jamf

How to Establish a Threat Hunting Team in Your Organization

So, you're thinking about threat hunting but don't know where to start.

Taking a proactive approach to finding active or persistent threats lurking within your organization requires a dedicated security team of skilled analysts, software tools, automation and the keen ability to think like an attacker.

In order to identify, mitigate and protect your organization against malicious threats, resources and time must be placed on hunting down the evidence and piecing together all the digital clues they leave behind.

While all this may sound overly complicated, we're here to help you develop a successful threat hunting effort.

In the world of IT, there are a number of different roles — some with statically defined purposes, others structured as a mash-up of two or more disciplines — creating opportunities for IT to meet the demands of dynamic environments, an organization's needs and a variety of risks. Arguably, one of the roles that fills some of the greatest needs affecting all organizations relying on computing devices, networks and the internet to perform business functions, is Threat Hunting.

Other names for this role include Threat Analysis and Malware or Security Research are essentially the same role: to investigate anomalies in applications, users, devices or data with the goal of identifying any unmitigated risks within the organization's computing infrastructure that is actively being exploited by a malicious actor.

You wouldn't be mistaken in thinking this role operates in a more specialized realm than some of the other IT positions out there. While it is considered a security-based role, in fact, it deals with a very specific sub-set of security with its own criteria and unique skill sets to be an effective and successful threat hunter.

In this paper, we'll discuss:

- ✓ Background on threat hunting
- ✓ Specialized skills and tools used
- ✓ Criteria for isolating threats within systems and apps
- ✓ Deep dive into a real-world threat hunting scenario

“If it bleeds, we can kill it”

Like the leading character, Dutch, from the first Predator movie mentions as his role transitions from being hunted by the titular antagonist to effectively becoming the hunter himself, a threat hunter assumes a similar role as they run systems, software and their processes through its investigative paces.

What are they looking for? Bugs in the code base, advanced persistent threats (APT), behaviors that are malicious or simply uncharacteristic with normal operating procedures. Essentially, anything that could represent vulnerabilities that could be potentially exploited by a threat actor to compromise equipment, account privileges, networks and ultimately, privacy and/or confidential data — or signs that an attack is already underway.

So now that we know the *what* behind threat hunting, you might be contemplating the *why*. Well, there’s an answer to that too. As threats become more advanced and attackers add techniques to their toolsets to stealthily slip past network and endpoint security defenses, organizations are increasingly seeing the value in establishing a team dedicated to extending threat analysis to detect instances of threats. Malware or indicators of compromise (IOC) and indicators of attack (IOA) for current, ongoing attacks may not be detectable using standard security tools and therefore need a human to help identify.

The process of proactively searching for undetected threats to your organization’s cybersecurity defenses focuses on indirectly protecting assets through direct analysis and investigation of malicious activities to locate hidden threats lurking within your network perimeter and/or targeting your endpoints. The goal being to “channel your inner Dutch” by taking an analytical, offensive approach to augment your defensive strategies. The quicker issues are detected, the faster and more efficiently risks can be mitigated. The limiting of “dwell time”, or the time between when an attack began and when it was detected, means an attacker has less time to cause damage and IT has more time to keep data secure and ensure endpoints remain compliant.

M-E-T-H-O-D, (man)

Traditionally, the role of a security analyst as it pertains to threat hunting involves sifting, sorting and categorizing various forms of data and logging information — and as importantly, combining that with intuition, security knowledge and an understanding of the devices involved — to determine how to best proceed in analyzing potential threats and mitigating real ones. While this still holds true, today's technological landscape looks different than it did five years ago — as it will invariably look slightly different five years from now — all thanks to advancements in technology. And while these advancements help security and IT, they also aide attackers and introduce new forms of risk...and you better believe they *will* take advantage of these advancements.

One of the biggest advantages to threat hunters comes by way of artificial intelligence (AI). Its contributions have paved the way for automation. In fact, AI lends itself particularly well to software that automates the detection of possible risks when leveraging machine learning and pre-defined or custom analytics — even metrics used to compare anomalies to known and unknown behaviors — to predict if the observed behavior types fall within a margin of error or are characteristic of a threat or attack.

Simply put, threat hunting is an exercise in finding patterns of unusual behaviors within a sea of data. Tools help us sift through the data to focus our ability to recognize unusual patterns.

A Typical Threat Hunting Process

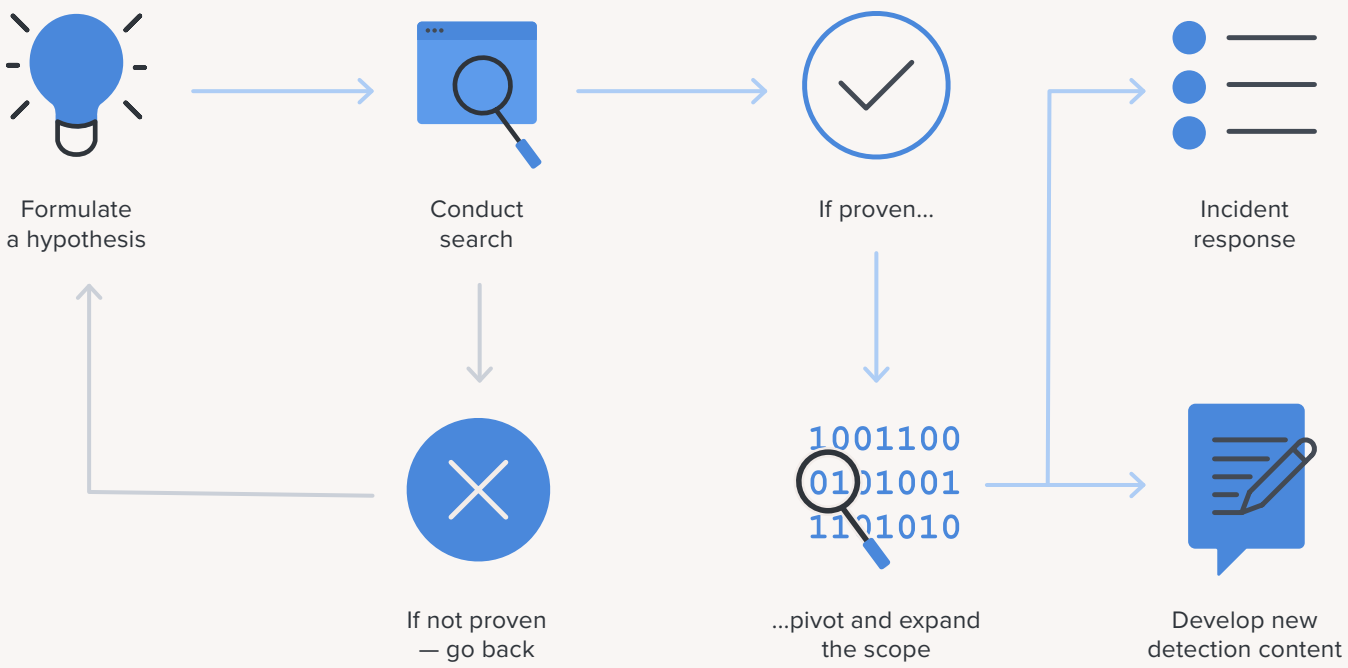

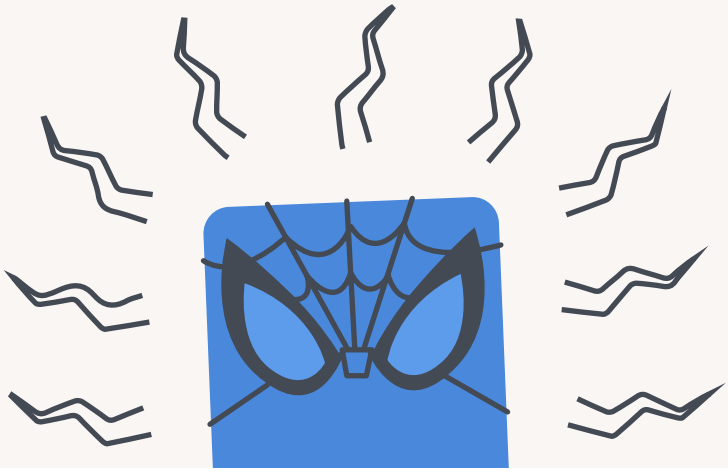

Before we dig into the use of frameworks, tools and skill sets, we need to start at the beginning by formulating our hypothesis. After all, every investigation – whether it's a capital offense or cybersecurity-focused one – begins with a belief about what has or will occur as a means to put together the evidence to support or contradict their initial assessment.

A hypothesis could be triggered by an anomaly detected in your SIEM, an email alert relating to a newly discovered malware strain or perhaps it's just your Spidey-Sense tingling that something may be slightly off. Whichever the trigger, the hypothesis is the jumping-off point to begin collecting and analyzing data.

Next, let's discuss the methodologies used in threat hunting, what each one target and the role each plays in guiding the overall investigative process, beginning with the creation of a hypothesis:

- **Analytical:** Driven by AI, User and Entity Behavior Analytics (UEBA), and/or custom metrics setup to develop scores that are aggregated and associated to specific risks.
- **Situational:** Provides ancillary data regarding an organization's critical systems and processes, often using the Crown Jewels Analysis (CJA) to determine asset criticality levels, then uses that data to determine risks and mitigating strategies.
- **Intelligence:** Incorporating data from myriad sources like reports, scans and analysis to assess and categorize threats using real-time, live data that is pertinent to the organization and relative to the assets currently in production.

Arguably, any process or collection of processes — ranging from relatively simple ones, like verifying file hashes to more complex tasks, like involving machine learning software to vet applications and network transmissions — may be grouped under the threat hunting umbrella. The challenge presented to organizations, however, comes not from the efficacy of each process to detect threats but rather the maturity level of each to effectively do so on a consistent basis and the ability to scale or adapt to meet current and future needs.

Couple that with the varying skill level of organizational security teams globally and it can be challenging to accurately detect these threats and investigate them appropriately. Threat hunting happens on a spectrum from simple to very complex processes. Some threats may evade even the most ardent practitioners for long periods of time.

One way to tackle this complexity is to adopt a formal framework that addresses the following concerns, among others:

- Centralized collection of known adversarial behaviors and detection analytics
- Advise organizations on how to improve threat hunting based on best practices
- United vocabulary to provide consistency in communications regarding attacks between departments and/or organizations
- Community driven, centralized effort to detect threats and better defense against existing, new and unknown attack types
- Increase the maturity level of the security and IT organizations by introducing processes and auditing mechanisms

A huge benefit to instituting a formal framework, especially for organizations with entry-level knowledge of security and staff, is an increase in the maturity level of the processes involved in detecting threats. The maturity level follows a model that classifies the level at which an organization's threat hunting program operates based on its capabilities to identify threats, procedures for data collection and level of analysis incorporated via automated vs manual methods.

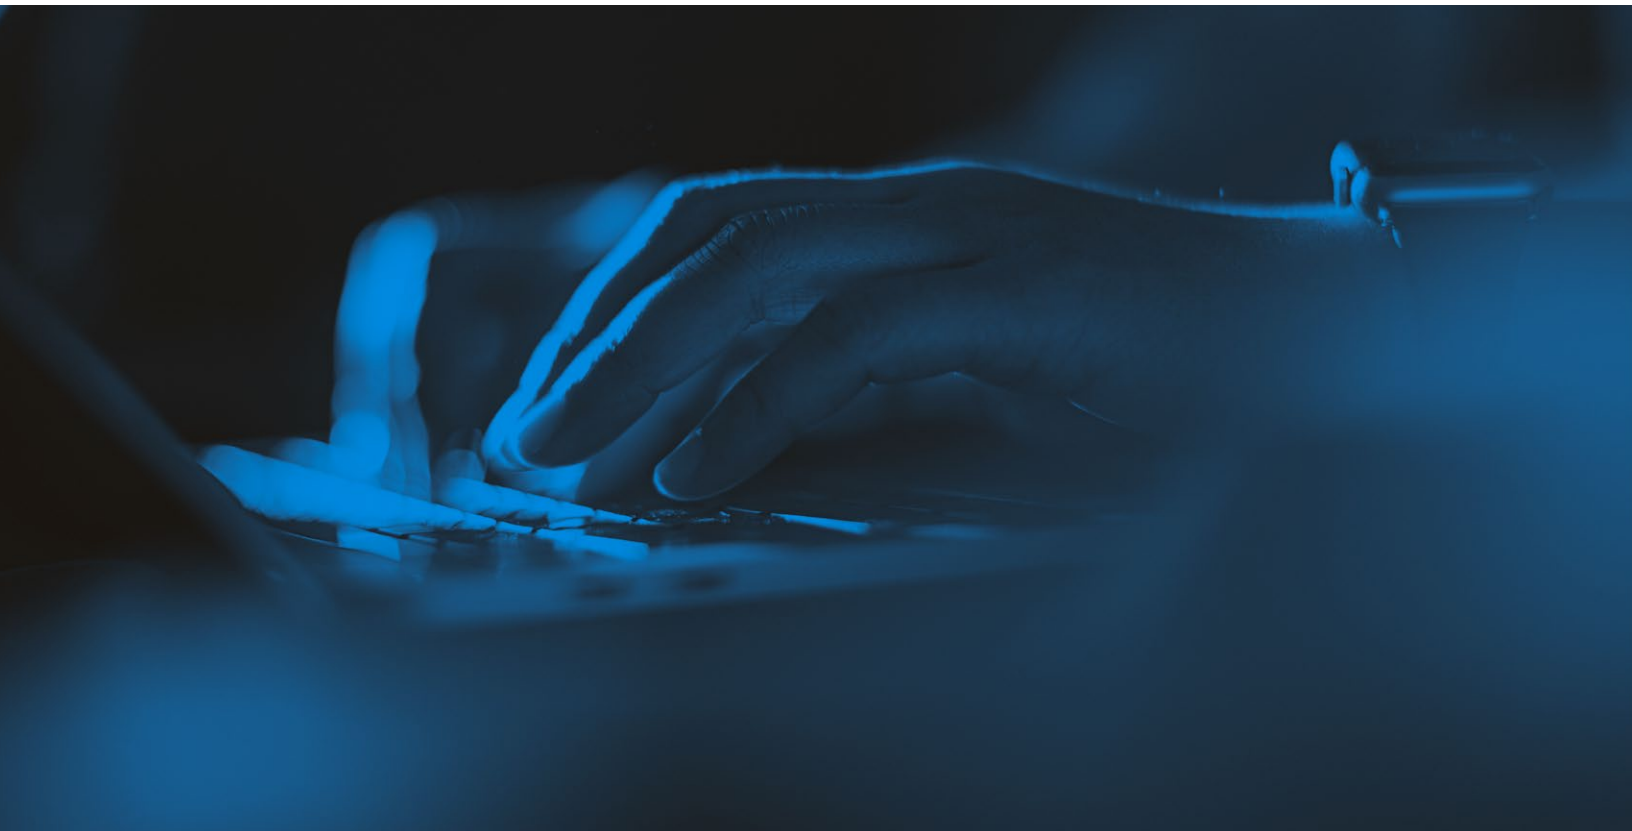

A little search into threat hunting frameworks will provide a whole slew of frameworks that could potentially be **a good match to address your organizational needs**. A few of the more notable **ones that have been around for quite some time** were developed by industry leaders or may even be aligned with Endpoint Detection and Response (EDR) products already in use within your organization like:

Cyber Kill Chain

By adopting the military concept of the "structure of an attack", **Lockheed Martin applies the *kill chain* term** to information security and introduces the *cyber kill chain* or "structure of a cyber-attack" in 2011. The latter relates to a breakdown in the way attacks occur, typically in phases, and how implementing the proper controls at each phase may disrupt attacks.

According to the model, there are seven levels in the cyber kill chain:

- Reconnaissance
- Weaponization
- Delivery
- Exploitation
- Installation
- Command and Control
- Actions on Objective

Each phase or level relates to actions undertaken by threat actors as they move through subsequent phases to complete their attack.

Armed with this information, security teams can prepare their organizations to protect against attacks, and threat hunters can analyze gathered data to effectively detect, deny, disrupt, degrade, deceive or contain threats in defense of the organization.

Diamond Model

This model was derived as an answer to the questions posed by several senior analysts in 2006, asking themselves open-endedly, how they do they perform their work. With this honest approach, the Diamond Model was birthed, based on quantifying the processes and principles at the core of collecting and analyzing data to better understand the mindset of threat actors and by extension, malicious activity with the aim of eradicating it.

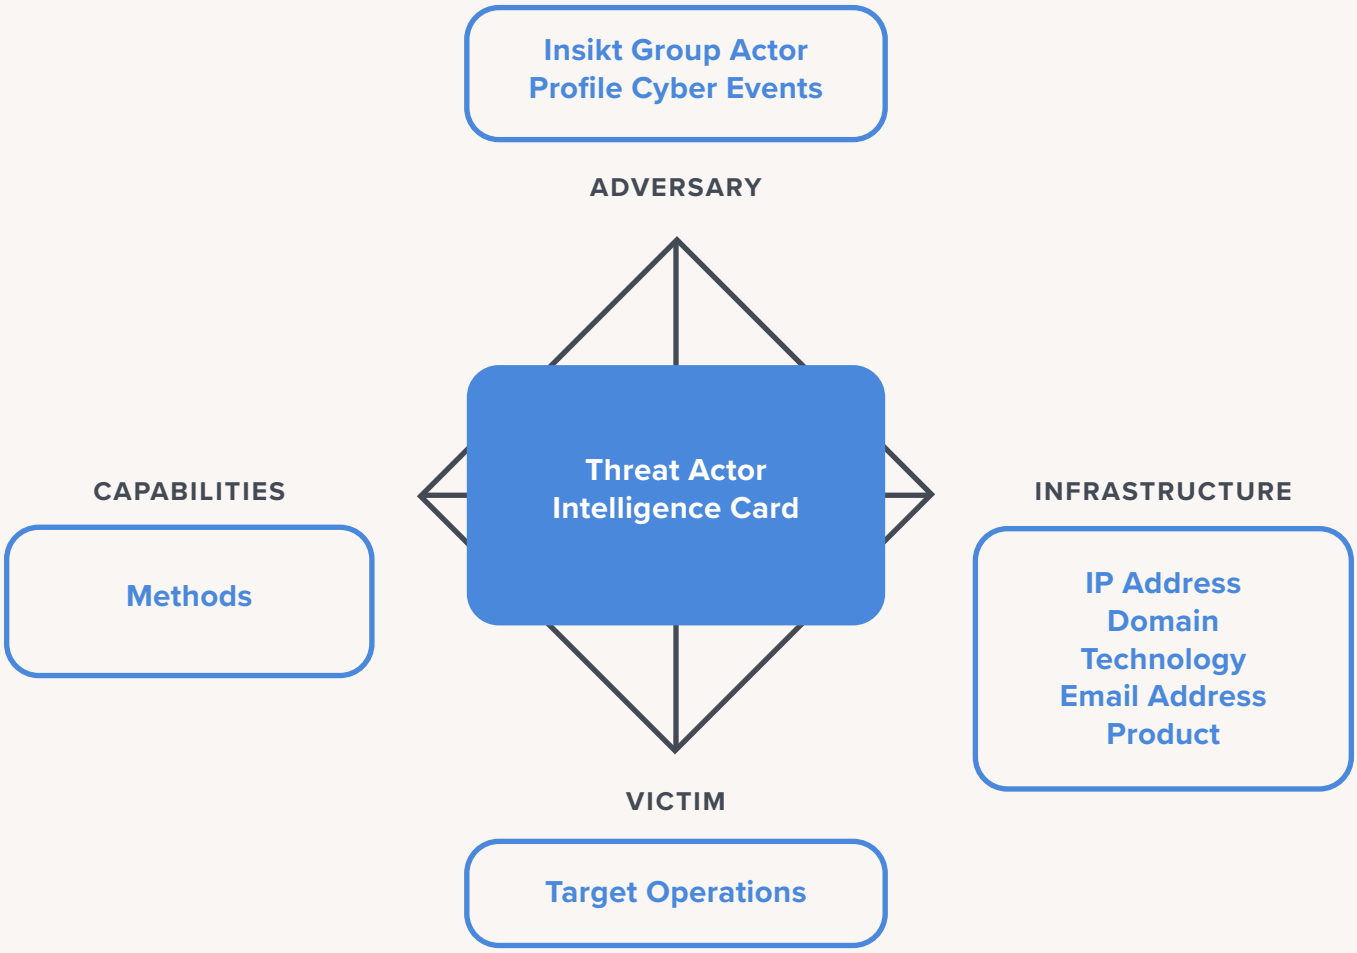

It gets its name from the four-pointed diamond shape, representing the core features of: Adversary, Capability, Infrastructure and Victim, present in each malicious event. By connecting the vertices, relationships between points can be drawn to produce an analytic that details how an attacker could potentially move from one feature to the other or pivot. This provides the analyst a holistic view to spot weaknesses and opportunities within the organization. By targeting those, analysts can obtain data necessary to protect the organization from having that feature exploited. The system is simple but can be tailored to match organizational needs, expanding core features and composing of sub-features as needed.

MITRE ATT&CK

Developed in 2013, the MITRE Corporation's framework is actually a set of three matrices, Enterprise ATT&CK, Pre-ATT&CK and Mobile ATT&CK – each of which targeting tactics and techniques that are unique to each matrix when threat hunting. Consisting of fourteen Enterprise categories:

- Reconnaissance
- Resource Development
- Initial Access
- Execution
- Persistence
- Privilege Escalation
- Defense Evasion
- Credential Access
- Discovery
- Lateral Movement
- Collection
- Command and Control
- Exfiltration and Impact

Each category, called **tactics, feature over 500 techniques and sub-techniques** contained within the Enterprise ATT&CK matrix alone. MITRE's ATT&CK framework provides a concise, organized description of each threat: what it does; how it works; what tactics it falls under; which sub-techniques are associated with each threat type; which data sources may be used to gather evidence; which systems defenses, if any, are bypassed during attack; what operating system platforms are affected by each threat and timestamp information, including last modified date and versioning details.

As touched upon previously, the **MITRE ATT&CK Framework is aligned with Jamf Protect**, the EDR software that provides threat hunting features alongside its malware detection and prevention capabilities for macOS.

And don't forget that security – like many other processes in IT – is iterative, meaning each phase, incident or stakeholder provides feedback that informs how process, tools and procedures should change in relation to the nature of risk mitigation over time.

For a detailed guide on **Building and Maturing Your Threat Hunting Program**, the SANS Institute has made this document available. It provides insight into the organizational models, metrics to collect, obtaining data sources, tooling considerations and levels of automation to use when developing a threat hunting playbook based on industry best practices for your organization.

Another consideration for assessing the maturity level of the Tactics, Techniques and Procedures (TTP) used within your organization is the **Hunting Maturity Model** (HMM). The HMM was developed by David Bianco to incorporate the characteristics of an organization's tools, security staff skill sets and processes used to effectively gather and analyze evidence in threat hunting.

Essentially, the HMM consists of five levels that determine organizational hunting capability ranging from Level 0 (least) to Level 4 (most). Here is a breakdown of the levels and their meanings:

Level 0 – Initial

Relying primarily on alerts generated from automated tools, such as antivirus, IDS and SIEMs, the incorporation of threat intelligence from pre-compiled and/or customized analytics are fed directly into monitoring systems with little to no routine data collection with the primary goal of manual alert resolution.

Level 1 – Minimal

While still relying primarily on auto-generated alerting, a limited amount of data collection routinely occurs. Additionally, threat intelligence is augmented by utilizing a combination of open and closed sources to track threats using the latest reports to extract key indicators and search through historical data. This is often considered the first level where the analyst conducts any real threat hunting.

Level 2 – Procedural

Leveraging data captures from sources toward the detection of a single type of malicious activity like analyzing logging data to detect malware programs, for example. At this level, organizations carry out gathering and analysis utilizing commonly available procedures developed by third-parties to hunt for threats but are not capable of generating their own analytics or wholly-new procedures.

Level 3 – Innovative

Expanding on the previous level to include targeting of multiple types of threat actor activities, while utilizing procedures generated themselves (as opposed to utilizing third-party created ones). Analytic skills are also linked to other advanced topics, such as data visualization and AI. At this level, procedures are performed routinely. They are well-documented and may even be published for others to review.

Level 4 – Leading

The final level in the HMM bears a striking resemblance to the previous level except for the inclusion of automation. At this level, the procedures vary in scope and scale. By finely tuning the procedures for automation, analysts are able to free themselves from running these procedures directly, allowing them to focus their efforts on the creation of new procedures to provide iterative improvement in a loop.

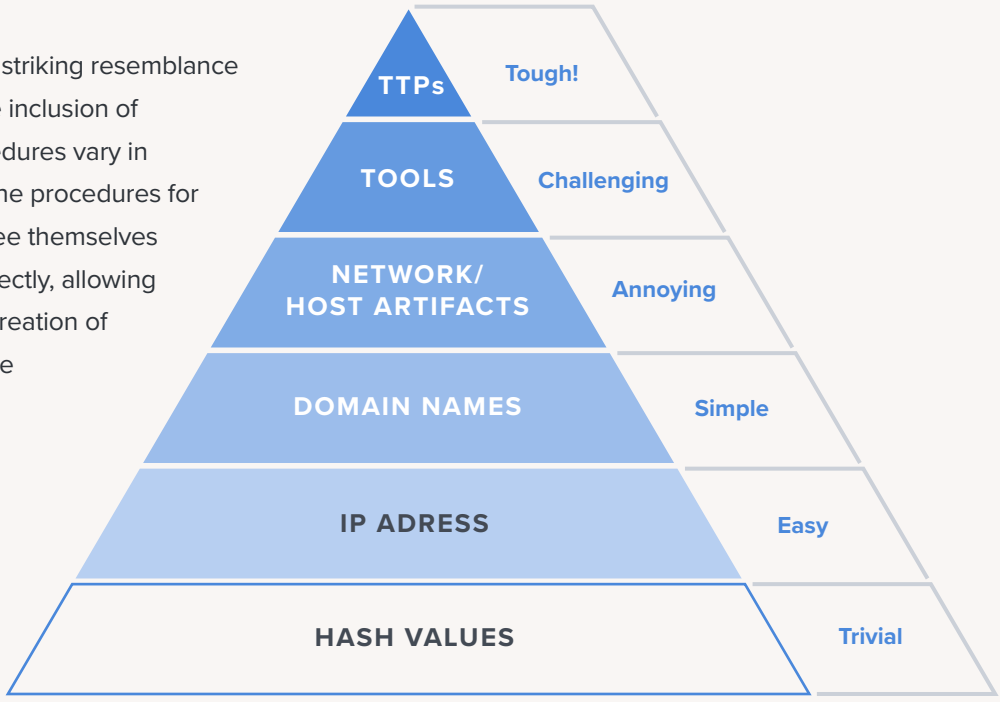

Skills to pay the bills

Unlike the Predator, the threats we've discussed don't bleed per se. But like all things digital, they do leave behind unique fingerprints that can be used to uniquely detect them. Therefore, they can be stopped. Lining up quite nicely with the mission of a security analyst, or threat hunter: **stop the threat through a series of objectives *before* it leads to something far worse**.

To be an effective threat hunter requires a mix of utilizing the right tools, experience and skills. It doesn't require you to be a veteran in the IT realm nor does it mean acquiring every certification available either. It does mean, however, that the following hard skills are shared among successful security analysts and, while all may not necessarily be a requirement, they are extremely beneficial to the role:

- Intimate knowledge of information and systems being protected
- Internal networks and how they communicate
- Endpoint management experience
- Ability to analyze and work with data
- Experience with performing data forensics
- Managing and collecting network traffic and analysis
- Ability to understand code and programming capability

Additionally, there are several soft skills as well that are also common to the success rate of security analysts:

- Recognizing patterns
- Deductive reasoning
- Effective communication
- Out-of-the-box thinking
- Overcome cognitive bias
- Ability to think like an attacker

"...A very particular set of skills. ...Skills that make me a nightmare for people like you."
– Liam Neeson

The above quote being the infamous lines spoken over the phone in the film Taken, summarizes the crux of how a security analyst armed with the proper set of skills is in stark proportion to the frustration level a would-be malicious actor experiences when having their threats thwarted each time.

Jamf's security team is dedicated to detecting and neutralizing Apple-specific malware threats.

Contact us today, or contact your Apple reseller, to put Jamf Protect's monitoring and prevention capabilities to work in your organization.

Happy hunting!