

Survey

Think Like a Hacker: Inside the Minds and Methods of Modern Adversaries

Written by [Matt Bromiley](#)

September 2022

Introduction

Protecting an enterprise environment is no easy feat, especially in today's complex digital landscape. A multitude of factors contribute to an overwhelmingly expansive attack surface that security teams must protect: hybrid cloud environments with multiple providers, a remote and disparate workforce, ever-increasing complexities of legacy, on-premises solutions. These are all part of a moving target that is increasingly difficult to defend.

Unfortunately, resources are limited, forcing organizations to make calculated investments from a defense-in-depth perspective across prevention, detection, response, and recovery. While the defensive approach is critical, it considers attack scenarios from a theoretical point of view often based on historical trends, threat intelligence, and internally calculated risk scenarios, all of which are often outdated. The defensive viewpoint is only one of many that must be taken into consideration. A comprehensive approach includes looking at things from the attacker's perspective.

In this inaugural 2022 SANS Ethical Hacking Survey, we aimed to understand the intricacies of how attackers think, the tools they use, their speed, their specialization, their favorite targets, etc. These insights are critical to investment decisions across an increasingly complex attack surface that is becoming more difficult to protect.

A few great insights and takeaways include:

- Approximately 37% of respondents indicate that they can break into an environment more often than not, if not **always**.
- Nearly 64% of respondents need five hours or less to **collect and potentially exfiltrate data**.
- **Social engineering and phishing** account for nearly half of all attack vectors that hackers cite as yielding the highest return on investment.

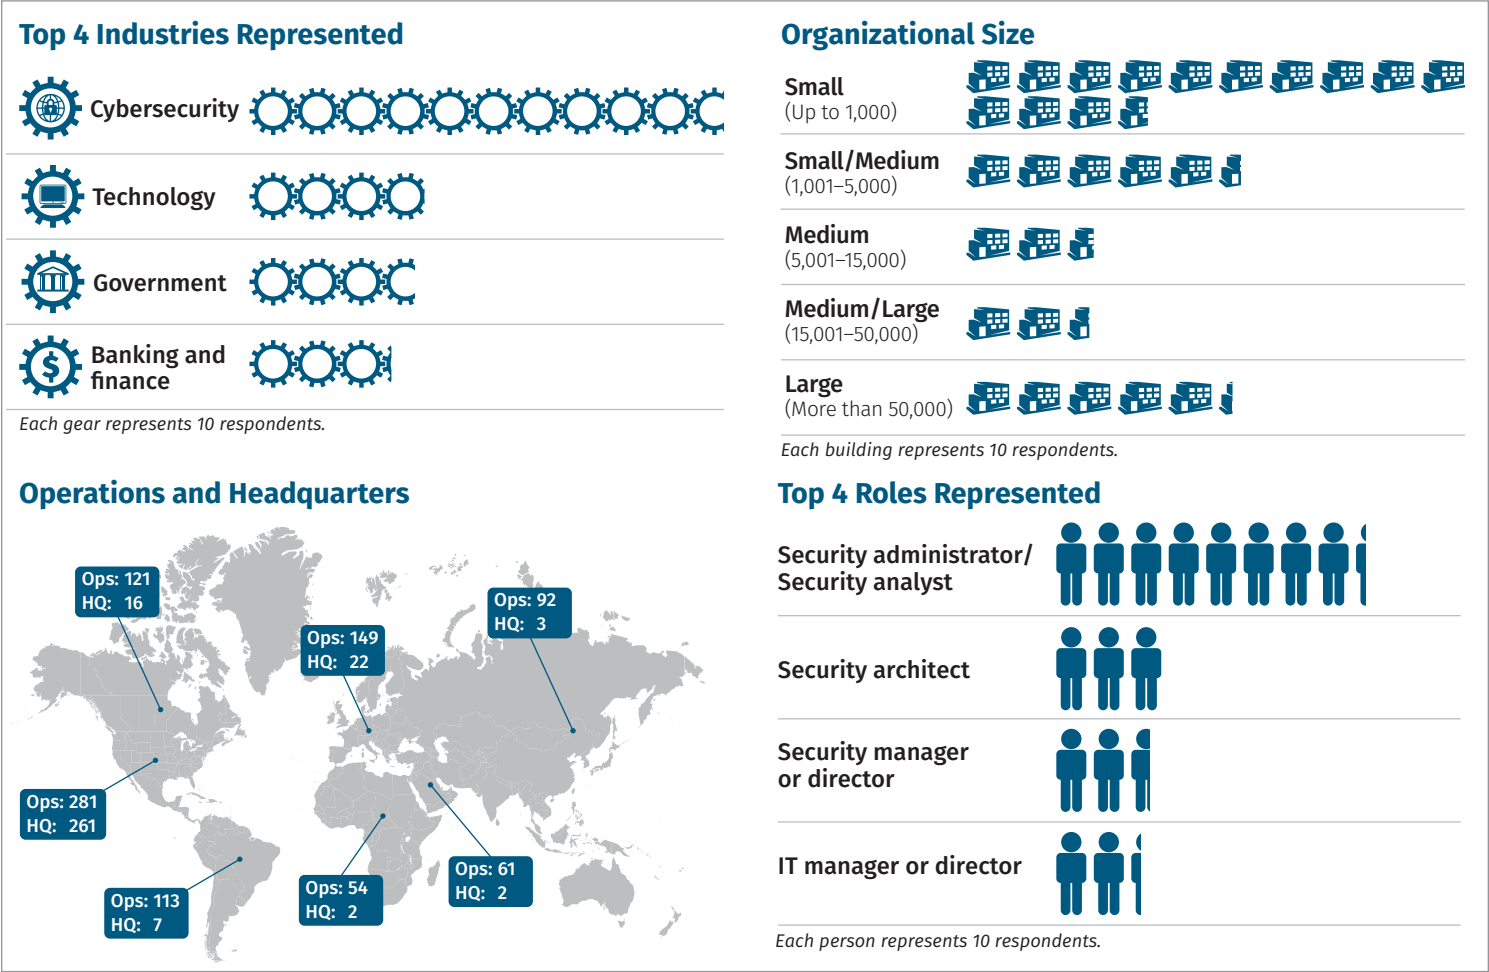
The bird's eye view of an adversary—whether sanctioned or not—can be a guiding light for security analysts and decision makers alike. Oftentimes, we see organizations that invest in security technologies that mitigate a wide range of threats leave commonly attacked ports and protocols wide open. Adversaries will choose the path of least resistance or the one they are most familiar with—and far too often, these are the same. Overlooked or assumed safety presents too much of a risk.

As you work your way through this paper, we encourage you to consider some of our respondents' answers in the same way we phrased the question: What is the return on investment on various attack vectors? What works and what doesn't? After you have absorbed the information in this paper, apply that knowledge to your defensive and offensive security investments. What works and what doesn't? How can your organization protect against what an adversary is likely to throw at you?

As with any SANS survey, we relied on our demographic questions to help depict what part(s) of the world our respondents are coming from, what industries they represent, and the size of their respective companies. A few notables from this year include:

- Companies with headquarters in the United States represent the lion’s share—a whopping 83.4%—of this year’s respondents.
- A little over a third, or 35.1%, of our respondents work in companies with 500 people or fewer.
- The largest segment of respondents (34.2%) work in the cybersecurity industry, and roles range from security administrator/analyst to CSO/CISO/VP of security or technology.

Figure 1 contains a brief breakdown of the respondents to this year’s survey.



An Examination of Ethical Hackers

When we devised the format and questions for this year’s survey, we wanted to ensure that we could capture the effort that adversaries—whether red teamers, penetration testers, or gray hats—put into breaking into their targets. However, we’re aware that welcome and unwelcome adversaries are very different. Some are present because we asked, while others we never want to see.

In this paper, we call out the difference between sanctioned and unsanctioned adversaries. Sanctioned adversaries, also commonly known as ethical hackers, are hired to attack their targets. Unsanctioned adversaries, often just called hackers, choose their targets based on motive, opportunity, financial gain, or intentional targeting (such as for corporate or state espionage). While the “why” may be different, the vectors, success rates, and observations share a nearly overlapped Venn diagram of tactics, techniques, and takeaways.

Adversaries, Welcome and Not
A *sanctioned adversary* (or an ethical hacker) is one that is hired to attack a particular network—think of red teamers or penetration testers. *Unsanctioned adversaries* represent the criminal element of attacks: black hats, criminal enterprises, and state-nexus espionage actors. The motives of sanctioned and unsanctioned adversaries may differ, but their approaches are often identical. In fact, 45% of respondents said if they used unethical measures, it would have a high or extremely high impact on their success.

Meet Our Hackers

It’s time to meet the hackers behind this year’s survey results. Our first set of questions focused on understanding the background of our respondents to better assess their responses and whether experience may guide an adversary differently. A large majority—nearly 84%—of our respondents have been ethically hacking for 10 years or less (see Figure 2).

In fact, the majority of respondents have been ethical hackers for one to six years. Because ethical hacking is conducted en masse by many security companies, this range of experience was expected and represents the commonly observed industry makeup. In addition, most of the respondents (87%) hold a security certification, with common ones including SANS GIAC Penetration Testing (44%), CISSP (40%), Security+ (28%), OSCP (26%), and CEH (26%).

Most of the respondents have served or are currently serving as internal security (76%), but many have also worked as a consultant, bug bounty hunter, or hacker for hire, all of which gives them experience from several environments (see Figure 3).

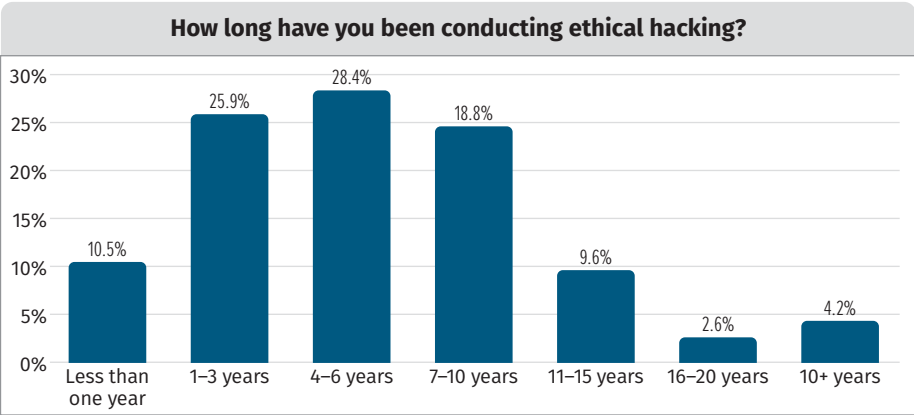


Figure 2. Years of Conducting Ethical Hacking

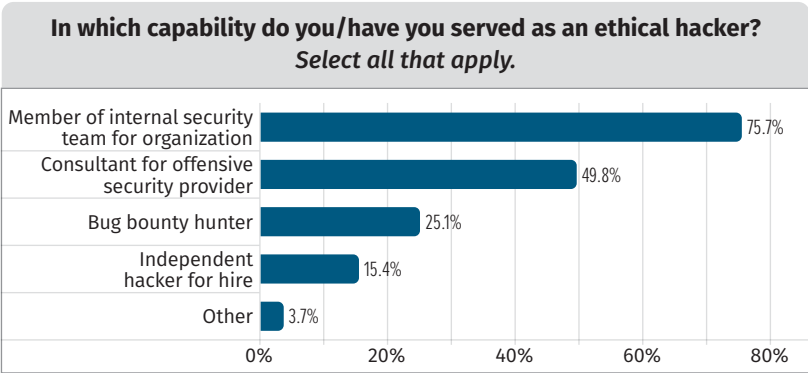


Figure 3. Experience as Ethical Hackers

We also asked respondents for their areas of specialty, again looking to see if these would guide subsequent results. The top five specializations, as seen in Figure 4, are network security, internal penetration testing, application security, red teaming, and cloud security.

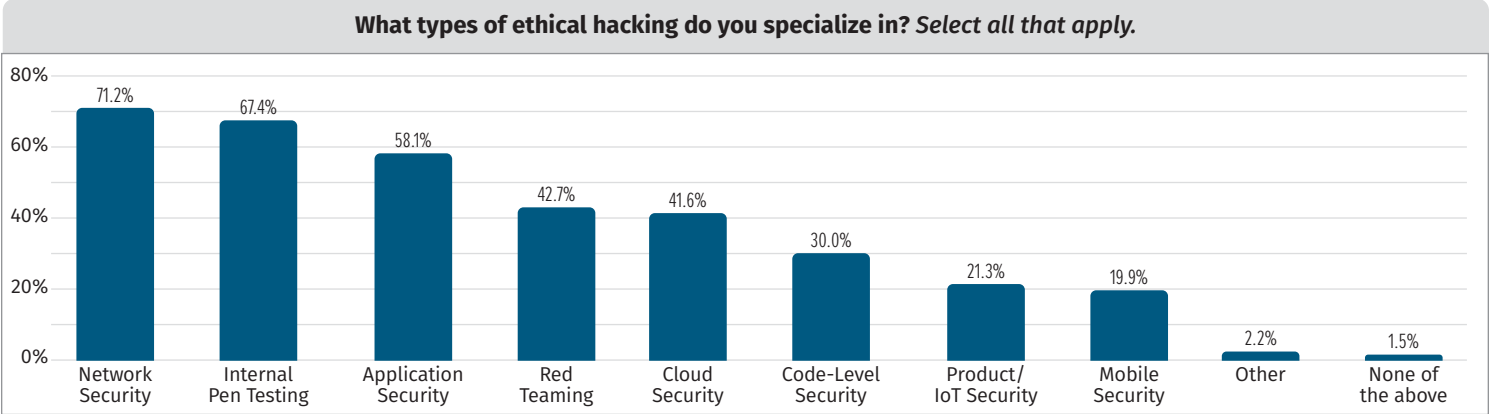


Figure 4. Areas of Specialization

These results come as no surprise, because they represent the architecture and testing designs of most organizations. For example, these days most organizations are deploying custom applications and have moved part of their infrastructure to the cloud. We would expect these areas of focus to grow year after year, because ethical hackers hone their craft to reflect their environment and/or their customers’ needs.

Speed

Our survey covered a wide range of topics, from asking how quickly ethical hackers could breach the perimeter of an environment to how quickly they could shift tactics if necessary. Speed became a central theme of questions, just as it is likely a top concern of security teams. Mean time-to-detect (MTTD) and mean time-to-contain (MTTC) speeds are often compared against adversary speeds, establishing acceptable time frames for detection and response and, when those time frames are unrealistic, making the case for significant investment in prevention.

More than half of respondents (over 57%) stated they could successfully discover an exploitable exposure in 10 hours or less (see Figure 5).

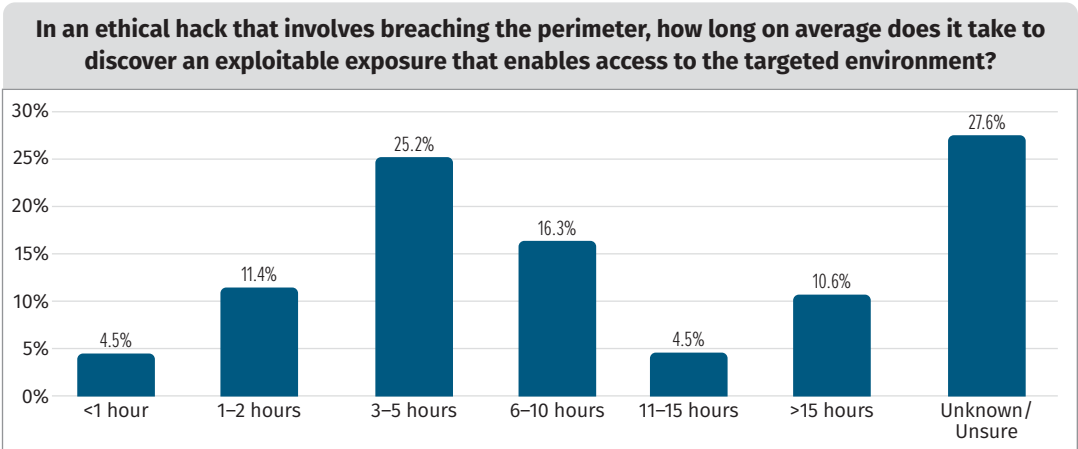


Figure 5. Time to Exposure Discovery

Reconnaissance and exploit *discovery* is one of the hardest phases for defenders to proactively identify—oftentimes they may not be aware of a weakness until it is exploited. Given the speeds reported by our respondents, this supports the need for proactive discovery and prevention and exploitable exposures, i.e., attack surface management, penetration testing, etc.

This question also surfaced an issue we see in many SANS surveys, and one worth calling out this early on. Nearly 28% of respondents to this question stated that it was unknown, or they were unsure of, how quickly they were able to identify an exploitable exposure. Possible explanations include:

- Ethical hackers do not keep track of or are unable to equate how much time perimeter discovery may take. This leads us to ask whether this a viable metric that organizations would benefit from.
- Is there an assumption that an ethical hacker will get in—therefore, measuring this metric offers no value?
- This is not the area where ethical hackers focus their time and efforts. We’ll examine Unknowns/Unsure in subsequent questions to see if this trend continues.

We posit another possibility here – outcomes from ethical hacking results may be measured in *overall outcomes*. An organization might ask, “Were you able to get in?” and expect a report of techniques, and not necessarily a breakdown of how quickly each step occurred (although many offensive security professionals do keep such metrics).

We were also curious to see if our results differed based on respondent specialty. Figure 6 provide a cross-section of this data.

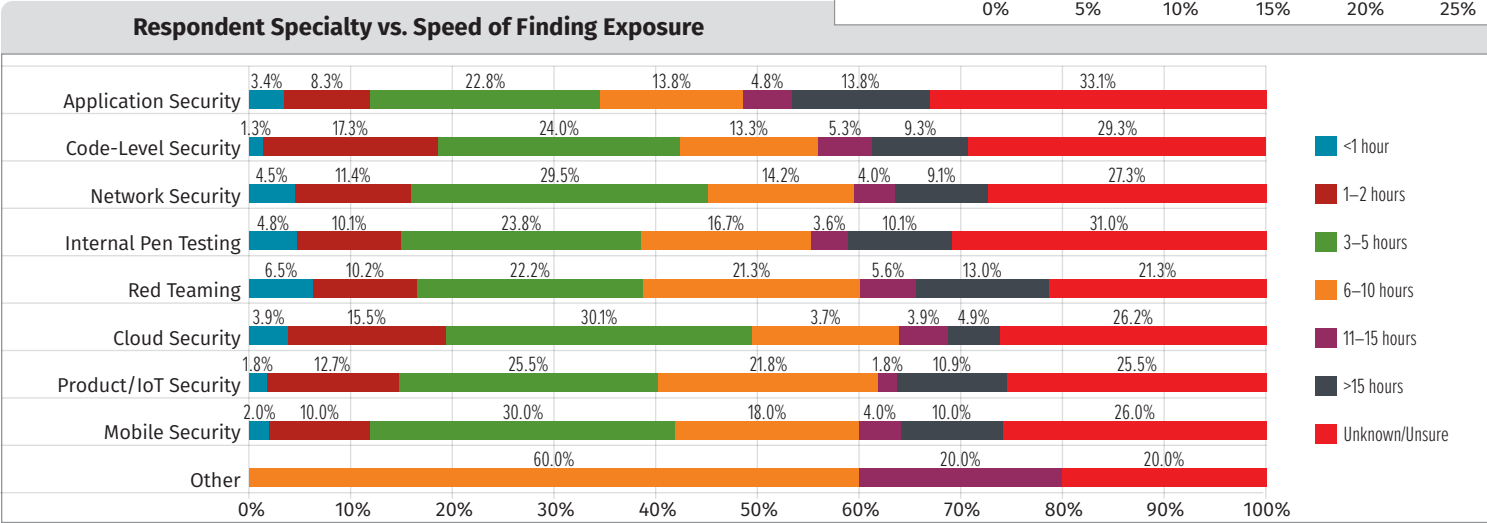
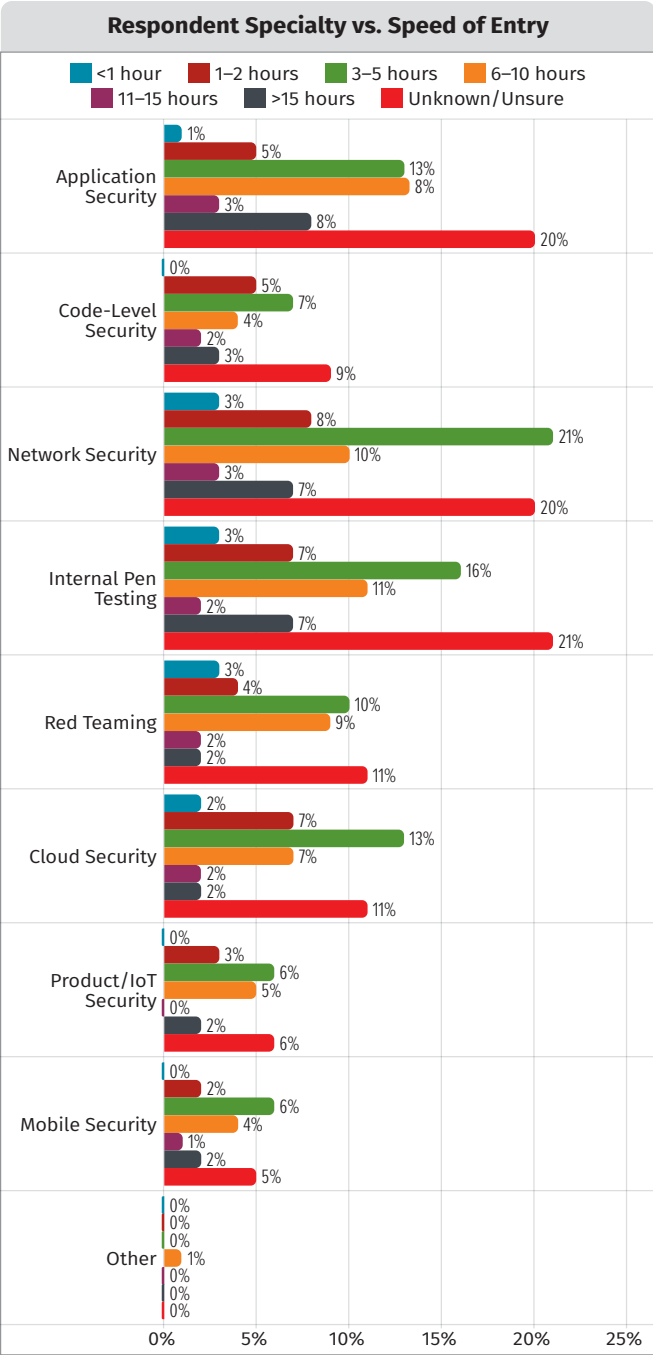


Figure 6. Respondent Specialty vs. Speed of Entry and Exposure

Our data shows that the majority of respondents with experience in application security, network security, and internal penetration testing were able to find an exploitable exposure in five hours or less. Regardless of specialty, few respondents were able to find an exploitable entry point in less than an hour.

The results in Figure 6 lead to our next question, which asks how long it takes an ethical hacker to move from discovery to exploit. Figure 7 again shows that an overwhelming majority (more than 58%) can exploit and break into an environment in five hours or less.

With regards to speed, “discovery of an exploit” and “exploitation leading to intrusion” are significant metrics that any organization should be tracking, especially in relation to the distribution of investment across prevention, detection, response, and recovery.

We also saw speed play an important factor in the survey in one way that many would not expect. We asked our respondents about the most significant factors contributing to vulnerable attack surfaces—what are organizations doing that open them up to compromise? As we can see in Figure 8, pace of application development/deployment and third-party connections are the top two factors for vulnerability exposure. Figure 8 expands on these findings a bit more.

Unfortunately, we find that these results align with what many in the industry have been struggling with from a visibility and an identification perspective.

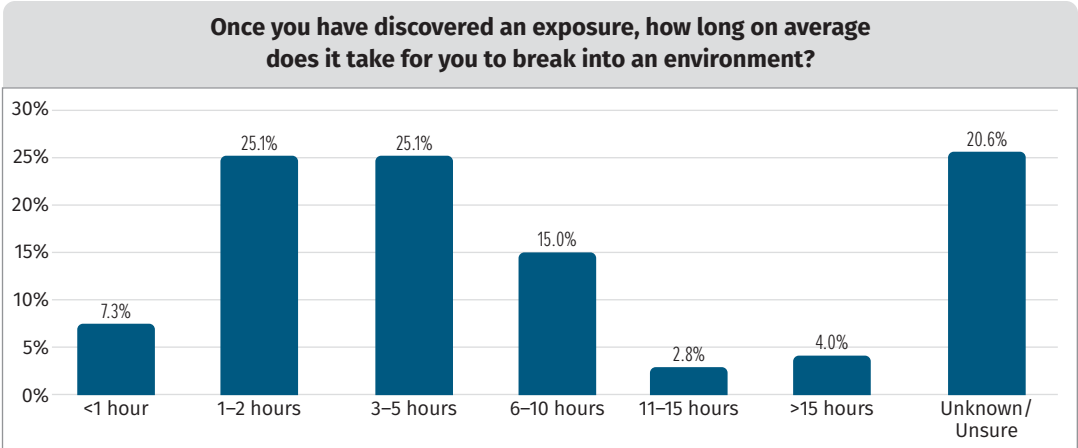


Figure 7. Length of Time on Average from Exposure Discovery to Breach

Key Defender Metric
Many respondents indicated that they could discover exposures within an environment and exploit them in 10 hours or less. That time becomes a benchmark measuring how adequately an organization is positioned to prevent exploits. Those that fail must be able to detect and contain exploits before the attacker escalates privileges, accesses the target, and exfiltrates.

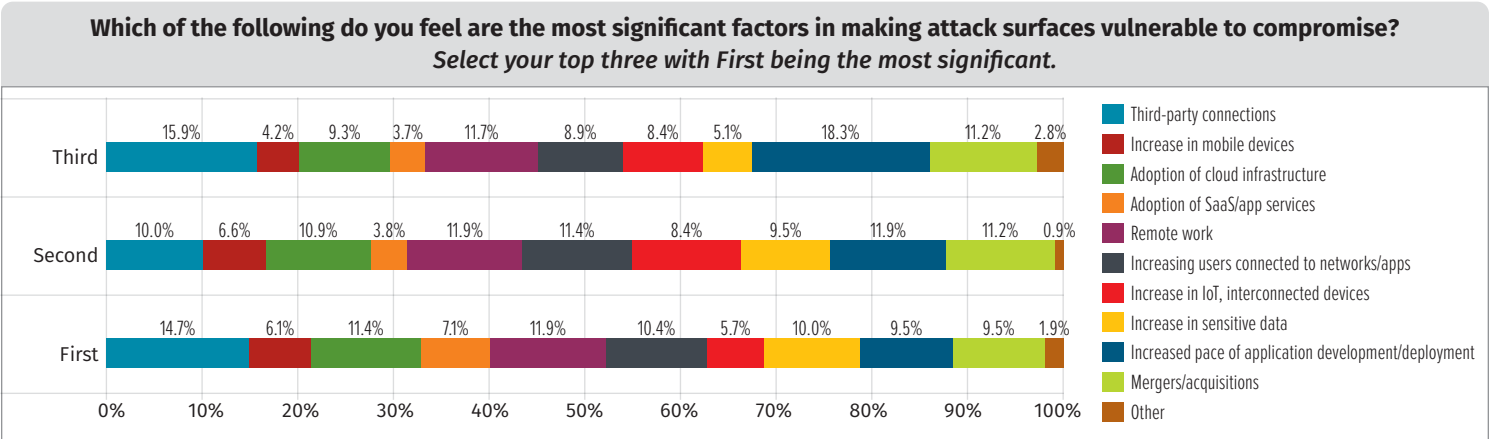


Figure 8. The Most Significant Factors in the Vulnerability of Attack Surfaces

On the topic of attack surface exposures, respondents were asked which attack surface exposures they most frequently run into. Supporting these findings, misconfigurations, vulnerable software, and exposed web services are highly correlative to increased pressure to deploy new applications that support business initiatives as seen in Figure 9.

Interestingly, we found there was little impact in relation to specialty and exposure prevalence. This is likely due to the fact that an adversary who is charged with compromising an environment will look for the path of least resistance (i.e., vulnerable applications, information disclosures, etc.) that is most prevalent across layers of the attack surface.

We asked respondents with cloud security experience how often they encountered improperly configured or insecure cloud/laaS assets. As seen in Figure 10, there’s an even split between “less than half the time” and “more often than not,” with small percentages that rarely see (4.6%) or always see (8.0%) misconfigured public cloud or laaS assets. These stats support an unfortunate truth that, as we see in previous figures, organizations develop and deploy applications that expose vulnerabilities, insecurities, and improper configurations for adversaries to take advantage of.

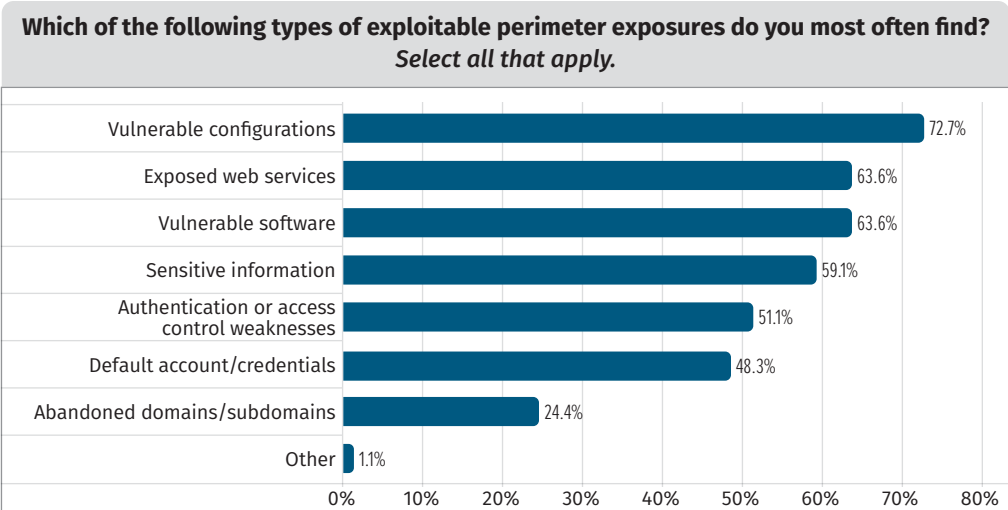


Figure 9. Most Common Exploitable Perimeter Exposures

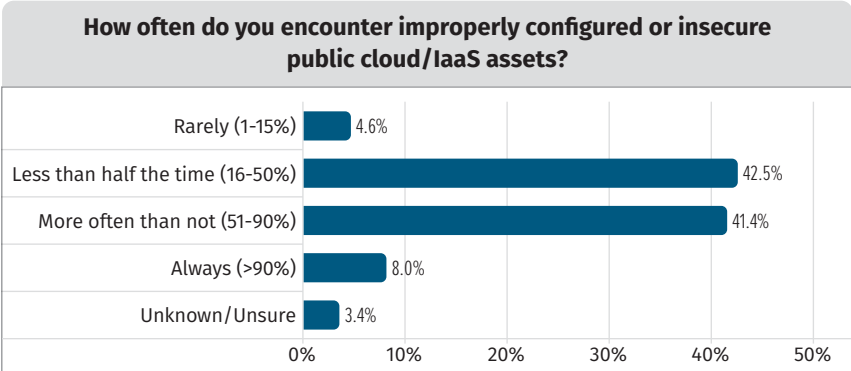


Figure 10. Incidence of Insecure Public Cloud/laaS Assets

Key Stages of an Attack

Once an adversary is inside an environment, execution of subsequent stages (gaining access to the target and exfiltrating data) occurs even faster than reconnaissance and initial exploitation. However, once initial intrusion occurs, an adversary is more open to detection.

Figure 11 looks at the time required to escalate privileges and/or move laterally among targets within a victim network, showing that 36% of respondents said they could escalate or move laterally within three to five hours and that a concerning 20% can do so in two hours or less.

We found it interesting that the majority of respondents consistently fell in the “five hours or less” time frame. Unsanctioned adversaries are rumored to move within *minutes* across networks, and many cybersecurity vendors claim that attacks occur within seconds. We were curious to see if we overlaid experience on these results, does that increase or decrease our respondents’ abilities?

Following up on the concept of speed of intrusion, we asked our respondents how quickly they were able to collect and potentially exfiltrate data. Figure 12 provides these results. Within these results we see the continuation of *faster* speeds, with nearly 64% being able to operate within the 5-hour window. However, as we can see above, ethical hackers are much more confident (to the tune of over 41%) that they can collect and exfiltrate data in two hours or less. This is a shift that we expected and could have predicted—as adversaries get further along in their attacks, they often either gain speed advantages due to lack of detection or become so familiar with the environment that exfiltration is simply another step in an already-established infrastructure.

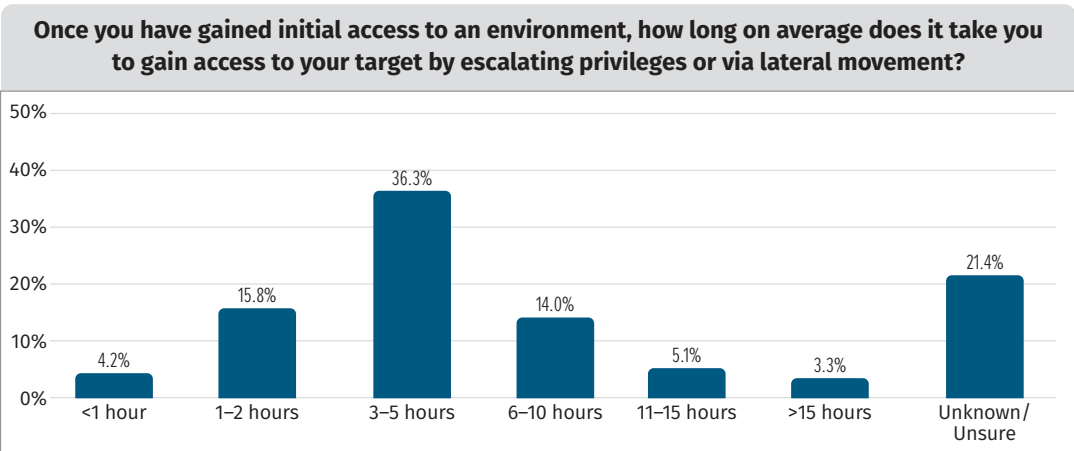


Figure 11. Time to Access Target

Key Defender Metric

We see adversaries consistently saying they are able to perform intrusion actions within a five-hour window. Whether it’s lateral movement, privilege escalation, or data exfiltration, security teams should be measuring their ability to proactively identify, and detect and respond as quickly as possible.

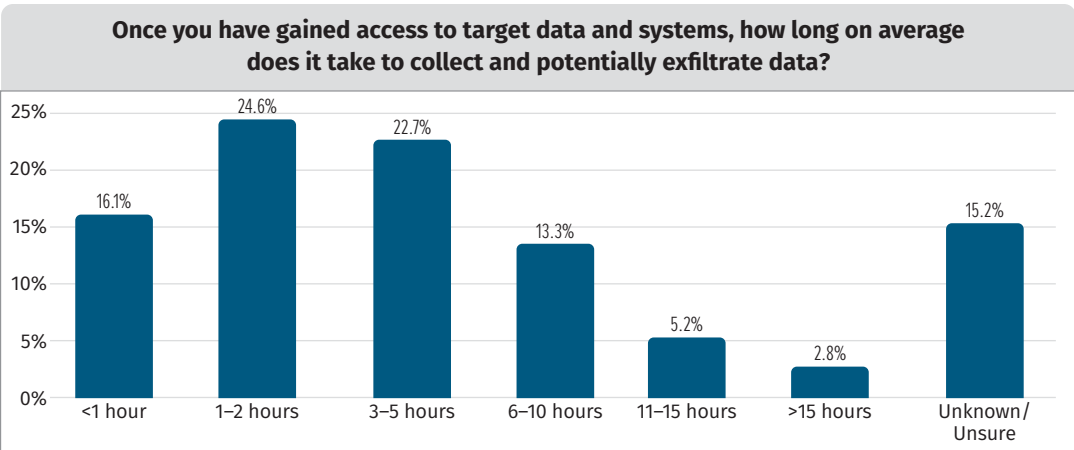


Figure 12. Time to Collect and Exfiltrate Data

The Acceleration of Speed

It is not surprising that adversaries take longer to enter a network and escalate privileges than they do to execute the final stages of an attack, such as data exfiltration. Those can be done quickly because the adversary has already established lines of communication, has access, and has identified key systems.

Finally, we asked our respondent adversaries to take a cumulative look at their intrusions and reveal how long an average end-to-end attack, incorporating all the stages above, takes. Figure 13 outlines those results.

Given the various *individual* time frames we examined above, the cumulative time frames are more dispersed. Again, we see a statistical outlier of ethical hackers who can complete an attack in less than five hours, but the results show a fairly even spread from 5–10, 11–15, 16–20, and 21–25 hours. One-fifth of the respondents said they needed 25 hours or more, but without details of those attacks (size of the network, scope of the intrusion, defense mechanisms, etc.), it’s tough to tell what types of hurdles they encountered.

It’s notable that another one-fifth of the respondents (more than 23%) responded with they do not know or are unsure how long an end-to-end attack takes them. We’ve already stressed the need to keep track of intrusion metrics, but this is a clear issue for offense and defense alike. A lack of metrics for determining how long intrusions take can create issues for benchmarks that security teams cannot rise to.

Ethical Hacking Speed Bumps

Of course, while we highlight that adversaries can move at speeds that defenders must *keep up with*, it does not mean that ethical hackers (or unsanctioned adversaries!) never run into speed bumps. In fact, it’s not uncommon for an adversary’s preferred method often is blocked or limited, meaning they must shift to other tactics. Figure 14 examines the results of this question.

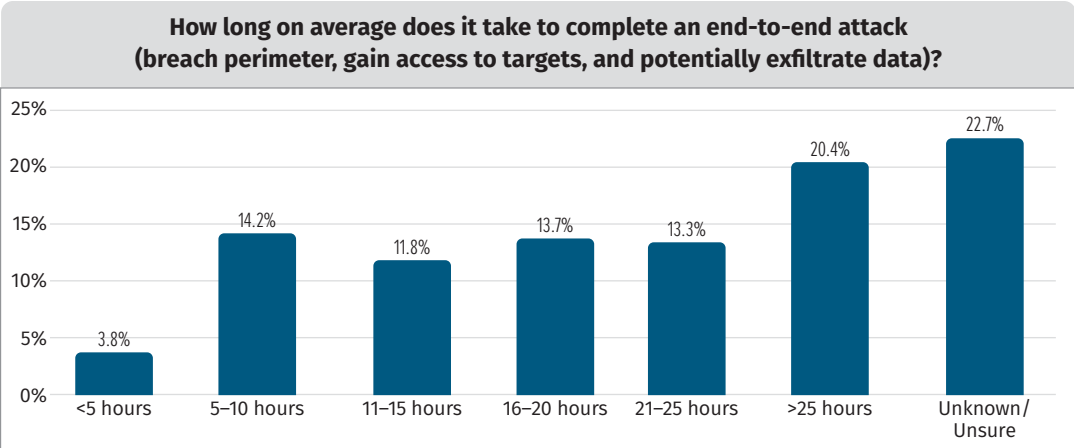


Figure 13. Time to Complete an End-to-End Attack

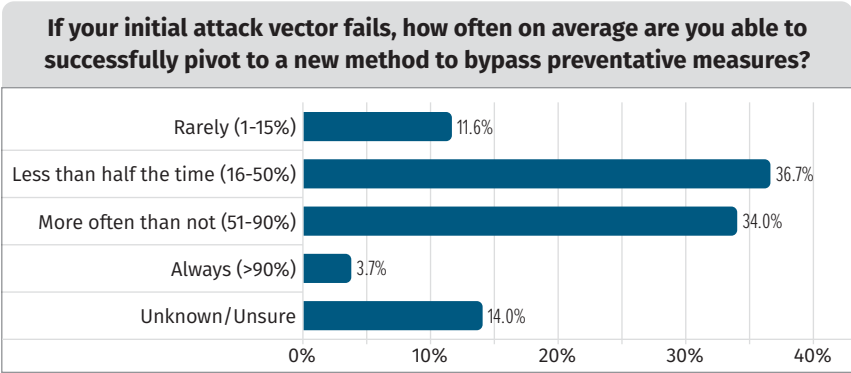


Figure 14 Frequency of Pivoting Attack Methods

Only 38% of respondents can pivot to new bypass methods more than half the time. Security teams get a welcome win when they can force adversaries to switch to unfamiliar tactics or techniques. We coupled these results with experience, to see how ethical hacking experience contributed to the ability to pivot and use new techniques. Figure 15 outlines these results.

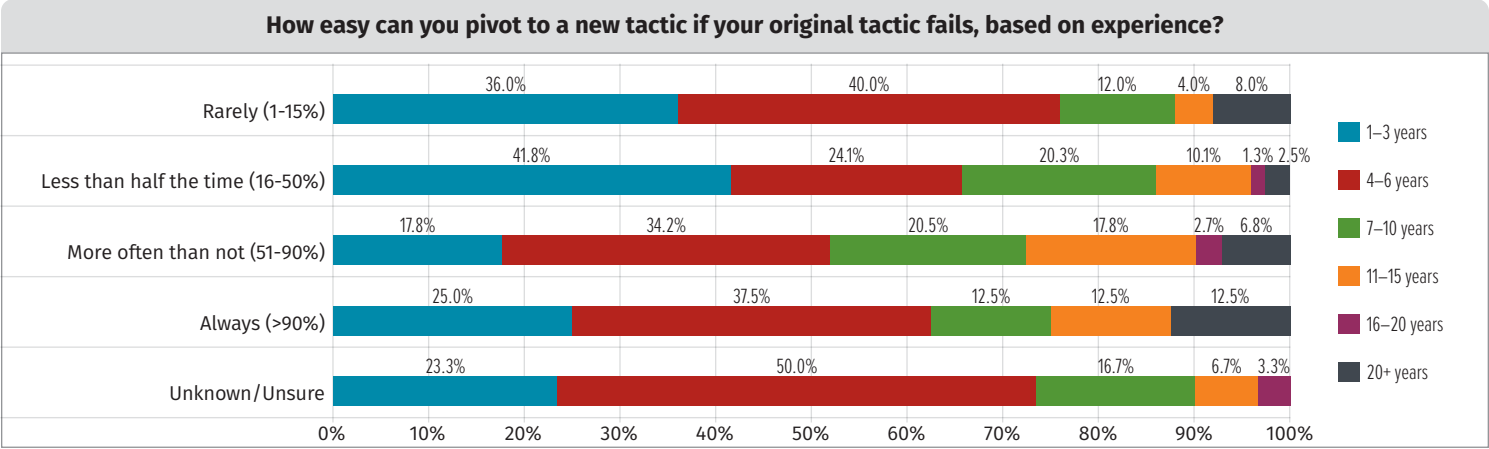


Figure 15. The Effect of Experience on Ability to Pivot

As we can see in Figure 15, experience is a contributing factor to an adversary’s ability to pivot to a new technique. Respondents with one to three years of experience were most likely to say that they are able to pivot less than half the time. Those with four or more years of experience are able to pivot more easily to bypass preventative measures, and those who said they could always pivot successfully were most likely to have four to six years of experience.

Tools, Tactics, and Techniques

An important consideration for adversaries is where they are likely to find the most success. Adversaries consider their business models just as any other business—and they will ask themselves where they will find the greatest return on investment, which technique will yield the most success, which should they pursue, and which should they abandon.

Figure 16 reveals the attack vectors that respondents said gave the highest return on investment.

As seen in Figure 16, it should come as no surprise that social engineering and phishing attacks are the top two vectors, respectively. We’ve seen this time and time again, year after year—phishing reports continually increase, and adversaries continue to find success within those vectors. The top five vectors are rounded out by web application attacks, password attacks, and ransomware. A close sixth place is Active Directory attacks, at 7.2% of respondents.

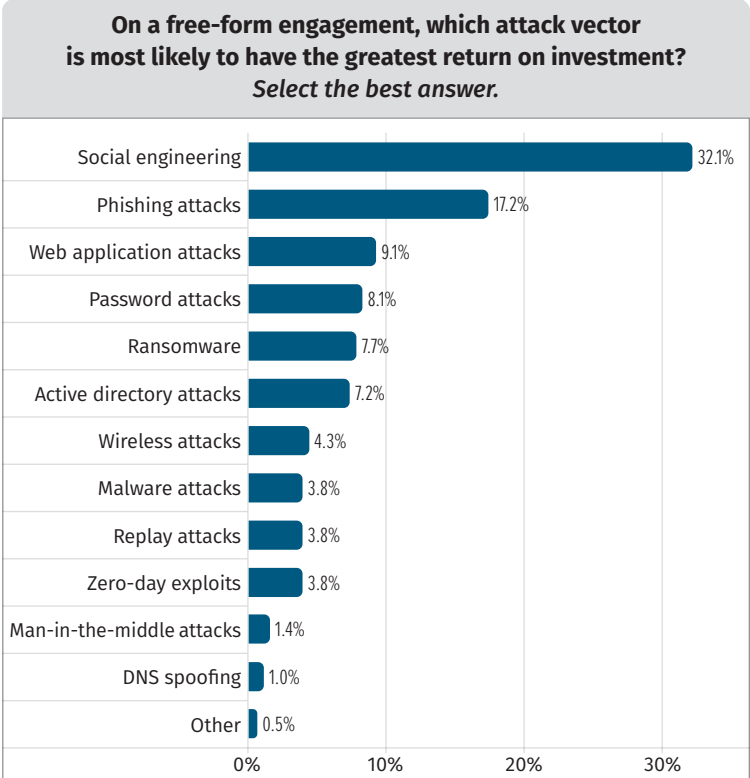


Figure 16. Attack Vectors with Greatest Return on Investment

It's worth noting that respondents had to choose one best answer for this question. However, we expect that many overlaps occur with these techniques/vectors. Ransomware intrusions (known for their monetary success) often involve Active Directory attacks and can be initiated via spearphishing. But intrusions that combine multiple tactics and techniques can get expensive either from a tooling, resource, or potential-detection perspective. Furthermore, many of the latter stages of an attack can be automated or driven by toolkits rather than by custom adversary tools/scripts.

We explored this topic further. Figure 17 looks at where respondents source their tooling.

As seen in Figure 17, it should come as no surprise that nearly 60% of respondents prefer open source tools. This is a hallmark of the hacking community (sanctioned and unsanctioned)—free and open-source tools, proof-of-concept code, and post-exploitation toolkits available to any and all. In fact, commercial tools were preferred by only 11.5% of respondents, clearly indicating that if hackers need tools, they prefer open-source.

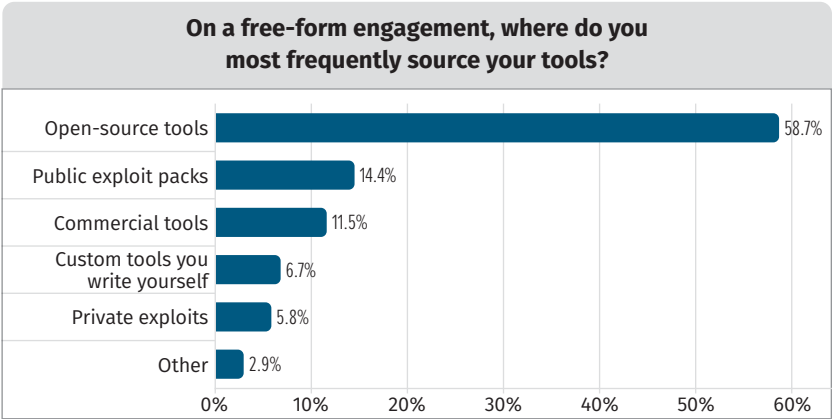


Figure 17. Tool Sourcing

Watching the Defenders

Finally, we asked respondents their opinion on how the defenders they encounter are performing. After all, their own success rates directly correlate to how *unsuccessful* the defense is. Figure 18 shows the results to our first question about adequate defenses.

The results presented in Figure 18 are astounding. Nearly three quarters of respondents indicated that organizations have only a few or some detection and response capabilities to effectively stop an attack. These are among the most revealing results from this survey. Adversaries realize that the ability to detect and respond is still significantly inadequate and use this to their advantage.

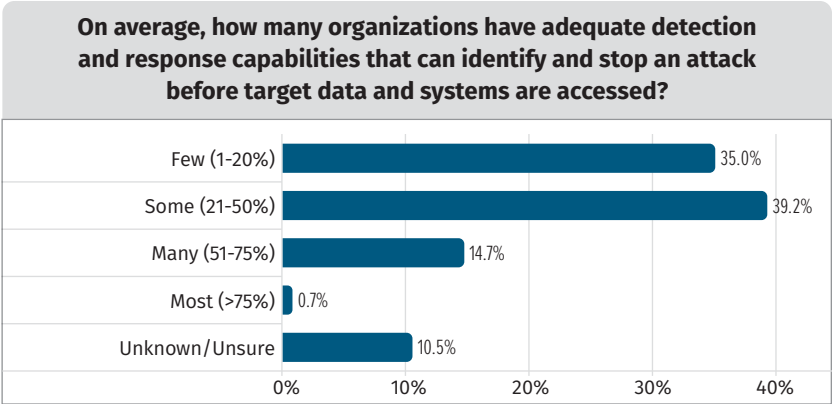


Figure 18. Defenders' Ability to Stop an Attack Before Systems Are Accessed

To see technology-specific detection and response capabilities, we drilled down deeper in our survey questions. For example, respondents were asked for their observations on the ability to prevent, detect, and/or respond to cloud- and application-specific attack techniques. Note: only respondents with specialties in cloud and application security were surveyed. Figures 19 and 20 have those results.

As seen in Figures 19 and 20, we arrive at the same place of inadequacy. Unfortunately, compared to our adversaries, organizations are ill equipped to respond to various types of attacks. One positive result was an uptick in capabilities of application-specific attacks, because we look for defense success wherever we can.

Closing Thoughts

This inaugural edition of our ethical hacking survey provides unique insights into the mind of today’s adversary. We set out to understand how attackers might approach an environment, where they find success in their attacks, and how easy it is for them to switch their tactics and techniques. We also wanted to get an adversary’s, rather than a defender’s, perspective as to whether organizations are detecting attacks or not.

Many of our surveys and whitepapers focus on a defensive perspective, often soliciting opinions from organizations defending against attacks. This survey yielded a new and welcome perspective. Hearing how adversaries had to change tactics and techniques or pivot in an environment can help organizations realize where they are making good investments and where they need to tighten up controls and policies. Remember, there are two sides to every story. Understanding how they work together can help you build resilient cyber defenses.

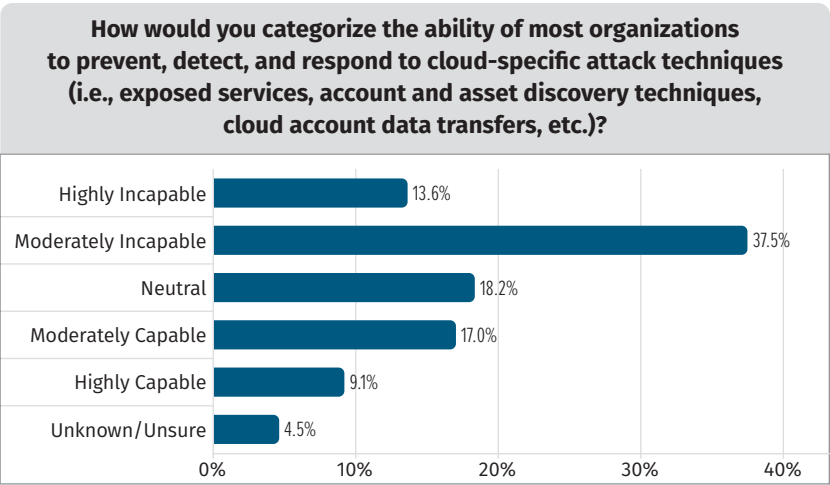


Figure 19. Ability to Prevent, Detect, and Respond to Cloud Attack Techniques

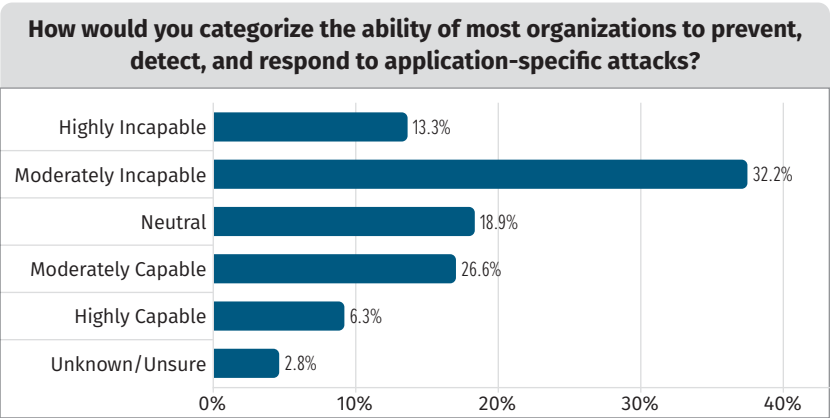


Figure 20. Ability to Prevent, Detect, and Respond to Application Attack Techniques

Sponsor

SANS would like to thank this paper’s sponsor:

