CHAPTER 5

Containing a Breach

Chris Pogue

INTRODUCTION

OK, so … now what? At this phase of the Breach, you are more than likely past the worst of it. I say, "more than likely" due to the fact that many Breach investigators make the mistake of myopically focusing on the Breach that they were brought in to investigate, discounting entirely that this particular compromise may be one of many. There may even be multiple attackers that have exploited multiple vulnerabilities to gain access to the target systems. In these instances, the respective investigation and remediation becomes exponentially more complex. Experience has afforded me the wisdom to confidently state that things are almost always much worse that they initially appear; so don't jump to conclusions too quickly. Research shows that it takes the average organization over 200 days to detect a Breach. That's a long window of time and considering organizations can be Breached multiple times within a close timeframe (I have seen this several times in my career), you could be brought investigating one Breach not knowing there's another one yet to be discovered that occurred months prior. An analogy I will draw is taking your car to the mechanic, paying a large bill to fix a problem with your engine, and a week later realizing something is wrong with your car again, you take it to the mechanic and they inform you something else has gone wrong in the engine, unrelated to what you fixed prior, and that it should be fixed. Sure it could be something completely unrelated but try explaining that to the angry customer. When dealing with Breaches the statements made by the victim organization will be scrutinized ad nauseam in the court of public opinion, by industry and government regulators, and potentially by opposing counsel. As their trusted advisor your guidance needs to be candid, technically sound, underpinned by the wisdom of your experience. The impact of a data Breach is staggering on many levels that include the loss of customer confidence, loss of market share, loss of company valuation; and are tremendously expensive as the victim will dole out potentially tens or hundreds of

CONTENTS

Data Breach Preparation and Response. http://dx.doi.org/10.1016/B978-0-12-803451-4.00005-8

millions of dollars in violations and fines, legal fees, and forensics fees. As the client marshals their collective efforts recovering from one incident, imagine the initial impact amplified exponentially should they be compelled to announce that "it happened again." However, presuming that you don't fall into that category, your organization or your client now needs to take steps to stop the bleeding.

The time frame from the initial Breach to containment of the Breach, is called the *Window of Compromise*. Now let's be careful with our terminology here, because this is important to understand. The *Window of Compromise* tends to get confused with the *Window of Intrusion*, which is the time frame from the initial Breach to final eradication of the attacker beach head. The important take away here is that you can contain a Breach while the attackers or malware are still present. Think back to the final stage of the Breach Breakdown introduced in Chapter 4—Exfiltration. If the bad guys lose communication with the target, and no longer have the capability to move data from the victim system(s) to their systems, the Breach is contained. Fig. 5.1 illustrates the window of intrusion versus the window of compromise.

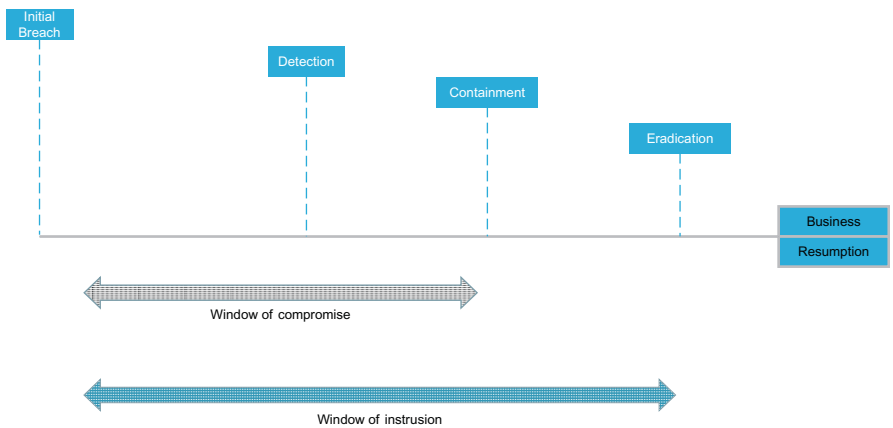

FIGURE 5.1 Illustration of the window of compromise and window of intrusion.

In our world of social media, security blogs, and "up to the minute news," the victim's ability to contain the security Breach is big news, and will travel fast. To say that you'd better be right is something of an understatement. I have said many times that during the *fog of compromise* and the push to make a statement, a large percentage of organizations have the propensity of overstating the complexity of the attack, while understating the impact. Pick an incident that has taken place over the past couple of years and read some of the media

coverage in chronological order for an illustration of precisely what I am talking about. The lesson learned from these misstatements is that any statement made by the victim can be syndicated to a global audience, and will, not may, end up in the hands of opposing counsel should the incident ever go to litigation, which it probably will. It is important to balance accuracy and speed when it comes to this type of communication which we discuss in further detail within Chapter 7. There is no room for error and there is no unringing the bell.

Finally, and arguably most importantly, the Breach actually has to be contained. Regardless of what you call it, or when it's communicated, it has to be done, and it has to be done right. There are two important aspects to this notion of "done right" that are important to introduce, implementation and validation. Very simply put, you have to fix what's broken, and make sure the fixes have their intended impact (these are frequently called *countermeasures*). I have witnessed many organizations either totally fail to remediate the vulnerabilities that the investigation identified to have been the cause of the Breach (remember, there may be multiple Breaches with multiple vulnerabilities), or implement a flawed plan with a road map that focuses on the wrong vulnerabilities (not the ones associated with the Breach or Breaches), in the wrong order (having no order of criticality), and at a later time that never arrives (daily business gets in the way). There are dozens of example of organizations that have been Breached multiple times over a one or two year period via the same vulnerabilities. Do not let your company or your customer be "that organization."

Now, assuming that you are, in fact, not "that organization," and the fixes or countermeasures after a Breach were deployed as planned, they have to be tested to make sure they are having their intended impact. This is a critical step that many victims fail to take, thereby providing them with a false sense of security. An external pentest team (yes, external—lots of reasons why, which will be covered later) should be engaged to try to exploit or circumvent the countermeasures, as well as help to identify any other potential attack vectors that may be present, and have been used in another Breach yet to be discovered.

Getting Breached is increasingly and sadly becoming commonplace. It's understood for the most part that defending every possible attack vector to include human beings, all day, every day, without ever missing a single one, is a bit of an unrealistic expectation. Therefore and rightfully so, there is a heightened emphasis on how an organization responds, and how quickly they can contain the Breach. Even if the Breach was not detected until 200+ days after the intrusion initially occurred, there will be a significant focused place on the speed of

response, ie…once you knew about it, how quickly did you act to stop the bleeding? Success here can very literally save the business millions of dollars; failure can mean closing the doors for good.

BREACH CONTAINMENT

The concept of the *Window of Compromise* was first introduced in the "Introduction" section to this chapter and will serve as the primary focus of the remainder of this chapter. Again, this window spans the time from the initial Breach, when a bad guy gains access to the target environment, to the victim containing the Breach and preventing further external communication with the attackers. To visualize this concept, think of a bank robber, if the measure of a successful bank robbery is getting away with the money and the bank robber gets trapped in the bank, he fails. Same concept here—if the attackers and any other malware or the malicious tools they brought with them are present yet their mechanism of communicating with those tools and/or exfiltrating data is disrupted or terminated, then the Breach has been successfully contained. It is important that you understand this concept and how it differs from the *Window of Intrusion*.

While the *Window of Compromise* can be achieved while the attackers and their malicious tools are still present, the *Window of Intrusion* cannot. The Window of Intrusion is the time frame from the initial Breach to complete eradication of the attacker. Understand that "closing the window" may (and in all likelihood will) take the form of a temporary fix that focuses on the prevention of further external communication and will not provide a long term solution. To completely remove this infiltration vector the associated vulnerabilities need to be identified and remediated, and the countermeasures tested to ensure that the attackers will not simply reenter the target using this same vector. At the conclusion of this window, the bad guys are gone, their tools are gone, the malware has been removed, and you are ready to resume business. These are two different windows with two different criteria which communicate two different things. It is important for you as the investigator to understand this, and be able to adequately explain it to a nontechnical audience. The example of the bank robbery is my favorite, and one I have used over and over in offices, Board room, and court rooms many times.

1. *Window of compromise*
   a. Bank robber has broken into the bank
   b. Stolen the money
   c. Failed to make a getaway

2. *Window of intrusion*
   a. Bank robber has broken into the bank
   b. Stolen the money
   c. Arrested by the police, he and his entire set of bank robbery tools have been hauled away, and possibly (under the best of circumstances) some or all of the money has been recovered

Now that you hopefully have a better understanding of the various windows present during this stage of the Breach, it's important to understand what "containment" does and doesn't mean.

What Are You Containing?

For an attacker to have gained access into the target environment he had to have taken advantage of a vulnerability or misconfiguration that provided him with that access. Potential attack vectors can include an unpatched server, a vulnerable plugin on a website, an improperly configured firewall, a default password, or human error. Whatever the case may be, something let the bad guys in and has been identified during the course of the investigation. There is where the focus of the initial containment steps should be.

It is a good practice to map out what you believe to be the Breach Breakdown in some sort of visual manner so that you can more clearly define your working hypothesis. You should also include a timeline of events that represents the chronological progression of the attack. This will be of particular interest to executives and general counsel as they prepare statements regarding what happened and when. In addition, you should also maintain a partner list of the impacted systems represented in the diagram. This list should include additional system details such as IP address, hostname, OS, system function (ie, webserver, database, workstation), and method of compromise. This diagram should depict which system the attacker initially used to gain access to the target environment, the systems that were used as he moved from the point of entry to the ultimate location of the targeted data, and the systems involved in harvesting and exfiltrating that data. In some instances, this process will be very short, as the Breach only involved a small number of systems, while in others you may have very large numbers. This is not entirely unlike Peter Chen's Entity Relationship Model,[1] or the string models used on police television shows. Whatever the case, you will benefit greatly from maintaining this diagram and partnering list as it provides a mechanism for you to track which systems were involved in the incident, and how. Trust me, while this may sound somewhat banal, it works (Fig. 5.2).

[1] https://en.wikipedia.org/wiki/Entity%E2%80%93relationship_model

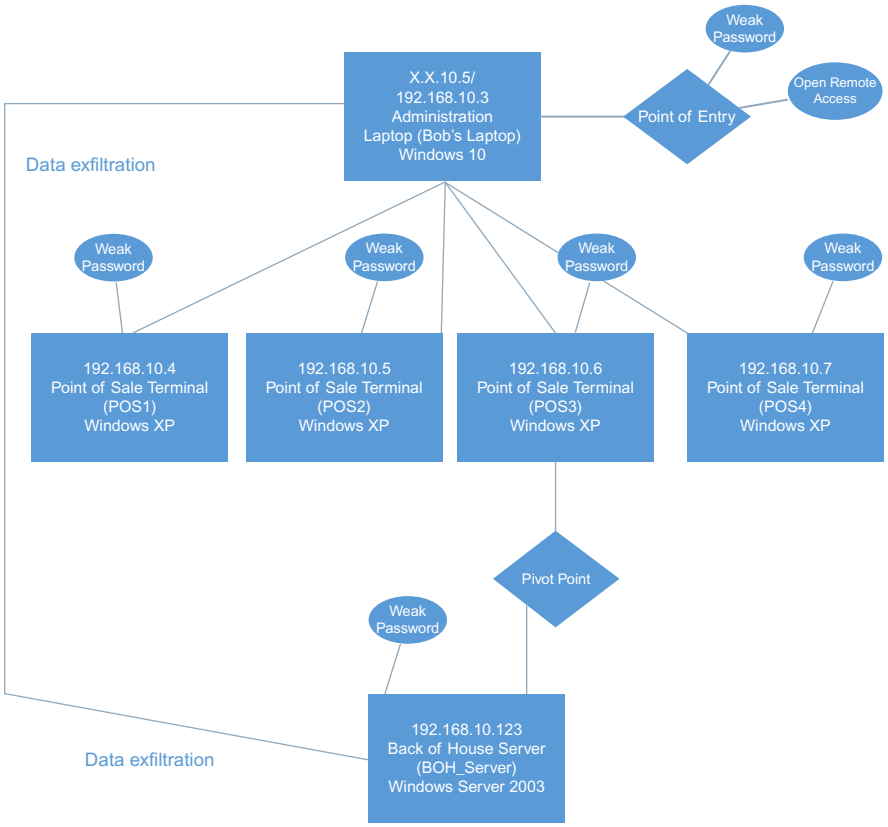

FIGURE 5.2 Sample Breach breakdown.

To effectively cut off the attacker's infiltration and exfiltration vectors (how the bad guys got in, and how they either got data out, or are maintaining communications with the target) it is critical that you fully understand how the Breach took place (hence the diagram). Many organizations that have been Breached are so overwhelmed by the gravity of what just happened that they panic (understandably so) and either lose focus or fail to gain it entirely. Their actions are erratic and disjointed instead of being coordinated and tactical. Senior management or the executives would want to fix every possible vulnerability and attack vector under the assumption that doing so will help them to save face with the Board, the customers, and the court of public opinion (movement for the sake of moving). However, the reality of the situation is that by not taking the time to formulate a logical response strategy based on the nature of the vulnerabilities that were present during the attack, they themselves are becoming the primary obstacle that will prevent them from achieving the very thing they are trying to accomplish. It's also important to mention if the Breach

ends up in litigation, they will want to be able to establish a defensible position of reasonableness; the foundation of which will be predicated on how they both planned to respond, and how they actually responded to the Breach. This is completely normal and to be expected especially considering the current post-Breach, litigation-infested landscape. However, this sort of "knee jerk" reaction creates the tendency to address symptoms rather than the root cause of the issues. It's haphazard, ineffective and should be strongly discouraged inasmuch as you are able to influence your organization. This is where being a seasoned investigator can be of tremendous value, having likely watched this play out literally hundreds or even thousands of times. You need to be that voice of reason and confidence. Remind them to take a deep breath, calm down, and proceed with forethought and logic.

The message that needs to be communicated is that while there may be multiple issues that have been identified, *only a very small subset (usually just one or two) needs to be addressed immediately in order to contain the Breach*. The focus of remediation efforts should be on these specific vulnerabilities, nothing more. That's not to say that fixing the other identified issues is not important; quite the contrary—it is vital that all of the vulnerabilities get addressed or they run the risk of being right back in the same mess in a couple of months. However, under the present circumstances, a more myopic approach is what's needed to effectively contain the incident. Additional vulnerabilities that were not part of the Breach should be noted, triaged based on criticality, prioritized, and put on a roadmap to be addressed later (so long as they are actually addressed). Many organizations actually do this during a Breach, and while I'm sure they have every intention of following through later, they become distracted by some other business driver and never actually complete the process. These are the ones that end up in the news several months later announcing that they have been Breached again.

Remediating Your Exposures

Once a thorough understanding of the components of the Breach Breakdown, the systems that were involved, and the vulnerabilities that were exploited have been established, remediation steps can begin. These steps should be prioritized based on their positioning within the network, the criticality of the vulnerability, and the complexity of the fix. Externally facing systems should be addressed first since they have the highest likelihood of being compromised again, followed by internal systems in and around the location of the targeted data. Using the diagram you presumably created that includes all of the systems involved in the Breach is a great mechanism for tracking system vulnerabilities, and determining the order of remediation. This is also true for deploying new countermeasures such as firewalls, encryption, or user-based access control mechanisms.

There is a common misconception that once the vulnerabilities on the affected systems have been addressed, they are now "safe" from attackers. Well, how do you know that the fixes and countermeasures that have been deployed are having the desired impact? Short answer is, "you don't." In many of the cases I have worked when I bring this exact point up, I get the "deer in the headlights" look. So, I happily (it's sort of fun at this point in my career) repeat my question, this time a bit more slowly, emphasizing every few words, "How do you know, that the steps that you have taken to secure the affected systems are having the desired impact?". The overwhelming majority of the time, the answer has been, "We don't."

Imagine the logic in that? There is a significant data Breach that is going to be expensive to investigate, will have a presently unquantifiable negative impact on the brand and company valuation to include stock price, and may very well end up being litigated for the next few years. You figure out how the Breach took place and focus all available resources on remediating the vulnerabilities that were exploited by the attackers. But, you don't see the need to get an external team of experts to VALIDATE that those fixes are actually doing what you think they are going to do. Why? I said they would work, and they will! That's good enough, right? Yet as crazy as that sounds, it is very much a reality in organizations all over the world.

Funny how in so many other aspects of life this is simply assumed but in computer security it's like a pink fluffy unicorn dancing on a rainbow. Would you let your plumber fix a water leak without testing it to make sure the pipe is not still leaking? Would you expect your mechanic to fix the air conditioning on your car without turning it on to make sure it's blowing cold air? Would you want your doctor remove a cast from a broken bone without taking an X-ray to determine if the bone had healed properly? No, no, and no. So why in the world, after suffering an expensive, damaging, data Breach would you not expect to test the remediation steps to make sure they are functioning properly? Hint: you shouldn't.

For this process, I recommend retaining external penetration test or "Red Team" services, which we discuss further in Chapter 8, to ensure that the specific vulnerabilities exploited by the attackers have been remediated, and to confirm that any countermeasures that have been deployed are having their intended impact. There are a few reasons for my recommendation to use an external team, rather than internal resources. One, they are experts in identifying and exploiting system, configuration, and application weaknesses. They will look at your systems from the *eyes of an attacker* and provide you with a candid view of your security posture; something you may not be willing or able to do. Two, they are not beholden to anyone within your organization and can therefore remain unbiased. Political pressures will exist from the executives, or IT manager (likely both) to provide a "clean bill of health" so that business can

resume. These pressures can also surface from individuals within the organization whose responsibility was to maintain the security of the impacted systems, who may very well be in jeopardy of losing that job. These pressures and the desire to get back to normal operations can lead to a premature or imprecise decision making that could very well do more harm to the organization than good.

An external tester that is not beholden to anyone within the organization is free (for the most part albeit not entirely) from these pressures. Three, they can also help to identify vulnerabilities that you may or may not have known about and help prioritize them based on their exploitability. Not all vulnerabilities are exploitable in the context of the current environment and its security controls. Understanding the impact of known vulnerabilities can help direct your remediation priorities. In addition, the likelihood of exploitability, due to its complexity, knowledge requirements, or the vector of attack also plays into this prioritization.

Are There More of You?

In many cases, malware and attacker tools play a significant role in a data Breach. Once initial access by the attacker has been achieved, these tools are utilized for everything from reconnaissance, privilege escalation, and lateral movement, to data harvesting and exfiltration. The good news about the presence of these utilities (if there has to be a silver lining) is that most of them have a signature and leave evidence of their existence or execution. There are some more advanced malware packages that live in memory and attack techniques that literally leave zero trace. In those cases, what is being outlined in this section would be of diminished value.

The hackers know they will probably at some point lose contact using their primary communication method with the installed malware. As such, there are typically at least one or more alternate "backup" command, control and communications mechanisms present. Some of the less technically advanced mechanisms can be as simple as installing a secondary remote access application (such as Bomgar, Logmein, VNC, or pcAnywhere), or they can be as advanced as having "phone home" activation triggers due to inactivity of the primary method. I mention this as Breach containment can't focus on the one tool, as many may exist. True containment and remediation can't occur until all potential infiltration and exfiltration vectors are analyzed and understood.

At one time, taking an MD5 hash of a binary and searching for a match within a corpus of evidence such as a forensic image was considered "advanced analysis." However, as forensic methodologies, technologies, and utilities have evolved in an effort to keep pace with attack vectors, this sort of activity now

falls into the category of "basic analysis." Today, we have other hashing algorithms that unlike simple hash comparisons that provide a strictly binary conclusion (the thing is either a 100% match or a 0% match), provide a percentage to which two files are similar; this is known as *Content Piecewise Hashing* or *Fuzzy Hashing*.[2] Using Jessie Kornblum's SSDEEP utility, you can compute fuzzy hash values for files, set a target percentage (eg, show all files that are 80% similar), and search other systems for potential matches. This is obviously exponentially more effective when searching for files that may not be an exact match for the known sample file.

Conversely, many types of malware are modified using a runtime compression program, yielding a file that is known as being "packed". Content piecewise hashing focuses on common sections in executable binaries such as malware (in this type of situation). Malware that has been "repacked" (reencrypted or repacked using the same or other packer) will of course change the sample file. While this is still an extremely valuable investigative tool, like any other tool in your toolbox, don't rely on it 100%.

There are also several mechanisms to track indicators of compromise such as OpenIOC,[3] STIX/TAXII,[4] CybOX,[5] and CRTIS[6] which we look at further in Chapter 8. The important thing to remember here is that you understand that the incidents have the potential, and even the likelihood, of being larger than they initially appear. Be flexible with your working hypothesis and make sure that the evidence remains the primary driver rather than the other way around. *Many investigators get into the bad habit of allowing their theory drive what evidence they choose to include and exclude any evidence that does not fit their theory*. Conducting a comprehensive investigation is not about being right or wrong the first time. It's perfectly normal to adjust your working hypothesis multiple times prior to completing the investigation. Your job is to be thorough and tell the full story of the Breach, so check your ego at the door.

REMOVING POSTED INFORMATION FROM THE INTERNET

In some Breach investigations, you may run into sensitive information that has been posted to the Internet on sites such as Pastebin and Github, as well as any number of Dark Web sites.[7] When this happens, the executive staff and legal

[2] http://ssdeep.sourceforge.net/

[3] http://www.openioc.org/

[4] https://stixproject.github.io/

[5] https://cyboxproject.github.io/

[6] https://crits.github.io/

[7] https://github.com/

counsel of the impacted organization need to discuss how they want to handle the situation. There will always be an impact to the brand that will have to be dealt with either through internal or external communication. For this reason, it is a good idea to retain a crisis communication firm that is experienced in handing these types of issues (see Chapter 7 for guidance on this). An improper response here, seeming aloof, or providing artificial value to the posted data (either too serious or not serious enough) can also have a negative impact on customer perception. There are literally scores of examples of Breached companies that have made statements regarding the status of an incident or the importance of the data that was compromised only to have to retract those statements at a later time. So, needless to say, allowing internal resources who may or may not have experience in such matters to handle them on their own, may not be the best option. Misstatements, or in some cases outright falsehoods about the incident, can absolutely crush executive, Board, customer, and shareholder confidence and should be avoided as much as possible··· and by "as much as possible," I mean never. Ever.

If the decision is made to attempt to have the posted data taken down, it is critical to engage both the appropriate law enforcement body and legal counsel. For the same reasons an organization would want to retain outside crisis management experts, Breached organizations should also look to retain outside counsel who specialize in data Breach litigation. While inside counsel may be tempted to handle these sorts of problems on their own, the likelihood is that they lack the knowledge and experience to handle the situation efficiently and effectively. There are many law firms that specialize in post-Breach activities that know what steps need to be taken to remove data from such public forums. Under such circumstances, these types of firms should be engaged early. Doing so will give the victim the greatest potential for success in removing their data. Similarly, law enforcement, based on the geographic region of the Breach as well as the geographic location of the site containing the posted data, may have the ability to demand the data be removed pursuant to criminal activity, or based on an ongoing criminal investigation. That being said if you are successful in removing posted information from the Internet, the success may be short-lived. The criminal who posted the information may learn the data has been removed and can repost it on the site or to an alternative site requiring your organization, legal team, and law enforcement to repeat the process. This in essence can resemble a game of whack-a-mole and continue until the good guys or the bad guys grow tired of the back and forth and stand down.

The key takeaway here is to engage the experts early on in the process. Their knowledge, like your own knowledge of incident response, is critical in positioning the victim organization in the best possible manner so that they can withstand the aftermath of the Breach. If left to their own internal capabilities

they run the risk of making mistakes that would otherwise be avoidable; mistakes that could prove to be extremely costly, and have a potentially devastating impact on the brand.

CONTAINING COMPROMISED SYSTEMS

There may be times, particularly in the instances where law enforcement is involved in the incident investigation, when the request is made to leave impacted systems alone so that attacker activity can be monitored. Requests of this nature are normally made by law enforcement agencies when the Breach may be part of an ongoing investigation involving other Breached entities. When this happens it is important that you engage the executive staff and legal counsel immediately. The business may have conflicting priorities with law enforcement that would prohibit such a request from being granted. This may include, but is certainly not limited to, Breach disclosure notification legislation, contractual obligations, industry specific compliance standards, international business requirements, and cyber security insurance policies. As badly as you may want to assist law enforcement, do not agree or commit to anything. This is decision that needs to be made by the executive staff and the attorneys, with the full understanding that under a worst case scenario, critical data that was not initially compromised becomes compromised while cooperating with law enforcement.

In other instances, you may be presented with the request to immediately contain the systems that are known to be compromised. There are four primary approaches in containing compromised systems, shutting them down, removing them from the network, patching them, or rebuilding them.

Shutting Systems Down

You will be involved in cases in which the victim may be compelled to shut down the systems affected by the Breach. Some examples of this include the theft of payment card information, the compromise and defacement of a webpage or the presence of malware with command and control (C2) capabilities on internal systems. Requests of this nature should also be fielded by the executive team and legal counsel as there will undoubtedly be significant business ramifications that will have to be carefully considered. Services being off-line could violate customer service level agreements (SLAs) or other contractual obligations, could halt the sale of goods or services resulting in significant revenue loss, or prevent customer payments from being processed.

In every Breach investigation you will face a myriad of complexities that will require input from a wide range of stakeholders. This underscores the point we have made multiple times within this book that incident response is no

longer strictly an "IT problem"; it is a business problem and needs to be recognized as such. The presence of data Breach legislation, the increased popularity of post-Breach litigation, cybercrime investigations, media attention, and cyber liability insurance have all added to the complexity of incident response. The seemingly simple decision to shut systems down, or leave them up can have a far reaching impact on multiple facets of the business. These decisions should not be taken lightly or made arbitrarily; rather they should be swiftly escalated to the appropriate parties so that a decision can be made that will hold the greatest value. It's important to note that this is not a "good choice" versus "bad choice" thing. In many if not most cases it will be more a matter of which decision will suck less.

Removing Systems From the Network

This option requires the least amount of technical capability—remove network cable—but also poses the greatest impact on the environment, since whatever function the system is performing is no longer going to be present.

Removing the system also can immediately prevent the incident from getting worse either by preventing the propagation of the attack, or abruptly halting exfiltration. This is a common decision made by smaller brick and mortar merchants whose cardholder data environment has been compromised. Pulling the network cable immediately "stops the bleeding" of credit and debit card numbers, and allows an investigation to begin without ongoing data loss. While this solution is easy to execute and only takes a couple of seconds, a much larger business impact exists that needs to be considered.

The challenge with this option is pretty obvious—the affected system is no longer going to be able to perform whatever function it was previously performing. So the business will either need to have an alternative solution in place, like rolling those services to a secondary or backup system, or they have to be able to absorb the loss of that functionality until a rebuilt system can come online. If the affected system is performing noncritical functions not associated with SLAs or other business critical capabilities, the impact should be relatively low. However, the more likely scenario is that the impacted system is doing something important, like processing payment card transactions, or allowing employees to send and receive email. In these cases, a business decision needs to be made in order to quantify the projected business impact in relation to the time it will take to transfer functionality to another system or deploy a temporary solution.

Patching Systems

Patching the affected systems is the least intrusive, so it will have the least likelihood of negatively impacting functionality and interoperability. I say least

because making changes on a system, even something seemingly positive like patching the OS or out-of-date applications and protocols will create the potential for something to stop functioning. I have seen this play out many times; an update is made and suddenly a dependent process stops working and everybody wonders, "What happened?" The responsible IT party or even the vendor will say something to the effect of, "The patch should have nothing to do with that"; which of course doesn't change the fact that everything was working properly prior to the patch, and after the patch was put in place, things are no longer working.

Patching also requires that you have a complete understanding of any tools or malware what were part of the incident and what their impact was on the targeted systems. While some binaries have very basic persistence mechanisms and are simple to remove, others can be much more complex making removal exponentially more difficult to the point of being time restrictive. It's important to make sure you precisely understand and execute proper removal steps or the threat will remain present. It is for this reason that the decision to patch requires the highest skill and greatest level of technical acumen.

Rebuilding Systems

There will be times when there is so much malware present that it's not time or cost effective to attempt to remove it all. In instances like this it's going to be easier and less expensive from an employee utilization perspective to simply rebuild the system from a gold standard image (provided one exists). Rebuilding the system will have an impact on business capability, but hopefully to a very limited extent (more to come on that). Also, the technical capability to do this is much lower than it is to remove malware and patch a system.

I have worked cases in the past where the target systems had literally hundreds of unique malware packages installed on it. When I was asked by the client which of the malicious binaries could have been responsible for data exfiltration, I was like, "Pick one. There are at least 50 here that could have been the culprit!" It wasn't funny at the time (well, it sort of was), but looking back on that situation, it underscores the point I am making here.

Now, rebuilding from a gold standard image presupposes a few things. First, that you actually have a gold standard image, second that the image is both up-to-date and is not going to introduce additional vulnerabilities or reintroduce old vulnerabilities (like the one that allowed the Breach to happen in the first place), and third, that it does not contain any malware (ie, the most recent image was taken *after* the Breach had occurred). Being able to say, "yes, yes, and no" is ideal, but in reality that likely won't be the case. Murphy's law is alive and well in the world of IR.

Gold standard images are only present is relatively mature organizations that are accustomed to frequent provisioning of systems. If that is the case, multiple images should exist based on functionality, such as basic end-user laptop, developer laptop, executive laptop, webserver, mail server, or backup server. Assuming that appropriate type of image is present, it's important to verify when that image was created as it could contain out-of-date applications or protocols that potentially could introduce additional attack vectors into the environment. Now it's not difficult to allow updates, but based on when the image was created and the volume of updates that have been published since that time, this process could take a significant amount of time, which is a luxury you may not have.

Also, based on when the image was created in relation to the incident, it may contain the same malicious programs or applications that currently are present. Likewise, the vulnerabilities or misconfigurations that allowed the Breach take place may also be present. If either of these possibilities is present, restoring from the image may actually do more harm than good.

The important thing to remember with any of the four options that have been covered is that they all have the potential of improving the situation as well as making it worse. I wish I could tell you that's going to be a cut and dry situation where there will be a clear right and clear wrong path, but there won't be. Like so many other aspects of a Breach response, decisions like these are more a case of which is going to suck less. This is why they need to be made with the key stakeholders, management, executives, and legal counsel. Your technical expertise should guide them, as well as provide them with an overview of the pros and cons of each option, but the final call is theirs. Just make sure they understand the potential impact of whatever choice they make. It's also a good idea to document your recommendations, and admonitions in an email. That way, if your advice is not taken, and things go horribly wrong, you can't be blamed. Bear in mind that unless your services have been retained by legal counsel thereby being protected as privileged, anything you put down in writing can become part of a post Breach E-Discovery request. It is a misconception to think that simply putting "Privileged and Confidential" on an email, or copying an attorney suddenly covers the entirety of that thread under privilege; it does not. While this concept is explained further in Chapter 9, you should always consult with legal counsel if you have any questions regarding legal matters. Sadly, I have learned these lessons from experience, so please take my advice, and make sure you cover your six.

SUMMARY

During the early stages of an incident you will almost always be asked, "What can we do to contain the Breach?" By now you should understand that providing an

answer to this question is not as easy as one might think. In the middle of a Breach most of the decisions you will have to make or will be called upon to advise on will be challenging both in terms of technology and the impact on the business, this one is no different. Solving one problem may end up creating another one, so there will almost always be competing priorities. For this reason, it is important that you engage the key stakeholders and legal counsel early and often when decisions need to be made that will impact the operations of the business.

Containing the incident may negatively impact a law enforcement investigation or the victim's ability to perform their core competency, therefore alternatives would need to be explored. Doing so may introduce the need for a phased or fragmented approach—containing certain aspects of the Breach immediately, while leaving others in place for a period of time, and addressing them later. This can be effective for short-term containment which may appear financially attractive, still requires a holistic, long-term solution. If post Breach remediation is not carried out effectively, and the lure of "good enough" being "good enough" overrides the wisdom of experience, this sort of corner cutting can lead to prolonged litigation, regulatory fines and violations, and loss of customer confidence and market share.

If the choice is made to implement short-term fixes (and it's more than likely will in some fashion), ensure that you make the recommendation that an external team of penetration testers is engaged to assess whether or not the deployed countermeasures are having the desired impact. Few things are worse during a Breach response than thinking or communicating that things have been "fixed," only to find out later that the "fixes" didn't work properly and the attack either continued or became worse.

In addition to business concerns, the victim organization needs to consider their legal and contractual obligations. Based on their geographic region and contract verbiage certain options may be off the table entirely, so make sure that legal counsel is updated regularly. If they don't have a firm grasp on the myriad of legal issues that emerge during a Breach make the recommendation that they retain outside counsel that specialize in Breach coaching and litigation.

Responding to and containing a Breach is tremendously complex. This fact necessitates a cross-functional team of experts representing multiple disciplines. As part of that team, it is your responsibility to have an understanding of how each unique skill fits together, what decisions will need to be made, the potential consequences, and how they are going to play into the overall Breach response. This is why good incident responders are in such high demand! So be thankful that you have chosen wisely, and keep reading!