



SANS 2019 Threat Hunting Survey: The Differing Needs of New and Experienced Hunters

Written by **Mathias Fuchs**
and **Joshua Lemon**

Sponsored by:
DomainTools

October 2019

Executive Summary

The 2019 SANS Threat Hunting Survey gathered current industry data from 575 respondents predominantly from small/medium to medium/large organizations that are working in the field of threat hunting or working alongside threat hunters. This year's report aims to help organizations understand what threat hunting is, why it is essential to protect their organizations, and how novice and experienced hunters can improve their processes.

Results demonstrate that confusion still exists about what respondents believe constitutes threat hunting and how to properly approach threat hunting. In addition to uncovering these areas of confusion, the report offers practical takeaways and action items that readers can use to strengthen their cybersecurity defenses within their organizations.

In this year's survey, we explore how threat hunting teams are tasked in an environment, where they hunt and how they hunt. More than half of the respondents use atomic indicators of compromise (IoCs) or an alert-driven approach to hunting. This year's survey results show that respondents have decreased their hypothesis-driven hunting over the past three years, which may pose some dangerous visibility gaps for organizations.

The results confirm that many organizations are still dual-tasking threat hunters, and very few have progressed over the past three years to standing up a dedicated team. It seems that threat hunting is still seen very much in its infancy for most organizations. This report explores how teams are structured, the priorities given to hunters along with other roles they are fulfilling in the organization, and how an organization resources a threat hunting team.

This report recognizes that organizations are still concentrating on technology as a key driver for increasing the capabilities of a threat hunting team. However, we question how useful a tool may be in the hands of an unskilled hunter, especially if training is not seen as a critical area to enable hunt teams.

Results indicate that organizations are still struggling to measure the benefits—or organizational impact—a threat hunting team can have. We suggest a process threat hunters can use to demonstrate to management why threat hunting is essential and how threat hunters can begin measuring the impact they are having in their organization.

This year's report provides several key takeaways and action items that readers should consider integrating into their threat hunting programs. We encourage anyone running a threat hunting team to start implementing change as soon as possible to ensure that your teams can keep pace with the ever-changing attack vectors and advances by adversaries.

Key Findings

- **35%** create hypotheses to drive their threat hunting efforts.
- **56%** use threat intelligence to hypothesize where attackers may be found.
- **34%** of hunters have major responsibilities for managing SOC alerts; **26%** perform IR and forensics of current breaches.
- **71%** indicate technology is the first or second focus of their threat hunting resources, followed by staffing (**47%**) and training (**41%**).
- **61%** report at least an **11%** measurable improvement in their overall security posture.

Survey Demographics

Survey respondents represented a global group of threat hunters from various organizations as well as management functions in security. Figure 1 provides a snapshot of respondents.




Figure 1. Key Demographic Information

Just under 8% of our sample (44 respondents) provided another job title we had not included in our standard list of related roles. They included such titles as security architect, threat intelligence analyst and threat researcher. Such roles and titles certainly offer great expertise in threat hunting teams.

Definitions of Threat Hunting

There are many misconceptions about what exactly qualifies as threat hunting. At its heart, threat hunting is a proactive approach to identifying signs of an attack, as opposed to the more reactive approach security operations center (SOC) analysts follow. Organizations with well-established hunting operations—especially proactive ones—have a better chance of catching an attacker early in the attack. But what exactly is threat hunting?

Description of Threat Hunting

Threat hunters leverage tools—and a whole lot of experience—to actively sift through network and endpoint data, always looking for suspicious outliers or traces of an ongoing attack. They consume threat intelligence to understand the tactics, techniques and procedures (TTPs) of attackers better. Most importantly, hunters create a hypothesis on how a potential attack might happen and search for data to prove or reject the hypothesis.

Is there a difference between threat hunting, incident response (IR) and SOC activities? For some, this might be an easy question to answer; however, for a large segment of respondents, these three areas still blur together. That blurred line is understandable, given that they are all interrelated and necessary. But they do have very defined roles, as illustrated in Figure 2.

The SOC is your “eyes on glass” team, continuously monitoring incoming security alerts from a SIEM and triaging them to determine whether they point to signs of malicious activity. The IR team is your firefighting team. Its members spring into action when your SOC has found something malicious and you need to determine how the incident occurred, how widespread it is, and where else the attackers are in the environment. Your threat hunting team covers the areas that your SOC is not watching and your detection mechanisms haven’t detected. Think of your threat hunters as a human, extremely intelligent SIEM solution that sniffs out evil in the environment based on the hypotheses it develops.

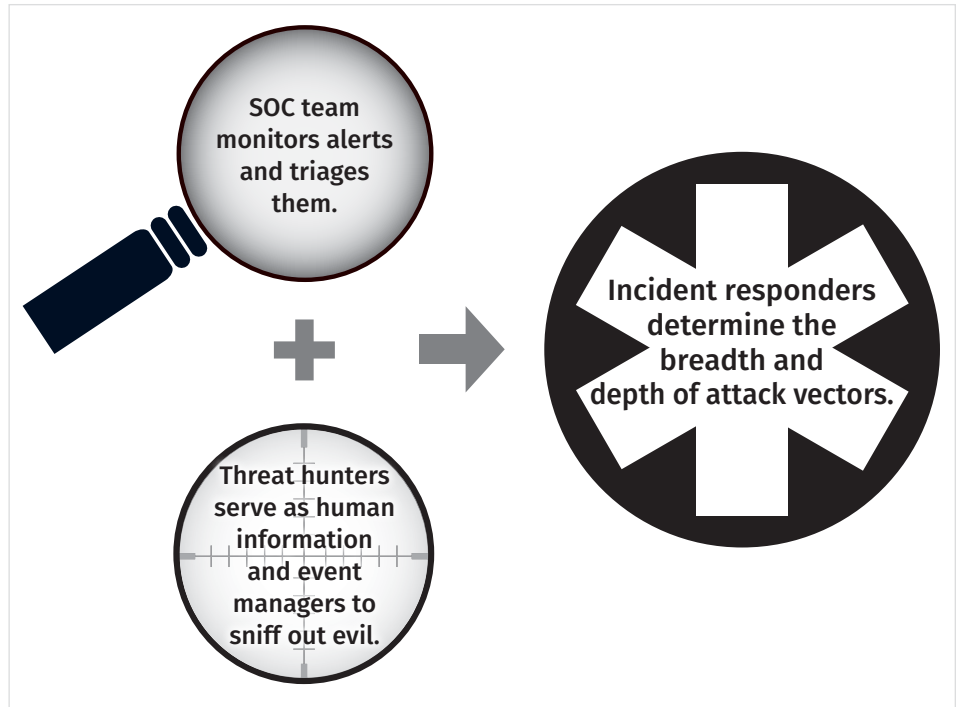


Figure 2. Interactions of Threat Hunting, SOC and Incident Response Teams

Examples from Respondents

Threat hunting is often misunderstood, or confusion arises about the differences in roles between SOC, IR and threat hunting. One way to understand the role of threat hunters is to really detail what they are accomplishing in their role as threat hunters.

Good Examples

The following examples are of processes defined by respondents that fit our definition of what threat hunting is. All are proactive steps that uncover a threat before an alert is sounded, or in the absence of an alert occurring.

- “Artifacts or TTPs are noted as a starting point for a typical hunt. The hunt is planned for a certain number of hosts based on intelligence. Normally a hunt uncovers some IOCs, result[ing] in a patching or triage effort and delivers interesting items for future hunts.”

- “The hunter creates a hypothesis to search on and uses a diversity of data source[s] to confirm their hypothesis.”
- “We do ad hoc threat hunting quarterly and also based on anomalies. We attempt to look for IoCs or other data related to a hypothesis and rule out the existence of certain threats/actors.”
- “Detect the undetected by understanding attackers’ methodology (TTPs) formulating new hypothesis asking, ‘Can we detect this?’”
- “Our great hunters usually find gaps in our detection or even deliver a use case that is used to create a new SIEM rule.”

To their credit, 35% of respondents create hypotheses to search for new threat activity. In addition, 56% use threat intelligence, such as adversary TTPs, to hypothesize where attackers might be found. These types of proactive approaches provide the best examples of what threat hunting teams should be doing.

Bad Examples

Unfortunately, some respondents still define threat hunting as reacting to or monitoring alerts. The following examples reflect that approach. In each of these examples, the hunters are simply reacting to alerts from tools or third parties. They are not proactively searching for potential threats within their environment.

- “Managing SOC alerts from SIEM and IDS/IPS sources”
- “Alert points to malware in file on machine or user account—hunter investigates alert, files involved, machines/users involved and engages as many others as necessary to verify that any threat(s) [have] been contained and completely remediated.”
- “Using alerts, start a hunt or grab IoCs from threat intelligence platform and hunt.”
- “As we [receive] alerts from different sources, we investigate the events. If there is a need to escalate the incident, we open the ticket.”
- “They are monitoring alerts from our clients and, starting from this, they are performing threat hunting.”
- “Checking logs and looking for something malicious then perform sandboxing”
- “We typically start with dashboards from our various technologies and start chasing things that seem unusual or odd when we have the time, which can be rare.”
- “Monitors alerts from EDR [endpoint detection and response] and investigates anomalies”

In fact, most respondents use a variety of reactive approaches, including using alerts (40%), or using IoCs via a SIEM or alerting system, to find adversary tools or artifacts (57%). Such approaches are valuable as supplements to proactive approaches, but should not be used as standalone threat hunting procedures.

Action Item

Take a moment to think about the threat hunting process your organization follows. Then ask yourself whether that process is proactive or reactive. If it is reactive, consider what you can do to become more proactive. For example, are your hunters watching alerts on a screen? Could they transition to covering the gaps that your security alerting solutions do not see? If they can, you will be utilizing your hunters to increase the defense of your organization by doing effective threat hunting.

Building and Running Threat Hunting Teams

Threat hunting is a complex yet rewarding task. Today's general conception of threat hunting is quite new. In the past couple of years, threat hunting teams have started to form in most larger organizations, but there is no norm of how exactly a threat hunting team should be put together.

Who Performs Threat Hunting?

While most organizations exercise some forms of threat hunting today, many threat hunting teams have additional major responsibilities, such as managing SOC alerts (34%), investigating incidents and breaches (26%) or designing security infrastructure (14%). Consequently, only 29% of the respondents conduct threat hunting as a formal program with specifically assigned staff. That is only a slight improvement of 0.8% compared with the 2018¹ survey and 1.8% compared with the 2017 survey.² Still, the results are moving in the right direction, as organizations recognize the importance of hunting.

The majority of the respondents either conduct threat hunting on an ad hoc basis (43%) or outsource it to a third party (7%). This year's numbers also show that organizations are not more likely to outsource threat hunting than in 2017.

Attackers become more versatile, and IT infrastructures grow more complicated every year. The borders between cloud and on-premises services start to disappear. Facing these challenges, threat hunting needs to find its way into all organizations that depend on IT to run their businesses, presumably close to 100% of all organizations. Threat hunting must be introduced as a structured approach supported by well-trained personnel, the right toolsets and a plan on how to mature over time. Even though SOC analysts hunting for threats in an ad hoc fashion is currently the predominant approach and could be construed as providing some proactive hunting, we believe it takes more to be successful and ready for future security challenges. Threat hunters require different skill sets than SOC analysts do.

Team Sizes and Structures

Today, the typical hunting team consists of one to four dedicated hunters (42%). Only 22% of threat hunting teams have more than five members. Because only 19% of the respondents plan to increase their investment in staffing by more than 25%, team sizes will not change by much with the addition of just one or two staff members. We don't see that as particularly bad, because threat hunting is supposed to be a qualitative rather than quantitative process.

¹ Data collected but unreported in "SANS 2018 Threat Hunting Survey Results," September 2018, www.sans.org/reading-room/whitepapers/analyst/2018-threat-hunting-survey-results-38600 [Registration required.]

² "The Hunter Strikes Back: The SANS 2017 Threat Hunting Survey," April 2017, www.sans.org/reading-room/whitepapers/analyst/hunter-strikes-back-2017-threat-hunting-survey-37760, p. 10, Figure 5. [Registration required.]

Profile of a Threat Hunter

Organizations value different traits in a threat hunter, with knowledge in baseline network communication and activity leading the way, valued by 75% of respondents as shown in Figure 3.

Technical skills such as IR (71%) and threat intelligence and analysis (67%) follow closely. In fact, of the list provided, only memory forensics was significantly less valued. Some respondents wrote in softer skills, such as curiosity, as critical traits for being a good hunter. We agree with that and would add persistence, observation and attentiveness to the list.

Threat hunters with an IR background are valued because modern threat hunting encompasses methodologies and techniques every incident responder should be familiar with. Additionally, 26% of the threat hunting team members are also involved in IR and forensics when not proactively hunting.

Key Takeaways

Threat hunting is still in its infancy. For that reason, team structures are not yet well-established for most organizations. Many are just starting with threat hunting, leveraging the resources they already have. Even though we strongly advise against using SOC analysts for hunting, it's essential to facilitate the exchange between SOC teams, incident responders and hunters. SOC operations might trigger an IR investigation, and incident responders contribute to daily SOC operations. In the same way hunting might be the trigger for an IR investigation, the results of the investigation feed back into how future hunting methodologies evolve by providing valuable intelligence. Figure 4 illustrates these interactions.

Action Item

The key to threat hunting success is to structure and professionalize threat hunting operations. On top of that, integration of threat hunting into your organization's security portfolio is vital.

Make sure that your SOC analysts, your incident responders and your threat hunters work together closely and openly share information as much as possible. That benefits all three and strengthens the detection capability of your organization.

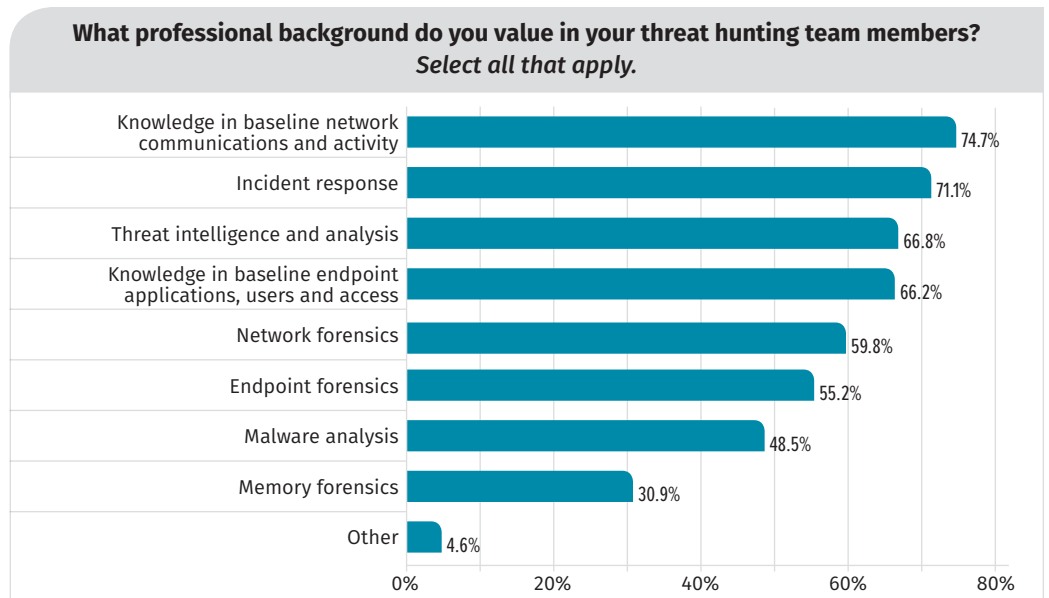


Figure 3. Professional Backgrounds of Team Members




Figure 4. Communication Between Teams

Methodologies for Performing Threat Hunting

Different organizations use different methodologies for threat hunting. Not all of them are equally efficient in catching traces of a breach. Let's look at what methodologies are most common and how well they work.

Threat Hunting Methodologies

Organizations report using a variety of methods for threat hunting activities. More than half of the respondents (57%) define using IoCs as a significant part of threat hunting, followed closely by threat intelligence about attackers' TTPs, as depicted in Figure 5.

This does not come as a surprise. Incident responders have been using IoCs that describe malware and adversary TTPs very successfully for years. While they constitute a viable approach to threat hunting, successful hunting requires a well-curated set of IoCs that provides high-quality IoCs—as opposed to a large number of IoCs.

Analysis of anomalies is included as part of the activities that fit their organization's definition of threat hunting by 54% of the respondents. No matter whether it is done manually, semi-automatically, machine learning-assisted or fully automatically, this is one of the most potent approaches to threat hunting; yet it is one that requires very experienced hunters to succeed. One of the most powerful techniques threat hunters use is finding anomalies by stacking artifacts.³

A real-world example of artifact-stacking would be hunting for running malware. In most organizations, infected systems are an anomaly rather than the norm. Threat hunters acquire a list of all running process executables from all machines, then they count every unique entry. Based on the hypothesis that malware is rare, the analysts identify the malware among the executables with a low number of occurrences. We call that principle the *least frequency of occurrence (LFO) analysis*. The huge advantage of anomaly-based hunting compared with IoC-based hunting is that the anomaly-based methodology is more suitable for identifying previously unknown malware and attacker techniques.

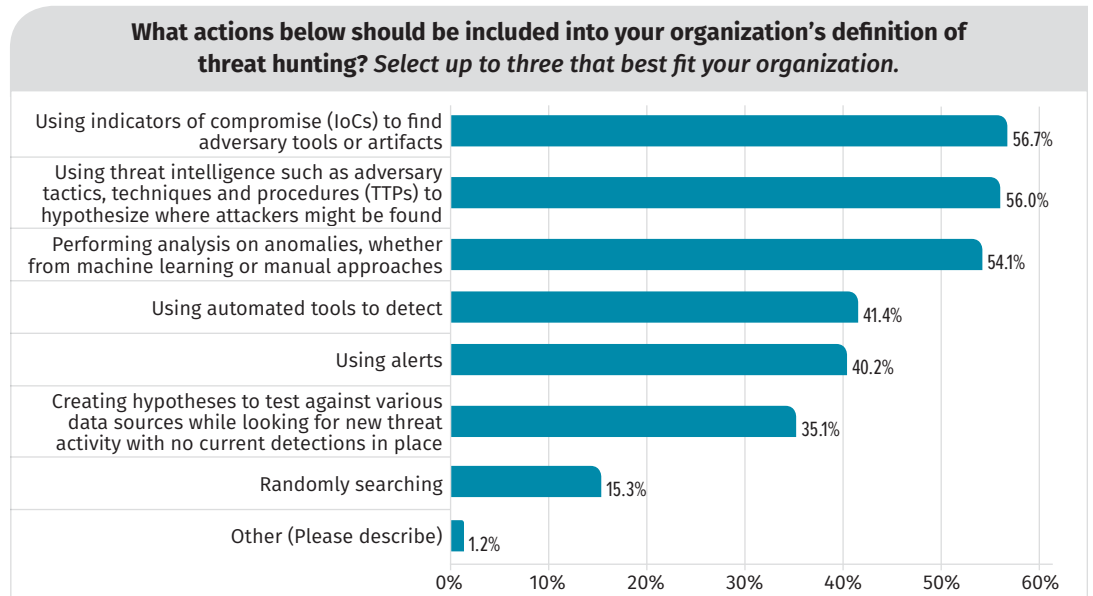


Figure 5. Hunting Methodologies

Action Item

IoCs have different levels of quality and life spans. Your detection ability does not increase just by adding more and more IoCs. Instead, you should review and quality-check every incoming new IoC carefully and test it against your environment.

Regularly check your IoC repository to retire outdated IoCs and tune your listings or remove overly noisy IoCs. Keep track of why you added or removed certain IoCs. In short, manage IoCs over their life spans.

³ www.youtube.com/watch?v=7q7GGg-Ws9s

Hunting activities should involve reacting to alerts, according to 40% of respondents. In our opinion, that qualifies as a standard security operation rather than as threat hunting. Threat hunting is part of nonstandard security operations. With that said, the alerts can certainly provide important information that can be used to create hypotheses for future hunts.

Still, only 35% of respondents would include hypothesis-driven hunting in their definition of threat hunting. That is very unexpected, given that such activities are at the forefront of the proactive hunting approaches recommended today. Hypothesis-driven hunting leverages threat intelligence to understand how an attacker might breach an organization. Assuming that the hypothetical breach has already happened, analysts hunt for evidence for and against the hypothesis. That way, analysts gain a better understanding of the threat landscape, the internal workings of the organization and attacker TTPs.

Furthermore, hypothesis-based hunting uncovers blind spots in an organization's security and investigative capabilities. Rather than leveraging only data that's already available, proving or rejecting a hypothesis additionally uncovers dangerous visibility gaps. Closing those gaps must be a high priority, because those gaps are in the exact spots where threat intelligence suggests an attacker might hit. Thus, we see those organizations that conduct hypothesis-based hunting as tending to better prioritize investments into visibility.

Only 1% of the respondents stated that they employ techniques we did not include in the survey in their definition of threat hunting. Some of the mentioned approaches included penetration testing and zero-day research. These techniques do not qualify as threat hunting. When organizations calculate risk, they usually base it on three major factors: threat, vulnerability and impact. Penetration testing and zero-day research fall under the vulnerability umbrella.

Key Takeaways

Using the right mix of methodologies not only provides better results but also saves money in the long run. However, good threat hunting requires experienced personnel, who are hard to get today. While we do endorse threat hunting using IoCs, ultimately we'd like to see everyone doing hypothesis-based threat hunting.

It takes time for threat hunting teams to mature. So if you are in the stage of using IoCs for hunting, your next step is to make sure that the IoCs you use are well-curated. For the following step, try to attract experienced incident responders who can introduce anomaly detection-using techniques such as stacking. At that point, you might realize that your tooling is inadequate for that task.

Action Item

Although hypothesis-based threat hunting requires quite a high maturity level, it is usually very successful and worth the effort. We strongly encourage all organizations to consider introducing hypothesis-based hunting.

A Closer Look at Hypothesis-based Hunting

Hypothesis-based hunting starts with generating a hypothesis. The first step in creating a hypothesis is to consume intelligence. You need to understand which kind of attack is likely to hit you. You can gather that type of information from the internet, from your country's national Computer Emergency Readiness Team (CERT), your peers and many more places. With intelligence about the types of breaches that might affect you, you can envision how an attacker might breach your organization.

Let's assume an attack group is known to mainly target organizations similar to yours. This makes it likely that the group would eventually hit you as well—or might have already hit you. You learned that this group usually tries to enter the network using vulnerable web applications on internet-facing web servers. Your hypothesis, then, is that the group has entered your network through a vulnerable web application on an internet-facing web server. So your next steps would be to ensure that you have the needed visibility into all of your exposed web servers. Given that visibility, you can start to look for evidence that either proves or disproves the hypothesis of a breach by that group.

Make sure you get the visibility you need. This approach still does not focus on likely attack paths into your network. So, to arrive at the top of the maturity pyramid, you must implement hypothesis-based hunting, which ultimately tells you if you see what you need to see, as shown in Figure 6.

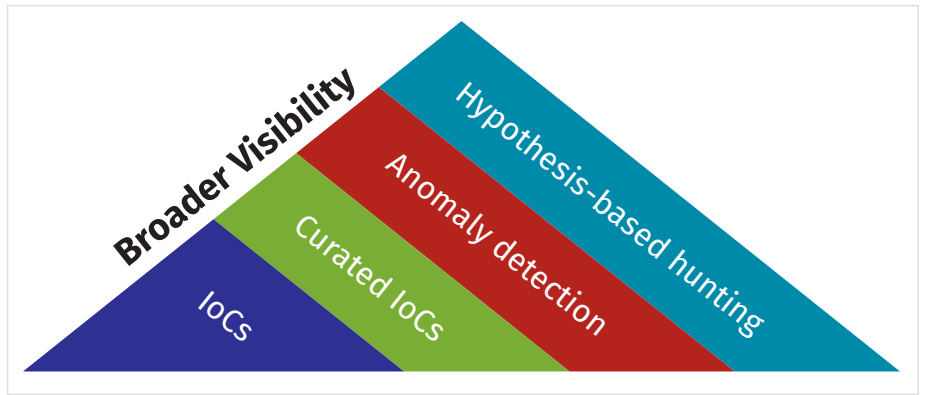


Figure 6. Hunting Maturity

Spending Priorities and Training Needs

Whether to invest in people, process or technology is always a much-debated topic among cybersecurity professionals. Where exactly should you be investing when it comes to threat hunting for your organization? Based on the results from this year's survey, we look at where organizations are currently investing and where they should look to invest in the coming 12 months.

Where the Money Goes

Investment in hunting operations is still an important area for organizations to ensure they are investing in the right places to aid their response and hunting missions. This year's respondents showed us that they are slowly increasing their spending on staffing and pulling back slightly on technology for threat hunting. While this is a good trend in the direction of less reliance on technology, it is also only a minor shift from the previous year, down 6%.

Respondents still appear to be focusing their overall spending priority for threat hunting on technology, as Figure 7 shows.

Focusing strongly on technology could be a costly mistake for some organizations. Technology is less likely to aid hunters in finding adversaries in their organizations than a skilled hunter is. While technology is the highest spending area, the response could represent an investment in visibility tools for respondents' environments. Either way, a fool with a tool is still a fool, so investing in knowledge development for hunters must become a priority for organizations to remain ahead of the curve in hunting adversaries.

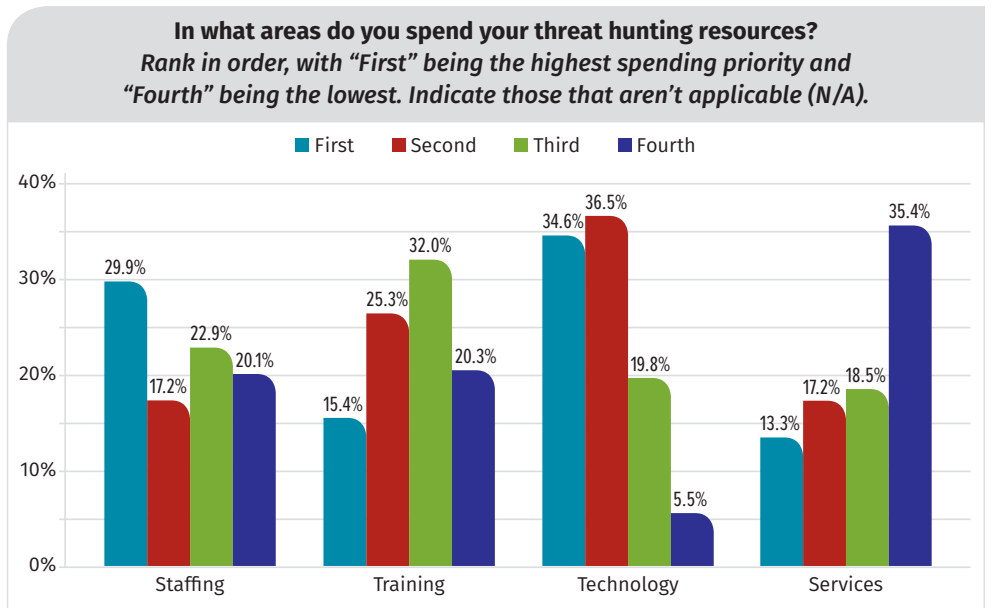


Figure 7. Spending Priorities

A fool with a tool is still a fool.

Action Item

Respondents are still seeing technology as a key requirement to aid hunters instead of skills and knowledge. While tooling is important, it cannot be a replacement for providing hunters with the right skills and training to be able to interpret information from technology that aids them, or to bend technology in more creative ways to catch adversaries. Invest in training your staff before you buy them shiny new tools. Who knows, maybe your trained staff will uncover creative ways to find evil that don't rely on a shiny new tool.

Building Up Skills and Capabilities

This year's survey results show a significant shift in organizations' mindsets around the type of skills required to perform threat hunting. In the 2018 Threat Hunting Survey, respondents were most likely to seek professionals with a background in log analysis and use of analytics tools.⁴ However, this year's respondents are focused on network-based skills for threat hunters, along with baseline knowledge of the network, endpoint applications and user behavior. Although it was not an option for respondents, only a few respondents wrote in log analysis. This may indicate that organizations are shifting away from the mistaken practice of using SOC and SIEM analysts as hunters, and are moving closer to having dedicated threat hunters in their environments.

Additionally, with the top four responses falling within nine percentage points, as evident in Figure 8, we are starting to see organizations value a wider spread of technical skills for threat hunters.

Threat hunting is a skill that requires deep knowledge and understanding of the footprints an

adversary leaves behind in an effort to compromise an endpoint, break out and pivot across the environment. In the past 12 months, organizations have been starting to recognize this as a requirement for the knowledge needed to perform threat hunting.

Unfortunately, this year we have seen a step backward, with respondents thinking that endpoint forensics and memory forensics are becoming less important in their threat hunters' skill set. This change could become a mistake for organizations that are not leveraging this type of knowledge to catch adversaries. Based on the *Verizon 2019 Data Breach Investigations Report*,⁵ the majority of breaches were due to phishing, which puts the endpoints squarely in the firing line for an adversary. Organizations need to recognize that endpoints are where an adversary attacks and moves onward in the kill chain. As a result, endpoint forensics is a rich source of evidence for hunting.

As we observed previously in Figure 7, only 15% of organizations are making training their No. 1 priority; however, 25% make it their second-highest priority. The need for well-trained threat hunters—in conjunction with organizations reducing their value on rich evidence areas, such as hosts and memory artifacts—might represent a gap in understanding the value of this type of evidence and the training that hunters are receiving to guide their hunting missions.




Figure 8. Team Member Backgrounds

⁴ "SANS 2018 Threat Hunting Survey Results," September 2018, www.sans.org/reading-room/whitepapers/analyst/2018-threat-hunting-survey-results-38600, p.8, Figure 4. [Registration required.]

⁵ <https://enterprise.verizon.com/en-au/resources/reports/dbir/>, p. 9

This year's survey shows respondents leaning strongly on assessing their current IT and business operations and baselining them as a crucial part of their threat hunting preparation. However, in the next 12 months, respondents have indicated that their most significant priority for planning is to develop a formal methodology to conduct threat hunting in their organization, as shown in Figure 9.

Looking back at the 2018 survey results, we observed 57% of respondents currently assessing and baselining their IT and business operations. Skip forward to this year's results, Figure 9, where we see an increase of respondents currently assessing and baselining their IT and business operations (61%). This increase in results between

2018 and 2019 may reflect the fact that organizations are struggling to move beyond baselining IT and business operations and have moved on to their intended next step: developing their own methodology (41% in 2019 and 42% in 2018). Organizations must move quickly to create a hunting methodology so they have a coordinated, measurable and repeatable approach. If organizations continue to delay their methodology approach, they will face being stuck in a perpetual baseline activity with their assets.

It is now also evident, from last year's and this year's results, that respondents have significantly moved away from the concept of using external paid services as their methodology for preparing or performing threat hunting, although compared with 2017 there is still an overall uptick from 26% of respondents who indicated they were currently using external paid tools. This trend can only be a positive step for the industry as organizations look to use their own threat methodology or information from their industry peers for hunting, rather than using paid services that may or may not be relevant to the threats they face.




Figure 9. How Organizations Are Preparing for Threat Hunting

Action Item

Shifting the focus to endpoints—and a deeper knowledge of the evidence available on them—still appears to be a shortcoming for a lot of threat hunting teams. This may also go hand-in-hand with knowledge about the amount of evidence that can be leveraged for hunting missions using evidence from an endpoint. Look to invest further in understanding the state of your endpoints quickly, and in developing a methodology to perform and measure your hunting activities.

Tools and System Data Needed for a Successful Hunt

The success of threat hunting relies on the quality of the information available to hunters. The more accurate that information is, the quicker and easier it can be to hunt for malicious activities in an environment. However, if data sources or the efficiency of tools to gather data or evidence is lacking, it affects a hunt team's ability to produce useful findings.

Tools Needed

Organizations are still placing a strong reliance on SIEM alerts as their current go-to tool for threat hunting, as they did in 2018, largely due to the ease with which information can be acquired (see Figure 10). While a SIEM may be the easiest source or tool for organizations to obtain, it generally provides low value from a hunting perspective. If an organization can alert on data in its SIEM, those alerts should inform their SOC and computer security incident response team (CSIRT) to run normal operations for active detection in an organization. Using a hunt team to perform the same role is not the best use of a hunter's skills. It won't help you uncover the threat actor you should be hunting for: the one who hasn't triggered one of your SIEM alerts.

On the plus side, endpoint security data and other endpoint-related information were rated just behind SIEM alerts as sources that organizations use to conduct threat hunting, again because of the ease of acquiring information. Endpoints are a data-rich source from which threat hunters can find evidence of malicious activity they wouldn't normally find in a SIEM due to data limitations or coverage. Threat hunt teams should be focusing their time on endpoints rather than existing SIEM alerts. While endpoints have been featured strongly as data sources, organizations have struggled to access endpoint user activity and forensics, followed by full packet capture. Full packet capture is understandably a hard source of evidence to




Figure 10. Information Needed for Successful Hunts

maintain and hold for extended periods of time. This is echoed by organizations finding full packet captures to be the most highly desired source for hunting that they are unable to attain.

Another interesting finding from this year’s survey is that respondents strongly agreed that information from security orchestration, automation and response (SOAR) alerts was not applicable to their threat hunting mission. Responding to alerts is, in general, not the task a hunt team should be spending its time on. It is time for the cybersecurity industry to take a positive step forward and recognize that SOAR alerts are best handled by automation, the IR team or a SOC.

Effective Data for Hunting

The value of the endpoint as a rich evidence pool for hunters is countered by the recognition that it is one of the hardest types of data to retrieve in an automated way, as illustrated in Figure 11. Respondents tell us that, while they strongly rely on endpoint data, it’s also the data source from which they must pull information manually. This is an area where we need to see

organizations apply their automation skills. Let’s face it: Endpoints are where the malicious adversaries spend the majority of their time, and automated security data is our best weapon.

The figure also suggests that respondents have moved heavily into automating the collection of network metadata, including atomic IoCs and DNS activity, to aid in their hunting activities. Due to the short-lived nature of atomic IoCs, being able to automate the collection and arrangement of this information is a worthwhile use of automation in the overall threat hunting methodology. Organizations need to also ensure they are cycling atomic IoCs so they aren’t reused well past their shelf life, which frankly is fairly short these days.

Key Takeaways

Organizations must pivot from the world of using a SIEM for hunting and move into more unstructured data for hunting to truly find that malicious actor sitting in your environment. If data can

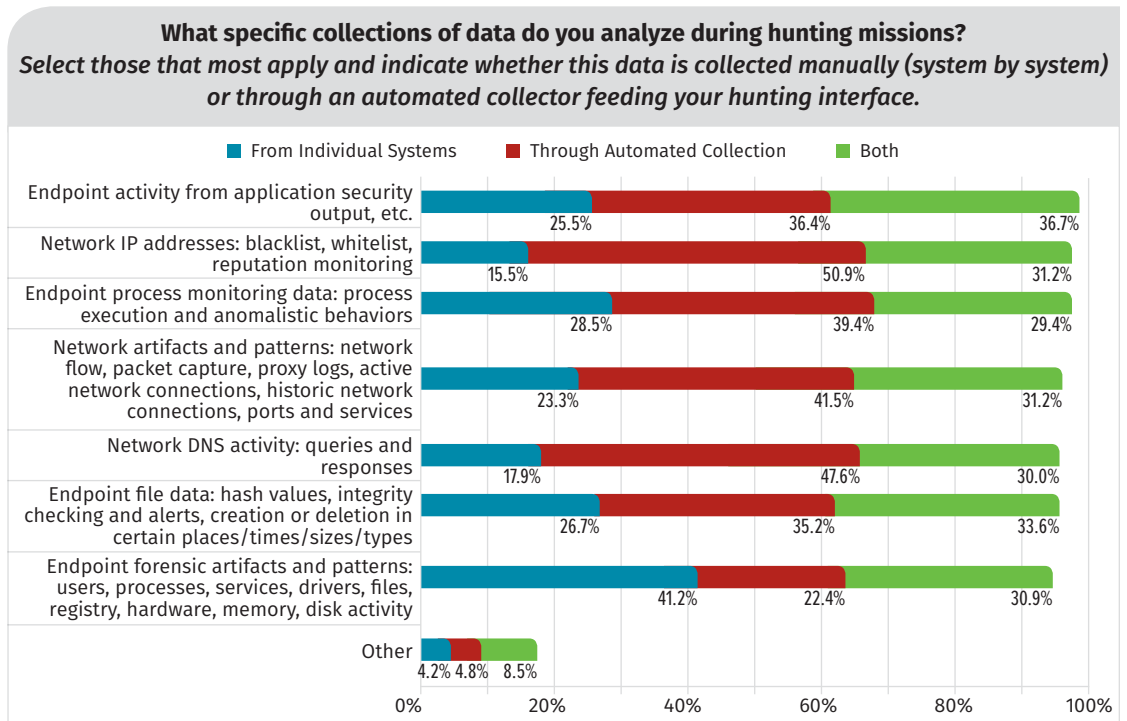


Figure 11. Manual and Automated Data Collection

Atomic IoC

An atomic indicator of compromise (IoC) is a single indicator of an adversary’s footprint following actions the adversary has taken on a system. These indicators can include IP addresses, URL addresses or text strings used by an adversary. Atomic IoCs are generally low-confidence IoCs and have a short window of use for hunters or IR professionals because they are easily changed by an adversary and can result in false positives when used to track an adversary. As an example, a compromised website (URL) might still be providing nonmalicious content alongside malicious content. Furthermore, it’s generally quick and easy for malicious actors to change their atomic IoCs, in the event a malicious actor determines the IoCs have been discovered by threat hunters.

be consumed by a SIEM, it should be used by the SOC to generate alerts. Some maturing organizations might have to use a SIEM to hunt. However, they should ensure they are converting their hunting into active detection so they start to increase the coverage in their environment. As organizations increase their maturity, they need to focus on automation or repeatable means to access endpoint data to aid hunt missions.

Effectiveness of Hunting Practices

Being able to measure and show the performance abilities of a threat hunting team is critical to the life of a team and its engagement by the rest of the business; it's a metric that can make or break a team, its funding or its objectives. Additionally, without some form of metric, how do you know whether you are achieving the objectives you set out to meet? Respondents this year indicate that measuring performance for a threat hunting team is difficult; however, there are ways to do this effectively.

Can You Measure the Performance of Threat Hunting?

Finding the right way to measure the effectiveness of threat hunting is something we are still seeing

professionals struggle with. This is evident in the large majority, 24%, of respondents telling us they are unsure if they have seen any improvement. However, in a positive light, the majority of the remaining respondents are seeing 11% or greater improvement to the overall security posture of their organization due to threat hunting. It's also reassuring to see that only 2% of respondents are seeing no improvement, which still shows that an active hunt team has an overwhelming impact to the posture of the vast majority of organizations. See Figure 12.

It's clear from our 2018 survey and again in this year's survey that there has been a large increase in the number of organizations that are struggling to understand the benefits of threat hunting or how to measure the impact of a threat hunting team. The 2018 survey identified 8% of respondents who didn't know whether they had measurable improvements. This may mean they don't know how to measure threat hunting performance. This year that number has grown to 24%. As an industry, we need to find a way to better measure and report the impact that threat hunting is having on securing our organizations.




Figure 12. Threat Hunting Improvement

Action Item

Organizations need to find measurable ways to show the benefits of threat hunting. Currently, one technique being adopted is to identify the coverage of your environment based on what you can actively alert on, via your SIEM or other security tools, and knowing what gaps you are left with. Using frameworks, such as the MITRE ATT&CK⁶ framework, easily uncovers those gaps.

Being able to identify where to hunt and why hunting is needed requires an understanding of the gaps you currently have. Using the MITRE ATT&CK framework to identify the gaps in your SIEM alerting coverage gives your threat hunters the ability to show benefits by eliminating coverage gaps.

This is only a simple example to start threat hunters off. However, frameworks such as MITRE ATT&CK provide a strong foundation to develop more strategic goals and measurable metrics to show effectiveness to an organization.

⁶ <https://attack.mitre.org/>

Debatable Expectations in Threat Hunting

One of the significant upsides that organizations have seen as a result of threat hunting is a marked improvement in more robust detections and better coverage across the environment, with 36% claiming significant improvement and another 53% realizing some improvement, as illustrated in Figure 13.

This is a fairly natural transition for organizations that are implementing effective threat hunting. Organizations should adopt the concept of “hunt once, detect forever” to ensure that hunting shows a return on the work that is done and teams aren’t continually hunting for the same technique from an adversary. Although this concept is generally adopted by more security-mature organizations, it is reassuring to see our industry finding this as its highest significant improvement.

Other key improvements are attack surface exposure/hardened networks and endpoints, with 35% seeing

significant improvement and 58% seeing some improvement, and more accurate detections and fewer false positives, at 32% significant improvement and 51% some improvement. Taken together, these improvements appear to be very valuable contributions to the overall security of respondents’ organizations.

As a by-product of hunting, we find that respondents are still spending a fair amount of time improving integration and normalization of multiple data sources needed to perform threat hunting effectively (see Figure 14 on the next page). These types of issues are common when you are hunting with data that is not commonly used or is not currently used for alerting, so it makes sense that hunting operations are forcing organizations to reassess their coverage, visibility and tailoring of their collections. However, this also results in lost time when hunting occurs.

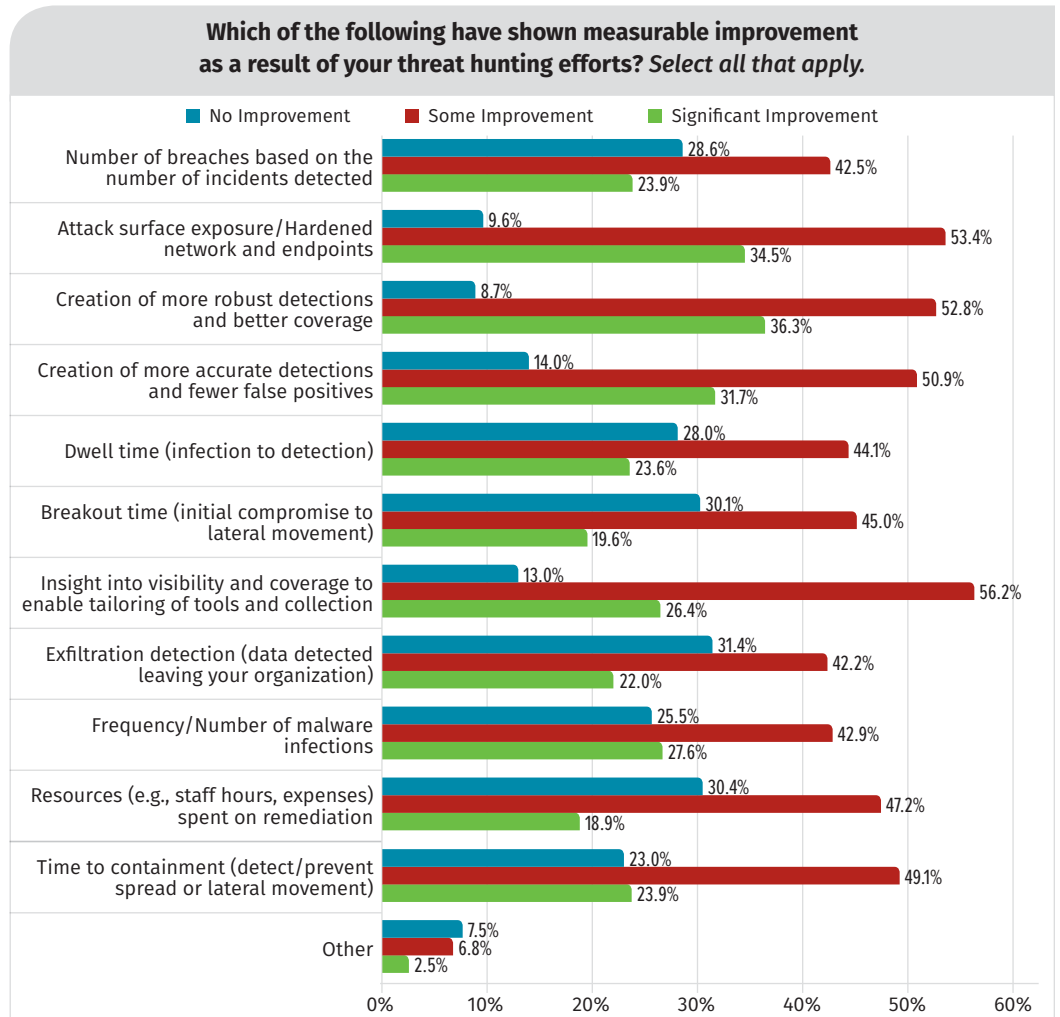


Figure 13. Improvements as a Result of Threat Hunting

Organizations should adopt the concept of “hunt once, detect forever.”

Conclusion

This year's survey shows that for many organizations, threat hunting is still very tool-focused and event-triggered, as illustrated by reliance on SIEMs and atomic IoCs for hunting. Only a small number of respondents see threat hunting as human-driven action that starts where automation ends. It is important for threat hunters to begin to use hypotheses to drive their hunts and add more proactive measures to their hunting toolboxes.

Even though spending now gravitates a little more toward people than in the past years, technology still prevails as the No. 1 investment for organizations. This can be a shortsighted approach. Hiring qualified personnel and providing them with the appropriate training to get the most from their tools is essential to hypothesis-driven hunting.

Network forensics and more general forensic skills lead the list of skills that organizations want to see in their threat hunters. Unfortunately, the survey shows that organizations put less value in endpoint forensics and memory forensics skills today. As visibility into networks decreases or gets more complicated with software-defined network stacks, that's not a smart move.

As in past years, organizations seem to struggle to measure the return on investment of threat hunting. Of those who found a way to measure the impact of threat hunting, the vast majority experienced at least an 11% increase in their overall security posture.

To continue to improve on the success of threat hunting and the value it provides to organizations, we need to see the trend going toward enhancing visibility on the endpoint combined with more specialized network-based methodologies. Organizations need to leverage threat intelligence even more to develop hypothesis-based hunting operations, employ experienced incident responders to lead threat hunting efforts and supply them with the right training and tooling to close identified visibility gaps.

What improvements do you need to make with respect to threat hunting tools and capabilities? *Select all that apply.*




Figure 14. Threat Hunting Improvements Needed To Be Successful

Action Item

When it comes to data used for threat hunting, organizations need to do better at ensuring that data is as consistent as possible. Respondents are showing us that the hunt team is spending time cleaning up log sources or data stores to enable hunting. Organizations must ensure that, as they on-board or acquire new log sources or data stores, they are parsed correctly and easily usable. Overall, this is not only a benefit to threat hunters, but also to SOCs and CSIRTs.

About the Authors

Mathias Fuchs, a certified instructor for [SANS FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics](#), is head of investigation & intelligence 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.

Joshua Lemon is a certified instructor for [SANS FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics](#). He is the director for strategic response and research in the Salesforce Security Response Center, providing research, development and identification of future technical capabilities for the center. Previously, he was computer security incident response team (CSIRT) manager for the Commonwealth Bank of Australia, leading one of the largest dedicated IR teams in the Australian commercial sector. Josh's experience in cybersecurity includes project management, threat hunting, IR, forensic analysis, reverse engineering, penetration testing, secure network design and software development. He holds GREM, GCFA, GDAT, GNFA, GCIH, GPEN, GPYC certifications.

Sponsor

SANS would like to thank this survey's sponsor:

