SANS

A SANS Survey

It's Awfully Noisy Out There:

Results of the 2018 SANS Incident Response Survey

Written by **Matt Bromiley**

October 2018

Sponsored by:
ThreatQuotient

SANS Analyst Program

Executive Summary

The past 12 months have been quite a time for those in the incident response (IR) field. We observed data breaches impacting millions of citizens from both financial and political standpoints. The current geopolitical landscape has unfortunately fostered an environment where hacking of any magnitude—including those actions that seek to undermine national elections—will suffer few, if any, tangible repercussions. Yet, we persevere and continue to defend our organizations.

External attacks were not the only hurdles that incident responders had to surmount during the past year. We saw the enforcement of privacy rules and regulations, such as the EU's General Data Protection Regulation (GDPR), and increased PCI security requirements. Given these factors and more, for this survey we settled on a theme of drowning out the "noise" and seeking to focus on the sounds that matter.

Our key takeaways from this year's survey include:

- Compared with 2017, IR teams are detecting, containing and remediating incidents much faster than before. For example, 10% of our respondents can detect within an hour of breach! **This is providing attackers less opportunity to cause damage and giving our teams more time to defend.**
- We're still seeing gaps in response capabilities, whether it's missed incidents, shortage of staff or simple lack of visibility into incidents or data breaches. Some 32% of our respondents were unsure of how many incidents they had not responded to. We cannot stress this enough: **Your IR team should be recording and reporting metrics to help hone its processes.**
- Respondents indicated difficulties in confidently identifying affected data and threat actors from breaches, which may lead to ineffective remediation and eradication. For example, 26% of our participants indicated they had been breached by the same threat actor more than once, with similar tactics, techniques and procedures (TTPs). **Without proper incident scoping, your remediation efforts may be for naught if the attacker can walk right back in.**

We'll look at these and many more results from this year's survey. Whether you're managing your own IR team or looking to implement a new team at an existing organization, we hope you'll find our takeaways impactful and actionable.

Changes in this year's survey allowed us to:

- Better understand how organizations are utilizing assessments of their IR programs
- Dig deep into IRs and determine whether our organizations are correctly classifying their incidents as true or false positives and adjusting procedures accordingly
- Determine whether an organization's efforts were preventing attackers from returning, or whether remediation efforts could be improved
- Uncover which technologies and practices our respondents truly feel would help them move past their current impediments

A Year on the Road

This Year's Respondent Base

Our total respondent population of 452 is drawn from security and other professionals who attested that they are engaged in cybersecurity IR either directly or indirectly (as a manager or security operations center [SOC] operative, for example). Together they represent a global pool of incident responders and organizations:

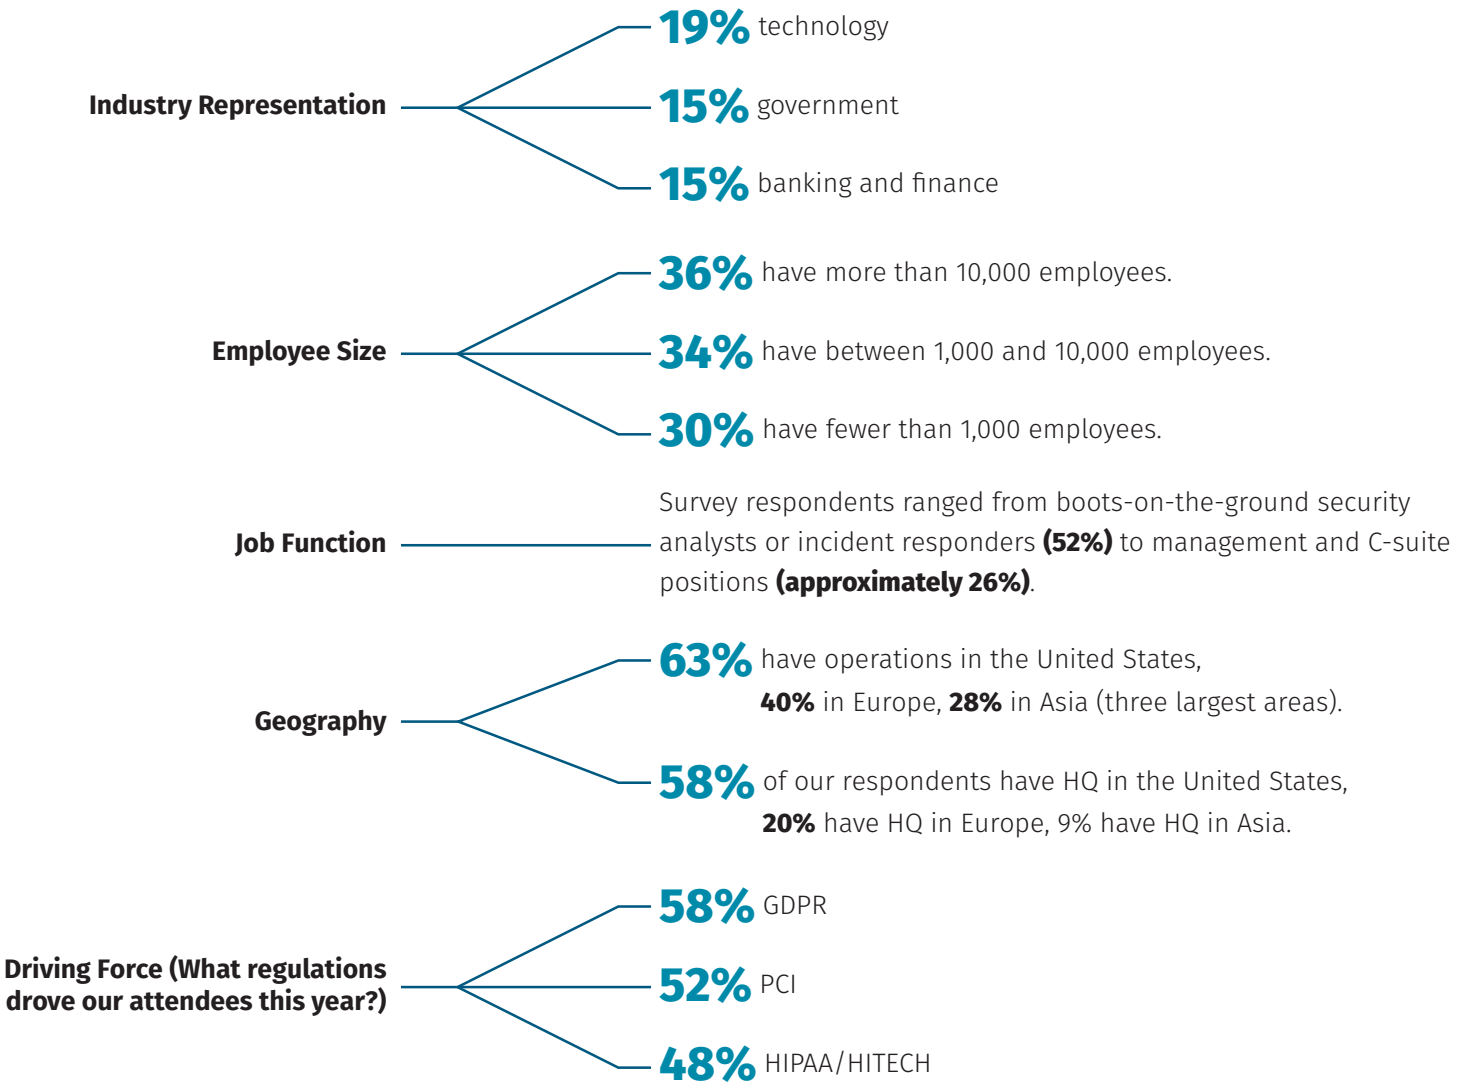

During the past 18 months or so, there were a number of significant breaches. It seems that every year the threat landscape gets worse. We saw Equifax, a major U.S. credit bureau, suffer a breach that put sensitive data of 143 million individuals at risk.[1] Financial data has been a consistently popular target, with multiple restaurant and hotel chains suffering breaches that put credit card data at risk. Between healthcare, higher education, retail and finance, very few sectors were left alone in the past year. Our survey respondents—who originate from many of the industries constantly under attack—were no different.

In this year's survey, our respondent middle ground hovered right around 25 incidents: Approximately 47% of our survey takers responded to between one and 25 incidents within their organizations, while 44% responded to more than 25. If those numbers seem light to you and your organization—don't worry, 24% of the latter group indicated that they responded to at least 100 incidents in the past year! With numbers of this size, we're hoping to see advanced teams and some form of integration and automation (we'll cover these topics later). We were happy to see that 4% of our respondents indicated that they had no incidents in the past 12 months of which they were aware.

We saw some organizations responding to more than 500 incidents in the past year, as well as sometimes 500-plus false positives. Incident classification—part of the process of helping your team fine-tune its tools and procedures—should be used on every incident to ensure you aren't wasting resources.

[1] www.ftc.gov/equifax-data-breach

Incident Metrics

Figure 1 provides a breakdown of number of incidents to which survey participants responded.

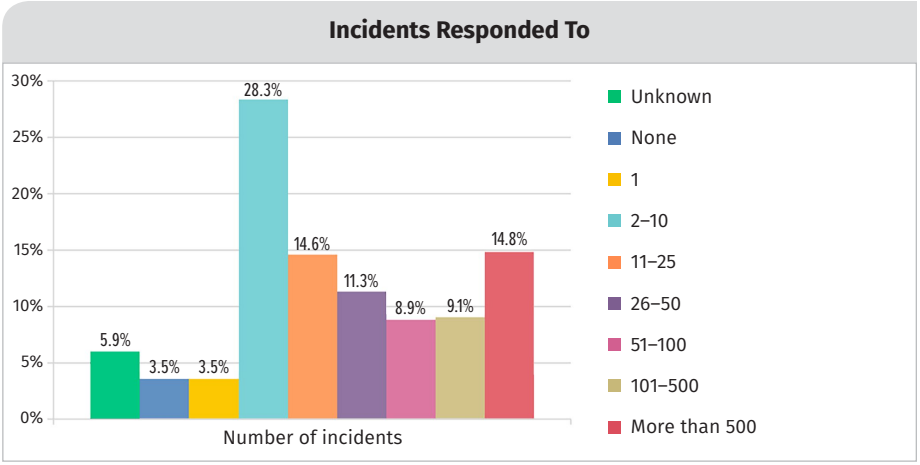

Figure 1. Incident Response Rates

This year we also asked how many of the incidents respondents deemed to be false positives. There's an important underlying question here: Is your organization spending too many resources (be it time, people or money) responding to incidents that are false positives? Does this speak to your alerting mechanisms, or your teams' abilities to correctly validate alerts? Approximately 74% of our respondents, as shown in Figure 2, were able to determine that at least one of their incidents was a false positive.

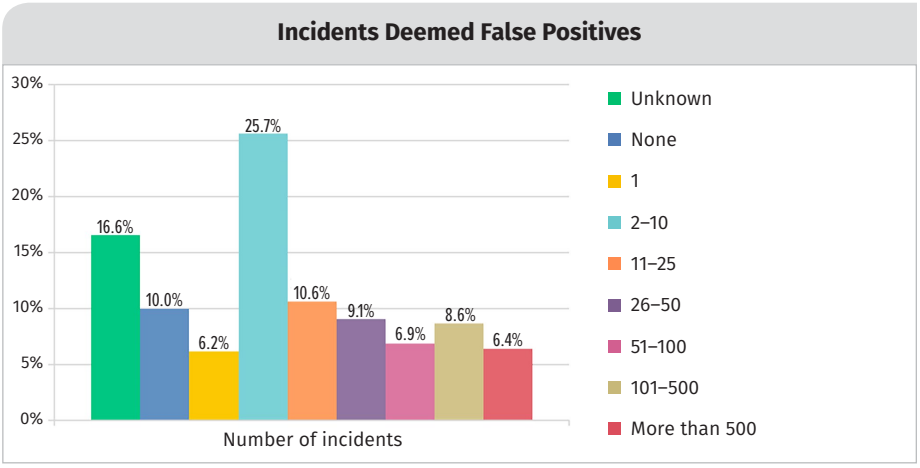

Figure 2. False Positive Rates

The higher we go in incident counts, such as greater than 50, we still have strong representation. In fact, 6% of our respondents indicated that they had *more than 500 false positives* in the past 12 months. That may indicate incident misclassification, which may mean that the team needs to hone its IR procedures.

Lastly, we also asked our respondents how many incidents they did not respond to. In the process of evaluating our teams, we must measure our successes and shortcomings to assess our procedures effectively. Our results on missed incidents were not dire, but they do indicate that we still have some work to do. There was an almost even split between respondents' organizations not missing any incidents and missing at least one incident, at 33% and 35%, respectively. See Figure 3.

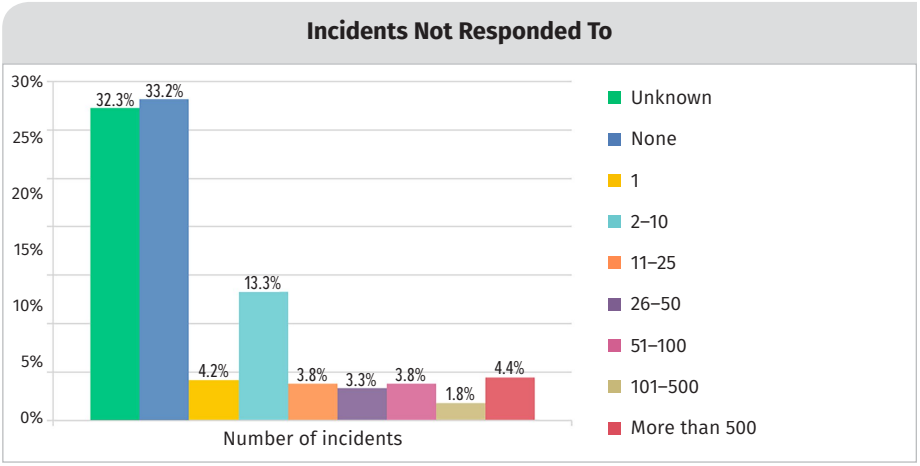

Figure 3. Missed Incident Rates

Unfortunately, the other third of our respondents (32%), as shown in Figure 3, responded "Unknown" to this sub-question. This is one area of concern where we think our IR teams still need improvement. Whether it's a lack of metrics or a number too high to fathom (too much noise, perhaps?), we're concerned when an organization is unsure of how many incidents it did or did not respond to. This may be leaving threat actors unchecked within the environment and providing a significant security risk to the organization.

No More Feedback: From Incident to Breach

While it's helpful to track and maintain metrics on the number of incidents that your team responds to, it is equally—if not more—important to track when those incidents turn into data breaches. As most IR teams know, data walking out the door may lead to breach notifications, involvement of external counsel, or regulatory-mandated investigations—all of which put extra stress on a team already working around-the-clock to get the business back to normal.

In this year's survey, respondents indicated that nearly a third (31%) of their incidents did not result in a breach of information, systems or devices. As seen in Figure 4, a small majority (54%) of incidents did materialize into such breaches; however, half of those affected *only* two to 10 systems (actually 26%).

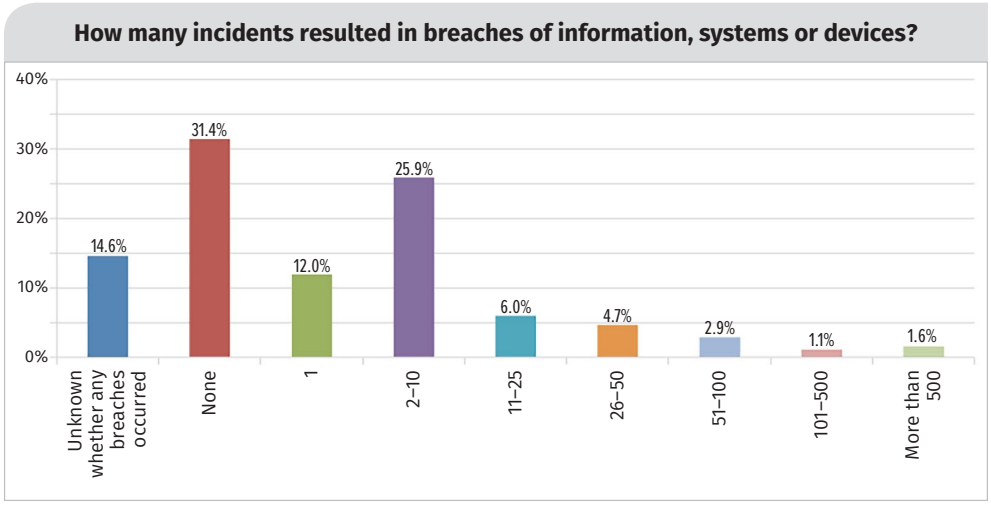

Figure 4. Incidents Resulting in Breaches

Now, that's not an argument to say that we *want* incidents to turn into breaches, but we are glad to see that when incidents do convert, they do so in small quantities. A very small percentage—approximately 3%—of respondents reported incidents breaching more than 100 systems. Finally, we still have issues with visibility, where nearly 15% of respondents reported not knowing whether incidents resulted in data breaches.

One of our traditions in this annual survey is to ask our respondents about the attacker components within their data breaches. As we've seen in previous sections of this paper, not all organizations are tracking and/or keeping metrics on their incidents. However, when they do, it's good to track trends of attack components so that they can determine whether their defenses are adequate and/or whether their organization is defending against the right threat.

We also like to compare these trends against what we are seeing in the field. A significant number of organizations these days are falling victim to account takeover or unauthorized access, particularly in the space of Office 365 and the well-known Business Email Compromise (BEC) land of data breaches. These attacks are passing through standard defenses, because they typically subvert email detection mechanisms and do not require malware to execute. In July 2018, the FBI announced that BEC scams had totaled approximately $12.54 billion in losses from October 2013 through May 2018.[2]

[2] www.ic3.gov/media/2018/180712.aspx

Our survey results lined up with what we are seeing in the field, with respondents indicating that 62% of their data breaches involved malware, while a close second included unauthorized account access (51%). Figure 5 provides further insight, showing that the top five attacker components within data breaches round out with theft of sensitive data (43%), APT/multi-stage attacks (35%) and insider threats (30%).

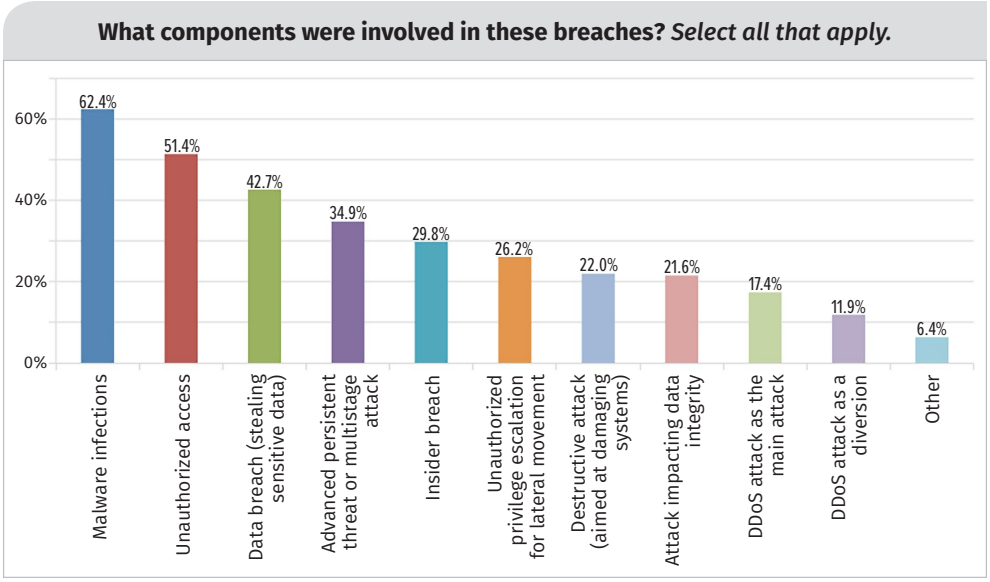

Figure 5. Breach Components

If Figure 5 confirms anything for your IR teams, it's that your organization is likely facing a multitude of threats from all angles and must remain vigilant wherever and whenever possible. Not all attacks require malware, and not all attackers will be nation-states. Understanding the type of systems that your organization has can be a crucial first step, and that segues perfectly into our next question.

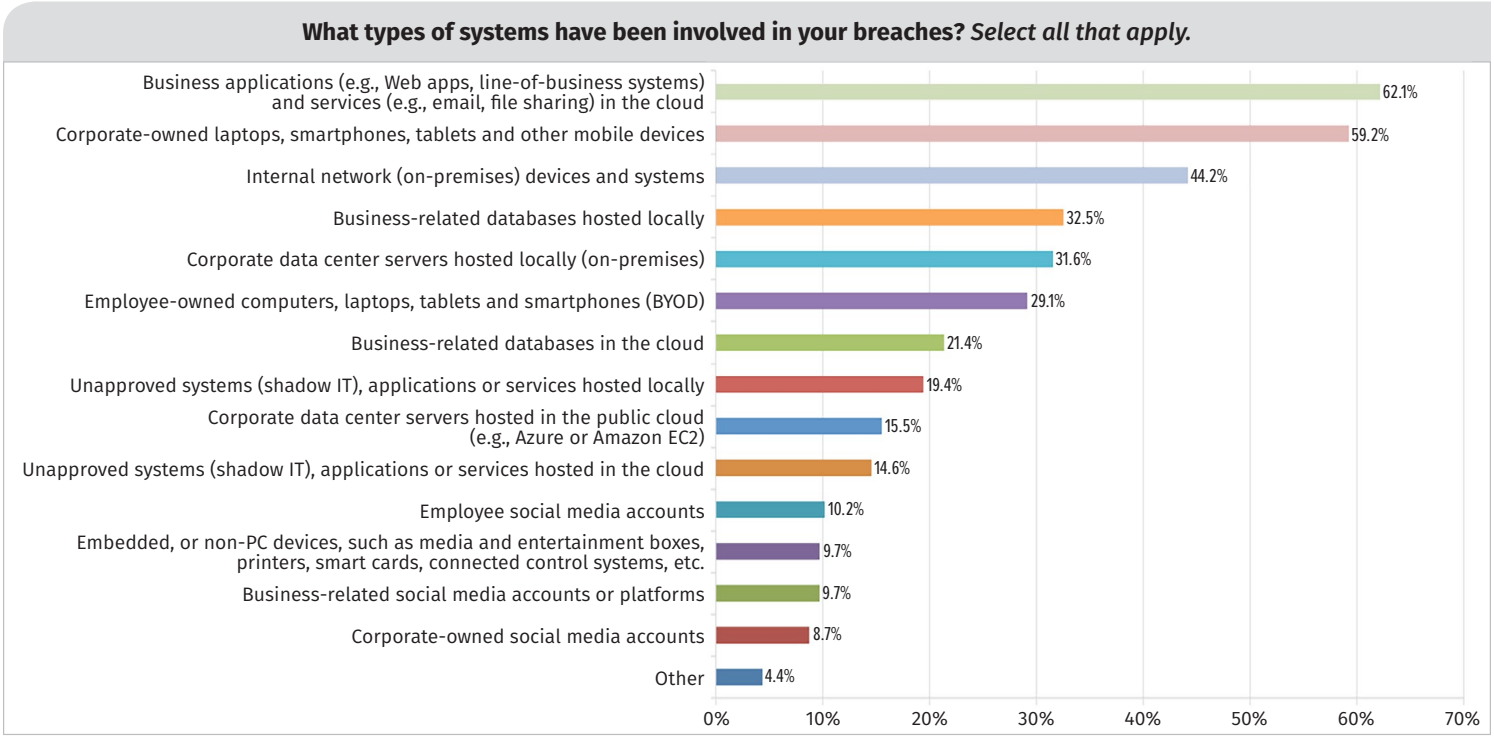

Figure 6. Types of Systems Involved in Breaches

To wrap up our analysis of incidents and breaches (a tough year on the road, it was!), we also asked our respondents what types of systems were involved in the incidents-turned-breaches to which they responded. The top seven categories are all business-related assets (see Figure 6).

Whether they are business applications, corporate devices or systems housing corporate data, it seemed that attackers had no issues moving throughout corporate networks. Rogue or unapproved systems accounted for only 19% when hosted locally, and 15% when hosted in the cloud.

TAKEAWAY

Visibility and IR metrics are absolutely critical to ensure that your organization is responding to the true incidents and not wasting resources (time, money, personnel) on the false positives.

Battling the Noise

Every year, we utilize this survey to identify whether we are seeing improvements in respondents' IR processes. The six-step incident process,[3] illustrated in Figure 7, is our go-to model for a repeatable, tried-and-true IR methodology.

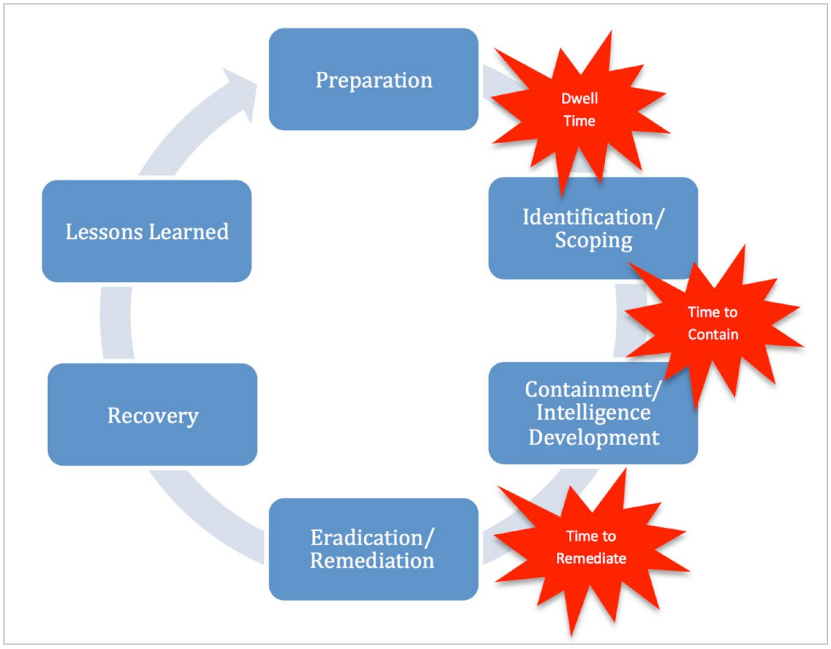

Figure 7. SANS IR Methodology

These steps identify the critical, high-level stages for almost any incident an organization or IR team of any size faces. The goal of the six-step IR process is to ensure that your organization is correctly scoping and containing incidents as they are occurring, successfully remediating the incidents to prevent additional attacker activity, and learning from each incident to improve the team's capabilities.

As highlighted in Figure 7, there are three critical time frames within IR that we hope to gain insight into:

- Time to Detect (aka the "dwell time"). The length of time between initial compromise and detection of a data breach or attack within an environment
- Time to Contain. The length of time between detection and containment
- Time to Remediate. The length of time between containment and remediation

If your organization needs to implement a tried and true IR methodology, consider the SANS six-step IR process.

These time frames represent some of the most important periods during any incident, as they are the critical periods when power shifts from the attacker to the IR team. With shorter dwell, containment and remediation times, attackers have less opportunity to cause damage to their victim organization(s). The three key time frames as reported by survey respondents are summarized in Figure 8. However, we will examine each in detail.

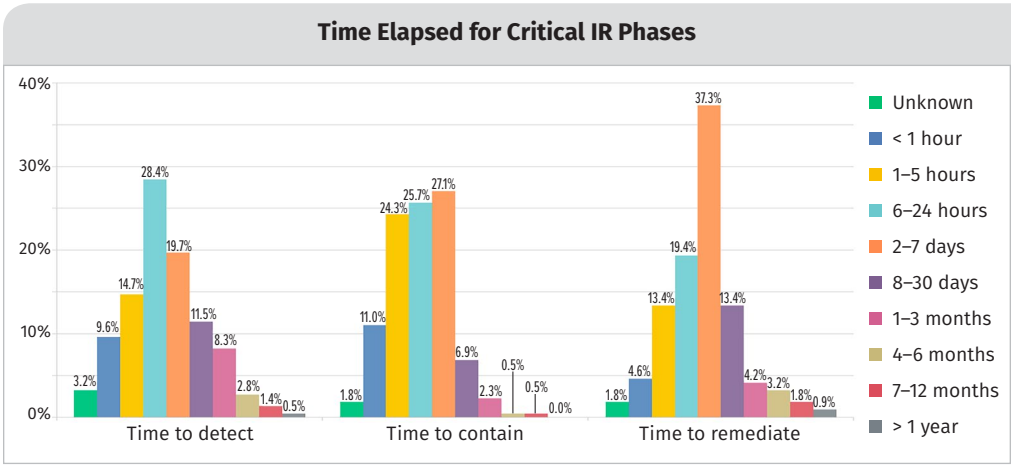

Figure 8. Detection, Containment and Remediation Time Frames

Time to Detect

This year we saw a much-welcomed shift to shorter detection times. Approximately 53% of our organizations are detecting incidents within 24 hours, an increase of three percentage points from last year. One of the more impressive changes is an uptick of two percentage points in organizations that can detect within one hour, which came in at 10% this year. We also saw a healthy shift in the "longer" (greater than 24 hours) end of detection, with organizations moving the

[3] "The Incident Handler's Handbook," December 2011, www.sans.org/reading-room/whitepapers/incident/incident-handlers-handbook-33901

needle to faster and faster detection. The largest change was an approximate three-percentage-point increase in organizations that detected incidents within six months.

We saw an appreciable drop in organizations that required at least a year to detect incidents, with less than half of 1% needing that long. Unfortunately, approximately 3% of respondents also indicated they were unaware of their dwell times—a percentage we always find troubling, because we would like to see organizations keep metrics detailing their incidents. Nevertheless, we were impressed with the shorter time frames all around.

Three of the most crucial time frames in any incident are:

- Time to detect. How long does it take to find a compromise?
- Time to contain. How long to limit or control the damage?
- Time to remediate. How long to remove the attacker for good?

Time to Contain

Containment, the step that seeks to limit or stop the damage as much as possible, is where the power begins to shift to the IR team. While this phase can sometimes take longer than detection, it is also significantly more comprehensive. This year we saw an impressive 8% growth, with approximately 61% of organizations containing incidents within 24 hours. This is huge! We love seeing our IR teams able to contain incidents quickly, which allows the business to recover from a breach faster.

We know our increase in containment times is directly attributable to improved processes, as this year only 36% of organizations required between 24 hours and three months to achieve containment—a 3% drop from 2017. We also saw an improvement in the number of organizations that did not know their time to contain metrics, with a drop of almost 1%. We hope our decrease in containment times was attributable to organizations improving processes, which we will explore later on. We know that attacks have been steady, so this is a welcome change and one we hope will become a permanent trend.

Time to Remediate

Our final time frame of interest is how long it takes for our respondents to convert containment into remediation. Remediation is a stage—typically well-planned but quickly executed—where the organization severs access, implements necessary blocks and removes the attacker(s) from the environment. Sure enough, where we independently saw improvements in dwell and containment times, we saw a noticeable improvement in remediation times as well. Approximately 75% of respondents were able to remediate within one week, compared with 65% in 2017. That's an impressive 9% increase in remediation speed.

Lastly, as we saw with other time frames, a percentage of our respondent base did not have an idea of their respective time frames. We've seen this number decrease across all three time frames; however, time to remediate received the most significant drop of approximately 3%.

As much as we like seeing detection, containment and remediation times decrease, this year we also wanted to know whether threat actors were returning to the same environment. Unfortunately, approximately 44% of our respondents indicated that a threat actor did return; however, in 10% of the cases, the threat actor returned with

different TTPs. Now, this may not mean the threat actor changed TTPs; it may mean the organization uncovered more of the attacker's playbook than was previously known (see Figure 9).

On the other hand, approximately 36% of organizations did not suffer another breach from the same threat actor. While we're glad that some organizations were able to successfully remediate their environment and eradicate the threat actor, we are deeply concerned to see nearly 44% of respondents suffering breaches from the same threat actor at least twice. There are some core reasons why this may present a considerable threat:

- A returning threat actor may not need to perform as much reconnaissance. It's likely you didn't change your entire organization during remediation.
- A returning, successful threat actor is now aware of your remediation techniques and can potentially augment his or her attack toolkit appropriately.
- A returning threat actor likely did not achieve his or her mission earlier, but may now work harder and faster to get there, which means your team needs to be twice as quick.

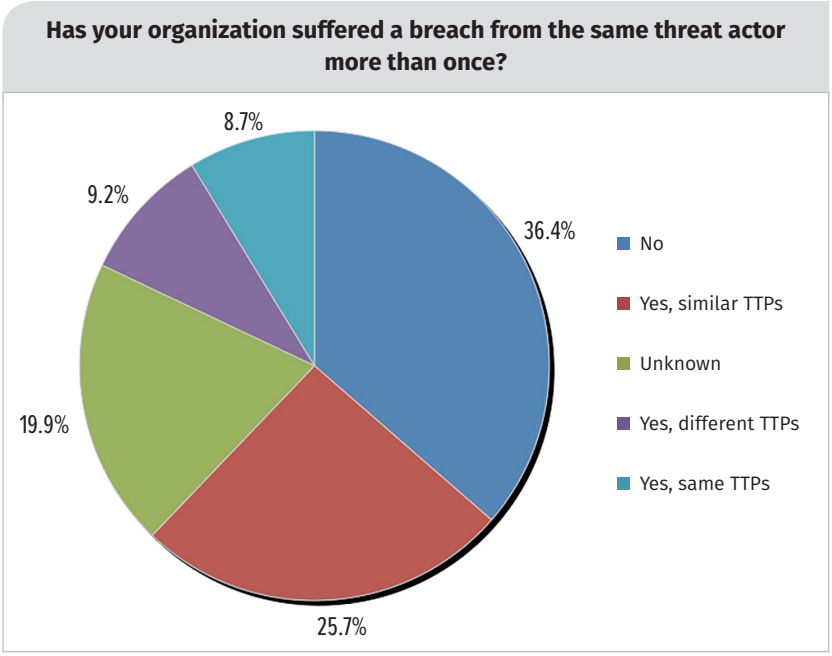

Figure 9. Statistics on Threat Actors Returning to an Environment

TAKEAWAY

Attackers can move quickly, but their activities can be limited by how quickly your organization can detect and respond. The faster you can detect, the less time they have to move deeper.

Next, we'll examine just how deep organizations are taking their investigations, and whether these levels are deep enough to ensure we are successfully remediating against skilled attackers.

Finding the Source

In the previous section, where we discovered that a good percentage of organizations had been breached by the same threat actor at least twice, we began to wonder whether remediation and eradication efforts were simply not enough. Then we began to wonder whether organizations were able to successfully gather as much data as possible about a breach, or if they were leaving holes open for the attacker to walk through as soon as the coast was clear (for example, was the entry vector of the intrusion effectively determined and/or patched?).

To begin, we asked our participants whether they were able to consistently and accurately discover the affected users, systems, data and threat actors involved. Without a healthy blend of each category, you run the risk of inaccurately scoping or misclassifying a breach. We also asked our respondents about the level of difficulty associated with the discovery of these data points. See Figure 10.

It's clear in Figure 10 where the strengths of our IR teams lie. An overwhelming majority of our participants expressed that they were able to identify both users and systems, either easily or with difficulty. To put it into perspective, respondents were able to identify users (56%) and systems (61%) with ease. On the other hand, our respondents indicated difficulties in identifying impacted data, with only 30% able to identify data easily. The most difficult category, without a doubt, was identification of threat actors. Approximately 40% of respondents said they were unable to consistently and accurately identify threat actor details. This gap in visibility might correlate directly to an organization's inability to successfully eradicate an attacker, leading to the same attacker re-compromising an environment.

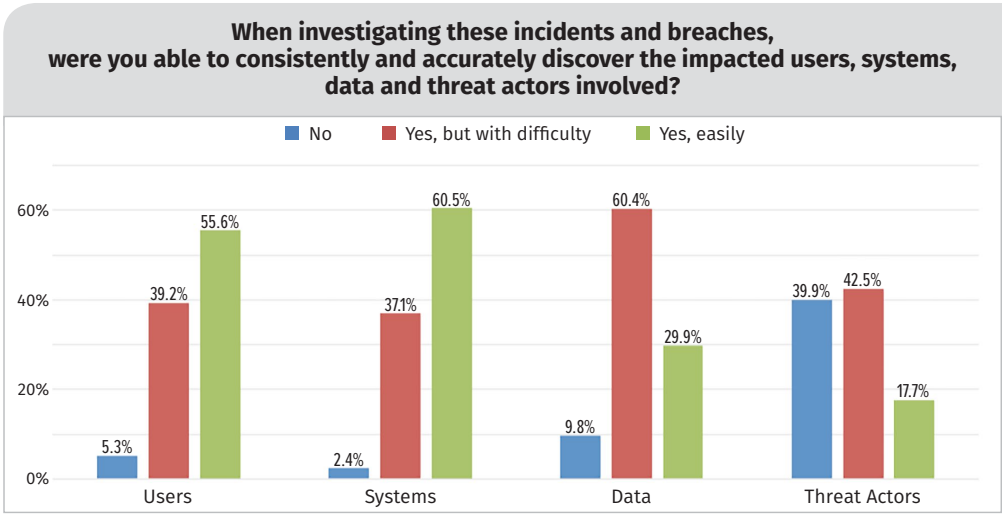

Figure 10. Difficulty of Data Recovery

EXPERT ADVICE

If your IR team is unable to confidently identify threat actor data, your remediation efforts may result in incomplete removal of the attacker. Couple enterprise visibility with threat-party enrichment (such as threat intelligence) to ensure that you aren't missing any obvious indicators.

It's All in the Malware

One mechanism we can use to determine how organizations are collecting threat actor data is by understanding how they analyze the tools the attacker leaves in the environment. Previously we indicated that not all attacks are malware-based; however, a significant number of attacks still rely on malware and their custom toolkits. We asked our respondents how they handled these bits of evidence when they encountered them. The results were promising.

Overall, 73% of our respondents collect malware that is destined for some sort of analysis, whether it's by a dedicated team or handed off to a third party. Approximately 30% of our respondents indicated that they have a dedicated malware analysis team or individual. Now, while we hope that one individual is not being overworked, we're glad to see that organizations are realizing the value of analyzing malware. Malware analysis can play an important role within any IR team, enabling the team to correctly garner indicators and actionable takeaways from their findings. Figure 11 shows our respondents' malware collection and analysis activities.

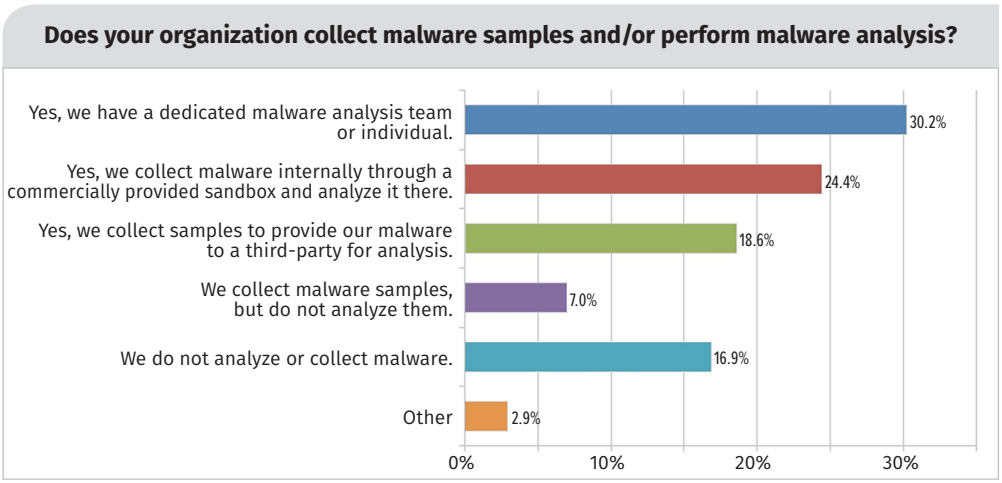

Figure 11. Collection/Analysis of Malware

Malware Analysis, IR and Then Some

Each year we like to look at the level of integration between the IR team and the SOC teams. This year, approximately 30% of our respondents indicated that the IR team is fully integrated with the SOC, and members are cross-trained. We love this! It's great to see that nearly a third of our respondents understand the benefits of having the SOC and IR teams work together so closely. While nearly 30% of our respondents have integrated IR and SOC capabilities, approximately 23% responded that the IR team is independent of the SOC. Interestingly, almost 7% of our respondents indicated that the IR team is completely outsourced, but does work with the SOC on investigations.

Figure 12 provides additional details on IR and SOC integration.

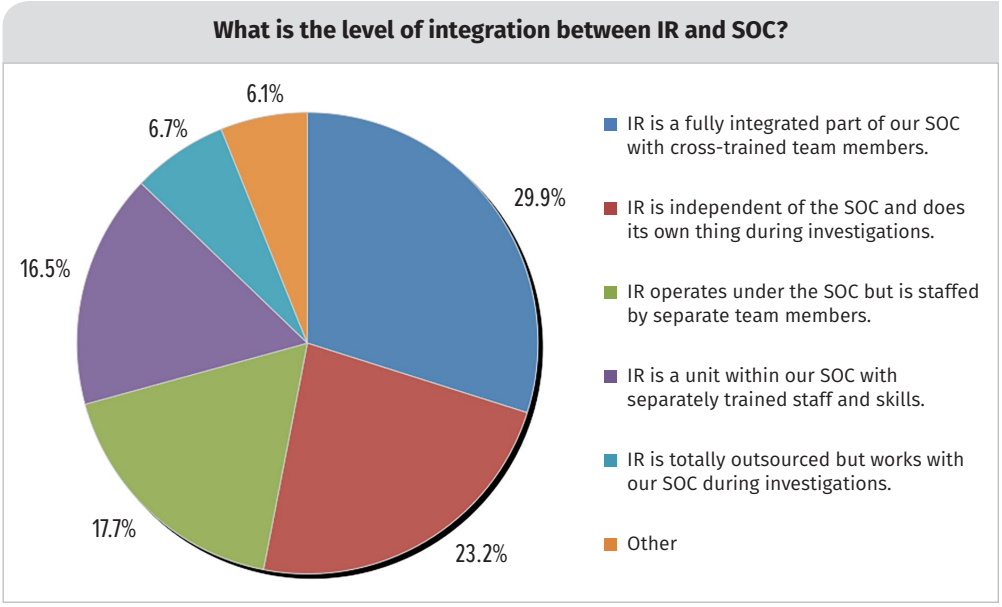

Figure 12. Integration Between IR and SOC[4]

TAKEAWAY

Malware analysis is a critical element to understanding the gravity of threats facing and within your organization. If you can't reverse-engineer yourself, get help from a third party that can give you actionable analyses and findings.

Planning for the Next Gig

It's time to look ahead—to the next gig. We typically save the final section of our survey analysis to ask our respondents to look ahead and let us know what the future of their IR programs looks like. Initially, we want to know what organizations are seeing as their top impediments. Every year, our top results remain the same: Organizations want more skilled people on their IR teams. Approximately 52% of our respondents indicated that their top impediment was a shortage of staffing and skills. Close behind is a lack of budget for tools and technology (48%) followed by poorly defined processes and owners (44%). See Figure 13 on next page.

Some of the additional impediments IR teams are running into include multiple levels lacking visibility, something that we've discussed as a key issue throughout this paper. We cannot stress this enough: Without proper visibility into the organization, it's going to be tough to effectively implement and act on each step of the six-stage IR process.

We've examined what's wrong, so let's look at how we're working to make it better. Particularly, we asked our respondents what IR improvements they planned to make in the next 12 months. Approximately 54% of our respondents indicated that they plan to allocate funds for additional training and certification of staff—a hopeful prediction, given the shortcomings we previously discussed (see Figure 14 on next page).

If your organization is running into impediments of poorly designed processes and owners, consider revamping using well-known methodologies and frameworks. There are plenty of examples out there, such as the SANS six-step IR methodology, which we covered earlier in this paper. Process owners are another story, but we can work on the processes themselves first.

[4] Respondents also had a response option of: "IR is totally outsourced but works with our SOC during investigations." No respondents chose that option.

What do you believe are the key impediments to effective IR at your organization? ***Select your top five choices, not in any particular order.***

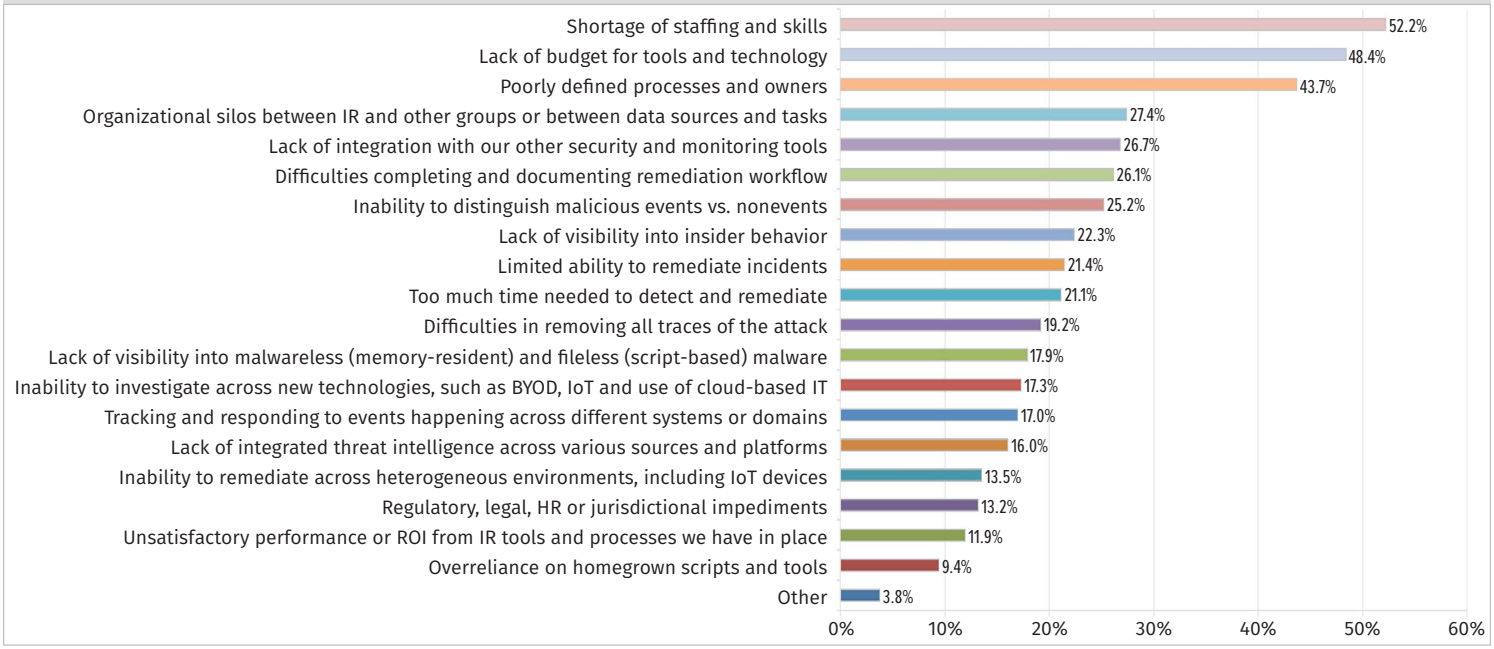

Figure 13. Top Impediments to Effective IR

What improvements in IR is your organization planning to make in the next 12 months? ***Select all that apply***

Figure 14. Planned Improvements/ Investments in IR

Respondents are asking for larger staffs, team members with knowledge and experience, and in-depth training, to name a few of the popular results.

Another pleasant surprise is that we see a planned increase in spending in automation and threat hunting. Approximately 50% of organizations said they plan to spend on improving automating the remediation processes (we need that!), and 45% are planning to spend on proactive threat hunting. Hearing that organizations are ready to move to proactive threat hunting capabilities is nothing short of promising—we just encourage you to ensure that your reactive game is tip-top first.

Staffing and skills are a consistent pain point. Be sure that your IR teams are receiving the training and support they need, and the other factors will receive the attention they need.

Conclusion

As with each iteration, this year's IR survey provided valuable insight into the current state of IR across a myriad of industries and organizations. While many reports focus on the severity of data breaches and threat actors' activities, we enjoy getting insight into how our IR teams have progressed within the past 12 months.

Our high points from this year's survey indicate that IR teams are still working hard to defend their organizations and are achieving measurable success in key areas. We're seeing drops in dwell, containment and remediation times—three of the most important time frames. However, the results concerning effective remediation and collection of threat actor data (with respect to a breach) indicate that organizations are potentially leaving gaps open.

Additionally, we identified issues in visibility this year, as multiple survey respondents were unable to answer questions completely due to lack of telemetry. One way to combat visibility issues may come from increased tool automation and integration, which will simultaneously allow the IR team to establish robust response processes. There's a win-win situation here: More visibility will also help reduce false positives, another issue we saw in this year's survey as well.

The worst thing an IR team can do is to create additional problems for itself, such as inaccurately scope an incident or refrain from utilizing all available evidence sources within the organization. These (and many more) issues could impede responders' ability to adequately protect the organization, and thus give the attackers a better chance at success. So, our final advice to our fellow incident responders is this: Keep fighting the good fight and causing disruptions for the attackers to work around. As tools and teams become nimbler, attackers will have a harder time keeping up. Utilize your tools, your environment and your people to build a formidable force, and attackers may think twice when your organization comes up in the crosshairs.

About the Author

Matt Bromiley is a SANS Digital Forensics and Incident Response instructor, teaching Advanced Digital Forensics, Incident Response, and Threat Hunting (FOR508) and Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response (FOR572), and a GIAC Advisory Board member. He is also a principal incident response consultant at a major incident response and forensic analysis company, combining experience in digital forensics, incident response/triage and log analytics. His skills include disk, database, memory and network forensics, as well as network security monitoring. Matt has worked with clients of all types and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor: