SANS | GIAC CERTIFICATIONS

WHITE PAPER

Is Your Threat Hunting Working? A New SANS Survey for 2020

Mathias Fuchs

SANS

A SANS Survey

Is Your Threat Hunting Effective?

Written by **Mathias Fuchs**

May 2020

Sponsored by:
Cyborg Security

Executive Summary

Last year's SANS threat hunting survey[1] revealed that staffing threat hunting teams adequately still proved to be a pain point in many IT security departments. The increase in organizations that execute threat hunting has led to a shortage of qualified personnel throughout the industry.

In this paper, we explore in more depth what exactly leads to the shortage of suitable personnel and how it affects security organizations' capabilities to utilize threat hunting teams. To grasp the impact of staffing challenges on threat hunting operations, we take a closer look at the metrics organizations are using to measure threat hunting effectiveness.

We also explore if and how security teams use threat intel to attenuate some of the adverse effects that a shortage of resourceful threat hunters has on organizations. We focus on the features that threat intel should exhibit to be useful, rather than a nuisance.

Analysis of the survey results indicates that even though some form of threat hunting has arrived in most organizations, there appears to be no consensus on exactly how threat hunting should look. Mainly, we still see some confusion about the daily tasks of SOC analysts versus the functions of threat hunters. The majority of our survey respondents rely heavily on tools, such as SIEMs and endpoint detection and response (EDR) tools. Even though both solutions offer the capabilities needed to support an adequate threat hunting operation, they usually do not come with batteries included. Many respondents asserted that employing the right experts to build up and maintain advanced threat hunting is challenging. First, the demand for experienced threat hunters appears to outweigh the supply. The second challenge our respondents face is the quality of threat intelligence. Even though the majority of respondents consume some type of threat intelligence for their hunting operations, only one of every three respondents said that they are highly satisfied with their sources.

The good news is that even though organizations are facing enormous challenges when introducing and running threat hunting operations, they still appear to benefit from them. Our results show that respondents are starting to put methodologies in place to measure the benefit of threat hunting. We believe that having these methodologies will lead to vast improvements in threat hunting operations, because measuring leads to more specific requirement definitions. These definitions help to shape threat hunting operations more precisely and make them more successful.

Ultimately, the survey results indicate that most of our respondents are on the right path and already seeing some success. Without a sufficient number of skilled staff, however, high-quality intelligence and the right tools to get visibility into the infrastructure success is limited. A world where we'll see a unified, widely accepted gold standard of threat hunting remains in the future, but we are headed in the right direction.

[1] "SANS 2019 Threat Hunting Survey: The Differing Needs of New and Experienced Hunters," October 2019, www.sans.org/reading-room/whitepapers/analyst/membership/39220

About the Respondents and Their Organizations

The 2020 survey had 255 respondents, with demographics as shown in Figure 1.

Top 4 Industries Represented

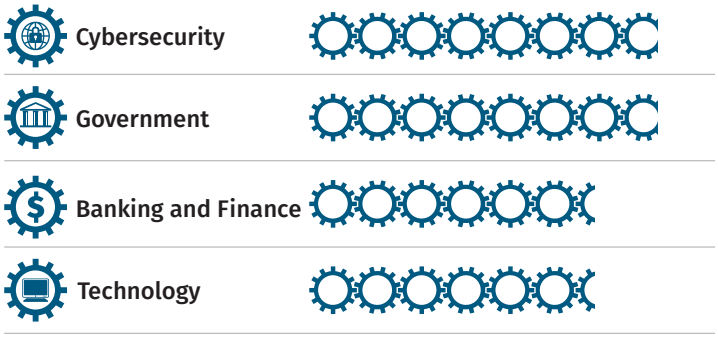

Each gear represents 5 respondents.

Operations and Headquarters

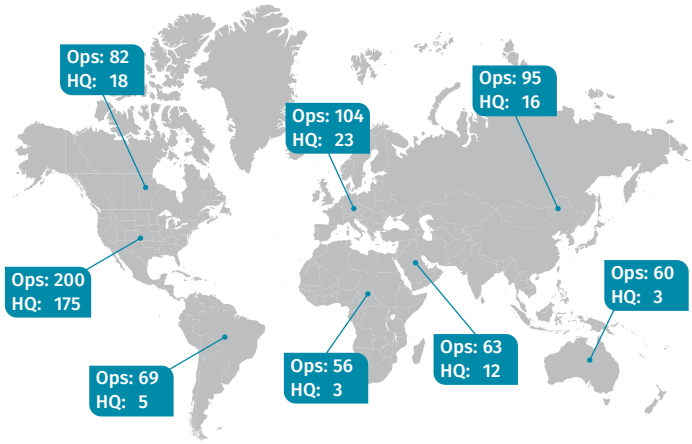

Organizational Size

Each building represents 5 respondents.

Top 4 Roles Represented

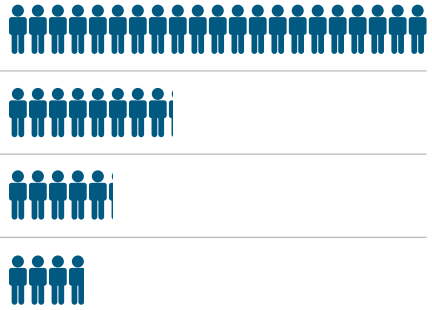

Each person represents 5 respondents.

Figure 1. Key Demographic Information

How Does Threat Hunting Work Today?

In their 2018 paper on threat hunting, authors Dan Gunter and Marc Seitz defined threat hunting as "the formal practice of threat hunting [which] seeks to uncover the presence of attacker tactics, techniques, and procedures (TTP) within an environment not already discovered by existing detection technologies."[2] That concise definition touches on most points that constitute threat hunting. First and foremost, organizations should conduct threat hunting in addition to using existing detection technologies to shed light on the dark spots in their detection methodologies and technologies. Threat hunting is a human-driven approach that operates outside the well-defined and controlled envelope of automated threat detection. Threat hunting is not only designed to detect adversaries, but also to uncover visibility gaps in detection mechanisms that are already in use.

> *Threat hunting is a human-driven approach that operates outside the well-defined and controlled envelope of automated threat detection.*

[2] "A Practical Model for Conducting Cyber Threat Hunting," November 2018, www.sans.org/reading-room/whitepapers/threathunting/paper/38710

In the SANS 2019 threat hunting survey,[3] we presented a model that introduced maturity levels for threat hunting operations, as shown in Figure 2.

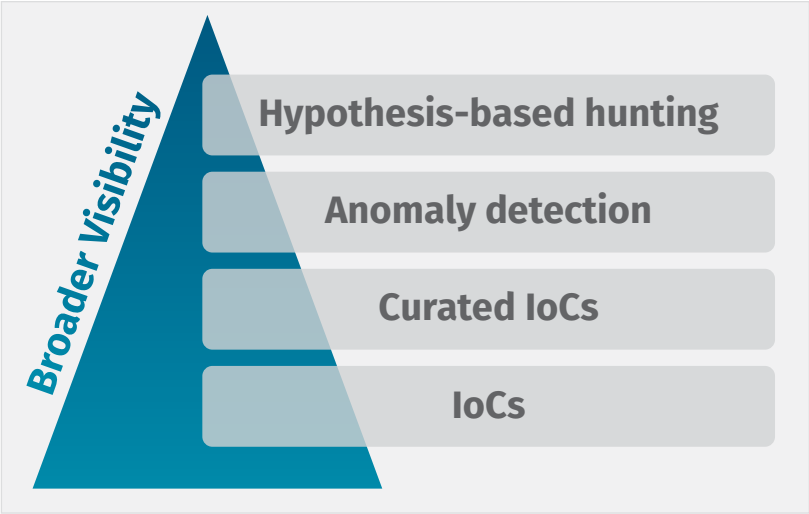

Figure 2. Threat Hunting Maturity Pyramid

Using indicators of compromise (IoCs) in an all-in approach forms the base of the pyramid. Organizations at this level consume threat intelligence in the form of IoCs and actively sweep their environment for them. These IoCs are rarely tailored for their environment. Thus, they often lead to a considerable number of false positives. This approach, however, does not uncover visibility gaps.

Evolving past the initial stage of threat hunting means curating IoCs. Instead of sweeping the environment with a large quantity of IoCs, hunters use a smaller number of high-quality IoCs that fit their environment. Curation not only requires selecting a subset from a vast pool of IoCs, but also putting context behind every single IoC to give them meaning. The process of curating IoCs is a task that organizations might outsource to their threat intelligence providers. However, not all vendors exercise the same level of quality control on the intelligence they release, so the task for staff threat hunters shifts from curating IoCs themselves to selecting the right vendors. The use of curated IoCs can be viable to detect adversaries, but usually does not help to identify visibility gaps.

The use of curated IoCs can be viable to detect adversaries, but usually does not help to identify visibility gaps.

Anomaly detection constitutes the next stage of the pyramid. Many detection technology vendors rely heavily on anomaly detection, which creates a baseline and detects outliers by using machine learning algorithms or implementing proprietary static algorithms. Organizations usually use these products to increase their detection capability. However, this is not the kind of anomaly detection that appears to be beneficial for threat hunting—only 14% of respondents have high confidence in those tools.

In a threat hunting context, anomaly detection is a more iterative and open-ended process. One example would be data stacking, where an analyst acquires a particular set of data, such as a list of all running processes within the environment. The analyst then counts the occurrence of every unique process throughout the environment to create a baseline. Because targeted malware is the exception rather than the norm, it will show a low frequency of occurrence. The power of stacking lies in the combination of different stacks that skilled hunters build dynamically based on what they find during the hunt. Even though this approach is well suited to detect adversaries, there is no guarantee it will also reveal visibility gaps.

Pro Tip: Stack RDP Logons to Uncover Visibility Gaps

An analyst extracts all type 10 logons from all Windows security logs in the environment. The analyst then extracts the source IPs for all logon events. In many situations, matching this list against the endpoints covered by security tools may identify visibility gaps.

[3] "SANS 2019 Threat Hunting Survey: The Differing Needs of New and Experienced Hunters," October 2019, www.sans.org/webcasts/2019-threat-hunting-survey-differing-experienced-hunters-111010

At the top of the pyramid is hypothesis-based threat hunting. This approach requires analysts to formulate a hypothesis on how a particular attack could happen. That hypothesis relies heavily on threat intelligence about the organization's specific risk profile. Analysts might still use anomaly detection and curated IoCs in this stage, but they do it very selectively and consciously. IoCs need context to be of any value at this point. Hypothesis-based threat hunting can meticulously identify ongoing attacks and even rule out specific attack patterns.

Additionally, hypothesis-based threat hunting does not only identify visibility gaps, but also the most dangerous visibility gaps. This form of threat hunting relies on well-curated IoCs that come with context, high visibility in the environment, a well-built hypothesis and management support.

How to Build a Hypothesis

Hypothesis-based threat hunting starts with building a hypothesis. Threat hunters begin that process by examining a multitude of factors. For example, let's say that industry peers report on a recent breach. Threat hunters then build the hypothesis that a similar breach happened to their organization. They compile everything there is to know about the recent breach, as well as a list of artifacts that they suspect they will see in their environment if a similar breach were to occur. That's where sole indicator-based threat intel is not enough—threat hunters heavily rely on context around the indicators to formulate the hypothesis and plan the hunt.

Now it's time to find ways to sweep the environment for those artifacts. These sweeps might be very targeted ("Get all error logs from all of our Tomcat servers running applications that use Struts 2," for instance). During the acquisition, analysts might recognize that they lack ways of getting data from some parts of the environment. It's their task to report these visibility gaps and, ideally, get visibility in due time. Many times, analyst inexperience with the existing technology causes visibility gaps. It takes time to repurpose detection equipment for human threat hunting.

After the threat hunters acquire and analyze all available data, they can either confirm or reject the hypothesis.

The Emotet PowerShell payload is a good example of the difference between simple detection queries based on IoCs and advanced detection queries.

The Emotet banking Trojan has been extremely prevalent in the modern threat landscape, often being the most observed malware. Although intended to bypass security controls and detection, the developers of the malware left a rather consistent method of detection behind in its Base64-encoded PowerShell command. Many researchers have noted this and created a simple signature (which is freely available) to detect its various strings, as shown in the example for Splunk in Figure 3. Before this Trojan is run in an environment, it can heavily be improved upon for accuracy and efficiency.

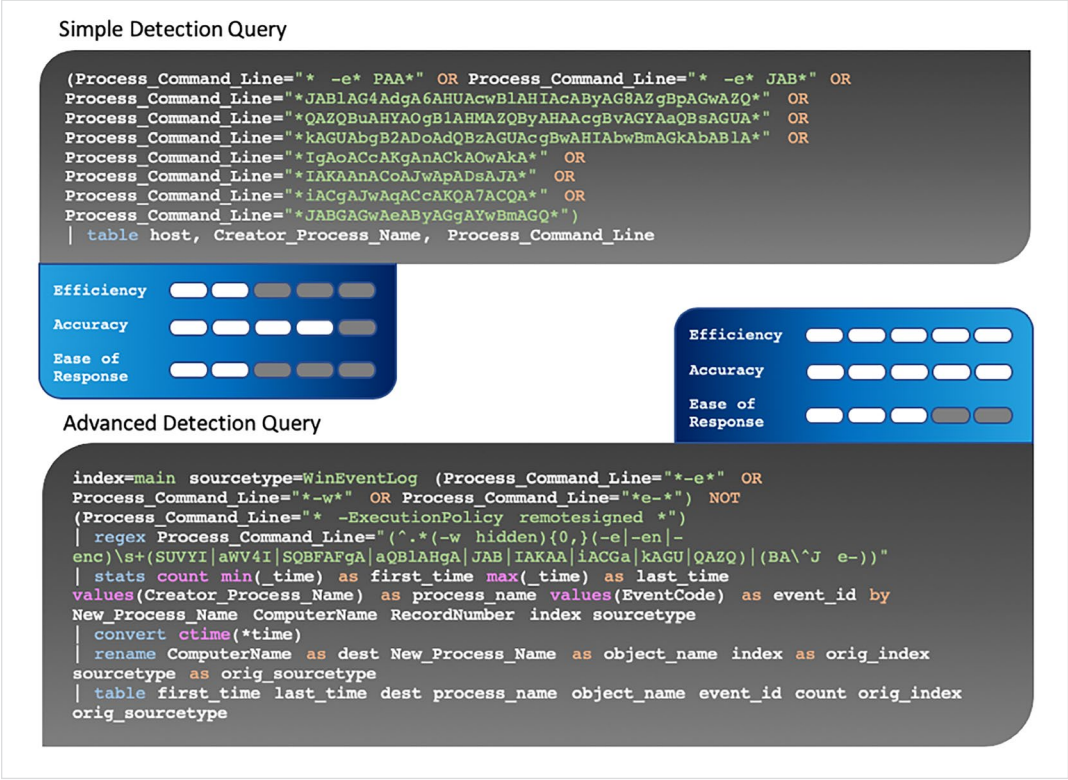

Figure 3. Splunk Example of Signature to Detect Emotet Banking Trojan

Has Threat Hunting Arrived in Most Organizations?

Challenges in IT security have changed massively over the past decade. Not so long ago, the main focus rested on protection, rather than detection. However, the industry eventually understood that investing in protective measures only was a losing game, and organizations have since shifted some of their spending to detection measures. But have organizations made the shift from passive detection to active hunting yet? Our survey results say yes. Sixty-five percent of respondents indicated that they already perform some form of threat hunting, and 29% plan to do so in the next 12 months. Only 2% of respondents claimed that they don't run threat hunting operations and don't intend to in the future.

Looking at the self-proclaimed maturity level, only 29% of respondents consider themselves mature or very mature when it comes to threat hunting. That's no surprise, because the concept of threat hunting is still relatively new and there is no blueprint for how to implement it in organizations. At best, low maturity levels result in inefficient hunting, and at its worst, ineffective hunting. (See Figure 4.)

Figure 4. Threat Hunting Maturity Level Assessment

When we asked respondents why they assessed their maturity in the way they did, their answers varied. While some respondents based their assessment only on their own opinion, others suggested more objective measures, such as the ratio between stopped and not-stopped attacks over time. From the responses, we observe a general trend that organizations are starting to build dedicated threat hunting teams, which in turn appears to lead to staffing issues because allocating the right talent becomes a greater challenge. There is also a strong disconnect in the numbers. While 70% of respondents have dedicated in-house staff doing the threat hunting, only 29% believed that they are mature or very mature when it comes to performing the task. That disconnect appears to relate back to the challenges in staffing qualified professionals as dedicated threat hunters.

Although many organizations struggle to attract qualified threat hunters, only 21% of respondents currently outsource their threat hunting activities. Despite that, the majority relies on externally produced threat intelligence.

Technical Aspects of Threat Hunting

Threat hunting may depend heavily on the human factor, but organizations also need the right tools in their tool chest to succeed. We asked respondents what tools and technologies their organizations are using today and what their collection of tools will look like in 12 months. Today, a broad majority (85%) relies on automated alerting tools, such as SIEM, IDS/IPS and EDR. (See Figure 5.)

Configurable, customizable, internally developed search tools come in at number two (61%). That percentage is an exciting development and indicates that threat hunting needs to go beyond what today's market tools deliver. Usually, creating these internal tool sets and using them to hunt goes back to a small number of highly skilled analysts. Keeping in mind that staffing appears to be one of the key challenges in the industry, the success with this approach is hardly scalable and only occasionally repeatable.

What tools/technologies do you currently use?
What tools/technologies are you planning to implement in the next 12 months?
Select all that apply.

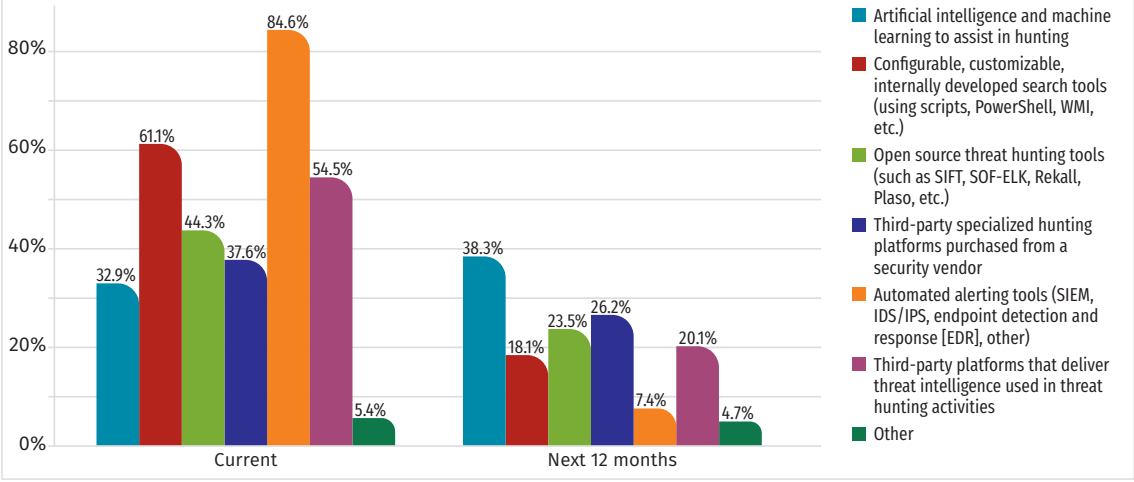

Figure 5. Tools and Technologies in Use Now and Planned for the Future

Even though only a third of respondents use AI and machine learning tools in their threat hunting activities, many other respondents have set their eyes on these tools for the future. Thirty-eight percent are looking to implement AI and machine-learning–based technologies in the next 12 months. However, when asked about their confidence in these tools, only 14% stated that they had a high confidence level.

Respondents are most satisfied with automated alerting tools, such as SIEM, IDS/IPS and EDR, with a high and medium combined confidence of 85%. Second place goes to configurable, customizable and internally developed search tools, with 57% indicating overall satisfaction with this technology. Interestingly, respondents' confidence in third-party platforms that deliver threat intelligence used in threat hunting activities is low, as demonstrated by the highest low confidence rating at 7%. (See Table 1.)

Table 1. Level of Satisfaction with Tools and Technologies

Tools/Technologies	Level of Satisfaction			
	High	Medium	Overall (High + Medium)	Low
Automated alerting tools (SIEM, IDS/IPS, endpoint detection and response [EDR], other)	41.7%	43.2%	84.9%	5.0%
Configurable, customizable, internally developed search tools (using scripts, PowerShell, WMI, and the like)	30.9%	25.9%	56.8%	6.5%
Third-party platforms that deliver threat intelligence used in threat hunting activities	23.0%	26.6%	49.6%	7.2%
Open source threat hunting tools (such as SIFT, SOF-ELK, Rekall, Plaso, and the like)	26.6%	18.7%	45.3%	0.7%
Third-party specialized hunting platforms purchased from a security vendor	18.7%	18.7%	37.4%	2.2%
Artificial intelligence and machine learning to assist in hunting	13.7%	14.4%	28.1%	6.5%
Other	2.9%	0.7%	3.6%	0.7%

The challenge in making threat intelligence usable in threat hunting depends on two factors. First, we need to get accurate, well-curated threat intelligence and understand the right way to deploy that threat intelligence to the environment. Second, we need threat intelligence providers to place greater focus on providing quality, rather than quantity, and we need the capacity to apply high-quality threat intelligence to organizations' environments.

To accomplish that, we need to understand where organizations are today and where they struggle in the process. As shown in Figure 6, 82% of respondents use open source intelligence (OSINT). The most significant advantage of OSINT is that it's usually free; the biggest disadvantage is that it's frequently up to the consumer to curate the data, which is especially difficult when it comes to adding context to technical indicators. As a result, many organizations consume free feeds but don't tweak them to meet their needs.

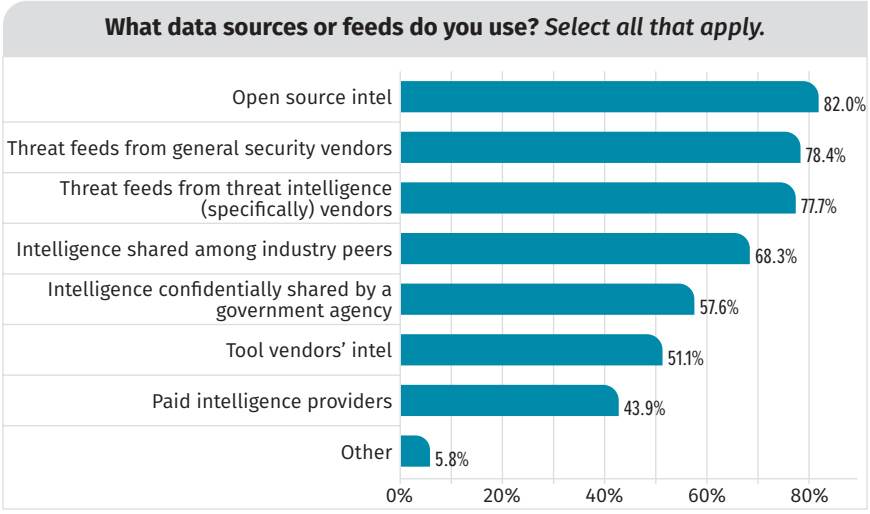

Figure 6. Data Sources or Feeds in Use

Seventy-eight percent use threat intel feeds compiled by either general security vendors or intelligence vendors. In our experience, general security vendors often focus more on quantity than quality, while specialized intelligence vendors usually wrap their intel with context. Unfortunately, organizations often base their buying decisions on quantitative measures, such as the number of indicators, because nontechnical people are involved and sometimes leading the purchasing process. ("Why should I pay more for an organization that gives me three pieces of intel a week versus one that gives me 100 a day?")

Sixty-eight percent of respondents share intelligence among industry peers. That intelligence is usually not publicly available and, thus, not visible to potential attackers. For that reason and because industry peers typically share the same risk landscape, we consider that valuable threat intelligence.

So how satisfied are organizations with the sources of their threat intelligence? The biggest complaint overall appears to be that the provided threat intelligence is too generic, reactive and biased, thus keeping the false-positive rates fairly high. Thirty percent of respondents have high confidence in threat intelligence shared among industry peers, while only 15% fully trust tool vendors' intel. Threat intelligence originating from specialized threat intelligence vendors ranks a close second place for the highest level of satisfaction, at 29%. (See Table 2.)

Table 2. Level of Satisfaction with Data Sources for Threat Hunting

Data Sources/Feeds	Level of Satisfaction			
	High	Medium	Overall (High + Medium)	Low
Threat feeds from threat intelligence (specifically) vendors	29.4%	41.9%	71.3%	3.7%
Open source intel	23.5%	47.8%	71.3%	7.4%
Threat feeds from general security vendors	17.6%	47.1%	64.7%	11.8%
Intelligence shared among industry peers	30.1%	33.1%	63.2%	3.7%
Intelligence confidentially shared by a government agency	25.7%	23.5%	49.3%	8.1%
Tool vendors' intel	15.4%	27.9%	43.4%	3.7%
Paid intelligence providers	18.4%	20.6%	39.0%	3.7%
Other	3.7%	0.7%	4.4%	0.0%

The Value of Threat Hunting

Threat hunting doesn't come cheap. Organizations need the right people, the right tools and time to get everything running smoothly. A proper grasp on how much threat hunting affects the security posture of an organization is critical to justify current and future expenses.

What Improvement Have Respondents Seen Over the Past Year?

We asked respondents how much threat hunting improved the overall security of their organization in the past 12 months. The median response lies at 40%, which is quite an impressive perceived improvement for an observation period of only a year. Even though we believe that threat hunting, if applied correctly, will have a positive effect on the security of any organization, we can't trust these numbers too much because only 37% of respondents claimed that they formally measure the success and effectiveness of threat hunting. Even though threat hunting is still a considerably new and hyped topic, justifying expenditures will become even more crucial when organizations face uncertain economic times.

So how do those 37% measure the effectiveness of threat hunting? The majority (58%) manually track threat hunting activities and outcomes; 33% do the same using automated tools. (See Figure 7.)

Interestingly, 43% considered the number of legitimate alerts generated by threat hunting as a measure of effectiveness. This measure, however, may actually be represented by the false-positive rate of applied intelligence, which—strictly speaking—is a metric that applies more to security operations than threat hunting. However, that measurement exposes the quality of the applied threat intelligence and, thus, is a valuable parameter for future buying decisions.

Figure 7. Methods Used to Measure Threat Hunting Effectiveness

Thirty-five percent of respondents measure the time to respond to alerts generated by threat intelligence sources. While response time is not a threat hunting metric, it shows how entangled the perception of threat hunting with security operation center (SOC) tasks still is. We can't emphasize enough that threat hunting activities and SOC activities have a different focus, as described in the "2019 SANS Threat Hunting Survey" referenced earlier.

Threat hunting metrics boil down to two significant groups of parameters: One is the amount of damage averted, and the other is visibility gained. As easy as this sounds, it isn't.

It's usually quite difficult to estimate the potential damage of an attack, especially when it's supposed to cover all adverse effects on the attacked organization. That means besides financial harm, you'd also need to factor in the loss of reputation and the like. When threat hunting activities catch the attacker early on, it's usually impossible to calculate the impact of what would have happened. The resulting paradox is that organizations only get good numbers when threat hunting fails and they experience the whole breach. Organizations that sport a mature risk management process might be better off because the various risks that support threat hunting metrics may already have price tags attached.

The second significant metric, visibility parameters, can also be tricky. The Hawthorne effect[4] describes that individuals and groups might behave differently when they are aware that they are being observed. In other words, measuring people affects what they do. So the danger in using gained visibility as the sole factor to measure threat hunting is that threat hunters might put more focus on identifying visibility gaps than hunting actual adversaries. The desired process of identifying visibility gaps is to have a hypothesis first and then figure out where and how to get the data needed to accept or reject the hypothesis. If that data is not available, the hunters have identified a crucial visibility gap. Only visibility gaps that come with a sound hypothesis are a sign of good threat hunting practice.

The danger in using gained visibility as the sole factor to measure threat hunting is that threat hunters might put more focus on identifying visibility gaps than hunting actual adversaries.

Based on these two parameters, we asked where our respondents saw the most significant improvement.

Ninety-three percent experienced a reduction in their attack surface. That metric relates to the identification of visibility gaps. Either respondents gained better insight into the vulnerability of assets or got first sight of the assets after remediating visibility gaps. For 89% of respondents, threat hunting improved detection creation and reduced false positives. Even though threat hunting and continuous monitoring are different activities, one might still benefit from the other. If a one-time threat hunting exercise creates detection capability that the SOC can adopt, it's great. Just be aware that it does not always work that way. (See Table 3.)

Table 3. Threat Hunting Improvements

	Level of Measurable Improvement				
	None	Some	Significant	Overall (Some + Significant)	Unknown
Attack surface exposure/hardened network and endpoints	3.8%	41.4%	51.1%	92.5%	3.0%
Creation of more accurate detections and fewer false positives	6.0%	41.4%	47.4%	88.7%	3.8%
Resources (e.g., staff hours, expenses) spent on remediation	11.3%	48.1%	27.1%	75.2%	7.5%
Exfiltration detection (data detected leaving your organization)	12.8%	36.8%	33.1%	69.9%	11.3%
Breakout time (initial compromise to lateral movement)	12.8%	42.9%	24.8%	67.7%	12.0%
Other	5.3%	3.8%	3.0%	6.8%	1.5%

[4] https://en.wikipedia.org/wiki/Hawthorne_effect

Sixty-eight percent observed an increased breakout time. That means that attackers took longer to pivot from patient zero to other systems. This can be a side effect of the first two points (smaller attack surface and better detection).

For 70% of respondents, exfiltration detection increased, and that improvement can be traced to fewer visibility gaps. You can only see something leaving your premises when you know it was there in the first place.

Lastly, 75% needed fewer resources for remediation activities, presumably because the attackers had less time in the environment. This number suggests that even though most organizations are still unsure about how to measure the effectiveness of threat hunting, they do have proper measures in place already (as indicated in Figure 6 on page 8).

What Is Holding Organizations Back?

We've now established that threat hunting has a positive effect on most organizations' security posture. So what prevents them from maximizing their success?

Unsurprisingly, the clear winner is the lack of skilled staff (72%)—the current approach of non-streamlined, very individual threat hunting relies heavily on professional input into hunting content and tools. That response aligns with budget constraints (51%) and the lack of defined processes (50%). Only 32% indicated that a lack of management support is a barrier. (See Figure 8.)

We also asked what needs to improve to mature threat hunting operations. In addition to looking for qualified staff to run the hunts (53%), respondents are looking for the enhanced contextual awareness that intelligence sources and tools provide (51%). Cloud-based hunting also poses a challenge to many organizations (46%), with a need to acquire tools and capabilities that can extend to the cloud.

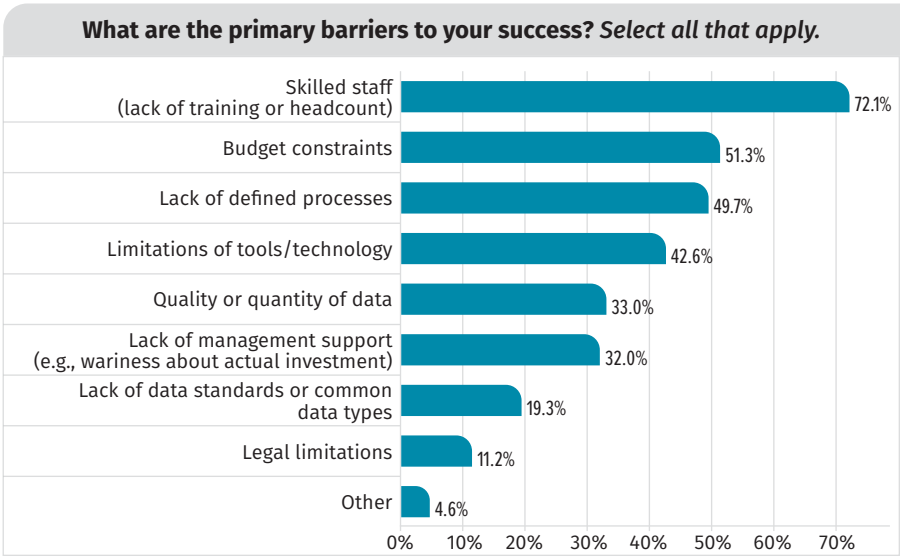

Figure 8. Primary Barriers to Threat Hunting Success

Measuring Threat Hunting

Measuring threat hunting is challenging, yet it's imperative to mature. Organizations need to implement precise requirements against which they can measure threat hunting. According to our survey, only 20% of respondents are documenting threat hunting requirements. At least 50% plan to define requirements eventually, and 24% indicated that they have requirements, but those requirements are ad hoc. While ad-hoc requirements can help in early stages, they are hard to measure over time compared to well-documented requirements. Undocumented requirements also tend to shift to whatever suits the involved parties, rather than what supports the cause. (See Figure 9.)

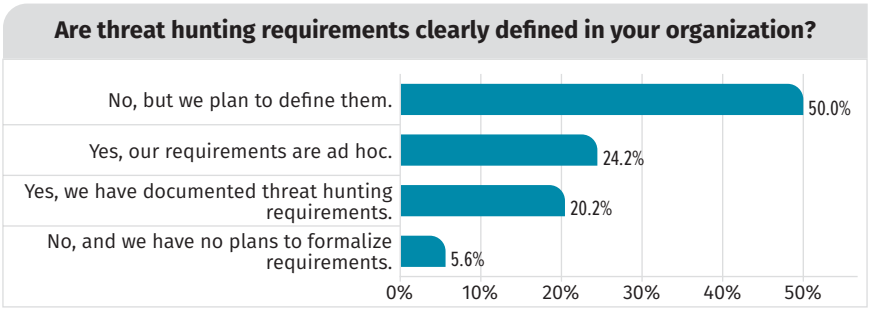

Figure 9. Clear Definitions of Threat Hunting

Who defines the requirements? The numbers show that requirements are mostly formulated from inside the threat hunting teams and rarely introduced by senior management. While it's generally a good idea to involve knowledgeable threat hunters and other security personnel in the process of defining requirements, the organization's executives must determine the strategic imperative. They decide how much risk the organization is willing to take and how much money to spend to minimize risk. On the flip side, executives are number-driven. Requirements need to be measurable, and it's always good to have some budget allotted to research that's not tied to strict requirements. (See Figure 10.)

In the end, requirements factor into staffing, vendor selection, build or buy decisions, and more.

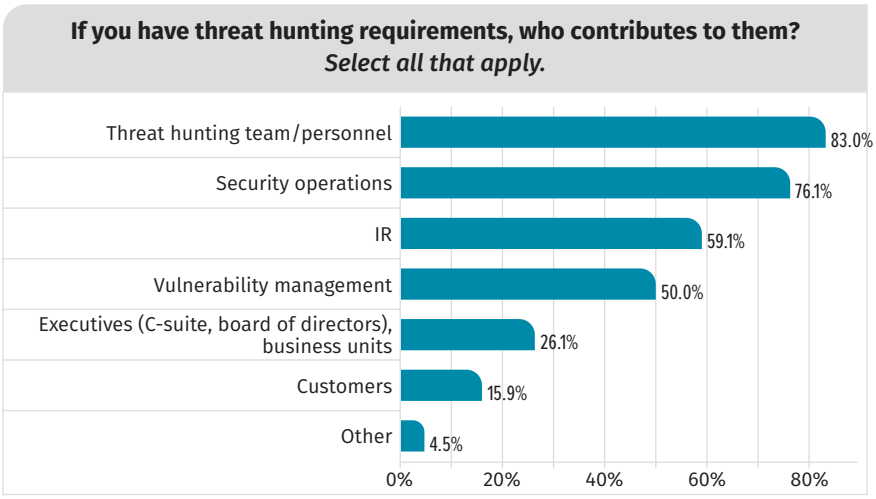

Figure 10. Who Contributes to Threat Hunting Requirements?

Even though most organizations don't have formal requirements, for those that do, 76% use them to assess the effectiveness of their threat hunting operations. That not only defines how organizations' threat hunting teams work, but also the quality vendors need to deliver to remain in the market. When the majority of organizations operate on precise requirements, the quality of tools and intelligence on the market will improve.

Conclusion

Threat hunting, in some form, has arrived in the majority of our respondents' organizations. Many others plan to introduce threat hunting within the next year. Threat hunters mostly rely on tools such as SIEMs, IDS/IPS or EDR and have high confidence in these tools. To run threat hunting operations, they need well-curated and accurate threat intelligence that comes with a particular context attached to every indicator. In addition, our respondents require skilled staff to design and run the hunts. Both good intelligence and qualified staff are hard to get, according to the responses we received.

Even though it's still a challenge to mount a top-notch threat hunting operation, our respondents saw significant improvement in their security posture because of threat hunting. They also identified better ways to measure the effects of threat hunting accurately, allowing them to formulate more precise requirements definitions, eventually leading to more streamlined and better-designed threat hunting operations. Requirement-driven threat hunting operations will help the community and vendors to develop even more useful knowledge and tools in the future.

There is currently no blueprint for how to run standardized threat hunting, which is why success depends so strongly on skilled and experienced hunters. Employing top-notch threat hunters is ideal, but to mature, organizations need to explore ways to equip new and unexperienced threat hunting teams with the right tools, the right content and the right mindset. That's the only way to scale.

About the Author

Mathias Fuchs, a certified instructor for SANS FOR508: Advanced Digital Forensics, Incident Response, and Threat Hunting, is head of cyber defense at InfoGuard AG, where he is actively engaged in building the incident response (IR) practice. In that role he uses his knowledge to shape his team; develop the necessary forensic, IR and threat hunting capabilities; and proactively mediate security vulnerabilities that would be more difficult to manage later. Prior to joining InfoGuard, Mathias was a principal consultant at Mandiant, where he led large-scale cybersecurity investigations. He also was the lead security architect at T-Systems and a security consultant for international clients in a variety of industries.

Sponsor

SANS would like to thank this paper's sponsor:

CYBORG SECURITY