THREAT HUNTING ESSENTIALS:

Part 1: Threat Hunting Defined

THE IMPORTANCE OF DEFINING THREAT HUNTING

The practice of cyber threat hunting continues to generate a great deal of discussion as organizations continue to seek out new ways to enhance their defensive capabilities. As the attack landscape continually grows, many security teams have found that the traditional approach of monitoring and responding to alerts is no longer scalable. Instead, a forward-thinking search for attackers and weaknesses is necessary to prevent compromises from escalating beyond recovery. For as long as humans have engaged in warfare, the strongest armies have not only focused on strong offensive capabilities, but also actively patrolled their defenses looking for signs of attack and methods of improvement. It has long been a common practice to continually assess your security posture for failures and opportunities to advance.

Now, as battlefields continually evolve into the modern cyber terrain we strive so hard to defend, this approach has also evolved with the current wave of cyber threat hunters. Security conscious organizations know that the strongest defenses can no longer position themselves as purely reactive . They must instead seek out the undetected, identifying the unpredictable before an attack can evolve beyond their control. Therefore, it is important that we move beyond the buzzwords and hype and instead set realistic expectations for what a threat hunter is and what they can achieve.

WHAT THREAT HUNTING IS NOT

WHAT THREAT HUNTING IS NOT

Defining threat hunting can at first appear to be a fairly easy task. Using the name itself we can craft a very simplistic definition of a cyber threat hunter as, "one who hunts for cyber threats." The problem with such a basic definition, however, is that it does not set the full parameters for what such an individual is actually capable of nor how their mission should be defined. Therefore, we must expand on this definition and further explain the nuances of such a role and how it fits into a modern security team.

Before we do this though it would seem to be more beneficial when crafting a detailed explanation of what a cyber threat hunter is by first eliminating what a cyber threat hunter is not. This removes the vague ideas surrounding "threat hunting," and instead allows us to set clear expectations on specific goals and skills. In this way we can separate what falls under the threat hunting umbrella and what instead aligns more with other groups.

Threat Hunting Is Not Incident Response

The first function that we can clearly define as not being part of threat hunting is one that it is often confused with – incident responders. It is true that many threat hunters will work closely with incident responders. In some situations, the direct outcome of a hunt is the identification of a breach that an incident response team would need to handle. Other times the outcome of an incident response engagement leads to newly discovered data that in turn feeds into new hunting hypotheses. Threat hunters and incident responders are organizational siblings both working in tandem towards a stronger defense, so it is easy to see how their duties can occasionally be confused. Their core missions however differ greatly, and we must understand this in order to better define a threat hunter.

We must recognize that first and foremost, the role of an incident responder is one that is highly reactive. Responders secure environments through engagements often initiated through the generation of an alert or notification that suspicious activity has been observed. As a result, an incident responder is tasked with responding to these identified threats with the goal of minimizing impact on the organization. This means that they are driven by a mission of business continuity above all else. Therefore, according to the standard PICERL (Preparation, Identification, Containment, Eradication, Recovery, and Lessons Learned) process, once a threat is identified the next steps are to contain and eradicate any malicious actors. Time is of the essence, and in many cases this process is very formalized, leaving little room for experimentation since the primary goal is rooted in returning the organization to a baseline of operation.

Threat hunting is instead much more proactive in its approach to securing an environment. Where incident response is driven by the detection (or evidence) of suspicious activity, threat hunting is derived from the assumption that suspicious activity is occurring but has yet to be detected. In this way it is rooted in hypotheses that must be tested and refined. Where incident responders see time as a luxury, many threat hunters have much more freedom to experiment since they are not constrained by the need to restore business operations. Incident responders also have a clear-cut way of defining what is considered a successful engagement since this can often be answered by asking the question, “Has business returned to standard operation?” If the answer is yes, then the incident response engagement can be seen as a success. Since many hunting engagements begin with a basic theory that can be expanded upon and evolve during the exercise, so too can the outcome of each scenario. Because of this discovery-based approach, what constitutes a successful hunt is then much more fluid in nature.

THREAT HUNTERS ARE NOT A RED TEAM

Penetration Testing, or the similar red team engagements, are the methods of securing an environment through controlled offensive exercises. Operators on these teams attempt to mimic the actions of malicious attackers as a way of testing and validating the security posture of an organization. This, like threat hunting, is a proactive method of identifying weaknesses and one that also generally begins with the question "Is the target suspectable to a compromise?"

It is the follow up questions where the missions of a red team and a threat hunter greatly diverge. Where a red team would follow up with "how would I compromise the target?" they would then formulate a plan to act on this theory. These offensive operators are focused on the "can I" aspect of the hypothesis, attempting to act as an attacker would in order to provide both evidence and recommendation for ways to further improve the security of the organization. A threat hunter would instead tend to ask, "has the target been compromised?" since they are driven by the search for evidence rather than the ability to execute and instead focus on the "have they" aspect of the hypothesis. In some situations, a red team may discover evidence of a compromise, but this is often a byproduct of their primary goals. Inverse to this, a threat hunter will rarely if ever execute an attack scenario in the execution of the goals. Instead, if a threat hunter is looking to recreate a set of attack events perhaps to contrast against a collection of data, they would generally engage a red or purple team to generate these events.

DEFINING THREAT HUNTING

Defining Threat Hunting

Now that we have established that threat hunting is a unique proactive security measure separate from incident response and offensive engagements, we can now set a definition for what threat hunting is. We do so however with a slight sense of caution. This is because threat hunting itself, being hypothesis driven and very open form, tends to be an interesting mix of science, art, and philosophy that requires a unique sense of intuition. Because of this there may always be debate on exactly what constitutes threat hunting. We will work from a generally agreed upon definition, however, and then work towards further breaking that definition down to better understand it.

Hunting Down Advanced Threats: 5 Threat Hunting Tactics that Work

Join our cybersecurity expert, Christopher Fielder, Sr Cybersecurity Strategist, for a discussion on how to begin implementing and developing a threat hunting program for your SOC or Cyber Fusion Center. We will examine threat hunting philosophies and the key tactics that will move your organization from a continually reactive posture to a proactive one that effectively keeps attackers from achieving their mission.

Key takeaways:

- Seek out the enemy using 5 threat hunting tactics
- How to generate a hypothesis-based threat hunt
- Defining the who, what and when and how to develop a threat hunting template

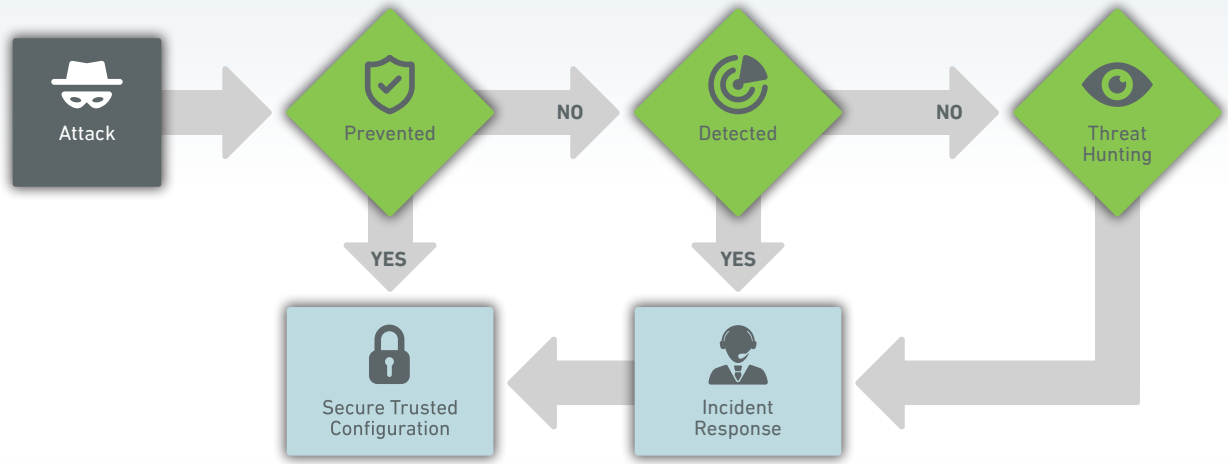

"Threat hunting is the proactive hypothesis driven discovery of artifacts, activity, or detection methods not accounted for in passive monitoring capabilities."

6 KEY ELEMENTS FOR EFFECTIVE THREAT HUNTING

1 PROACTIVE

This is still a relatively simple definition, but one that greatly expands on the primary goals of a hunt while setting specific expectations. The first being that every threat hunt is a proactive discovery exercise. This is an important differentiator since many security practices tend to focus more on detection rather than discovery. Detection is absolutely an important aspect of a security architecture but focusing solely on detections leads to a security posture that is always rushing to clean up incidents after they have occurred rather than a balanced posture where incidents may be identified before they reach their final objectives.

2 HYPOTHESIS DRIVEN

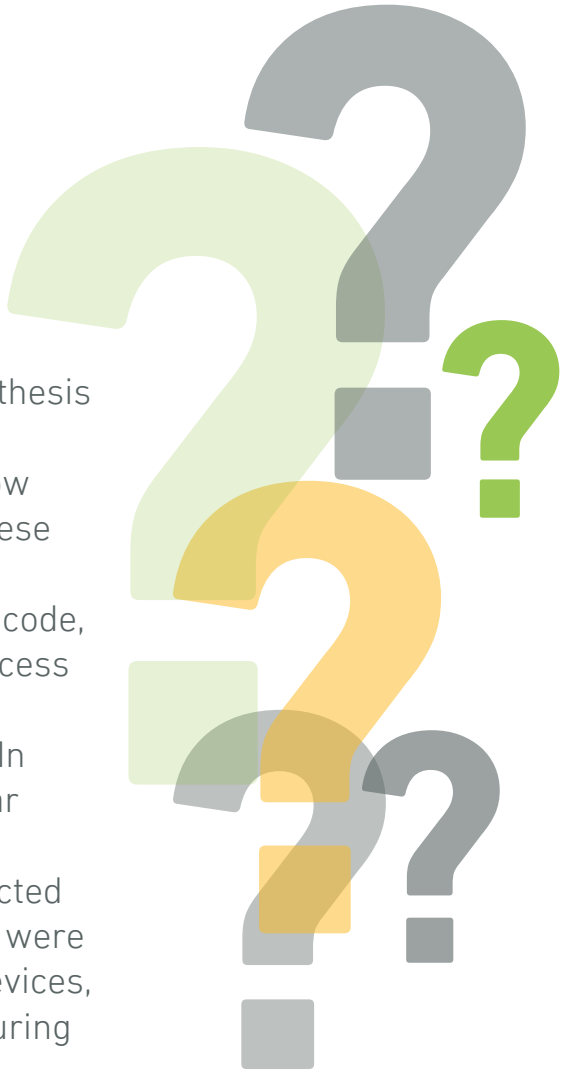

The next important statement in our definition is that threat hunting is primarily hypothesis driven. This means that quite often each hunt will begin with a series of questions or theories. Generalized questions could include, "If I were to attack this environment, how would I do it? What would I attempt to gain access to? What would be my targets?" These types of questions then lead to hypotheses that can be tested. If a hunter believes a particular machine may be targeted, such as an engineer's machine hosting valuable code, they may then begin formulating their theories around how an attacker might gain access to this device. They may begin looking for odd services, unusual network connection, abnormal behaviors, or anything that seems unusual for this device or environment. In many cases the hunt may return no results, but that does not mean that this particular theory was incorrect. These theories should not be seen as rigid because they can be further expanded upon and refined based on evidence gained during the view of collected data. Perhaps no unusual behaviors were found when only the engineering machines were in scope for the hunt. When the same queries where applied to an expanded set of devices, abnormal data was then found. This is an example of how a hypothesis can change during testing.

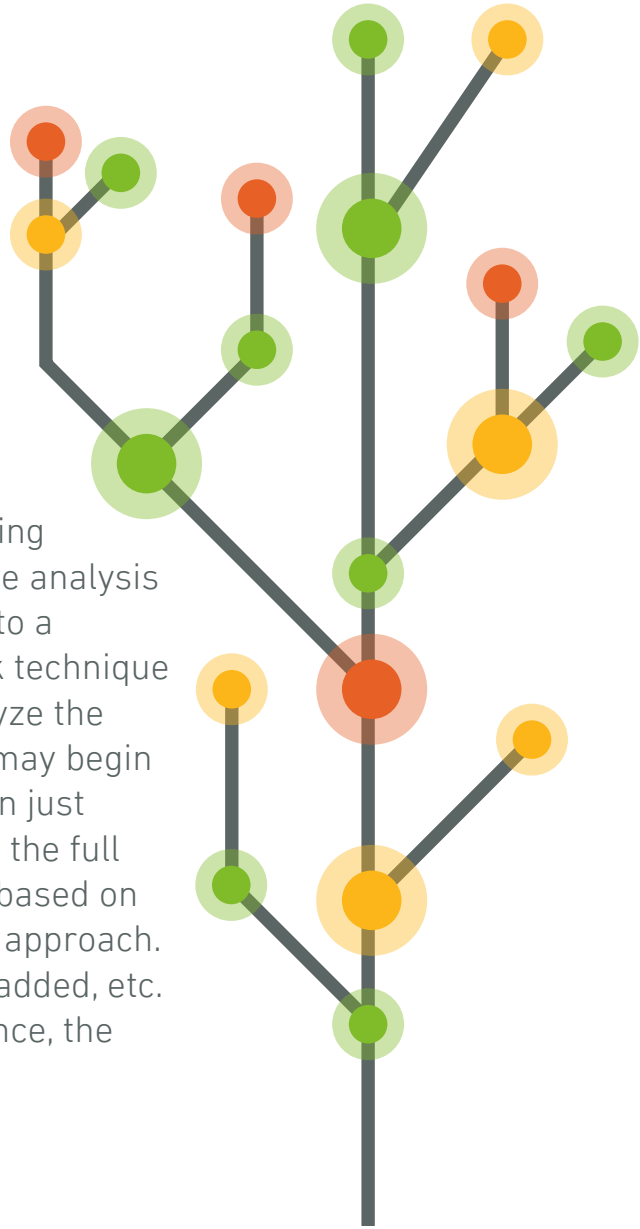

3 A WORD ON RETROSPECTIVE ANALYSIS

Before we go further, it is necessary to point out that although a majority of hunting attempts are driven by the testing of theories, there are cases when retrospective analysis may be what initiates a hunt. This is when newly discovered evidence is applied to a preexisting set of data. For example, the publication of a newly discovered attack technique that was found in another organization's environment. A threat hunter may analyze the process of such an attack and then use this data to drive their search. The hunt may begin by following the publicized steps for the discovered attack in a very linear fashion just to find any exact matches. This strict Point A to Point B testing does not make up the full scope of the hunt though, as the steps may then further evolve. This tends to be based on any suspicions the hunter may have of how different attackers may modify their approach. The initial vector may change, the target could migrate, additional steps may be added, etc. This exercise acts as a further example of how even when driven by static evidence, the actions of a threat hunter should remain fluid and open to experimentation.

4 DISCOVERY

Another key element of threat hunting is that it finds its roots in discovery rather than detection, an element that ties into its proactive nature. These are two terms that on their surface may appear interchangeable even though they certainly are not. Instead, what we can say is that while detections are the key building blocks of security, they cannot exist without the foundation of discovery. For every detection that exists, meaning the passive identification of malicious events (i.e., alerts), there must first have been a discovery of such events in order to build the detection criteria. For example, these criteria could be signatures, hash values, deviations from standard operation, indicators of compromise (IOCs), etc.; all of which are produced once the initial compromise is discovered. Prior to the recent trend of threat hunting, many of these elements were discovered postmortem. A compromise would occur and during the forensic analysis phase of an incident the evidence would be collected and shared with the larger security team to identify and prevent similar incidents from happening. Now, with the adoption of threat hunting, discovery is still the basis for detection methods, but it does not have to be extracted once an attacker has completed their objectives. Instead, many indicators can be harvested much earlier in the kill chain and aid in expelling an attacker before they are successful. In some situations, such as the creation of behavioral based detections and anomalous activity, detection methods can even be developed from a hunter's hypotheses.

5 ARTIFACTS AND ACTIVITY

So far, we have defined ***who*** a threat hunter is. A defense analyst proactively searching for evidence of a possible compromise from a large collection of data. We understand ***how*** they operate, working from a theoretical approach and basing their searches off a series of hypotheses. We know ***where*** they execute these hunts, from an environment they are tasked to defend under the assumption it has already been compromised. Finally, we know ***why*** they do so is to discovery evidence that has yet to be detected passively and further their defensive measures. What we haven't yet defined is ***what*** exactly a threat hunter is looking for. We know the ***what*** can broadly be described as "evidence," but we can further break this down into two categories as described in our definition: artifacts and activity.

Artifacts can best be described as something left behind during the course of malicious activity. If we were to use the metaphor of a threat hunter as an archeologist, then artifacts would be the physical evidence that is extracted during an excavation. This could be many different things, for example an unauthorized outside resource that was brought into the environment, such as a script or an attack tool. It could also be something generated within the environment in a malicious manner, such as an archive created to be exfiltrated or an unauthorize network mapping for lateral movement. An artifact could even be an something internal that was modified in a malicious manner, like an exploitable update being applied to approved software so it may be leveraged as an attack vector. Whatever the artifact may be, its discovery generally results in something the threat hunter is able to obtain. Forensically, this type of discovery can yield valuable results such as hash values, YARA rules, code analysis, and more that could then be used to not only aid in future hunts, but also enhance future detections.

5 CONTINUED

Activity is something that can also be discovered but is less tangible. These are unusual behaviors or actions that may never use unauthorized artifacts. Unusual logins, odd connection times, interesting commands that run outside of standard parameters, and more could all be considered activity that a threat hunter may find useful. They are valuable pieces of data but not ones that may result in the discovery of an actual artifact. An example may be a user account from the finance department accessing a payroll server to modifying timecards. This may at first appear to be legitimate activity, but when a knowledgeable hunter crafts additional filters to refine this data they may yield further information about a compromised user account. Additional data could include connections being made from a remote login, occurring outside of normal business hours, or perhaps even executing more modifications than usual. These are all examples of potentially malicious activity that didn't result in a standard artifact being discovered. However, that does not mean that this identification did not provide valuable results. Depending upon the capabilities of an organization's security stack, behavioral based detections could be extrapolated from this exercise and used to identify future malicious activity of a similar nature. These types of behaviors might also feedback into a threat hunter's theories to further allow them to identify unauthorized activity.

6 ENHANCING DETECTION METHODS

Despite our focus on these concepts, the discovery of unauthorized artifacts and activity is not the sole purpose of a threat hunter and it should not be considered as such. A common misconception that can often discourage junior threat hunters is the belief that in order to classify a hunt as successful, the outcome should always be the discovery of something malicious. This is absolutely not the case. In fact, it may be more accurate to say that when every attempted hunt results in the discovery of malicious activity, that is not a signifier of success, but rather a sign that an organization has very poor defenses. That is because in an ideal environment, a strong defense would block all attackers, or at a minimum passively detect them early on. Such mechanisms would be so strong that every hunt would return no results since compromises could not linger undetected. Unfortunately, none of us work in the ideal environment and instead we must balance our defensive posture with usability, which can result in compromises occasionally occurring. That is why the goal of the threat hunter should be a blend of both discovery of artifact and activity, and the continual improvement of passive defenses.

6 CONTINUED

Consider the scenario of a threat hunter working under the hypothesis that a targeted attack will occur against the internal human resources department. The attacker's purpose could be obtaining employees' personally identifiable information (PII). Once the hunt begins, data is collected, organized, and reviewed; activity is monitored and vetted; queries are written, and results are analyzed; and the outcome of all this work is then the determination that nothing malicious has occurred. To a novice threat hunter this may appear as a wasted effort; but a seasoned professional knows there is greater value in this exercise. That is because the results are not the only outcome of these events. Instead, look at the hunt as a whole. The queries that were written based on this theory could now be automated to run against future data sets, thereby creating a new passive detection mechanism that may yet yield results. The collected set of data could be reused for additional variations of the original hypothesis, perhaps replacing the goal of obtaining PII instead with access to company data. Observed activity that was cleared as legitimate work actions may then be used as a baseline to compare against abnormal activity later. All of these scenarios result in incremental increases in the overall defensive posture of the organization even when nothing malicious was discovered. And that is mission of the threat hunter: search through the past, discover in the present, and prepare for the future.

CONCLUSION

How to Accelerate Threat Hunting with Fidelis Elevate

Threat hunters thrive on data, and the most successful threat hunters are the ones with the highest level of visibility into their environment's data. At Fidelis Cybersecurity, our mission is to give every security team the visibility they require to defend. We do this through holistic visibility that leverages integrated network traffic analysis, digital forensics and incident response, and deception technology - a proven trifecta that allows your organization to successfully secure your environment while proactively hunting out unauthorized activity.

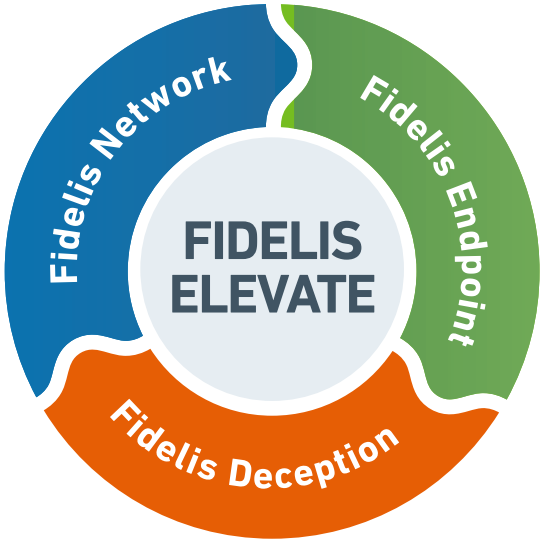

NETWORK TRAFFIC ANALYSIS

With Fidelis Network, threat hunters obtain an enhanced view of network activity through deep inspection and analysis of all forms of content, including unpacking and extracting deeply embedded files. Fidelis Network bi-directionally scans all traffic, regardless of port or protocol to reveal the network and application protocols, files, and content. With direct, internal, email, web and cloud sensors, you gain full network coverage and visibility.

ENDPOINT DETECTION AND RESPONSE

Endpoint Detection and Response is built with visibility as its core component. Monitor and query endpoints in real-time and retrospectively while both on and off network using new or saved advanced Boolean logic. Fidelis Endpoint allows threat hunters to see all process actions, user activity, registry events, file system activity, memory data, and more on Windows, Mac and Linux endpoints. Review full system software inventories for known CVE and KB vulnerabilities and record key events with playback analysis that automatically delivers an incident timeline, along with prioritized alerts.

DECEPTION

With Deception Technology, a threat hunter gains deeper visibility with the ability to classify all networks and assets, communication paths, and network activity to profile your users, services, and systems. They obtain full view of servers, workstations, enterprise IoT devices, legacy systems and shadow-IT, while ensuring an always-current profile as changes occur within your environment to automatically adapt deception layers. Enhance this further with the ability to actively test your hypotheses by generating breadcrumbs and decoys to lure attackers out and monitor their activity while protecting your production environment. Fidelis lets you take threat hunting to the next level.

Fidelis Cybersecurity is a leading provider of threat detection, hunting and response solutions. Fidelis combats the full spectrum of cyber-crime, data theft and espionage by providing full visibility across hybrid cloud / on-prem environments, automating threat and data theft detection, empowering threat hunting and optimizing incident response with context, speed and accuracy.

By integrating bi-directional network traffic analysis across your cloud and internal networks with email, web, endpoint detection and response, and automated deception technology, the Fidelis Elevate™ platform captures rich metadata and content that enables real-time and retrospective analysis, giving security teams the platform to effectively hunt for threats in their environment. Fidelis solutions are delivered as standalone products, an integrated platform, or as a 24×7 Managed Detection and Response service that augments existing security operations and incident response capabilities. Fidelis is trusted by Global 1000s and Governments as their last line of defense. Get in the hunt. **www.fidelissecurity.com**

Fidelis® Cybersecurity