1 Failure: The Most Common Option

As security professionals, we simultaneously hear platitudes about how users are our best resource, as well as our weakest link. The people contending that users are the best resource state that aware users will not only *not* fall prey to the attacks, they will also respond to the attacks and stop them in their tracks. They might have an example or two as well. Those contending that the users are the weakest link will point to the plethora of devastating attacks where users failed, despite their organizations' best efforts. The reality is that regardless of the varying strengths that some users bring to the table in specific circumstances, users generally are still the weakest link.

Study after study of major data breaches and computer incidents show that users (which can include anyone with access to information or computer assets) are the primary attack vector or perpetrator in an overwhelming percentage of attacks. Starting with the lowest estimate, in 2016, a Computer Technology Industry Association (CompTIA) study found that 52 percent of all attacks begin by targeting users (`www.comptia.org/about-us/newsroom/press-releases/2016/07/21/comptia-launches-training-to-stem-biggest-cause-of-data-breaches`). In 2018, Kroll compiled the incidents reported to the UK Information Commissioner's Office and determined that human error accounted for 88 percent of all data breaches (`www.infosecurity-magazine.com/news/ico-breach-reports-jump-75-human/`). Verizon's *2018 Data Breach Investigations Report* (DBIR) reported that 28 percent of incidents were perpetrated by malicious insiders (`www.documentwereld.nl/files/2018/Verizon-DBIR_2018-Main_report.pdf`). Although the remaining 72 percent of incidents were not specifically classified as resulting from an insider mistake or action, their nature indicates that the majority of the attacks perpetrated by outsiders resulted from user actions or mistakes.

Another interesting finding of the 2018 DBIR is that any given phishing message will be clicked on by 4 percent of people. Initially, 4 percent might sound extremely low, but an attack needs to fool only one person to be successful. Four percent means that if an organization or department has 25 people, one person will click on it. In an organization of 1,000 people, 40 people will fall for the attack.

NOTE The field of statistics is a complex one, and real-world probabilities vary compared to percentages provided in studies and reports. Regardless of whether the percentages are slightly better or worse in a given scenario, this user problem obviously needs to be addressed.

Even if there are clear security awareness success stories and a 96 percent success rate with phishing awareness, the resulting failures clearly indicate that the user would normally be considered the weakest link. That doesn't even include the 28 percent of attacks intentionally perpetrated by insiders.

It is critical to note that these are not only failures in security, but failures in overall business operations. Massive loss of data, profit, or operational functionality is not just a security problem. Consider, for example, that the WannaCry virus crippled hospitals throughout the UK. Yes, a virus is traditionally considered a security-related issue, but it impacted the entire operational infrastructure.

Besides traditional security issues, such as viruses, human actions periodically result in loss of varying types and degrees. Improperly maintained equipment will fail. Data entry errors cause a domino effect of troubles for organizational operations. Software programming problems along with poor design and incomplete training caused the devastating crashes of two Boeing 737 Max airplanes in 2019 (as is discussed in more detail in Chapter 3, "What Is User-Initiated Loss?"). These are not traditional security problems, but they result in major damage to business operations.

History Is Not on the Users' Side

No user is immune from failure, regardless of whether they are individual citizens, corporations, or government agencies. Many anecdotes of user failings exist, and some are quite notable.

The Target hack attracted worldwide attention when 110,000,000 consumers had their personal information compromised and abused. In this case, the attack began when a Target vendor fell for a phishing attack, and then the attacker used the stolen credentials to gain access to the Target vendor network. The attacker was then allowed to surf the network and inevitably accomplish their thefts.

While the infamous Sony hack resulted in disaster for the company, causing immense embarrassment to executives and employees, it also caused more than $150,000,000 in damages. In this case, North Korea obtained its initial foothold on Sony's network with a phishing message sent to the Sony system administrators.

From a political perspective, the Democratic National Committee and related organizations that were key in Hillary Clinton's presidential campaign were hacked in 2016 when a Russian intelligence GRU operative sent a phishing message to John Podesta, then chair of Hillary Clinton's campaign. The resulting leak of the email was embarrassing and was strategically released through Wikileaks.

In the Office of Personnel Management (OPM) hack, 20,000,000 U.S. government personnel had their sensitive information stolen. It is assumed that Chinese hackers broke into systems where the OPM stored the results of background checks and downloaded all of the data. The data contained not just the standard name, address, Social Security number, and so on, but information about their health, finances, mental illnesses, among other highly personal information, as well as information about their relatives. This information was obtained through a sequence of events that began by sending a phishing message to a government contractor.

From a physical perspective, the Hubble Space Telescope was essentially built out of focus, because a testing device was incorrectly assembled with a single lens misaligned by 1.3 mm. The reality is that many contributing errors led to not only the construction of a flawed device but the failure to detect the flaws before it was launched.

In an even more extreme example, the Chernobyl nuclear reactor had a catastrophic failure. It caused the direct deaths of 54 people, another approximately 20,000 other people contracted cancer from radiation leaks, and almost 250,000 people were displaced. All of this resulted from supposed human error, where technicians violated protocols to allow the reactor to run at low power.

These are just a handful of well-known examples where users have been the point of entry for attacks. The DBIR also highlights W-2 fraud as a major type of crime involving data breaches. Thousands of businesses fall prey to this crime, which involves criminals pretending to be the CEO or a similar person and sending emails to human resources (HR) departments, requesting that an HR worker send out copies of all employee W-2 statements to a supposedly new accounting firm. The criminals then use those forms to file fraudulent tax refunds and/or perform other forms of identity theft. Again, these attacks are successful because some person makes a mistake.

NOTE If you are unfamiliar with U.S. tax matters, W-2 statements are the year-end tax reports that companies send to employees.

Other human failures can include carelessness, ignorance, lost equipment, leaving doors unlocked, leaving sensitive information insecure, and so on. There are countless ways that users have failed. Consequently, sometimes technology and security professionals speciously condemn users as being irreparably "stupid." Of course, if technology and security professionals know all of the examples described in this section and don't adequately try to prevent their recurrence, are they any smarter? The following sections will examine the current approach to this problem and then how we can begin to improve on it.

Today's Common Approach

There are a variety of ways to deal with expected human failings. The three most prevalent ways are awareness, technology, and governance.

Operational and Security Awareness

As the costs of those failings have risen into the billions of dollars and more failings are expected, the security profession has taken notice. The general response has been to implement security awareness programs. This makes sense. If users are going to make mistakes, they should be trained not to make mistakes.

Just about all security standards require that users receive some form of awareness training. These standards are supposed to provide some assurance for third parties that the organizations certified, such as credit

card processors and public companies, provide reasonable security protections. Auditors then go in and verify that the organizations have provided the required levels of security awareness.

Unfortunately, audit standards are generally vague. There is usually a requirement that all employees and contractors have to take some form of annual training. This traditionally means that users watch some type of computer-based training (CBT) that is composed of either monthly 3- to 5-minute sessions or a single annual 30- to 45-minute session. CBT learning management systems (LMSs) usually provide the ability to test for comprehension. Reports are then generated to show the auditors to prove the required training has been completed.

As phishing attacks have grown in prominence, auditors started to require that phishing simulations be performed. Organizations also unilaterally decided that they want phishing simulations to better train their users. Phishing simulations do appear to decrease phishing susceptibility over time. These simulations vary greatly in quality and effectiveness. As previously stated, this optimistically results in a 4 percent failure rate.

In general operational settings, training is provided, but there are few standards or requirements for such training. There may or may not be a safety briefing. There are sometimes compliance requirements for how people are to do their jobs, such as in the case of handling personally identifiable information (PII) in certain environments covered by regulations or requirements, such as the Healthcare Insurance and Portability and Accountability Act (HIPAA) and the Payment Card Industry Data Security Standard (PCI DSS). The PCI DSS even requires that programmers receive training in secure programming techniques. NIST 800-50, "Building an Information Technology Security Awareness and Training Program," even attempts a more rigorous structure in the context of the Federal Information Security Management Act (FISMA).

Unfortunately, awareness training, security-related or otherwise, is poorly defined and broadly fails at creating the required behaviors.

Technology

Independent of awareness efforts, IT or security technology professionals implement their own plans to try to reduce the likelihood of humans falling for attacks or otherwise causing damage. For the most part, these are preventative in nature. For example, a user cannot click

on a phishing message if the message never gets to the user. For that reason, organizations acquire software that filters incoming email for potential attacks.

There are also different technologies that can stop attacks from being completed. For example, data leak prevention (DLP) software reviews outgoing data for potentially sensitive information. An example would be if a file attached to an email contains Social Security numbers or other PII, DLP software should catch the email before it goes outside the organization.

The purchase of these technologies is generally random to the organization. While awareness and phishing simulation programs are generally accepted as a best practice, there are no universally accepted best practices for many specific technologies, with a few notable exceptions such as for anti-malware software, which is a staple of security programs.

Cloud providers like Google and Microsoft are becoming increasingly proficient at building effective anti-phishing capabilities into their platforms like Gmail and Office 365. As a result, many organizations are considering whether purchasing third-party solutions is even necessary. Either way, every software solution has its limitations, and no single tool (or collection of tools) is a panacea.

Governance

Although we discuss governance in more detail in Chapter 13, "Governance," for an initial introduction it is sufficient to know that governance is supposed to be guidance or specification of how organizational processes are to be performed. The work of governance professionals involves the specification of policies, procedures, and guidelines, which are embodied in documents.

These documents typically reflect best practices in accordance with established laws, regulations, professional associations, and industry standards. In theory, governance-related documents are expected to be living documents and used for enforcement of security practices, but it is all too common that governance documents only see the light of day during a yearly ritual of auditors reviewing them for completeness in the annual audit.

In an ideal world, governance documents should cover how people are to do their jobs in a way that does not make them susceptible to attacks and in a way that their work processes do not result in losses. This includes how specific actions are to be taken and how specific decisions are to be made in performing job functions.

That ideal world represents the embodiment of a system. A good example of this is McDonald's. Generally, McDonald's expects to hire minimally qualified people to deliver a consistent product anywhere in the world. This involves specifying a process and using technology to consistently implement that process. Although people may be involved in performing a function, such as cooking and food preparation, technology is now driving those processes. A person might put the hamburgers on a grill, but the grill is automated to cook the hamburgers for a specific time at a given temperature. The same is true for french fries. Even the amount of ketchup that goes on a hamburger is controlled by a device. Robots control the drink preparation. McDonald's is now distributing kiosks to potentially eliminate cashiers. Although a fast-food restaurant might not seem to be technology-related, the entire restaurant has become a system, driven by governance that is implemented almost completely through technology.

We Propose a Strategy, Not Tactics

We described in the book's introduction how the scuba and loss prevention industries look at the concept of mitigating loss as a comprehensive strategy. When organizations fail to do this, they attempt to implement random tactics that are not cohesive and supporting of each other. For example, if you think the fact that users create loss is an awareness failing and that the solution is better awareness, you are focusing on a single countermeasure. This approach will fail.

A comprehensive strategy is required to mitigate damage resulting from user actions. This book provides such a strategy. This strategy is something that should be applied to all business functions, at all levels of the organization. Wherever there can be a loss resulting from user actions or inactions, you need to proactively determine whether that loss is worth mitigating and then how to mitigate it.

NOTE Implementing the strategy across the entire business at all levels doesn't mean that every user needs to actively know and apply the depth and the breadth of the entire strategy. (The fry cook doesn't need to know how the accounting department works, and vice versa.) The team that implements the strategy coordinates its efforts in a way that informs, directs, and empowers every user to accomplish the strategy in whichever ways are most relevant for their role.

In an ideal world, you will always look at any user-involved process and determine what damage the user can initiate and how the opportunity to cause damage may be removed, as best as possible. If the opportunity for damage cannot be completely removed, you will then look to specify for the user how to make the right decisions and take the appropriate actions to manage the possibility of damage. You then must consider that some user will inevitably act in a way that leads to damage, so you consider how to detect the damaging actions and mitigate the potential for resulting loss as quickly as possible.

Minimally, when you come across a situation where a user creates damage, you should no longer think, "Just another stupid user." You should immediately consider why the user was in a position to create damage and why the organization wasn't more effective in preventing it.