Best Practices for Effective Threat Hunting

A PROTECTWISE WHITEPAPER

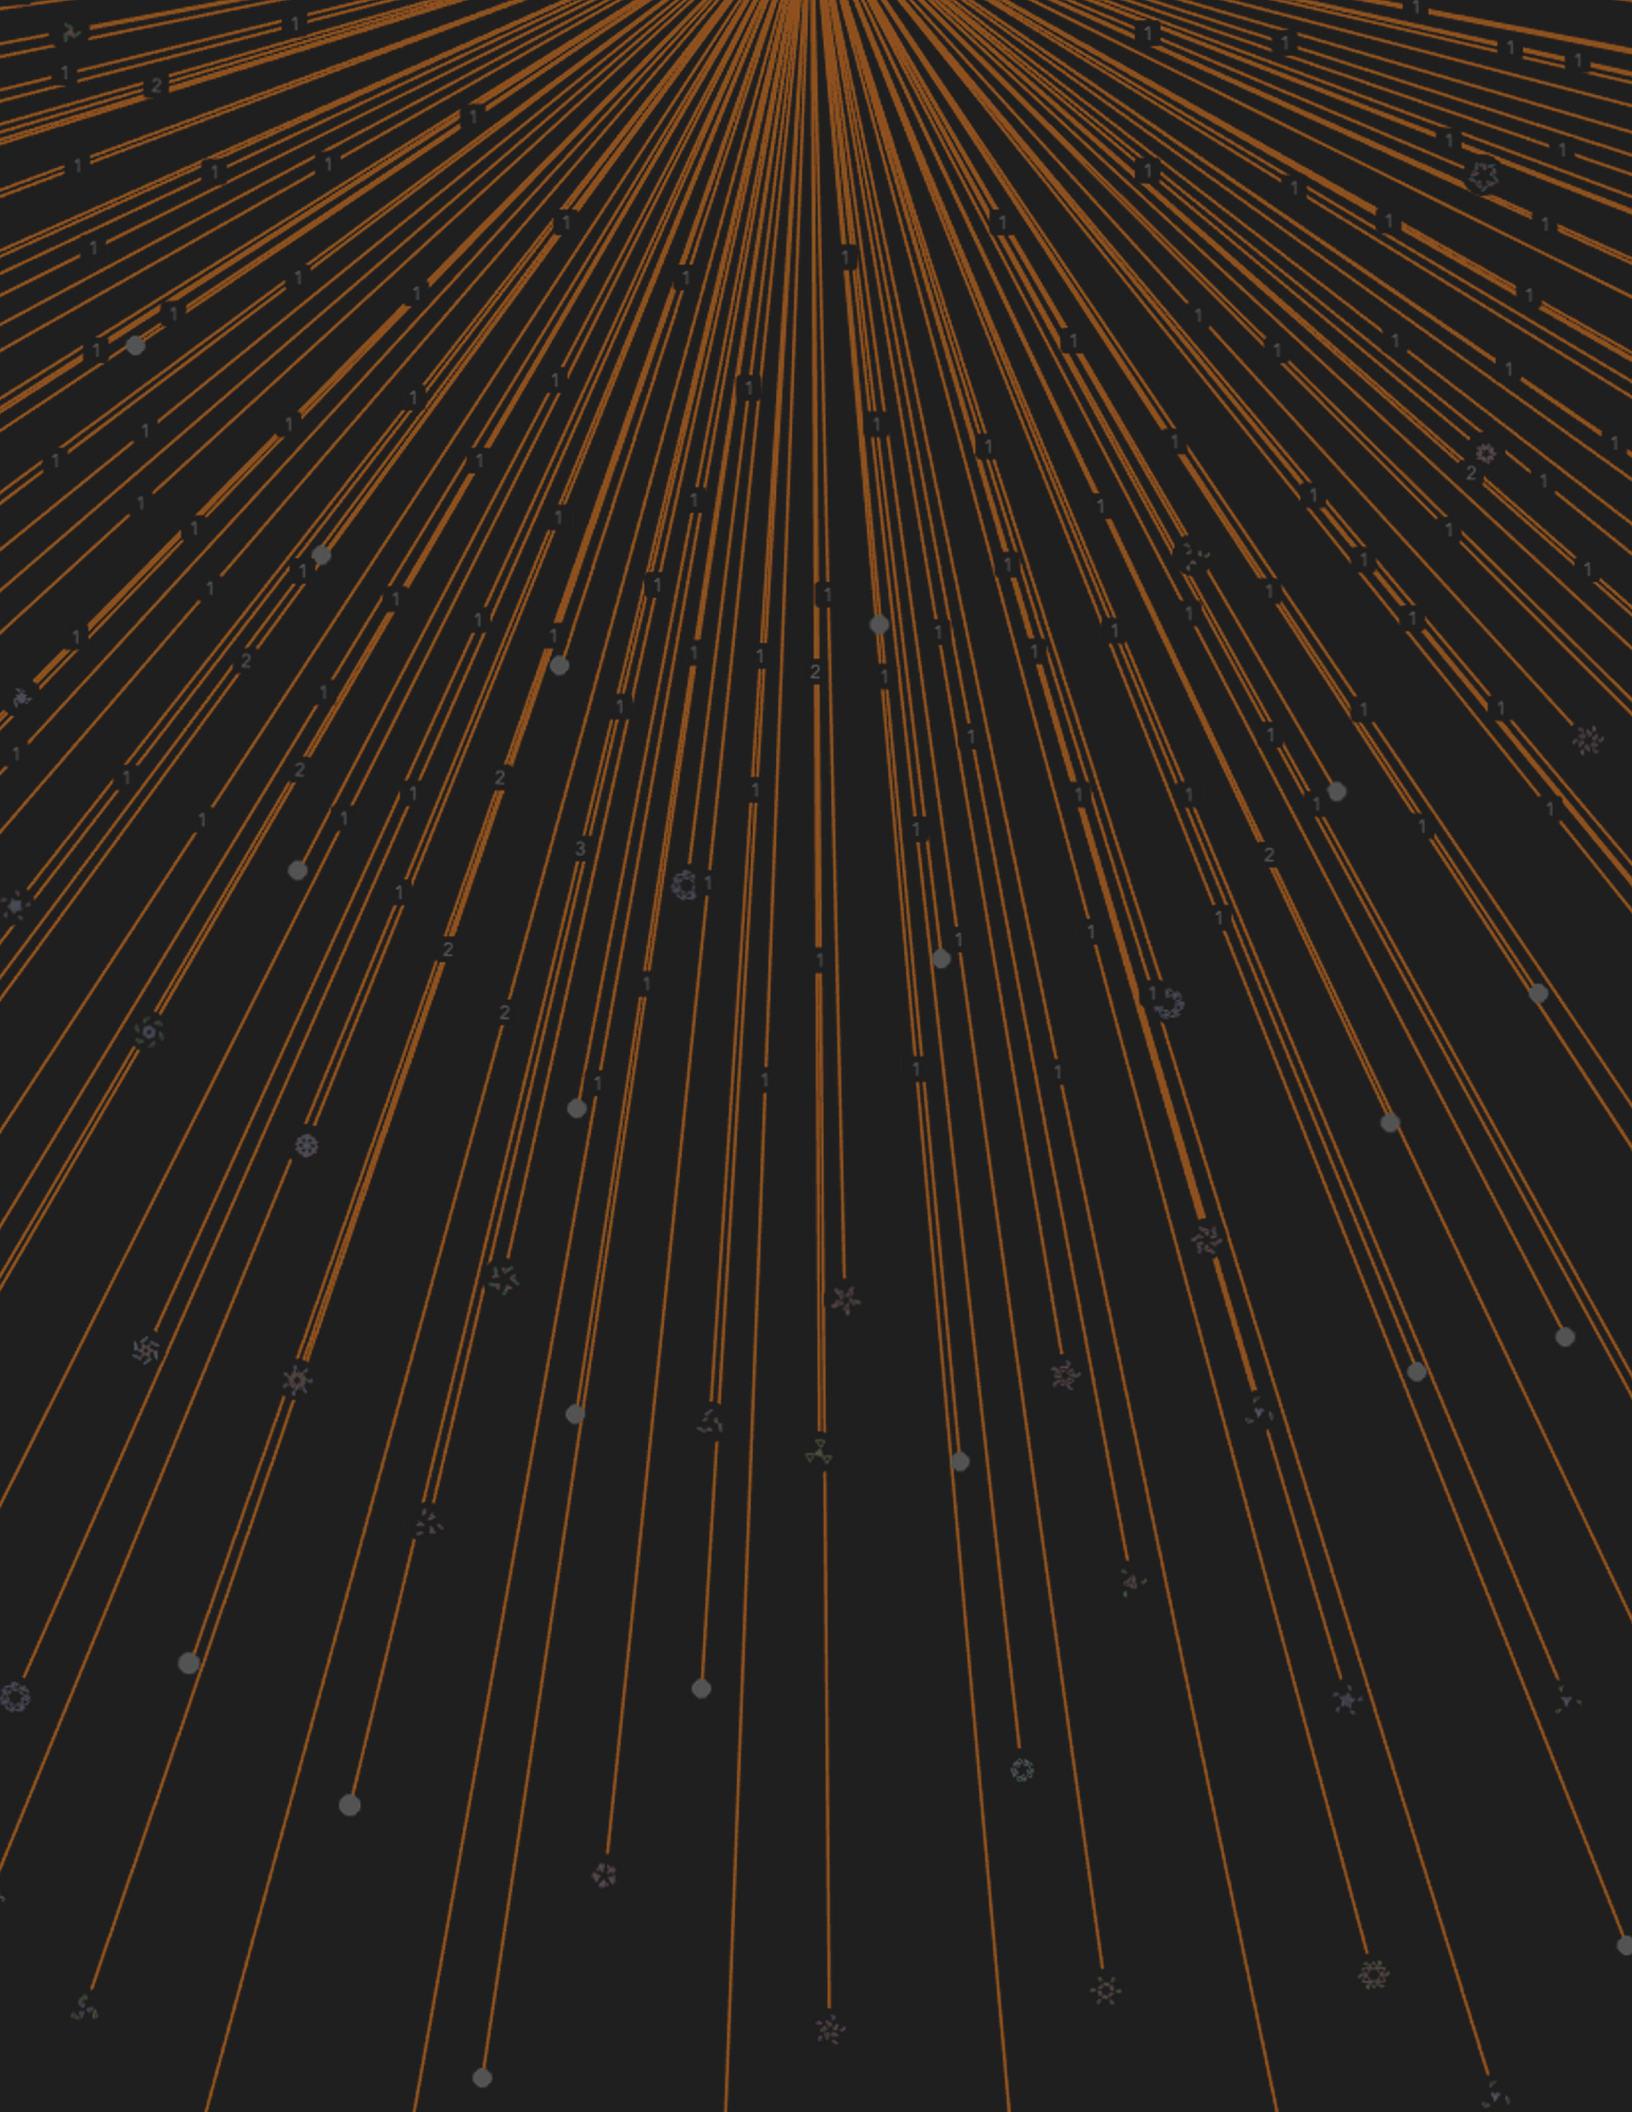

THREAT HUNTING IS DIFFERENT THAN THREAT DETECTION

Detection is reactive in nature and involves receiving alerts or looking for known threats or signs of attacks. Detection primarily happens in real-time, largely using threat intelligence and signatures. There are additional real-time use cases for automated techniques like anomaly detection and machine learning algorithms.

Threat hunting is the proactive approach of searching out malicious activity that may have evaded detection mechanisms by either being new or generally unknown. So, instead of focusing solely on real-time alerts, threat hunters may look into the past to discover suspicious trends or indicators of compromise (IoC). Hunts may begin as hunches or theories with leads based only on an analyst's intuition, new intelligence, or previous experiences.

When hunting, analysts shouldn't always focus on finding the next big APT. Instead, they should focus their efforts on discovery and learning to help them and their team hone their skills. Hunting also has the added benefit of driving higher satisfaction and retention since analysts are being challenged rather than performing mundane or repetitive tasks commonly associated with threat detection.

When analysts discover evidence of attacks that were previously unknown, they can use that information to improve automated detection. This could be new malware, a new type of attack, or a new vulnerability. But it could also be something less malicious like a broken business process, a misconfiguration, or an unapproved network configuration.

Building a threat hunting function in your security team is no easy task, though. There are many considerations, including the maturity of the existing security team, the depth of analyst's experience, and the systems and technologies that are available. This guide offers some best practices to keep in mind when scoping out whether or not your organization is ready to set up a threat hunting team.

1. PREPARE FOR SUCCESS

Threat hunting is most successful when dedicated resources can remain focused on their goals. It will be difficult to build out a threat hunting function if a team finds itself constantly putting out fires or without the right technology to get through the reactive nature of day-to-day activities efficiently. Before the team starts hunting, consider the following:

- **ORGANIZATIONAL READINESS:** Like many business projects, it may be best to consider a top-down approach to building a threat hunting function so that everyone from the top floor to the shop floor is ready to support the team. Brief executives why a threat hunting function is necessary, and ensure that they understand the value the team brings to the organization. Then, educate groups within the organization on what they can expect from threat hunters and the team's benefits to them.

- **SKILLSETS:** Threat hunting requires a unique skillset and a different frame of mind. Threat hunters are typically skilled and tenured security professionals who are highly proficient with incident response processes and capabilities. They then marry that knowledge and experience with intuition and an investigative mindset to seek out known threats proactively.

- **TECHNOLOGY:** Products in an organization's security stack should work together so hunters can remain focused on seeking out threats. Intrusion Detection Systems (IDS), SIEM, firewalls, and endpoint security products should work together to automate detection and mitigation workflows. Information from these products should feed a unified body of correlated forensic evidence for better context, and should be retained for periods of time that exceed breach windows.

---

2. UNDERSTAND THE THREAT LANDSCAPE

It's often said that the best defense is a good offense, and a good offense means knowing what you're up against. For that reason, it's important for threat hunters to have a deep understanding of the threat landscape. Threat hunters should stay abreast of publicly known external threats, trends, and remain connected with peers about new vulnerabilities, actor groups, and malicious tooling.

On top of that, threat hunters should have an intimate knowledge of their organization's networks. Whether it's a public or private cloud, an on-premises network, or an industrial control system network, analysts need to know what type of users and devices are on their networks. Endpoint security products can help provide visibility into laptops and mobile devices on which vulnerabilities could pose a dangerous risk of hacks or data leaks.

Security teams also need to consider the vertical or industry in which their organization operates to include any unique vulnerabilities or threats they might face. For example, hospitals may face life-threatening impacts if they become victims of ransomware attacks. Although ransomware attacks aren't unique to healthcare alone, it's important for security professionals in this industry to focus on immediate and potentially life-threatening incidents that could compromise a hospital's ability to serve its patients.

3. AUTOMATE DETECTION AND RESPONSE

Time spent responding to alerts from various security products takes away from time teams could spend proactively hunting for threats. There are many approaches to making sure detection and response capabilities have been automated as much as possible. Whatever approach is taken though, consider API-driven products for their ability to work seamlessly with off-the-shelf and with home-grown security products.

A good place to start is making sure security products can integrate in some manner. Data used for threat hunting should be enriched by additional context made available by a wide range of other products in the ecosystem. For example, correlating network and endpoint data will help more easily identify devices and users causing network vulnerabilities.

4. COLLECT THE RIGHT DATA

Effective threat hunting requires collecting the right data, including the network artifacts needed to build a searchable repository. Collect network data at ingress and egress points on internal and external segments owned by the organization and on those that are not, like in the Cloud, when possible.

Artifacts collected internally are also valuable when considering the lateral movement of threats or searching for suspicious behaviors. The most pertinent information for analysts, though, can be taken from network packet captures, or PCAPs. While PCAPs are a key element for effective threat hunting, also keep in mind the importance of supplementing these with data like logs and device information from endpoint security products.

Artifacts that could be extracted from retained PCAPs, then indexed and made searchable may include:

- HTTP headers and transactions
- DNS queries (to understand which hosts were talking to on the network)
- Complete flow for searching and narrowing down PCAPs in a deep-dive analysis
- SSL certificates, which can say a lot about the nature of network traffic
- Files from protocols like FTP, HTTP, and SMB
- Email protocols and headers

If this data has been collected over a long period, there needs to be a mechanism in place that allows analysts to pivot from one data point to the next so they can work through multiple leads efficiently. These leads might take them to dead ends, but others might take them somewhere else in time. For that reason, it's important to make sure threat hunters can make sense of and pivot through massive amounts of data quickly.

5. EMBRACE THE CLOUD

On-premises storage used to be a burden for its high hardware and maintenance costs which limited scalability. Today, however, cloud storage costs are calculated in fractions of cents, and cloud computing prices are lower than maintaining proprietary data centers. As a result, enterprises can cost-effectively retain data, including full-fidelity PCAPs, for as long as necessary (weeks and even years) enabling analysts to perform an infinite amount of threat hunting in the Cloud.

Storage is just one part of this new dynamic, though. Organizations can rent cloud servers with incredible amounts of power, capable of running complex statistical analysis, or of indexing massive amounts of stored data quickly. That makes handling these massive amounts of data and correlating that data across multiple systems easier and faster, and provides threat hunters with a more contextual understanding of what's going on inside their networks. For example, rather than saying an event happened at a specific point in time, threat hunters can see instantly a 360-degree view of network segments, users, devices, and even external parties involved in security events.

6. DEVELOP AND PURSUE LEADS

Threat hunters have to think like detectives which means they're going to develop and pursue leads. If something doesn't quite add up or make sense, keep asking why until a root cause is found. Don't just focus on picking up those pieces that only prove a hypothesis. If that doesn't work, take a step back to reexamine the circumstances around the gathered body of evidence - and then start asking questions again.

Keep in mind the importance of collecting evidence that both supports and disproves a hypothesis. False positives aren't always unacceptable, and a small quantity of them may be okay as long as they're not overtaking analysts workload.

Security teams should use false positives to tune detection systems and processes, and to educate themselves on things they're learning from their networks. They should also follow trends about new malware or actor groups. Security teams should learn about how these function, then test this threat intelligence in-house to identify detection gaps or existing malicious activity.

7. AVOID TUNNEL VISION AND CONFIRMATION BIAS

Analysts are going to spend a lot of time hunting, so they need to know how to manage their time effectively. It's easy to get stuck focusing on a single artifact in a sea of artifacts. Spending time focusing on that one piece of evidence is okay, but not at the expense of discovering something else later that disproves what the analyst was looking to prove in the first place. For example, a hunter may have found an HTTP request that seems malicious and become too focused on this one event that he or she fails to see obvious signs that determine whether it's a true- or false-positive.

It's also important to avoid, confirmation bias, or the tendency to interpret new evidence based on an existing belief. What this usually means is an analyst starts out looking for malicious activity, and when something suspicious is found, he or she might say the discovery immediately proves a hypothesis. This could be damaging because a lot happens on a network that's just anomalous, not dangerous. Instead, turn that impulse on its head and prove that something is legitimate rather than malicious.

8. DOCUMENT EVERYTHING

Threat hunters aren't going to remember everything, or to memorize every little detail they come across. Everything should be documented because what appears to be malicious one day might not be the next.

Case notes from previous hunting engagements give threat hunters additional context and history about things they'll discover in the future, and may help them avoid wasting time pursuing a chain of evidence that, ultimately, reveals that a series of apparently suspicious actions really isn't malicious at all. It also creates a record that the entire team can use in the future, even after the analyst has left the organization.

CONCLUSION

Threat hunting isn't a failure if nothing is found. Simply going through the process of hunting and tuning the organization's detection systems and processes is of great benefit to the entire security team. Hunters might spend hours or even days without finding anything significant, but that doesn't mean their time has been wasted. Hunting also results in an improved knowledge of the organization's network, and can even uncover broken business processes that need to be fixed.

Don't be afraid to embrace the analyst mindset. Test theories, hypothesize about something learned from public reports, or focus on something curious. Chasing after these ideas is an effective way to learn and to improve individual skills, and the skills of the overall security team. As threat hunters spend more time going through both malicious and non-malicious traffic, they start to develop better intuition when things look a little strange.

And remember, threat hunting should never be a one-man show that relies on a single "InfoSec rock star." Share documentation and grow together as a team. Chances are other team members will know something about what another hunter is hunting, or will find something else that helps with bias.

About ProtectWise

ProtectWise™ is disrupting the security industry with The ProtectWise Grid™, its enterprise security platform that captures high fidelity network traffic, creates a lasting memory for the network, and delivers real time and retrospective alerting and analysis in a rich, innovative visualizer. By harnessing the power of the cloud, The ProtectWise Grid provides an integrated solution with complete detection and visibility of enterprise threats and accelerated incident response. The ProtectWise Grid delivers unique advantages over current network security solutions, including an unlimited retention window with full-fidelity forensic capacity, the industry's only automated smart retrospection, advanced security visualization, and the ease and cost-savings of an on-demand deployment model. For more information, visit www.protectwise.com.

20170519

PROTECTWISE™

1601 WEWATTA ST • SUITE 700 • DENVER, CO 80202 • 1.303.625.6802
WWW.PROTECTWISE.COM • INFO@PROTECTWISE.COM