment. However, in a proactive environment where Active Defenders are regularly engaging with the data being collected by their tools and actively looking for signs of attack, they will typically recognize what portion of that intel is relevant to their environments. Furthermore, they may also be able to articulate to other areas within their organization to explain how and why what they found matters from a business perspective.

Understanding Existing Threats

Generally speaking, because they are engaging with, learning from, and potentially training with offensive security professionals, Active Defenders also have a better understanding of existing threats. The majority of traditional attacks use the same well-known exploits, which can also be used by the offensive security community for testing purposes.[xviii] Therefore, the Active Defender engaging with this community is likely to be familiar with them as well.

Furthermore, by understanding how legitimate services in their networks can be leveraged by attackers for nefarious purposes, Active Defenders are in a stronger position to help their organizations. For example, PowerShell, which is a cross-platform framework that contains a command-line shell, scripting language, and configuration management capabilities, is used regularly by defenders to automate tasks, access certain settings, and configure systems. Unfortunately, the ubiquity of this framework makes it particularly easy for attackers to use it as well, hiding their activities in legitimate traffic. This technique is called a living off the land attack because it specifically abuses legitimate services that are native to the operating system or the user environment. We'll see more examples of how living off the land attacks can be used in Chapter 5.

Attacker Behavior

Most importantly, because of their relationship with the offensive security community, as we've already mentioned, the Active Defender can more easily discover attacker behavior. They have insight into what it means for an attacker to live off the land, what kind of tools they're likely to need and upload for post-exploitation, and what legitimate things on the network can be leveraged in malicious ways.

When alerts do fire that make references to tools well-known to be used by attackers, or they observe certain patterns of malicious behavior, the Active Defender will recognize the potential for real risk rather than mentally lumping these alerts in with all the others they receive. For example, because they are immersed in the offensive security community, Active Defenders are also more likely to recognize any similar tactics, techniques, and procedures (known as TTPs) used in offensive security engagements. In order to understand why this matters, let us explore how TTPs fit into a defender's detection capabilities.

Pyramid of Pain

In 2013, security professional David J. Bianco created the Pyramid of Pain, which ranks the indicators used to detect an attacker's activity by the "pain" it would cause them if they are discovered by the defender.[xix] The higher the indicator, the more pain for the attacker, because these higher indicators are more difficult for the attacker to alter and still continue their nefarious activity. Likewise, the higher the level, the more challenging it is for the defender to uncover and block an attacker. For example, at the bottom of the pyramid are hash values, because it is trivial for an attacker to change the hash value of a file but fairly easy for the defender to detect. At the very top are the TTPs, because these steps are how the adversary actually accomplishes their goals, which requires the defender to understand attacker behaviors, not just the tools they use. Obviously that also means that detecting these TTPs is the most difficult challenge for the defender.

MITRE ATT&CK

The MITRE ATT&CK framework built on this idea in 2015, offering a way to catalog the tactics and techniques used by attackers as well as more specific techniques known as sub-techniques, such that it became possible to strategically track specific groups of attackers.[xx] In particular, understanding an attacker's technique of choice can be useful because it provides the types of data a defender needs for detection purposes.

TTP Pyramid

Taking this idea a step further, in 2022 Christopher Peacock breaks down the top of the original Pyramid of Pain by further separating tactics and techniques from procedures in something he calls the TTP Pyramid.[xxi] He explores each of these levels, explaining how they differ. Tactics are the bottom of this pyramid and include what the attacker is strategically attempting to achieve, such as credential access.

While useful to know, this information provides little actionable detail. At the next level we find techniques, which are at the heart of the ATT&CK framework previously described. They provide an overall insight into the tactical goal of the procedure in question—for example, doing credential dumping via LSASS memory. However, what the techniques by themselves do not provide is exactly how the attacker meets this objective. There are a number of methods to dump LSASS memory, such as by using tools such as Mimikatz or ProcDump among others. The specific way in which the adversary carries out their techniques is known as a procedure and is found at the top of the TTP Pyramid.

To further clarify the idea in nontechnical terms, we can think of the tactics as any desired outcome. Take, for example, climbing a mountain. The technique is the general plan for achieving that outcome. In this case, that might involve putting on the appropriate gear and moving upward in elevation. The procedure is specifically how one goes about implementing this technique. For example, in terms of putting on the gear, the procedure might include putting the right foot into the right boot, tying laces in a counterclockwise bow knot, and then repeating that process with the left.

Toward a Deeper Understanding

If a defender truly hopes to build effective detections, they need to understand not only the technique but the procedures being used during an attack. While traditional defenders are often familiar with the tactics used by attackers, they typically only have limited familiarity with attacker techniques. While possible, it is even less common for them to be acquainted with the various procedures the attackers use.

By virtue of their immersion in the offensive security community, the Active Defender is not only aware of more of these techniques, but understands them and how they work as well as the distinction between techniques and procedures. The Advanced Active Defender may even have a working knowledge of these procedures. This knowledge is a key feature of the Active Defender, making them significantly more valuable than their passive counterparts.

Return to the Beginning

Let's return now to the scenario we saw at the beginning of the introduction in which the attacker sends a phishing message that leads to the download and detonation of malicious software. Once the user unknowingly launches the malware in question, it pops up a cloned system authentication dialog box that asks for their username and password and relays them back to the attacker. At the same time, the malware calls back to the command-and-control (C2) framework that

the attacker is using such as Cobalt Strike or Metasploit. Using that connection, the attacker opens a command shell on the victim's computer and attempts to run the password stealing software Mimikatz to gather additional passwords, but the antivirus software stops them in their tracks. Completely undaunted, they instead choose to use `rundll32.exe` with `comsvcs.dll` to dump memory and capture passwords.

The attackers do some basic reconnaissance next to determine what legitimate processes running on the local system can be used to provide camouflage, such as drivers, OneDrive, or `explorer.exe`. The selection might be based on the legitimacy of the process, communication patterns, and/or stability. Once this decision is made, the attacker injects the same malware into this process, which helps their activity blend into the normal traffic and behavior of the local machine.

The attacker focuses next on initial network reconnaissance and finds a terminal server on the network. They are able log into it using the initial credentials stolen from the user. Once logged in, the attacker downloads and executes additional malware, giving them a second C2 via a different domain. Although this new malware triggers a Network Security Monitoring (NSM) alert, it does not prevent the malware from running. As a result, the attacker was able to exfiltrate passwords and other company data stored on the terminal server via the second C2 environment they've established.

The security analysts do a good job of initially responding to the case. Because they had a basic understanding of well-known attack tactics, such as gaining initial access, and techniques, such as phishing, they went about deleting the phishing message and looking for others. They also noticed the malicious traffic coming from the initial C2 domain used by the attacker and took appropriate steps to block it.

Unfortunately, because they are traditional defenders, they made some common mistakes and missed some key evidence. For example, they were unaware that there was a pivot to a second machine because they reimaged the initially infected system to expunge the attacker from it without further investigation.

In addition, although alerts for the attempted Mimikatz install were observed, they assumed because it was blocked, no further investigation was warranted. While an NSM alert did fire against the second machine, indicating that elements of Cobalt Strike were being used, the analysts thought it was a false positive. Furthermore, because the attacker switched C2 domains and used long sleep times, the additional C2 was missed—and so the intrusion on the second computer was never detected.

Being unfamiliar with some of the attacker techniques or procedures, they also did not recognize the connection between the attempt to use Mimikatz and the framework of Cobalt Strike, which might have led them to uncover the second C2 implementation. Furthermore, because there is no visibility into

file exfiltration, and since the file was downloaded over the C2 channel, there aren't any artifacts to determine what files might have been taken. By the time the initial user's credentials were rotated, the attacker had already compromised other accounts and systems, working their way toward Domain Admin.

Had there been an Active Defender in the house who understands how attackers operate and could recognize more of the techniques and procedures being used, they would most likely have known that attackers commonly attack one system with the goal to pivot to others. Therefore, they would not only know what techniques to look for that indicated evidence of lateral movement but be able to start actively hunting for it.

In addition, the Active Defender would expect that just because the procedure of running Mimikatz did not work for the technique of credential dumping, the attacker would likely try another procedure. As a result, they should be looking for evidence of other procedures that could achieve the same goal. They would also examine other systems for possible related strange behavior, such as unusual process execution or additional services added. They might also know to seek out potential evidence of additional C2 channels or evidence of data exfiltration.

Summary

The Active Defender, through their immersion in the offensive security community, is someone who cultivates the hacker mindset. They use this mindset to think differently about attacks and attackers, focusing on relationships and connections rather than lists. We've seen how the "active" part of Active Defender involves an agility in monitoring, responding to, and learning from the threats on their networks rather than some form of attacking back as well as why being active is better than being passive. Furthermore, we've discussed how Active Defenders often participate in proactive activities such as threat modeling, threat hunting, and attack simulations. Finally, we discussed what an Active Defender is likely to understand compared to their traditional counterparts and how that could be useful in a particular scenario. Now let's turn to Chapter 2, where we'll begin to discuss how to become an Active Defender.

Notes

i. `http://tmrc.mit.edu/hackers-ref.html`

ii. Casey Smith. "Hackers Teaching Hackers keynote," Columbus OH. November 4, 2022.

iii. `https://medium.com/@johnlatwc/defenders-mindset-319854d10aaa`

iv. www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/systems-engineering-for-mission-assurance/crown-jewels-analysis

v. https://docs.microsoft.com/en-us/azure/security/develop/threat-modeling-tool-threats

vi. www.threatmodelingmanifesto.org

vii. Ibid.

viii. Neil Wyler, "Practical Threat Hunting: Straight Facts and Substantial Impacts" presentation, 11/19/20.

ix. Robert M. Lee, *The Sliding Scale of Cyber Security*, 2015. www.sans.org/white-papers/36240

x. www.jcs.mil/Portals/36/Documents/Doctrine/pubs/dictionary.pdf

xi. Ibid.

xii. ACM course slides, 2019.

xiii. J. Strand, *Offensive Countermeasures: The Art of Active Defense. 2nd Ed*. 2017.

xiv. www.fortinet.com/resources/cyberglossary/active-defense

xv. https://go.illusivenetworks.com/hubfs/Platform%20Overview%20-%20Illusive%20Active%20Defense%20Suite.pdf

xvi. www.acalvio.com/us-federal/#:~:text=Active%20Defense%20is%20%E2%80%9CThe%20employment,and%20hoping%20for%20the%20best

xvii. www.sans.org/white-papers/sans-2021-survey-threat-hunting-uncertain-times

xviii. www.bleepingcomputer.com/news/security/90-percent-of-companies-get-attacked-with-three-year-old-vulnerabilities

xix. http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html

xx. https://attack.mitre.org

xxi. www.scythe.io/library/summiting-the-pyramid-of-pain-the-ttp-pyramid