IRA WINKLER AND DR. TRACY CELAYA BROWN

YOU CAN STOP STUPID

STOPPING LOSSES FROM ACCIDENTAL AND MALICIOUS ACTIONS

WILEY

Table of Contents

List of Tables

Chapter 8

Chapter 11

Chapter 15

List of Illustrations

Chapter 4

Chapter 8

Chapter 9

Chapter 12

Chapter 17

Chapter 19

You CAN Stop Stupid

Stopping Losses from Accidental and Malicious Actions

Ira Winkler
Dr. Tracy Celaya Brown

WILEY

Introduction

We believe that the title of a book is perhaps its most critical characteristic. We acknowledge that the title, *You Can Stop Stupid* is controversial. We had considered other possible titles, such as Stopping Human Attacks, but such a title does not convey the essence of this book. Although we do intend to stop attacks that target your users, the same methodology will stop attacks by malicious insiders, as well as accidents.

The underlying problem is not that users are the targets of attacks or that they accidentally or maliciously create damage, but that users have the ability to make decisions or take actions that inevitably lead to damage.

That is the fundamental issue this book addresses, and it makes a critical distinction: The problem lies not necessarily in the user, but also in the environment surrounding the people performing operational functions.

What Is Stupid?

Managers, security specialists, IT staff, and other professionals often complain that employees, customers, and users are stupid. But what is “stupid”? The definition of “stupid” is having or showing a great lack of intelligence or common sense.

First, let's examine the attribute of showing a great lack of intelligence. When your organization hires and reviews people, you generally assess whether they have the requisite intelligence to perform the required duties. If you did hire or retain an employee knowing that they lacked the necessary intelligence to do the job, who is actually stupid in this scenario: the employee or the employer?

Regarding a person who shows a great lack of common sense, there is a critical psychological principle regarding common sense: You cannot have common sense without common knowledge. Therefore, someone who is stupid for demonstrating a great lack of common sense is likely suffering from a lack of common knowledge. Who is responsible for ensuring that the person has such common knowledge? That responsibility belongs to the people who place or retain people in positions within the organization.

In general, don't accuse someone in your organization of being stupid. Instead, identify and adjust your own failings in bad employment or training practices, as well as the processes and technologies that enable the “stupidity.”

Do You Create Stupidity?

When people talk about employee, customer, and other user stupidity, they are often thinking of the actions those users take that cause damage to your organization. In this book, we refer to that as *user-initiated loss* (*UIL*). The simple fact is that a user can't initiate loss unless an organization creates an environment that puts them in a position to do so. While organizations do have to empower employees, customers, and other users to perform their tasks, in most environments, there is little thought paid to proactively reducing UIL.

It is expected that users will make mistakes, fall for tricks, or purposefully intend to cause damage. An organization needs to consider this in its specification of business practices and technological environments to reduce the potential for user-initiated loss.

Even if you reduce the likelihood for people to cause harm, you cannot eliminate all possibilities. There is no such thing as perfect security, so it is folly to rely completely on prevention. For that reason, wise organizations also embed controls to detect and reduce damage throughout their business processes.

How Smart Organizations Become Smart

Consider that large retail stores, such as Target, have a great deal to lose from a physical standpoint. Goods can be physically stolen. Cashiers can potentially steal money. These are just a couple of common forms of loss in retail environments.

To account for the theft of goods, extensive security controls are in place. Cameras monitor areas where goods are delivered, stored, and sold. Strict inventory control systems track everything. Store associates are rewarded for reporting potential shoplifters. Security guards, sometimes undercover, patrol the store. High-value goods are outfitted with sensors, and sensor readers are stationed at the exits.

From a cash perspective, cashiers receive and return their cash drawers in a room that is heavily monitored. They have to “count in” the cash and verify the cash under the watchful eyes of the surveillance team. The cash registers keep track of and report all transactions. Accounting teams also verify that all cash receipts are within a reasonable level of expected error. Also, as important, the use of credit cards reduces the opportunity for employees to mishandle or steal cash.

Despite all of these measures, there are still losses. Some loss is due to simple errors. A cashier might accidentally give out the wrong change. There might be a simple accounting error. Employees might figure out how to game the system and embezzle cash. Someone in the self-checkout line might accidentally not scan all items. Criminals may still be able to outright steal goods despite the best controls. Regardless, the controls proactively mitigate and detect large amounts of losses. There are likely further opportunities for mitigating loss, and new studies can always be consulted to determine varying degrees to which they might be practical.

An excellent example of an industry that intelligently mitigates risk is the scuba diving industry. Author Ira Winkler is certified as a Master Scuba Diving Trainer and first heard the expression “you can't stop stupid” during his scuba instructor training. The instructor was telling all the prospective instructors that there will always be some students who do not pay attention to safety rules. It is true that scuba diving provides for an almost infinite number of ways for students to do something potentially dangerous and even deadly.

Despite this, scuba diving is statistically safer than bowling. When you consider how that may be, you have to understand that most scuba instruction involves safety protocols. Reputable dive operators are affiliated with professional associations, such as the Professional Association of Diving Instructors (PADI). PADI examines how dive accidents have occurred and works with members to develop safety protocols that all members must follow.

For example, when Ira would certify new divers, all students had to take course work specifying safe diving practices. They also had to go through a health screening process and demonstrate basic swimming skills and comfort in the water. They then had to demonstrate the required diving skills in a pool.

When it comes to certifying people in open water, all equipment is inspected by the students and instructors prior to diving. The potential dive location is chosen based upon the calmness and clarity of the water and limited depth so that students don't accidentally go too deep. Before the dive, there is a complete dive briefing, so students know what to expect, as well as safety precautions and instructions about what to do if a diver runs into trouble. The instructors are familiar with the location and any potential hazards. The number of students is limited, and dive master assistants accompany the group as available to ensure safety. Additionally, instructors are required to ensure there is a well-equipped first aid kit, an emergency oxygen supply, and information about the nearest hospital and hyperbaric chamber.

To become an instructor, Ira went through hundreds of hours of training, especially including detailed training about how to handle likely and unlikely problems. This training includes extensive first aid training. From a risk mitigation strategy, instructors maintain personal liability insurance. Similarly, the sponsoring school maintains liability insurance while also paying for supplemental insurance to cover potential injuries to students. The dive facilities, be they pools, boats, quarries, or so on, also maintain liability insurance.

Essentially, PADI and other professional associations have proactively examined where potential injuries may occur and determined how to prevent them as best as possible. Although some accidents will inevitably occur, there is extensive preparation for those incidents, and the result is that diving is a comparatively safe activity.

Not All Industries Are as Smart

Retail loss prevention and dive instruction have clearly created comprehensive strategies for preventing and mitigating loss that accounts for human error and malfeasance. Unfortunately, many industries, and ironically even many practices within the same industries that are otherwise relatively secure, are not dealing with human error well. For example, Target, which generally has an outstanding loss prevention practice, failed when it came to a data breach where 110,000,000 credit records were stolen.

When an organization fails to account for humor error and malfeasance, and fails to put in sufficient layers of controls, the losses can be devastating. When organizations fail to implement an effective process of risk mitigation to account for user-initiated loss, there is a great deal of blame to go around, but organizations tend to point to the “stupid user” who made a single error.

No case is more notorious for this than the massive Equifax data breach. When Richard Smith, former CEO of Equifax, testified to Congress regarding the infamous data breach, he laid the blame for the data breach squarely on an administrator for not applying a critical patch for a vulnerability in a timely manner. Not immediately applying a patch is not uncommon for organizations the size of Equifax. However, a detailed investigation showed that there was a gross systemic failure of Equifax's security posture.

After all, not only did Equifax allow the criminal in, the criminal was able to explore the network undetected for six weeks, breach dozens of other systems, and download data for another six weeks. The attack was detected only after Equifax renewed a long-expired digital certificate that was required to run a security tool.

This type of scenario is common in computer-related incidents. Whether it is the failing of an individual user or someone on the IT team, a single action, or failure to act, can initiate a major loss. However, for there to be a major loss, there has to be a variety of failures to allow an attack to be successful.

Similar failures happen in all operational units of organizations. Any operational process that does not analyze where and how people can intentionally or unintentionally cause potential loss enables that loss.

The goal of this book is to help the reader identify and mitigate actions where users might initiate loss, and then detect the actions initiating loss and mitigate the potential damage from the harmful acts.

Just as the diving and loss prevention industries have figured out how to effectively mitigate risk arising from human failures, you can do the same within your environment. By adopting the proper sciences and strategies laid out in this book, you can effectively mitigate user-initiated loss.

Deserve More

When we consult with organizations, we find that one of the biggest impediments to adequately addressing user-initiated loss is not getting the required resources to do so. The underlying reason is that all too frequently, people responsible for loss reduction fail to demonstrate a return on investment. In short: You get the budget that you deserve, not the budget that you need. You need to deserve more.

If people believe scuba diving is dangerous, the scuba industry will collapse. If accounting systems fail, public companies can suffer dire consequences. These industries recognize these dangers, and they take steps to demonstrate their value and viability. However, many other professions do not adequately address risk and prove their worth.

The common strategy of dealing with user-initiated loss is to focus on awareness and letting people know how not to initiate a loss. Clearly, this fails all too frequently. Therefore, money put into preventing the loss appears wasted. There is no clear sense of deserving more resources.

It is our goal that you will be able to apply our strategies and show you are deserving of the resources you need to properly mitigate the potential losses that you face.

Reader Support for This Book

We appreciate your input and questions about this book. You can contact us at www.YouCanStopStupid.com.

How to Contact the Publisher

If you believe you've found a mistake in this book, please bring it to our attention. At John Wiley & Sons, we understand how important it is to provide our customers with accurate content, but an error may occur even with our best efforts.

To submit your possible errata, please email it to our Customer Service Team at wileysupport@wiley.com with the subject line “Possible Book Errata Submission.”

How to Contact the Authors

Ira Winkler can be reached through his website at www.irawinkler.com. Dr. Tracy Celaya Brown can be reached through her website at DrTre.com. Additional material will be made available at the book's website, www.youcanstopstupid.com.

I
Stopping Stupid Is Your Job

While professionals bemoan how users make their job difficult, the problem is that this difficulty should be considered part of the job. No matter how well-meaning or intelligent a user may be, they will inevitably make mistakes. Alternatively, the users might have malicious intent and intend to commit acts that cause loss. Considering the act “stupid” assists a malicious party in getting away with their intent.

Fundamentally, you don't care about an individual action by a user; you care that the action may result in damage. This is where professionals need to focus. Yes, you want to have awareness so users are less likely to initiate damage. However, you have to assume that users will inevitably make a potentially harmful action, and your job is to mitigate that action in a cost-effective way.

Part I lays the groundwork for being able to address the potential damage that users can initiate. The big problem that we perceive regarding the whole concept of securing the user—as some people refer to it, creating the human firewall—is that people think that the solution to stopping losses related to users is awareness. To stop the problem, you have to understand that awareness is just one tactic among many, and the underlying solution is that you need a comprehensive strategy to prevent users from needing to be aware, to create a culture where people behave appropriately through awareness or other methods, and to detect and mitigate loss before it gets out of hand.

Any individual tactic will be ineffective at stopping the problem of user-initiated loss (UIL). As you read the chapters in Part I, you should come away with the holistic nature of the problem and begin to perceive the holistic solutions required to address the problem.

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.