SANS

A SANS Survey

SANS 2018 Threat Hunting Survey Results

Written by **Robert M. Lee**
and **Rob T. Lee**

September 2018

Sponsored by:
Anomali

SANS Analyst Program

Executive Summary

Threat hunting is a focused and iterative approach to searching out, identifying and understanding adversaries who have entered the defender's networks. Results from the SANS 2018 Threat Hunting Survey show that, for many organizations, hunting is still new and poorly defined from a process and organizational standpoint. Unfortunately, most organizations are still reacting to alerts and incidents instead of proactively seeking out the threats. Threat hunting itself cannot be fully automated. The act of threat hunting begins where automation ends, although it leverages automation heavily. That said, many organizations are finding success by focusing on core continuous monitoring technologies and relying on more security automation in their environments to make hunting more effective.

The survey of 600 respondents reveals that most organizations that are hunting tend to be larger enterprises or those that have been heavily targeted in the past. The survey uncovers some other interesting data points, including the fact that, of the organizations that achieve measurable improvements in their security, most measure improvements in speed and accuracy, while the same percentage report that the use of hunting reduced their exposures. The survey also shows that threat intelligence and hunting must go hand in hand to work effectively. Responses indicate intelligence is key to effective threat hunting and that focusing on people and training are paramount for that effectiveness.

This paper looks at the state of threat hunting and suggests approaches that organizations can take to enhance their threat hunting programs.

Top Survey Findings

- Threat intelligence leads threat hunting, and survey results demonstrate that organizations are investing more in cyberthreat intelligence (CTI) than before.
- Trained staff are key to running threat hunting engagements.
- Hunting is starting to show that organizations are using intelligence properly to identify threats instead of solely relying on traditional alerts and alarms.
- Threat hunting is helping organizations find threats more effectively.

What Is Threat Hunting?

Threat hunting is aptly focused on threats, and to be a threat, an adversary must have three characteristics: the intent, capability and opportunity to harm. Threat hunters focus their search on adversaries who have those three characteristics and are already within the networks and systems of the threat hunters' organization, where hunters have the authority to collect data and deploy countermeasures.

Many security personnel likely think that they have been doing this type of activity, at least in part, since long before the term threat hunting emerged; in many cases, that is true. The recent focus on threat hunting is not about rebranding what many defenders have endeavored to do over the years; rather, it is about placing an appropriate, dedicated focus on the effort by analysts who purposely set out to identify and counteract adversaries who may already be in the environment. Threat hunting requires some specific analytic skills, such as familiarity with the enterprise and the ability to generate and investigate hypotheses. Hunting benefits from analysts using automation to make these hunts faster, easier, more frequent and more accurate. (Automation will be discussed later in the paper.)

Threat Hunting

A focused and iterative approach to searching out, identifying and understanding adversaries who have entered the defender's networks

Why hunt? Threats are human. It is the adversaries themselves, not just their tools (such as malware), that interest threat hunters. These adversaries are persistent and flexible and often evade network defenses. The threats are often identified as advanced

persistent threats (APTs), not just because of the capabilities that the adversaries wield, but also because of their ability to initiate and maintain long-term operations against targets. Focused and funded adversaries will not be countered by security boxes on the network alone.

For their part, threat hunters do not simply wait to respond to alerts or indicators of compromise (IoCs). They actively search for threats to prevent or minimize damage. Additionally, threat hunting does not need to find threats to be measured as successful. The act of threat hunting should essentially test an organization's capability to reliably detect and respond to threats. Consider threat hunting a hypothesis-driven approach to validating the collection, detection and analysis of data ahead of an incident.

Threat hunting is a hypothesis-driven approach to validating the collection, detection and analysis of data ahead of an incident.

One of the most notable highlights of the 2018 survey is that it demonstrates a more accurate use of threat hunting in many organizations. This change in threat hunting practices has increased since the last survey in 2017,[1] which showed many organizations typically were hunting through traditional intrusion detection. In this year's survey, many more organizations were using proper threat intelligence to help identify the best locations inside an organization's network to look for anomalistic behaviors that are direct indicators of threats.

Good Examples of Threat Hunting

Respondents provided brief descriptions of their threat hunting processes. Here are some examples of good processes:

- "Starting with Tools, Techniques or Procedures (TTPs) or a vulnerability, develop hypotheses to determine whether our infrastructure is impacted, and then test those hypotheses."
- "First, baseline the environment for normal activity. Create a hypothesis based on the kill chain. Utilize ATT&CK framework for TTPs. Run IOC sweeps from threat intel reports."
- "Gather intel, develop a hypothesis, create a scope and execute the hunt."
- "Form a hypothesis or use evidence from intel, then determine the best way(s) to find activity on the network or hosts, both for the current point in time and for future events."
- "Identify a hypothesis of what to hunt for, review documentation of past hunts, peer review the proposal, notify the team and begin work, collect and normalize data, analyze data, identify findings, take immediate action for any detected intrusions and declare the incident, and determine non-immediate adjustments to controls and detection mechanisms."
- "A threat hunting process starts with generating hypotheses (assuming we have been breached in a given way) and then verifying the hypothesis by hunting for the related indicators in all relevant data sources using log analysis and then marking the hypothesis as true or false in the end."

[1] "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

Bad Examples of Threat Hunting

These are examples from respondents that are just performing intrusion detection—not threat hunting:

- "Notice an alert in the system and slowly tear it apart from endpoint to endpoint."
- "Spend a lot of time reviewing logs from the SIEM and formulating custom queries in the SIEM."
- "Analyst watches logs and endpoint events. Non-baseline behavior or triggered events create a potential incident. Analyst reviews network traffic and isolates potentially affected systems. Standard IR rolls from there."
- "Our entire operation is constantly monitoring the environment to establish its baseline. As soon as we detect something odd or we are made aware of something risky in our environment (such as a malicious IP address communicating with us), we start an analysis on that resource: network behavior, processes behaviors, logs and possible strange evidence through the filesystem and registry. If we confirm something 'evil' we move on with the process of containment, eradication, recovery and then the lessons learned."
- "We have antivirus deployed on most endpoints. The signature that has been triggered the most is investigated and hunted for its root cause, and we try to reduce its count by the next week."
- "Threat hunting is triggered by SIEM alerts or AV alerts."

Takeaway: Begin Consuming Intelligence

Threat hunting is part of nonstandard security operations. It is a good combination of threat intelligence and hypothesis generation based on likely and probable locations of intrusion into your network.

We advise organizations that consider hunting as reacting to alerts to continue to find ways to increase visibility into threat intelligence capabilities. Once an organization begins consuming threat intelligence, natural hunting begins to take place. It is similar to knowing that the latest burglary technique in your neighborhood involves people trying to steal cars by entering through garages. You might put a camera in your garage and monitor it a bit more closely.

Modernizing Hunting Operations

Threat hunting is key to detecting adversaries in a variety of environments. However, through the years many in security operations have directly associated threat hunting with intrusion detection. This year's survey shows the beginnings of a move away from that mindset.

In the 2017 survey,[2] many respondents indicated that hunting was centered completely on reactionary indicators instead of proactive threat intelligence and predictive analytics about where adversaries are likely to be. Threat hunting based on targeted inspections of likely locations of bad activity is useful, but the key challenge is that most traditional security operations center (SOC) indicators usually fail to detect the very threats they are tuned to detect. Most of these indicators focus on malware-based indicators instead of the behaviors or strategies of the adversaries.

One of the key indicators that threat hunting is growing in scope and need is the fact that 43% of respondent organizations are now performing continuous threat hunting operations. In 2017, the number was only 35%, which shows that many organizations are now adopting methodologies that are key to reducing adversaries' overall dwell time. Figure 1 illustrates the various methodologies respondents use for threat hunting.

Figure 1. Threat Hunting Methodology

As more organizations continually perform threat hunting, we hope that adversary dwell time, which averages above 90 days , will start to fall in the next few years. As recently as 2013, the average dwell time was over six months. The decline since then shows that the adoption of threat hunting and stronger analytical techniques has had a significant impact on reducing the overall dwell time of adversaries across most networks.

Where are most of the respondents' organizations obtaining the threat intelligence used in threat hunting? Almost 58% of that intelligence is created internally based on previous attacks, and 70% originates from third-party sources, as shown in Figure 2.

Most organizations use traditional alerts and alarms to identify threats. This is not threat hunting in any way, but we offered it as an option to show that organizations have a difficult time transitioning between the traditional approach of intrusion detection via

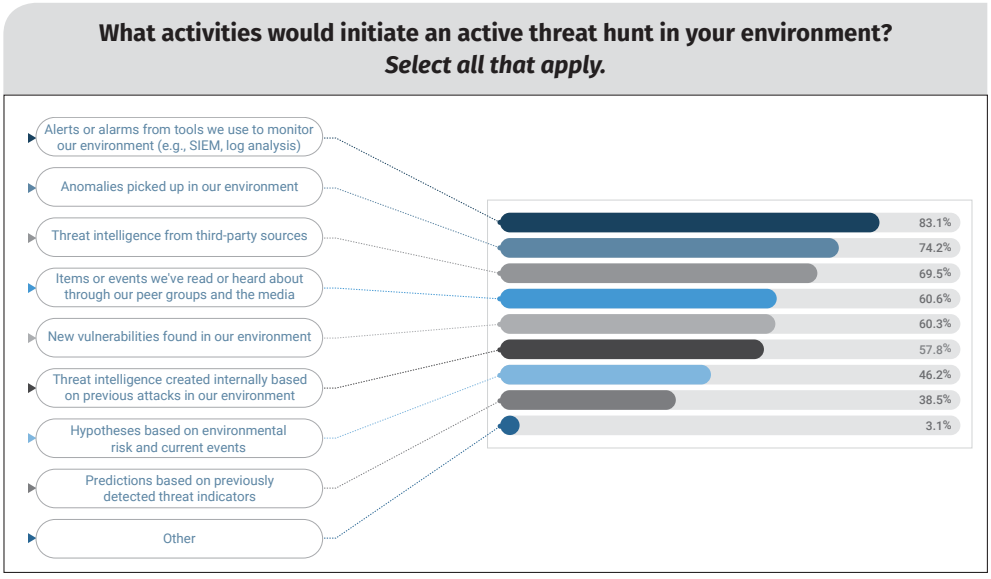

Figure 2. Activities that Initiate a Threat Hunt

[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

alerts and a forward-leaning, proactive hunting engagement led by threat intelligence. Over time, as organizations become better at threat hunting, more of them use internal threat intelligence to drive their hunting operations. Nothing is more valuable than correctly self-generated intelligence to feed hunting operations. Many organizations do not have a mature threat intelligence capability, however, so they tend to initially rely on third-party intelligence (which can include antivirus signatures) to feed their security operations and to hunt to detect recurring or new adversaries in their environments.

Takeaway: Blend Internal and Third-Party Intelligence

The 2017 survey showed that not enough organizations were creating or ingesting modern threat intelligence feeds to tune sensors to initiate hunting operations. In this year's survey, more organizations are creating and ingesting intelligence feeds. We recommend that organizations continue to focus on transitioning their hunting operations from reacting only to signatures provided by third-party intelligence capabilities. A solid blend of both internal self-generated intelligence augmented with third-party feeds will continue to reduce overall adversary dwell times across organizations' networks.

Hunting Still Seen as a Technology Solution

Hunting operators use technology to help increase the speed and accuracy of their operations, but many organizations are still prioritizing buying tools and technology over developing a well-versed staff. This might be a mistake, because hunting is akin to special operations: You need highly trained and skilled personnel to lead and execute core hunting operations. Among survey respondents, 41% said technology was most important, compared with 30% who said staff was most important (see Figure 3).

In what areas do you spend your threat hunting resources?
Rank in order, with "First" being the highest spending priority and "Fourth" being the lowest.

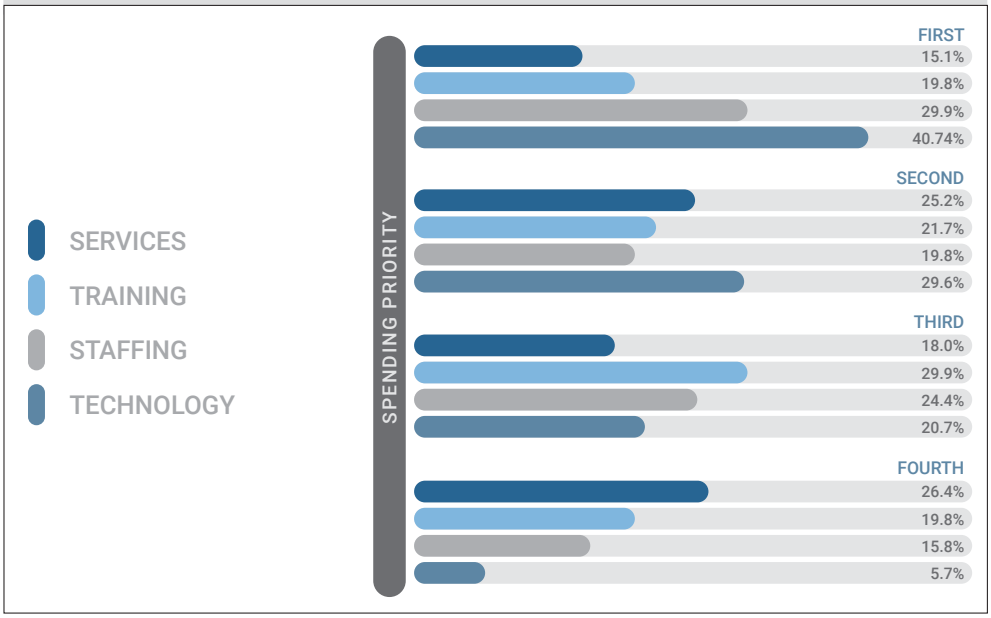

Figure 3. Spending Priorities

This perspective should likely be more balanced, because fully automated threat hunting doesn't exist. Threat hunting automation is similar to spell-check in a word processor. While it can help to identify mistakes, it is, by its nature, largely human-driven and is more of a tool than true automation. Ironically, while the staff isn't rated as high as technology, training of staff does get high rankings. This shows that respondents do, in fact, believe that staff is important, because you can't train a device, but you can train a person. See Table 1 for a comparison of training rankings from 2017 and 2018.

Table 1. Training Rankings: 2017 vs. 2018

TRAINING RANK	2017	2018
FIRST	11.9%	19.8%
SECOND	28.8%	21.7%
THIRD	33.8%	29.9%

Takeaway: Prioritize Threat Hunting Training

Organizations should begin placing a higher priority on developing extremely skilled staff to lead and execute threat hunting operations across their organizations. The right team might be more important for hunting than tools/technology. The more that organizations understand that well-trained hunters will likely create the tools they need, the more likely those organizations will reap the benefits from that skilled staff. Many organizations tend to hire cheaply, assuming that these threat hunting skills are found in many hires. However, there is evidence that organizations that have hired the best staff tend to be the best at detecting threats. The same cannot be said for organizations that spare no expense on tools and technology, unless they spend equally or more on highly skilled operators.

Threat Hunters: Does Your Team Have What It Takes?

During the past several years, the skills and tools necessary to be considered an effective hunter have been hotly debated in security operations circles. This year's survey, similar to last year's, shows that the core skills needed to hunt effectively are core information security baseline skills.

There is a clear pyramid of skills and tiers that must be attained in order to be adept at threat hunting. Most respondents cited network, endpoint, threat intelligence and analytics skills as baseline skills. Of the respondents, 73% selected threat analysis as a key skill needed, second only to log analysis and analysis skills at 83%. See Figure 4 on the next page.

Why Can't Threat Hunting Be Fully Automated?

Automation is such a misunderstood word, especially in the context of threat hunting. Hunting needs capabilities to help enhance speed, accuracy and effectiveness. The best hunting teams heavily leverage automation to aid in increasing the scale and efficiency of hunts across the enterprise. However, by its definition, hunting is best-suited for finding the threats that surpass what automation alone can uncover. Threats are, after all, moving targets. Still, it is important to recognize the intertwined nature of automation and the human process of threat hunting.

Tools and capabilities that aid threat hunting are driven by SOCs. Traditional information security architecture such as SIEM analytics, log file analysis, intrusion detection and antivirus are largely automated capabilities based on signature-based rules fed and maintained by analysts. Hunting concepts using these capabilities often record and identify, but then possibly ignore, small anomalies that are the barely visible tracks of advanced adversaries. Ignoring these trivial anomalies is easy because there are too many to properly vet in even a modest-sized network.

After discovering an adversary, security teams often realize that their sensors did, in fact, record the adversaries' activities. At the time those alerts occurred, however, the teams were too overwhelmed to pay any attention to them. These early warning capabilities can be enhanced greatly by utilizing threat intelligence effectively. With proper intelligence, additional threat indicators of compromise and the right analysts using properly tuned tools, some seemingly benign alerts can be identified as major events. In other words, threat hunting, threat intelligence and security operations can move together in harmony.

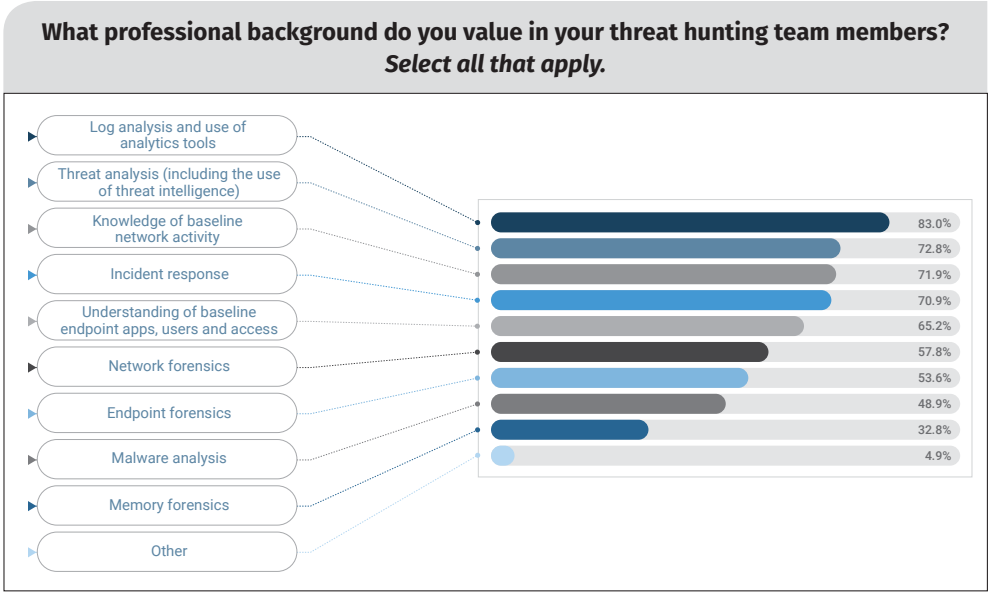

Figure 4. Professional Background for Threat Hunting Team Members

The responses were nearly identical to those in the 2017 survey, as shown in Table 2.

Table 2. Critical Baseline Skills for Threat Hunting, 2017 vs. 2018 Surveys

BASELINE SKILLS	2017	2018
ANALYTICS	79.5%	83.0%
KNOWLEDGE OF BASELINE NETWORK ACTIVITY	77.6%	71.9%
UNDERSTANDING OF BASELINE ENDPOINT APPS, USERS AND ACCESS	66.2%	65.2%
THREAT ANALYSIS (INCLUDING THE USE OF INTELLIGENCE)	69.9%	72.8%

Digital forensic and incident response (DFIR) skills make up the next tier of attainment on the skills pyramid. After operators have mastered baseline and threat intelligence skills, they then move forward with mastery of core DFIR skills that cross over endpoint, network, malware analysis and memory forensics. Again, responses to the 2018 survey closely matched the 2017 results (see Table 3).

Table 3. Critical DFIR Skills for Threat Hunting, 2017 vs. 2018 Surveys

CRITICAL DFIR SKILLS	2017	2018
INCIDENT RESPONSE	66.2%	70.9%
ENDPOINT FORENSICS	49.8%	53.6%
NETWORK FORENSICS	57.5%	57.8%
MALWARE ANALYTICS	49.3%	48.9%
MEMORY FORENSICS	38.4%	32.8%

The final tier on the pyramid involves using all the skills that inform respondents to make the best guess, hunch or outlier detection. This skill evolves over years of experience, because it sometimes involves something that even the best hunters can't place their finger on: the feeling that something is "off" in a location on the network. This is directly correlated to the fact the experienced analysts can see beyond just what

the data are telling them. In the future, more machine-learning tools might also help enable the operator to develop these intuitive skills, making this tier of the pyramid more accessible. For now, the ability to see what no one else can is limited to a distinct few, thus making this level of hunting extremely valuable for an organization that can recruit top-tier hunters.

Takeaway: Value Staff and Hunting Skills

Trained staff must be valued more highly, especially because customized abilities are used so frequently in environments. For threat hunting, baseline security skills are critical, and DFIR skills augment those skills.

The ability to use threat Intelligence likely needs to be valued higher among the baseline skills and tools needed for effective threat hunting. Organizations need to invest more in CTI to obtain greater leverage in threat hunting. Our advice when trying to hire skilled hunters is to keep in mind that the extent to which hunting is a science or methodology has yet to be exactly determined. We continue to expect rapid advances in the coming years as organizations share more best practices and tools that enable lesser-skilled hunters to home in more quickly on threats without the help of top-tier hunters.

The Hunting Armory: Choose Your Desired Weapon/Hunting Tool

Which tools are used in hunting? Most (90%) survey respondents indicate that they use existing infrastructure tools for hunting. Many staff in organizations are developing their own customizable home-developed tools: 62% of survey respondents note the importance of having a properly trained hunting team to create these capabilities. Augmenting homegrown solutions are open-source capabilities integrated alongside standard SOC capabilities. For hunting, open-source solutions are used more frequently (48%) than purchased commercial third-party hunting platforms (33%), as illustrated in Figure 5.

What tools do you utilize to perform hunting? *Select all that apply.*

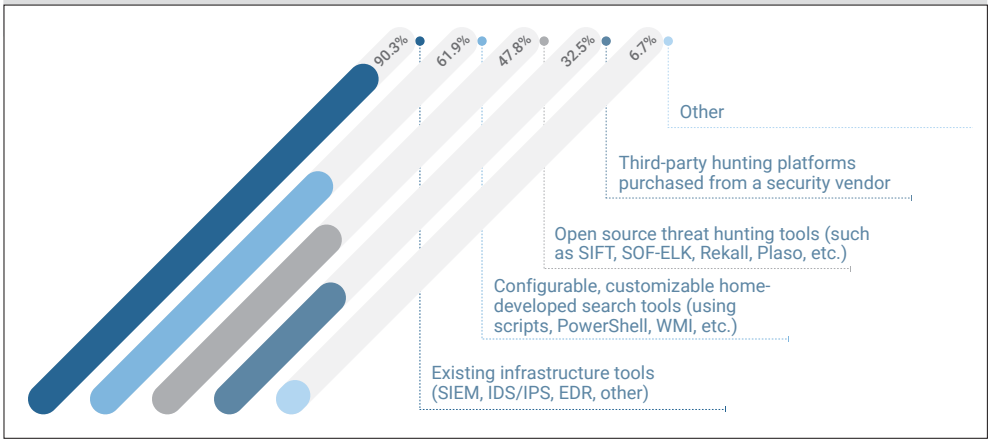

Figure 5. Tools Used to Perform Threat Hunting

Two questions arise. First, are these tools providing enough of a view, given that most tools are detection- and rules-based? Second, where are the rules coming from? Based on the survey results, it is clear that most organizations are treating hunting as an aggressive SOC exercise using detection. For this to result in any type of success, most SOC-based operations must be baselined and tuned specifically to their environment. We know this is unlikely based on the current average dwell time of adversaries for most organizations, which hovers around 90 days. It is unlikely that a typical

These results are interesting considering that technology was ranked as more important than staff in an earlier part of the survey. The key takeaway is that qualified staff can more easily implement homegrown and open-source solutions. Without that key staff, it is unlikely that these solutions will be implemented properly.

organization, even with a formally tuned SOC and tools, would still be able to detect a determined adversary breaching its environment. A tool provides a baseline for hunting operations by providing a decent "horizon view" of the entire network environment that can be used to detect historical anomalies during a hunt. However, it is extremely rare that a SOC using automated tools alone will detect all of the adversaries breaching its environment. This is the entire point of hunting, which draws on threat intelligence to formulate the most likely locations in the environment where an adversary will appear based on prior behaviors, attacks and objectives.

Takeaway: Hire the Best, Then Choose Tools

Across the board, tools help augment properly trained staff. Qualified staff can create and implement their own solutions, in many cases, using open source and at times commercial platforms. Tools help increase efficiency but should not replace the importance of hiring the best hunters an organization can afford.

Organizations therefore should hire the best people for their hunting operations and have them figure out what tools they need. Prioritizing tools before hiring a key hunting team, a common occurrence, might arm your team with the wrong hunting tools and leave the team in a situation where it needs to create its own solutions anyway. The survey data back up that most organizations are relying on homegrown solutions over commercial capabilities specifically for hunting operations.

A properly trained hunter is not an easy hire, and it may well be an IT organization's most expensive hire. Having said that, the best people will help reduce overall costs by not making sloppy purchases. Time and again, the smartest individuals are usually very conservative with what they purchase. They don't want to waste time managing expensive tooling platforms that don't aid them in their specific hunts across the organization.

Endpoint Hunting: The Elusive Target

Hunting has been gaining ground in recent years in showing its effectiveness. Among respondent organizations, 27% found one to three threats, and 21% found four to 10 threats, as noted in Figure 6.

So, did respondents find known threats? There is no real pattern to dictate whether the threats respondents found were known/unknown/evolved. (An evolved threat is one that continually is improving itself through experience of hacking many targets over an undefined period of time. The evolved threat is one that is growing in maturity.) However, if you combine the known and evolved groups on Figure 7, it shows that the threats came back, and that hunting

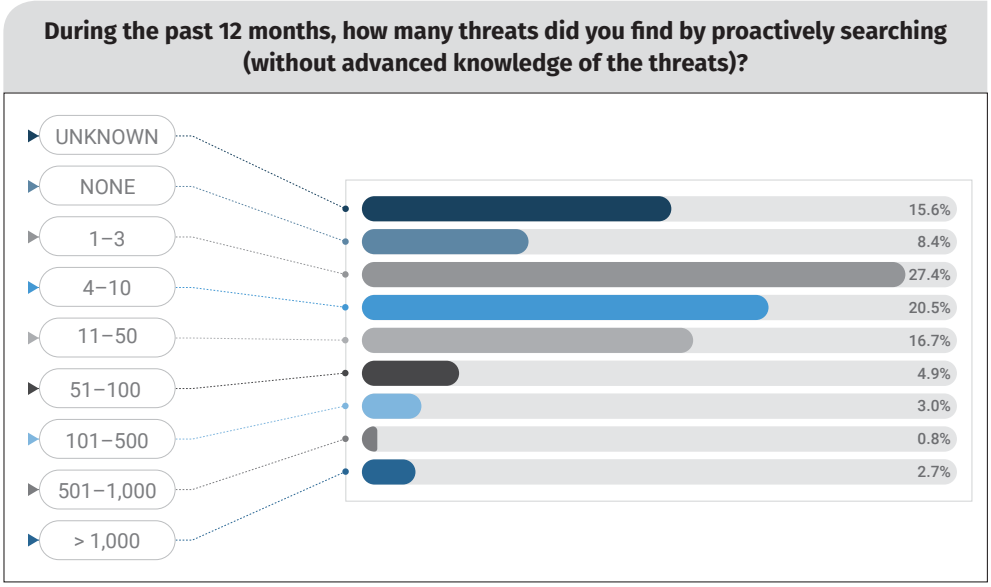

Figure 6. Threats Found by Proactive Searching

aided in detecting the threats that did come back.

As in the 2017 survey, this year's survey found that one of the more important sources of data used for hunting is still fairly hard to collect. Endpoint detection and response (EDR) collection for subsequent analysis is still fairly new in the world of information security. While many groups are opting for simple endpoint collection utilities such as sysmon from Microsoft, organizations are finding it hard to identify tools and capabilities to help process and analyze endpoint data. This is challenging partly because of the sheer volume of possible data collected by a single endpoint, including event logs and registry, disk and forensic artifacts. Correlation capabilities in many cases are extremely limited, which has reduced many hunters to examining across systems instead of across the network. Table 4 provides a look at the types of data organizations need and the difficulty they have in acquiring that information.

Of all threats detected through hunting, what percentage were new and unknown to your security systems? What percentage were known to existing security systems? What percentage evolved from other threat patterns already present in your environment?

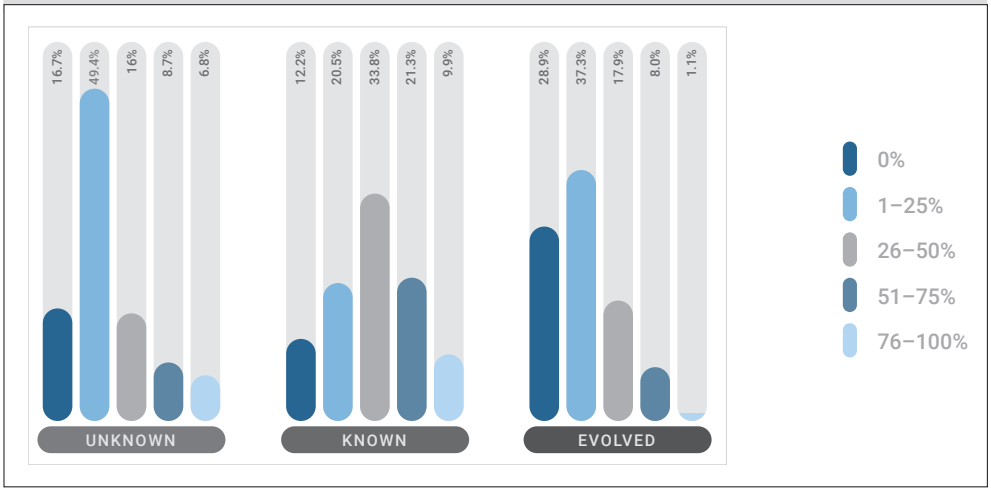

Figure 7. Threats Detected Through Hunting

Table 4. Difficulty in Acquiring Needed Information

DIFFICULTY ACQUIRING NEEDED INFORMATION	NEED BUT UNABLE TO ACQUIRE	ABLE TO ACQUIRE WITH DIFFICULTY	ABLE TO ACQUIRE EASILY
SIEM ALERTS	7.8%	14.9%	74.3%
ENDPOINT SECURITY DATA (ANTIVIRUS, ENDPOINT PROTECTION SUITES)	5.2%	22.4%	70.5%
NETWORK IDS/IPS FEEDS	6.7%	23.9%	67.5%
WEB PROXY LOGS	10.4%	22.0%	63.4%
OPEN SOURCE THREAT INTELLIGENCE	8.2%	28.7%	59.7%
ENDPOINT SECURITY AND SYSTEM EVENT LOGS	7.5%	34.0%	56.7%
EMAIL LOGS	8.6%	33.2%	55.6%
ENDPOINT PROCESS ACTIVITY	11.6%	33.2%	53.4%
DNS	10.8%	33.6%	53.4%
NETWORK TRAFFIC FLOW	13.4%	32.8%	52.2%
THIRD-PARTY CUSTOMIZED THREAT INTELLIGENCE	18.7%	29.9%	46.3%
INTERNALLY GENERATED THREAT INTELLIGENCE	17.5%	32.5%	44.8%
ENDPOINT USER ACTIVITY AND FORENSICS	9.0%	47.0%	41.4%
FULL PACKET CAPTURE	27.6%	36.6%	32.5%
DECEPTION AND DECOY SYSTEM DATA CAPTURE	54.5%	19.8%	19.0%
OTHER	10.4%	6.0%	6.0%

The easiest data to obtain are core security data typically offered by standard baseline operations. The data are mainly grouped by SOC-related automated IDS information (SIEM, endpoint security agents, network IDS/IPS and weblogs). This leads to the perception that baseline data are the easiest to acquire. However, acquiring DFIR data is still a skill that requires a lot more work. Both are needed in hunts.

What specific collections of data do you analyze during hunting missions?
Select those that most apply and indicate whether these data are collected manually (system by system) or through an automated collector feeding your hunting interface.

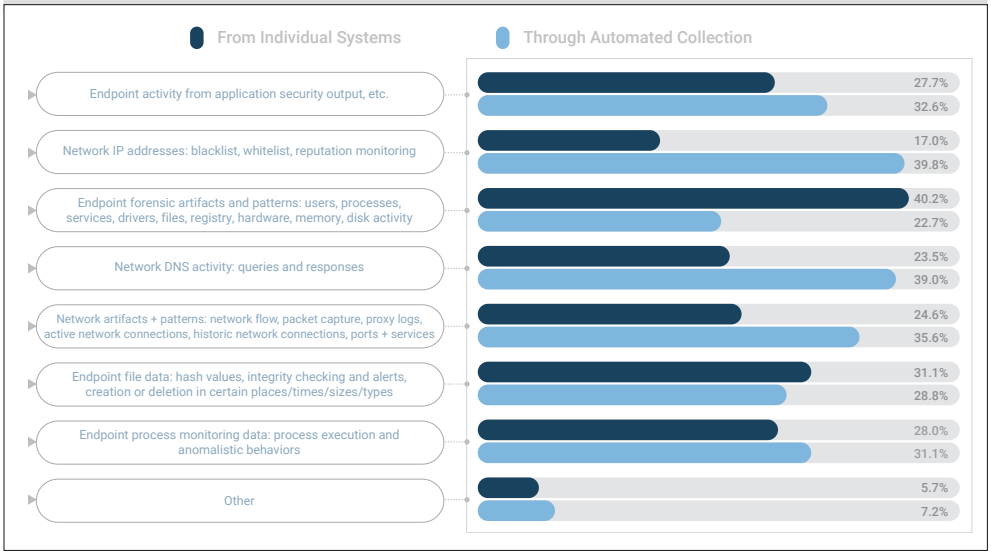

Figure 8. Collections of Data Analyzed During Hunting Missions

While network data were rated high on the scale, endpoint analysis is still a gaping hole in most hunting operations. The 2018 survey showed little change from the 2017 survey in this regard. Not all endpoint data have been challenging to collect. Generally, endpoint security data such as anti-malware are fairly simple to collect. Event logs and file system data are the most difficult elements for analysis in hunting operations.

Most organizations state that endpoint data (including event logs) are the best source of hunting telemetry used to identify malicious behavior. See Figure 8.

Does this mean that network data are less important? Absolutely not. One of the key elements not generally available across all hunts is full-packet captures. Full-packet captures are desired in subnets and network segments targeted by adversaries. Hunters should place network sensors in specific locations with full-content packet interception enabled to add additional depth to network data collected and to provide for additional containment once found during incident response. Threat intelligence and prior attacks will tell the hunter the best locations to monitor fully.

Survey data indicate that endpoint monitoring data are mainly collected via individual system access, while network data are collected through automation.

Takeaway: Collect Full-Packet Captures

Endpoint data are still relatively difficult to acquire, as reported in both the 2017 and 2018 surveys, and organizations haven't seen a noticeable improvement in collecting or analyzing data from these sources. EDR data are stronger, showing that organizations that adopt the capabilities of these systems are having an easier time collecting difficult-to-obtain DFIR endpoint data. Organizations should consider dropping in "wiretaps" at locations on specific endpoints or network enclaves with a high rotation frequency to collect full-packet captures.

Measuring Success Drives Improvements

It is difficult to improve on what you cannot measure. In this year's survey, 48% of respondents noted that they measured the improvements made to the organization from their threat hunting activities. See Figure 9.

Figure 9. Measuring Improvements

Of the 48% who measured hunting improvements, only 3% found that their threat hunting efforts did not improve the organization. However, an additional 8% did not know. The primary finding here is that organizations should be measuring the improvements that they see from threat hunting. In doing so, they can show a return on their investment to the organization. Even determining that threat hunting efforts are not improving anything serves a purpose, because it may change how you hunt or your investments in the practice entirely. See Figure 10.

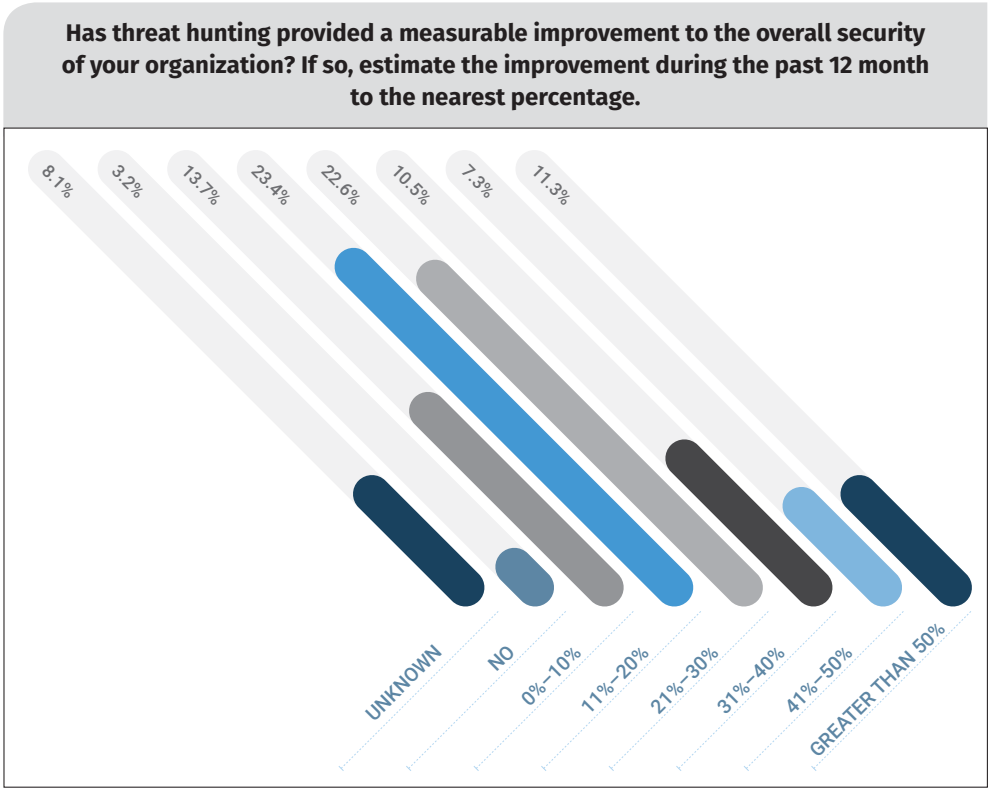

Figure 10. Measurable Improvements from Threat Hunting

Organizations have limits as to how, where and to what level they can invest their time and resources in security. Therefore, it is vital to determine a return on investment and prioritize the most significant improvements that can be made to the organization. Threat hunting is an intensive process and should not replace areas such as continuous monitoring and network security monitoring. However, when used and measured correctly, threat hunting can add significant value to security programs to help keep the organization in a proactive instead of reactive stance.

Tips for Measuring Success

Simply attempting to measure threat hunting is not going to drive value. It is important to create a structured approach to qualitative or quantitative measurements. Survey respondents had a number of options to choose from on where they saw improvements. The most significant area of improvement was time for containment (88%). Additionally, significant improvement was made in attack surface hardening (48%) and decreasing adversary dwell time (40%). When combining "significant improvement" and "some improvement," 74% of the respondents noted improved efforts. See Table 5.

Table 5. Measurable Improvements as a Result of Threat Hunting Efforts

MEASURABLE IMPROVEMENTS AS A RESULT OF THREAT HUNTING EFFORTS	NO IMPROVEMENT	SOME IMPROVEMENT	SIGNIFICANT IMPROVEMENT
TIME TO CONTAINMENT (DETECT/ PREVENT SPREAD OR LATERAL MOVEMENT)	10.0%	55.8%	32.5%
AMOUNT OF BREACHES BASED ON THE NUMBER OF INCIDENTS DETECTED	13.3%	54.2%	28.3%
RESOURCES (E.G., STAFF HOURS, EXPENSES) SPENT ON REMEDIATION	22.5%	49.2%	25.8%
ATTACK SURFACE EXPOSURE/HARDENED NETWORK AND ENDPOINTS	3.3%	48.3%	47.5%
FREQUENCY/NUMBER OF MALWARE INFECTIONS	14.2%	45.8%	37.5%
DWELL TIME (INFECTION TO DETECTION)	11.7%	45.0%	40.0%
EXFILTRATION DETECTION (DATA DETECTED LEAVING YOUR ORGANIZATION)	19.2%	45.0%	31.7%
BREAKOUT TIME (INITIAL COMPROMISE TO LATERAL MOVEMENT)	21.7%	44.2%	30.8%
OTHER	3.3%	9.2%	4.2%

These are all great options for measuring when looking at a threat hunting program. Additionally, organizations should consider measuring collection of data. As an example, a sample threat hunt could have a hunter testing a hypothesis on how newly discovered adversary tradecraft might be found in the organization. The hunter would have to pull data from various collection sources around the network to test the hypothesis. The process of testing the hypothesis might reveal that collection is exactly where it should be (a great validation that you could reliably detect threats using that tradecraft), or that collection is not where it should be or as has been reported to be (i.e.., you do not have the collection you thought you did, or entirely new collection efforts are needed). Determining that your collection is good, that it needs to be tuned or that you need new collection efforts altogether all constitute important outputs of a threat hunt that can be measured and improved over time.

Improvements Required for Continuing Success

More than 40% of respondents marked options to survey questions that noted improvements they need to make. The least significant improvement options marked were storage (29%) and less intrusiveness on the host (28%). The most frequently marked options were better investigative functions (59%) and more staff with investigative skills (also 59%), as shown in Figure 11. Both of the top options relate to the effectiveness and efficiency of staff on hand as well as to an increasing need for skilled personnel.

Threat hunting is not a generic skill set. Rather, it is focused on hypothesis-driven efforts to uncover adversary activity and test the organization for threats in advance of an incident. Thus, we can't suggest a single discipline. Skills that relate to current data collection in organizations, the amount of collection that is desired, and the analysis of that data are all appropriate investments. Training should be considered in universally needed skill sets such as incident response and threat intelligence functions, as well as more specific skills involving common data sources such as network forensics, memory forensics, and intrusion detection systems. Most security teams often feel they need more security analysts, when in fact the harsh reality is that the industry is still struggling to find appropriate talent due to a lack of resources and candidates. Therefore, training and ensuring the effectiveness of analysts already on board is particularly critical.

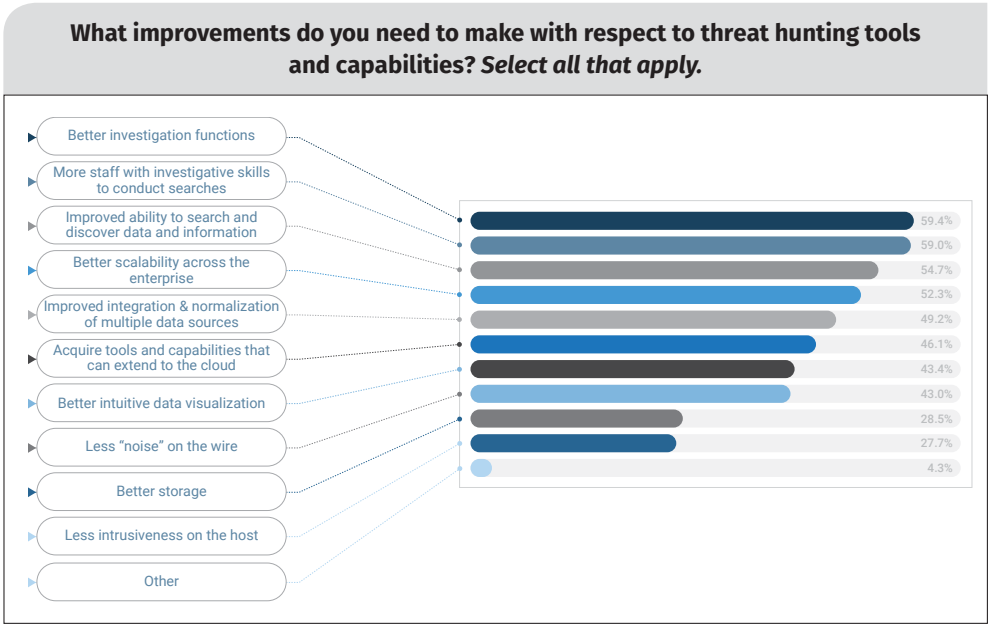

Figure 11. Threat Hunting Improvements Needed

Takeaway: Test Against Tradecraft

One recommendation for greater effectiveness is to use better CTI. Organizations must seek to improve their investigative functions and move past indicators to focus on threat behaviors such as adversary tradecraft. Adversary tradecraft is a far more scalable, transposable and long-lasting form of detection. Testing against adversary tradecraft and generating hunts based on new tradecraft types can lead to significant improvements in investigation capabilities and efficiency.

Threat Hunting: A Growing Necessity

Threat hunting is seen as a consistently growing area of investment in organizations, but it should not replace existing security efforts. Instead, it should seek to complement them. More than half (55%) of survey respondents expect to see investments increase in staffing, and 65% expect increased investment in tools related to threat hunting. Interestingly, these investments largely mirrored each other and were seen to go hand in hand. See Figure 12.

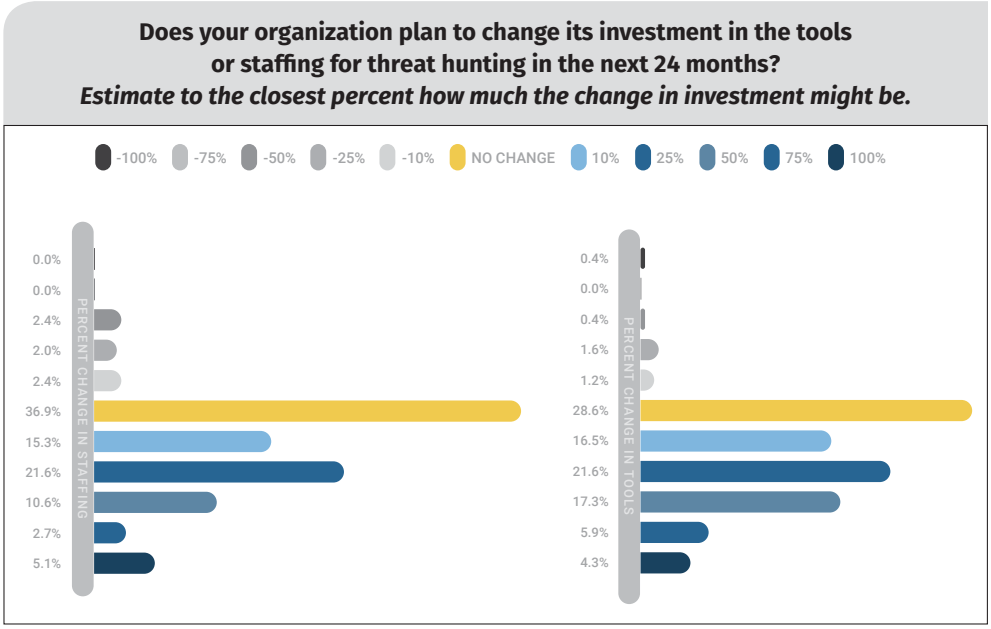

Figure 12. Plans to Change Investment in Threat Hunting

It is a common adage in security that it is easier to get technology than it is to get more skilled analysts. If this holds true for organizations, it will be important to look at technology choices and ensure that the choices help make analysts more effective and efficient, as opposed to introducing entirely new capabilities. Technologies that introduce new capabilities often have a people cost associated with them; if staffing or training are not factored in when assessing technology investments, then the technology investment could run the risk of becoming shelfware.

Takeaway: Prioritize Staffing and Training

We recommend that organizations prioritize finding new staff and training existing staff to ensure that they are ready and able to make use of their technology investments. Additionally, organizations should seek technology that makes it easier for existing staff to test their hypotheses in the organization. Threat hunting is a human-driven process, and thus tools should complement those efforts instead of seeking to replace them. Threat hunting cannot be fully automated, but automation can significantly increase the effectiveness of hunters.

Conclusion

A clear theme from this year's survey responses is that threat intelligence is core to threat hunting. More respondents this year were consuming threat intelligence, ranking threat analysis and threat intelligence as a baseline skill required for threat hunting, and noting the effectiveness of using intelligence to drive their processes. This should come as no surprise, because threat hunting consists of generating a hypothesis and testing that hypothesis in the environment. One of the three primary methods of generating hypotheses is the intelligence-driven method. Thus, having a core intelligence skill set is likely to increase the number and effectiveness of the hypotheses generated and tested.

Additionally, endpoint collection still lags behind network collection and is seen as a difficult data source for most organizations to obtain. Yet respondents rank it as a valued skill set and note that memory forensics, incident response and log analysis are the core types of correlation required for threat hunting. Another key finding is that the No. 1 investment area of threat hunting is still technology, although respondents indicated that the lack of trained staff in numerous areas is an important reason why they did not perform threat hunting or why they did not perform it as effectively as they should.

Too many respondents are trying to continuously hunt or are waiting to hunt based on triggering events. Continuous monitoring and incident response more appropriately map to continuous process and triggering events. Security operations that are proactively finding new hidden threats in the environment are not necessarily performing threat hunting—they simply constitute proactive security.

Threat hunting can be a resource-intensive process, and should be an analyst-focused, hypothesis-driven process. To accomplish this, it is effective to schedule hunts and not overwhelm the organization. Even a few hunts per year, when done correctly, can be highly effective for the organization.

The threat hunting process depends on the structure imposed by hypothesis-generation and testing. That structure leads to repeatability, measurability and success that are not bound to immediately finding threats. Threat hunting is not simply a compromise assessment or continuous security monitoring. Ultimately, threat hunting is an approach that drives security benefits across the organization by making sure that human adversaries are met by human defenders who are taking full advantage of the environment that they defend.

About the Authors

Rob Lee is the curriculum lead and author for digital forensic and incident response training at the SANS Institute. With more than 20 years of experience in computer forensics, vulnerability and exploit discovery, intrusion detection/prevention and incident response, he provides consulting services in the Washington, D.C. area. Before starting his own business, Rob worked with government agencies in the law enforcement, defense and intelligence communities as a lead for vulnerability discovery and exploit development teams, a cyber forensics branch, and a computer forensic and security software development team. Rob was also a director for MANDIANT, a company focused on investigating advanced adversaries, and coauthored *Know Your Enemy: Learning About Security Threats*, 2nd Edition.

Robert M. Lee, a SANS certified instructor and author of the "ICS Active Defense and Incident Response" and "Cyber Threat Intelligence" courses, is the founder and CEO of Dragos, a critical infrastructure cybersecurity company, where he focuses on control system traffic analysis, incident response and threat intelligence research. He has performed defense, intelligence and attack missions in various government organizations, including the establishment of a first-of-its-kind ICS/SCADA cyber threat intelligence and intrusion analysis mission. Author of *SCADA and Me* and a nonresident National Cyber Security Fellow at New America, focusing on critical infrastructure cybersecurity policy issues, Robert was named EnergySec's 2015 Energy Sector Security Professional of the Year.

Sponsor

SANS would like to thank this survey's sponsor:

ANOMALI®