CHAPTER 1
Understanding Risk

If you're reading this book, I'd hazard a guess that you've read some of the doom-and-gloom cybersecurity books out there as well. There are many, and many are great (see the bibliography for suggestions). What's more, I am sure you have had your fill of statistics. Dreadful statistics showing how cybercrime is increasing by the day. I'll include some of those, too, just to satisfy any morbid curiosity left in you, but they are essentially useless. By the time the ink is dry on these pages, the numbers have changed. For the worse.

A BRIEF SAMPLING OF DREAD

- Hacker Attack Rate: 39 Seconds

  Assistant Professor of Mechanical Engineering Michel Cukier at the A. James Clark School of Engineering conducted the study that profiled the actions of hackers using brute-force methods to gain access to a set of exposed computers. The results showed that the computers were attacked about 2,244 times per day.
- More than 33 percent of United States consumers have experienced a cyberattack.

  This was reported in a survey by Zogby Analytics commission for the Hartford Steam Boiler Inspection and Insurance Company (HSB), with the most likely victims being between 18 and 24 years old. Moreover, the associated incident costs ranged from $500 for 56 percent of the cases to between $1,000 and $5,000 for 23 percent of the cases.
- According to the "Internet Security Threat Report—Symantec 2017" (Volume 22, April 2017):

- It takes on average two minutes for an Internet of Things (IoT) device to get attacked.
- The average ransom amount for a ransomware attack went from $373 in 2014 to $1,077 in 2016.
- Over the last eight years, more than 7.1 billion identities have been stolen as a result of data breaches.
- In 2016, the United States was number one both in number of data breaches (1,023) and in identities stolen (791,820,040).

- According to the "2017 Data Breach Investigations Report" (Verizon):
  - 75 percent of the breaches are perpetrated by outsiders, versus 25 percent involving insiders.
  - 62 percent of breaches featured hacking, of which 81 percent leveraged stolen or weak passwords.
  - 66 percent of malware was installed through malicious email attachments.
  - 73 percent of the breaches were financially motivated; 21 percent were espionage-driven.
- According to the "Small Business Trends" website (https://smallbiztrends.com):
  - 43 percent of cyberattacks target small business.
  - Only 14 percent of small businesses rate their ability to mitigate cyberrisks vulnerabilities and attacks as highly effective.
  - 60 percent of small companies go out of business within six months of a cyberattack.
  - 48 percent of data security breaches are caused by acts of malicious intent. Human error or system failure account for the rest.
- According to Juniper Research's study titled "The Future of Cybercrime & Security: Financial and Corporate Threats & Mitigation" (Juniper Research, Ltd.):
  - Cybercrime is expected to cost businesses over $2 trillion by 2019.

  - Although North America has seen the lion's share of these breaches (60 percent in 2015), the proportion will level off as global digitization levels the playing field.
- According to "Cybersecurity Ventures' Predictions for 2017 through 2021":
  - The cost of cybercrime damages worldwide is estimated to be $6 trillion annually by 2021.
  - In 2016, the cybersecurity unemployment rate dropped to zero percent, and it is expected to remain at that level through 2021, with a projected job-to-skills shortfall of 1.5 million positions by 2019.
- ISACA's *2016 Cybersecurity Global Data Snapshot* lists social engineering, insider threats, and advanced persistent threats as the top-three threats facing organizations.
- According to Barkly Protects, Inc.:
  - One-third of the IT professionals surveyed by Barkly reported their security had been bypassed by a cyberattack in 2016.
  - 71 percent of organizations targeted with ransomware attacks were successfully infected.
  - Over half the organizations that suffered successful cyberattacks in 2016 are not making any changes to their cybersecurity posture in 2017, with budgetary constraints cited as the main block to improved cybersecurity.

How Much Is It Worth to You?

In the misty past, a person's most valuable possessions were things he could see: his castle, gold, tapestries, even his heirs! Their value was tied to their physical existence.

Today, the concept of *value* has expanded beyond tangibles to include intangibles such as data, intellectual property, and reputation. As a matter of fact, many intangibles hold more value than tangibles. Consider, which is more important: an artisanal pizza or the *recipe* for the artisanal pizza? It's no accident that the phase following the Industrial Revolution has been nicknamed the Information Revolution.

The rise of intangible valuables affects every individual as well as businesses of all sizes. These things of value that individuals and businesses create are—like all things of value—coveted by others and therefore warrant your protection. So, just like you would protect your valuable jewelry, you must protect your valuable data. It's a simple concept.

What is interesting in this analogy is the assumption that we all share a common understanding of what is *of value*. You certainly have no problem intuiting that a set of diamond earrings is valuable and should therefore be stored in a secure place. Which place and how secure? That, too, is straightforward to understand. We have an innate sense of value to guide us in these decisions—something that tells us that a $300 pair of earrings is safe in the jewelry box in the apartment while a $30,000 pair of earrings is best protected in a bank's safe deposit box. Easy to understand and easy to make a value judgment on.

We make these types of judgments every day, and we're very good at it. We understand what's of value, and we understand the risks to this value:

Earrings? Theft!
Property? Fire!

It ends up that we are very good at making complex risk management decisions on a daily basis. Who knew?

Risk! Not Just a Board Game

Consider this situation: It is 11:00 at night, and you just finished dinner with friends at your favorite restaurant. Walking to your car, you reach an intersection and see that the walk signal is red. You look left. You look right. You see a car down the block, shrug it off, and cross the street. No problem.

Now, let's change this scenario a little. Same story, only this time you are pushing a stroller with your baby in it. What's the decision now? Do you cross the street or wait for the signal to change? My bet is you wait.

We just stumbled on the concept of *risk acceptance*, which will prove to be of real importance in the pages that follow. The bottom line is that we all live with risk every single day of our lives. We constantly make decisions about risk and, when we're done evaluating, we take action signifying our acceptance of this risk.

In the example just suggested, in one case you accepted the risk that you can cross the street against the light, and in the other case, when you had your baby along, you did not. How does this translate to the cyberworld? In some cases, we accept the risk of having our information available out there (e.g., when using Facebook, Instagram, Swarm, and the like), and in others we do not (e.g., when we are using our credit card or revealing our medical records).

Studying risk is taking a trip down a fascinating, complex, and intricate labyrinth. It is hard-core science—involving complex mathematics, ethics, and philosophy—with potential life-and-death implications (e.g., the risk of reprisals when we attack a terrorist group, the risks that first responders take every day, etc.). This is certainly not a book to start you on this type of journey, although I have included a selected bibliography for you to consider at the back of this book. The purpose of this book is to expose you to some risk management and tech concepts so that we can develop a common language when discussing how to protect your things of value from cyber-based threats. With that in mind, let's start with a simple definition. What is risk?

Risk is the combination of the likelihood of an event and its impact.

What's the risk of a hurricane in Miami?

Well … how likely is it, and what will its impact be when it hits?

Why do you need know this? Because, for starters, the answer determines whether you want to move there, if you want to start a business there, if you want to send your kids to school there, and how much your insurance will cost you to protect you from this risk, and so on.

How can you determine the likelihood that a hurricane will strike Miami? You have tons and tons of statistical data that give you a good sense of the frequency of hurricanes hitting the area over the past couple of hundred years.

What's the impact? There is the cost of rebuilding, the cost of business losses, environmental damage, and the potential of loss of life, among many others. You get the picture. Who is good at keeping these types of statistics? Insurance companies. But this alone is not the complete picture. Let's fill in the blanks.

First, let's assume you have decided that you want to live in Miami. That's key. That decision (like the road-crossing example discussed earlier) implies some degree of risk acceptance out of the gate. You know that hurricanes strike Miami, yet you choose to live there. Fine. (Who am I to judge? I live in New York!)

But you're not simply living in Miami. Knowing that hurricanes may hit Miami, you have chosen to live in a "hardened" house, meaning a house that is as hurricane-proof as you *choose* to make it. Before buying the house, you did your research, compared options, and decided to buy a house that can withstand a Category 3 hurricane. That was your choice. It was a very important one: You didn't just choose any house. You chose to get a hardened one. This hardened option? That's a *control*. Controls act against risks. Your control was to buy a house that can withstand a Category 3 hurricane. That *mitigates* your risk: If a Category 3 storm hits, you are protected.

But what happens if a Category 4 storm hits? Well … you have to deal with that risk.

So, how do you deal with the Category 4 risk? You call your insurance company and tell them that you need hurricane insurance. They quote you an outrageous amount, to which you respond with "Wait! I have a Category 3 hurricane-proof home!" They go recheck their numbers and come back with a much more reasonable amount, but they tell you they will cover only Category 4 storm damage and higher. Nothing below. You sign on the dotted line.

You have now *transferred the risk* of a Category 4 storm and higher to the insurance company. Now, you're breathing more easily. What did you accomplish?

First of all, you *assessed the risk* of your desire to live in Miami and decided that *this was worth it to you*. You *applied a control* to *mitigate this risk* by buying the hardened house. You then *transferred the rest of the risk* to the insurance company.

Did you eliminate the risk? Nope! You can still get hurt if a tree falls on you during a Category 3 storm and the insurance will not pay for your medical bills. In other words, no matter what you do, there is always some risk that is left over. Always. *Risk is never zero*. This leftover risk is called *residual risk* and that's the risk that you choose to accept.

Or not.

If you don't accept the risk, you can do more to make it more to your liking. You can apply more controls. For example, you can invest in a Category 4 hurricane-proof home and renegotiate with the insurance company to cover more or different kinds of risk. You can buy your own weather station to predict hurricanes before anyone else does and evacuate the area. You can build a bunker. You can do all sorts of things insofar as the things you do to mitigate your risk don't exceed the value of what you're trying to protect. In risk management terms: The cost of controls cannot exceed the value of the asset you're trying to protect. That would be silly, and clearly, you're very smart, so you wouldn't be doing this anyway.

One last thing: There are all sorts of controls. Broadly speaking, they fall into one of the following categories: *preventative*, *detective*, *corrective*, and *compensatory*. Your weather station is a detective control, the hardening of your home is a preventative control, and your bunker is a compensatory control. What's the corrective control? The insurance inspector who comes once a year to confirm that the house remains hardened. (Notice that your insurance policy is not a control. It doesn't *do* anything to reduce risk; it only accepts some of the risk you transfer to it. You're still the one incurring the risk; it's just that you're not on the hook for the damages that may result.)

What happens if you employ all these controls at once? Well—you risk strategist, you!—you just developed a *defense-in-depth* strategy to mitigate your

risk. You rolled out your controls in a way that they complement one another; if the first fails, the next kicks in, and so on.

Okay, now we're talking the same risk management language. Enough for our purposes (secretly, we haven't even scratched the surface of the risk management field, but let's accept this and move on). But wait, you say: How does the hurricane example apply to my data in New York? Believe me, it does, and I'll show you how. For now, just note those definitions and examples. We'll apply them to your data universe shortly.

But I need to come clean about something before we go further down the cyberpath: Remember all these statistics about likelihood that the insurance company had? How useful they were in determining risk and so forth? Well … we don't really have any of those for cyberattacks. Not enough to build a very robust statistical model about likelihood. At least not yet.

We certainly have a good sense of the *impact*: Your life (or your business) will be in shambles if someone steals your identity (or your business data), so we're all pretty clear on impact. The rest, *you* are going to estimate. Notice the change from "we" to "you"? Excellent! I am betting that you, of all people, know best when it's okay to cross the street. I trust you.

This all makes sense when the risk is personal, right? You get to make the call on what risk you're willing to accept. But what if the risk is to a business? The same concept applies. The people who decide whether or not to accept business risk can only be the owner(s) of the business. No one else. If it is a small business, the owner (usually the president or CEO) is responsible for accepting or rejecting risk. For larger firms, the board of directors is ultimately responsible for risk management decisions. In the absence of a board, then it is the CEO or the executive appointed by the shareholder(s) to run the company. If you are not "them," then you need to trust them that they are making the right decision. That's why they are there. The responsibility is theirs, and theirs alone, to decide whether it is safe to "cross the street." It is your responsibility to advise them one way or another.

To summarize our risk definitions:

- *Asset*: Anything of value
- *Risk*: Likelihood of an event, multiplied by its impact
- *Mitigated risk*: Existing risk after controls have been applied
- *Residual risk*: What's left over after risks have been mitigated or transferred as much as possible
- *Accepted risk*: Residual risk that has been accepted, aka *the risk of doing business*
- *Controls*: Active countermeasures, be it processes, systems, or applications, that prevent, detect, correct, or compensate against risk