10

PROACTIVE DEFENSE TECHNIQUES

If you tell the truth, you don't have to remember anything.

—Mark Twain

Now that we've covered the fundamentals of social engineering and OSINT collection, it's time to talk about how an organization can minimize the impact of these attacks or even prevent them altogether. Although you'll rarely be able to stop all attacks, you can take steps to reduce an attack's success rate and lessen its harm if it does succeed.

This chapter covers three such techniques: awareness programs, reputation monitoring, and incident response. We'll discuss the elements of a successful awareness program, explain how to implement OSINT monitoring and technical email controls, provide integration with incident response, and finally, produce threat intelligence.

Awareness Programs

Awareness programs are company initiatives designed to provide guidance to users in situations when they encounter—or, unfortunately, fall victim to—a social engineering attack. These programs are essential because they expose users, who are probably already receiving phishing emails, to tactics that malicious attackers may use without the potential negative outcome.

One approach to conducting these trainings is to teach users about common trends in the security industry. Offering this kind of general advice is rarely enough. Hopefully, the previous chapters of this book helped you understand that traditional security guidelines—such as looking for the green padlock in the address bar of your web browser, paying attention to spelling and grammar in emails, and checking link addresses—is no longer enough to prevent phishing attacks. Sure, some attackers still make those mistakes. But the sophisticated ones capable of doing catastrophic harm to an organization are not.

A better approach is to inform users about the specific problems that the organization faces as a result of phishing. For example, if the organization is experiencing, say, an influx of Nigerian Prince emails, or an aggressive business-email-compromise campaign that spoofs the CFO, letting users know the details will better equip them to resist these attacks. Users are likely to encounter one of these specific attacks, so they should know to look out for them.

How and When to Train

Though training should occur often to keep users aware of current trends, you also shouldn't detract from their assigned duties. Offering awareness activities frequently enough for users to retain the lessons without becoming a nuisance is a delicate balance. I recommend providing training at least quarterly. Although monthly trainings provide more security, they can be cumbersome to both the users and those managing the program.

During this periodic education event, you should provide examples of phishing emails that the organization received since the last training. If your organization performed any testing, you could also distribute statistics like those discussed in Chapter 9. Most importantly, you should tell users the steps they must take if they receive a phishing email and the steps they must take if they fall victim.

When discussing the example phishing emails, point out any clues that indicate that the emails are fake. Do this from a logic, language, and technical standpoint. Draw attention to any requests that violate standard operating procedures or reason. For instance, raise the question of why the CFO on vacation in Thailand needs you to release $45 million, send it to a PayPal account, and then text a number in Belize. Point out the grammar errors, which could include missing key phrases, different spelling conventions (like *organize* in the United States instead of *organise* in other parts of the world), or using the wrong term for employees (*associate* for Walmart and *cast member* for Disney). Teach users to hover over links to see the page to which the email is trying to send them. Encourage them to forward suspicious emails to the security team. Discourage them from forwarding chain letters or responding to suspicious emails without first talking to the security team.

In one type of successful program, "Security Thought of the Month" training, the security team discusses one concept related to social engineering, or any security-relevant topic, for that matter, per session. These concepts can coincide with previous or upcoming engagements. They can also complement current events. For example, in the United States, employers provide employees with tax forms in January, allowing them to meet the April 15 tax-return deadline. Coincidentally, many successful W-2 phishing attempts occur in the early weeks of January. Similarly, September or October tend to be the best months to talk about identity theft and e-commerce, because they come right before the holiday shopping season.

Nonpunitive Policies

One of the main reasons people fail to report falling victim to phishing emails—whether it be a click, a download, or information they entered into a web form—is that they are embarrassed they did so or fear for their job. But an unreported successful phishing attempt could cause significant downtime, or if the organization fell victim to ransomware, the purchase of Bitcoin or Apple gift cards to unlock it.

Employees should know that it is acceptable to report that they've fallen victim to a social engineering attack. Although doing so may mean that they must complete additional training, they shouldn't have to update their resumes as a result. Many social engineering firms put provisions in the contracts they sign with their clients that prevent employees from being fired as a result of the testing. (To my knowledge, this language has not been tested in court.) I have not been personally party to any litigation regarding such a provision nor am I intimately aware of anyone who has. Consult with your legal counsel before attempting to enact this in any contracts.

In rare cases, employees may need to be let go because they cannot grasp the concept of security awareness. This occurs if the employee becomes a greater liability than an asset. Still, terminating an employee's contract should be the last resort. Begin by exhausting every attempt to train the employee, including going beyond the typical awareness programs. Also attempt to implement additional technical controls.

Incentives for Good Behavior

Though it's not in your interest to punish people's mistakes, it's helpful to reinforce good behavior. Once again, however, doing so properly is a delicate balance. The reason it's delicate is that, occasionally, people will try to game the system.

To provide an example of how offering incentives could go wrong, consider what happened to Wells Fargo in 2016. Between 2009 and 2015, Wells Fargo had set unrealistic sales goals for its staff. The bank later discovered that these goals had incentivized 5,300 employees to create fake Wells Fargo accounts, in some cases for family and friends, but in other cases for strangers. Management discovered this when the strangers started incurring fees.

To prevent employees from gaming the system, avoid offering incentives for reporting the *most* phishing attempts. Those incentives would encourage employees to get their emails on phishing lists, creating more work for the security team. Instead, you could incentivize reporting *clever* or *unique* phishing emails, passing all phishing simulations or reporting them, or something along these lines. The idea is to reward reporting in general, especially of clever or unique phishing attempts, rather than rewarding the largest quantity of reports. Should an organization base rewards on quantity and employees purposely get on phishing lists, eventually one may come through that looks real and the users fall victim; meanwhile, the security team is busy analyzing all the other emails they forwarded.

Here are a few free or low-cost prizes that you could offer:

- Getting 15 minutes off for reporting a unique or widespread phishing email
- A $10 Amazon or Starbucks gift card
- A parking spot for a week
- Entry into a drawing for a big prize
- Free lunch for a week

Providing anything of perceived value to employees for doing a good job will help reinforce the good behavior you seek. This is a bit of social engineering in itself, but it aims to bring about positive outcomes for employees and the organization.

Running Phishing Campaigns

Though controversial, running phishing campaigns as part of your training efforts can reliably expose your employees to realistic phishing attempts and allow you to test the organization's response as a whole.

The first decision you should make after deciding to simulate a phishing campaign is whether to conduct the engagements internally or hire a third party to do so. To choose the best option for your company, ask yourself how often you plan to run such engagements and what your budget is. If outsourced, phishing engagements may take 4 to 24 hours of billable work per engagement, depending on the SOW, scope, and complexity you desire. If you'd rather test internally, you must figure out who will conduct the engagement, what their other duties are, and what impact on your security posture their time away will present. If you have the budget to purchase a phishing simulation service from a non-consulting company such as Proofpoint, Cofense, or KnowBe4, you could take that route as well.

Reputation and OSINT Monitoring

Proactive OSINT monitoring is just as essential as proactive social engineering. *OSINT monitoring*, or the practice of periodically conducting OSINT on oneself or one's clients, is also sometimes called *brand and reputation monitoring* or *dark web monitoring*. The benefit of OSINT monitoring, in any form or flavor, is that it allows the organization to see what potential attackers can see. This allows the organization to act appropriately, ahead of an attack, whether by removing the data if possible, increasing monitoring, or implementing disinformation or deception.

Since OSINT is largely passive, you can't really run simulations to condition users in ways that prevent attackers from collecting OSINT. There are few opportunities for detection. In many cases, the OSINT may come from user accounts, and the organization can't force a user to remove something from social media, unless it interferes with intellectual property through the Digital Millennium Copyright Act (DMCA) or some other legal criterion.

Implementing a Monitoring Program

When implementing an OSINT monitoring program, focus on finding information that might pose risks to the business. Don't use this as a means to spy on or pry into employees' personal lives. One easy way to ensure that your testing remains ethical is to outsource the OSINT monitoring (discussed in "Outsourcing" on page XX).

If your organization chooses to implement its own OSINT and reputation-monitoring program, it has to decide the parameters within which to operate. In doing so, it must define what to test for, when to test, and how to test. Since employees can post anything at any time, monthly or quarterly testing is a good practice. Otherwise, many of the considerations required to set up a phishing campaign are applicable here as well. Will the testing be automated or manual? What is your budget? How in-depth is the engagement expected to be? What is the scope? How will you ensure that you respect your employees' privacy and agency to post to their social media?

Determining the amount of manual testing to conduct will drive the budget conversation. Having someone actively search for OSINT about an organization requires paying the investigator (and possibly the investigator's employer). Automated code doesn't require this, but the owner of the code may charge a fee for using the service.

Define a scope similar to the one used for social engineering engagements. This is essential to avoid violating your employees' privacy. While the organization should care about what vendors, employees, contractors, visitors, and partners post publicly, avoid looking for content shared privately or between friends. Don't force employees to connect with you on social media or try to join their friends lists by using fake accounts.

Outsourcing

From my experience, it's often best to have third parties handle the OSINT and reputation monitoring. When third parties collect OSINT on your employees, you can alleviate concerns that the organization spies on employees. It also keeps your organization's security team away from

personal accounts belonging to other employees, which reduces the chance of stalking or harassment accusations. Finally, it keeps actual abuse from occurring under the banner of security.

In addition to avoiding harassment claims, having third parties conduct OSINT monitoring allows the investigators to operate with minimal bias. They're more likely to act as a sieve rather than a pump, meaning that they filter out the extraneous information, irrelevant to security, often by using automated web scanners with little to no malicious intent, and provide the organization with only relevant information.

Incident Response

Incident response is the set of predefined actions that an organization will take if an adverse event meets criteria that classifies it as an incident. Part of being prepared is thinking through what may happen if social engineering is successful. The remainder is studying how your systems interact so that you can take steps to prevent widespread or catastrophic impact. As stated earlier in this book, the time to decide which actions to take is not during an active social engineering campaign.

The SANS Incident Response Process

The *SANS Institute*, an organization for security research and education, defines a cohesive incident response process that includes the following steps: preparation, identification, containment, eradication, recovery, and lessons learned (Figure 10-1). Also known as *PICERL*, this incident response standard takes into account the entire life cycle of an incident, beginning before it is classified and ending after it is resolved. At each stage, you define the steps necessary to minimize the impact of the attack, restore services as quickly as possible, and fix the root cause of the event.

[F10001.EPS]<<CIRCULAR ARROW DIAGRAM SHOWS STEPS OF THE SANS INCIDENT RESPONSE PROCESS: PREPARATION, IDENTIFICATION, CONTAINMENT, ERADICATION, RECOVERY, AND LESSONS LEARNED.>>

Figure 10-1 The SANS incident response process

In the *preparation* phase—which often grows out of the *lessons learned* phase that follows an incident—you anticipate future incidents by running awareness programs, performing phishing simulations, and monitoring your OSINT.

The *identification* phase begins the moment the organization becomes aware of an event that the organization classifies as an incident. In a social engineering context, the following could trigger this phase: a user self-reporting clicking on a phish; a user self-reporting that they provided information or access to a caller over the phone; an alert about ransomware from malware prevention tools; server logs that indicate spidering, unusually high access, or downloads of public files; any emails received on a honeypot email address (which serves no purpose other than to be discovered by spidering tools); suspicious email alerts from email-filtering software like Proofpoint or Mimecast; or custom alerts based on internal and external threat intelligence.

Once you've identified and classified an incident, you enter the *containment* phase, in which you take steps to ensure that the threat cannot spread further. When it comes to social engineering, containment action could include sinkholing a domain or removing an email from the email server and queue. You could also directly block a domain, IP address, or email address from being accessible from the organization's network and systems. Finally, you might isolate a computer system from the rest of the network or put it in a sandbox, force a user password reset, or even deactivate affected user accounts. Once you've taken an action, you should send a mass email to users to prevent them from falling victim.

In the *eradication* phase, you solve the problem. You remove any malware that was installed. You analyze the incident's root cause and begin to identify actions to help you recover. Unless

malware is involved, this is usually a very short phase.

In the *recovery* phase, you take any actions necessary to recover from the incident entirely. This might mean reenabling user accounts, changing network configurations to remove sandboxes or other segregations enacted as a result of malicious activity, and reverting changes made by the malicious actor.

After everything is back to normal, the *lessons learned* phase is when you analyze the root cause of the incident to determine existing gaps in your knowledge and execution. You'll then remediate these gaps as part of the preparation phase. This is the time to be introspective and decide what could have been done better.

Responding to Phishing

Now that you understand the basics of incident response, you need to define what users should do if they fall victim to the various kinds of social engineering attacks. Let's start with phishing. A phishing email that contains links or files might allow an attacker to gain access to your systems. Your goal, then, should be to expedite containment and eradication.

One tip for making rapid responses possible is to choose a single nonblack color for all networking cables. This will allow you to instruct employees to unplug that colored cable from the back of the computer or the wall. Bear in mind that systems may still be connected to the network wirelessly, and you should define behavior for disconnecting wireless devices as well.

Ask users to report the approximate time at which the incident occurred. This detail helps the security team locate the logs they should comb through, rather than leaving them with no clues as to where to look. Upon falling victim to the attack, users should either log out, lock their screen, power off, unplug from the network, or hibernate their system. The actions that a user should take depend on the organization's capabilities and its potential response to a given incident. For example, if no forensic analysis will occur, there is no reason to hibernate the system. Alternatively, if the organization is going to rebuild the system from known, good media, the correct action may be to just shut down the system after gathering the artifacts.

The user should also report the source email address or website of the phish, any windows and applications that were open, and whether anything unusual happened on the screen.

You should print out these guidelines, laminate them, and put them in each user's physical workspace. At any given time, a user should be able to reference the sheet and perform the required actions.

Responding to Vishing

While similar to phishing, vishing presents unique challenges. In the absence of monitoring all phone calls, the ability to identify and act upon vishing calls relies upon the staff reporting them, as well as knowing the actions to take in advance. No widespread or accurate intrusion detection systems (IDSs) or SEIMs encompass phone calls. Companies can monitor the internet traffic of any phones connected to the corporate Wi-Fi network, but not ordinary phone calls or their context. Fortunately, an attacker cannot (immediately) log in and take over a network from a phone call. Even if someone who is vishing gets information, technical controls might prevent them from causing further harm. Regardless, you must define actions for responding to vishing attempts.

First, upon noticing that they are experiencing a vishing attack, users should either hang up, ask to call back, try to get information out of the caller, or lie to confuse the caller. Security management within the organization will need decide which actions to train employees to perform. A level of risk is associated with encouraging users to solicit information or lie; the user may give up valuable and accurate information inadvertently. Users should then contact information security and provide the approximate time at which the incident occurred, any actions they took, the phone

number that called them, the information they were asked, and the information they provided. They should also provide any additional information about the caller, such as their accent, dialect, tone, or mood, or the presence of background noises.

Responding to OSINT Collection

Detecting OSINT collection is hard, because platforms like Shodan, Censys, and Have I Been Pwned don't automatically allow you to set alerts on queries, though you could write your own code to set alerts for whenever the organization's assets appear on such platforms. Have I Been Pwned allows organizations that can demonstrate ownership of a domain to set up such alerts, but will not share any breached credentials belonging to accounts on the domain. But since OSINT collection typically takes place in the reconnaissance phase of an ethical hacking engagement, it's common to see it occur alongside scanning and enumeration, which are detectable.

The first layer of detection lies within a CDN like Cloudflare or Amazon CloudFront, if used. The next layer is within the web server logs or the application logs of the web applications. These sources will educate the organization as to who is scanning and what is being scanned. Often this will lack the context needed to differentiate between web scanning en masse and an actual adversary attempting to collect OSINT or working through scanning and enumeration.

Decide what actions you should block. Examples can include blocking users after a certain number of 404 errors caused by spidering; blocking or rate-limiting spidering to a certain number of events per second; blocking anyone who downloads a certain number of public files, using a specific user-agent string in the browser or script; and blocking users who navigate to a honey page.

Handling Media Attention

Depending on the severity of the attack, the publicity it receives, and other events happening in the news cycle, the media may seek to speak with people from your organization during an incident. While doing so should not be your top priority, failing to respond to media inquiries can send a worse message than if you just admit that you don't know all of the details at this time. While the media should reach out to the organization's public relations team, some may attempt to contact and interview any employee.

To control the message being conveyed to the public, enforce a media blackout on all employees except those defined in your incident response plan. Provide unauthorized employees with a template response for handling such inquiries. This can be as simple a statement as "I am not authorized to discuss the details of the subject of your request" or a redirect to the designated media representative at the company.

The parties who are authorized to talk to the media must understand what tone to take, how to decline to answer, and whom to speak with in order to learn the facts they will share with journalists. Also define a person or committee to review and approve any messages that the public relations person will provide the media.

I also highly recommend consulting with your organization's internal PR team and any external PR consultants that your organization uses. They will be able to speak to your organization's specific policies and procedures, whereas I am speaking in more general terms.

How Users Should Report Incidents

If you don't tell users how they should report suspected incidents, they may bombard a gate guard who has no means of resolving the problem. One strategy is setting up a phishing@organization.com email address to collect the phishing emails that users receive. You might also set up cyber@organization.com or incidents@organization.com as catchall addresses that forward emails to the appropriate parties.

When a user has fallen victim to a phish and may have introduced malware to the environment, reporting via email may not be the best approach. That's because it's possible that the entire email system has been compromised and that attackers can read, block, or alter the message. Depending on the size of the organization and whether a user is onsite or remote, you could direct users to report the incident face-to-face, call a phone number, or send a message using private chat or a secure texting platform like Signal, Wickr, or Wire.

Technical Controls and Containment

When it identifies a phishing email, the security team should collect the email, and any relevant information about it, while adhering to organizational policies and procedures. This will pay dividends when producing threat intelligence, which may be necessary depending on which industry the organization is in and whether it belongs to any Information Sharing and Analysis Centers (ISACs).

From the email itself, you should collect the source email address, details about whether the email address was spoofed (which is discussed in Chapter 12), and the source IP address. You should also block the source addresses and check logs to see if any other users or hosts communicated with the source addresses. To analyze the attack further, use the tools discussed in Chapter 12.

If the phishing scheme involves a malicious file, you can upload it to the malware-detecting website VirusTotal to get a quick answer about the file's content, assuming it is known malware. Additionally, take a cryptographic hash of the file, and check systems on the network for files that produce the same hash. Tune the SEIM to alert you about incoming instances of this file as well.

Sinkhole any IP addresses or URLs associated with the phishing email to redirect users to a benign page and set up alerts on workstations for those who attempt to access the malicious site. Once the email is sinkholed, the security team can contact all users, advising them to avoid the email. Once the incident is in the recovery phase, transition the information collected to threat intelligence and follow the organization's guidance on publication and distribution.

Conclusion

Part of keeping your organization safe is keeping users informed, aware, and alert. By applying a nonpunitive policy coupled with incentives for positive behavior to reinforce desired actions, you can drastically improve the security posture of your organization while empowering employees to make good decisions. Once users are trained, know what to look for, and understand what to do when they fall victim to social engineering attacks, it is time to engage them by using internal or external testers to measure their adherence to the organization's guidance.

Even after you train users, it's still necessary to test them via phishing simulations and OSINT monitoring. People sometimes want to publicly share occasions like promotions, their last day on the job, or their first day on the job—and as discussed in Chapter 5, they don't think about other information they're including in the frame of their photos. Similarly, people want to share their work experience so that they can demonstrate their competencies to hiring managers, but this makes the technical details listed in their resumes searchable.

Integrating your training with your incident response process is a crucial aspect of defending against social engineering. Awareness programs help the organization avoid invoking incident response in the first place, but they should also help the incident response process work more smoothly in instances where users fall victim.

Explicitly telling users what to do if they find themselves the victim of social engineering also dramatically improves the organization's security posture. If nontechnical people are left to their own devices, asked to perform technical actions without direction, it's likely that they'll ignore the

problem or try to cover it up. Defining steps for users and the security team to take when something goes wrong will save the organization a lot of headaches and allow it to focus on restoration.