The New School of Information Security

Adam Shostack • Andrew Stewart

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The views and opinions expressed in this book are wholly those of the authors and do not represent those of their employers or their employers' clients or customers.

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact:

U.S. Corporate and Government Sales
800-382-3419
corpsales@pearsontechgroup.com

For sales outside the United States, please contact:

International Sales
international@pearson.com

Visit us on the Web at informit.com/aw.

Editor-in-Chief
Karen Gettman

Acquisitions Editor
Jessica Goldstein

Managing Editor
Gina Kanouse

Project Editor
Anne Goebel

Copy Editor
Gayle Johnson

Indexer
Cheryl Lenser

Publishing Coordinator
Andrea Bledsoe

Cover Designer
Adam Shostack
Andrew Stewart

Composition
Jake McFarland

Library of Congress Cataloging-in-Publication Data:

Shostack, Adam.
The new school of information security / Adam Shostack and Andrew Stewart.
p. cm.
Includes bibliographical references.
ISBN 0-321-50278-7 (hardback : alk. paper) 1. Information technology—Security measures. 2. Computer security. 3. Computer security equipment industry. 4. Business—Data processing—Security measures. I. Stewart, Andrew, 1975- II. Title.

HD30.2.S563 2008
658.4'78—dc22

2007052580

ISBN-13: 978-0-321-50278-0
ISBN-10: 0-321-50278-7

Text printed in the United States on recycled paper at RR Donnelly in Crawfordsville, Indiana.
First printing March 2008

Preface

"I didn't have time to write you a short letter,
so I wrote a long one."

—Mark Twain

We've taken the time to write a short book, and hope you find it enjoyable and thought-provoking. We aim to reorient security practitioners and those around them to a New School that has been taking shape within information security. This New School is about looking for evidence and analyzing it with approaches from a wide set of disciplines. We'd like to introduce this approach to a wider audience, so we've tried to write this book in a way that anyone can understand what we have to say.

This isn't a book about firewalls, cryptography, or any particular security technology. Rather, it's about how technology interacts with the broader world. This perspective has already provided powerful insights into where security succeeds and fails. There are many people investing time and effort in this, and they are doing a good deal of interesting research. We make no attempt to survey that research in the academic sense. We do provide a view of the landscape where the research is ongoing. In the same spirit, we sometimes skim past some important complexities because they distract from the main flow of our argument. We don't expect the resolution of any of those will change our argument substantially. We include endnotes to discuss some of these topics, provide references, and offer side commentary that you might enjoy. Following the lead of books such as *Engines* *of Creation* and *The Ghost Map*, we don't include endnote numbers in the text. We find those numbers distracting, and we hope you won't need them.

Some of the topics we discuss in this book are fast-moving. This isn't a book about the news. Books are a poor place for the news, but we hope that after reading *The New School*, you'll look at the news differently.

Over the course of writing this book, we've probably written three times more words than you hold in your hands. The book started life as *Security Decisions*, which would have been a book for managers about managing information security. We were inspired by Joan Magretta's lovely little book, *What Management Is*, which in about 200 pages lays out why people form organizations and hire managers to manage them. But security isn't just about organizations or managers. It's a broad subject that needed a broader book, speaking to a wider range of audiences.

As we've experimented with our text, on occasion removing ideas from it, there are a few fascinating books which influenced us and ended up getting no mention—not even in the endnotes. We've tried to include them all in the bibliography.

In the course of writing this book, we talked to a tremendous number of people. This book is better for their advice, and our mentions are to thank them, not to imply that they are to blame for blemishes that might remain. If we've forgotten anyone, we're sorry.

Simson Garfinkel and Bruce Schneier both helped with the proposal, without which we'd never have made it here. We'd both like to thank Andy Steingruebl, Jean Camp, Michael Howard, Chris Walsh, Michael Farnum, Steve Lipner, and Cat Okita for detailed commentary on the first-draft text. But for their feedback, the book would be less clear and full of more awkward constructs. Against the advice of reviewers, we've chosen to use classic examples of problems. One reviewer went so far as to call them "shopworn." There is a small audience for whom that's true, but a larger one might be exposed to these ideas for the first time. We've stuck with the classics because they are classic for a reason: they work. Jon Pincus introduced us to the work of Scott Page. We'd like to apologize to Dan Geer for reasons that are either obvious or irrelevant. Lorrie Cranor provided timely and much appreciated help in the academic literature around security and usability. Justin Mason helped correct some of the sections on spam. Steven Landsburg helped us with some economic questions. We would like to thank Adam's mom and Andrew's wife.

We'd also like to thank the entire community contributing to the Workshop on Economics and Information Security for their work in showing how to apply another science in broad and deep ways to the challenges that face us all in security.

It's tempting in a first book to thank everyone you've ever worked with. This is doubly the case when the book is about the approaches we bring to the world. Our coworkers, managers, and the people we have worked with have taught us each tremendous amounts, and those lessons have been distilled into this book.

Adam would like to thank (in roughly chronological order) cypherpunks Eric Hughes, Steve Bellovin, Ian Goldberg, and others too numerous to name, for fascinating discussions over the years, Ron Kikinis, coworkers at Fidelity, Netect (Marc Camm, David Chaloner, Scott Blake, and Paul Blondin), Zero-Knowledge Systems (Austin and Hamnett Hill, Adam Back, Stefan Brands, and the entire Evil Genius team), my partners at Reflective, and the Security Engineering and Community team at Microsoft, especially Eric Bidstrup and Steve Lipner. In addition, everyone who I've written papers with for publication has taught me a lot: Michael J. Freedman, Joan Feigenbaum, Tomas Sander, Bruce Schneier, Ian Goldberg, Austin Hill, Crispin Cowan, and Steve Beattie. Lastly, I would like to thank my co-bloggers at the Emergent Chaos Jazz Combo blog, for regularly surprising me and occasionally even playing in tune, as well as the readers who've commented and challenged us.

Andrew would like to thank Neil Todd and Phil Venables for their help and guidance at the beginning of my career. I would also like to thank Jerry Brady, Rob Webb, Mike Ackerman, George Sherman, and Brent Potter. Please note that my mentioning these people does not mean that they endorse (or even agree with) the ideas in this book.

Finally, we'd both like to acknowledge Jessica Goldstein, who took a chance on the book; Romny French; our copy editor, Gayle Johnson, and our project editor, Anne Goebel.

Chapter 1

Observing the World and Asking Why

In December 2006, Turkish authorities announced the arrest of Ali Y'nin and nine accomplices for bank fraud. They accused Y'nin of leading a gang that sent millions of virus-laden emails. About 11,000 of the recipients opened the email message and unknowingly infected their computers. Then when the victims used online banking services, the gang captured the passwords for those bank accounts and drained them using false identification, fake ATM cards, and Western Union money transfers.

How have we found ourselves in a world in which a small Turkish gang can drain bank accounts on such a massive scale? The police state that Y'nin and his accomplices sent 3.4 million emails and compromised about 11,000 bank accounts. That is a success rate of only 0.3%, but it is hard to imagine that Y'nin was disappointed at being able to access the bank accounts of "only" 11,000 people.

Part of the answer is that because the interconnected world of computers and the internet provides many advantages to criminals, they are drawn to electronic crime. Attacks can be automated and carried out in large numbers. Imagine Y'nin attempting to perform the same fraud, but in person at bank branches. If each member of his gang tried to walk into the same bank branch claiming to be a different person each time, even a bored security guard would catch on after a while. If the gang spent all day traveling to different banks and spent one hour per account, they would be doing nothing but going from bank to bank eight hours a day for over six months. The internet makes everyone more efficient, even criminals. Perhaps especially criminals.

Although Y'nin and his gang were eventually caught, it is much harder to catch an electronic thief than a robber in the physical world. Investigating a burglary might take the police an hour or perhaps a day. An electronic break-in executed across international borders might require months or years of investigation. Only a few national police agencies take on cases that require such an investment of time and effort, whereas anyone connected to the internet can now attack computers around the world. In some of these countries, laws about electronic crimes might not be clear, or there may be no effective local law enforcement to make an arrest. Is it illegal to send email spam from China? What happens if an attacker launders his attack through a computer in Nigeria? Some large companies are dedicating resources to helping police forces investigate attacks that matter to them, but it is not clear if this strategy is a good investment. Another challenge for law enforcement is that the skills required to investigate computer crime quickly go out of date because of the rapid advance of technology. If an officer learned to develop latent fingerprints thirty years ago, that knowledge is still valuable in investigating crimes. In contrast, the ability to perform a forensic investigation of a computer that runs Windows 95 is of little use today.

Because attackers can carry out attacks in a highly automated way and because they are unlikely to ever be caught, online crime is attractive to criminals not just in Turkey, but everywhere. American brokerage houses have found themselves losing millions of dollars to schemes in which criminals use other people's money to "pump and dump" the stock market. The scheme starts when a thief buys some thinly-traded penny stock. The thief then breaks into the victim's bank account and uses the person's money to buy up that stock. The stock rises in price, and the thief then sells his holdings in the now-inflated stock, leaving him much richer and the victim much poorer. (If the thief is clever, he might even set up automated sale orders. The link between the thief and the automated selling of the stock is hard to prove, as is the fact that someone gained illegal access to the victim's account.)

When confronted with computer crime, it is hard to shake the impression that information security is failing. It can seem that these failures are everywhere, filling our electronic world with spam, computer viruses, and identity theft. Even worse, these problems seem to increase even as we spend more time and money on security. We might expect that a rise in electronic crime is a natural result of the world's becoming increasingly electronic. As money and influence move online, so do crime and vandalism. But as crime and vandalism move online, so must security. Ideally, security shows up first and allows us to preempt problems, but that seems to be a rare occurrence. It is often easier to experiment with and build software without security features, so they tend to get added later or not at all. The design of security measures can also cause frustration by getting in the way of the wrong things, so people seek to minimize such features.

But information security matters; it *is* important. It matters to companies and their shareholders. It is of great importance to the general public, whose personal data is stored by the companies and organizations with which they interact (and by some with which they don't). We all hope our private files and email correspondence remain secure. The security industry and security professionals are the guardians of that personal information. They seek to frustrate bad guys such as Y'nin and his ilk by employing standard ways of working and by deploying security technologies. Unfortunately, these efforts have not always been successful.

This chapter delves into some of the most apparent failures of information security. These topics often have a nuanced history. By discussing them in detail, we lay the groundwork for the first half of this book, in which we analyze the myriad factors that have allowed such failings in information security to occur. In the second half, we build on the sum of these observations to reveal what we believe must happen to improve the state of information security in the world, how those changes can be made, and who is in a position to make them. Everyone will benefit from these changes, from multinational corporations to individual consumers.

Many books about information security focus on an idealistic notion of what security *should* be, or they approach security problems from a purely mathematical or technological perspective. Our approach is to begin by looking at the state of the world and trying to understand why it is the way it is. We believe that only through a balanced, well-rounded understanding of the nature of problems can we begin to design solutions that are both effective and efficient. We begin our discussion with a widely visible failure of information security.

Spam, and Other Problems with Email

The flood of unsolicited email flowing into our mailboxes seems to get worse each year, despite more antispam software, more laws, and more email lost to spam filters. In 1994, a law firm decided that the internet would be an ideal way to advertise its legal services. The firm sent a message to thousands of discussion groups, advertising its services. This was widely seen as having opened the floodgates to today's deluge of spam.

Sending an email message is so inexpensive that it makes sense to send one to every email address that can be found, rather than trying to pick specific recipients. Imagine if companies didn't have to pay anything to deliver paper catalogs.

Everyone's mailbox would be stuffed full of catalogs from every company in the world! After all, they can't make money unless people know about their revolutionary product. The United States today doesn't have a general-purpose privacy law that forbids the secret harvesting or sale of most types of personal information, so email addresses are not protected. Privacy laws in other countries vary, but strong privacy laws don't seem to inhibit spam.

There are two types of spammers. The first are companies you did business with once, which then send you emails forever. Even if you ask them to stop, the mail keeps coming. Consumers see this as spam. However, these companies have real products to sell. They're not outright fraudsters. The second type are criminal spammers who send spam about things such as sex pills, stocks, or quick fixes to your credit. These criminals often break into computers and use them, along with their network connections, to send spam.

As spam was rising, so was a new problem—*adware*. Adware companies called themselves "affiliate marketers." They claimed that people *chose* to install software that displays pop-up ads to the user. Sometimes this was and even still is true, but often the adware is embedded in other software and installs itself without the meaningful consent of the PC's owner. (By meaningful consent, we mean that the person installing the software understands what he is getting into.) Adware can also piggyback on a program that a user wants. Sometimes this is done with the cooperation of the author of the desirable program, who takes part of the revenue and earns a living by giving away his software. Other times, this is done as an unauthorized repackaging of innocent software. The adware industry has been creative in devising new ways for its software to surreptitiously install on people's computers. Adware uses innovative means to ooze into the obscure corners of a computer so that it can't easily be removed. Today, some experts say it can be more cost-effective to reinstall a computer than to remove a bad adware infection.

Another attack that uses email is *phishing*. Phishing is the art of sending fraudulent emails designed to look like they are from a company such as a bank. The phisher's goal is to lure people into visiting a web site that *looks* like their bank's real web site. The phisher (or an associate) then uses the fake but authentic-looking web site to convince people to provide personal information such as usernames, passwords, or mother's maiden name. The attacker then takes that information and uses it to access the victim's real bank account. Unpleasantness ensues.

At its root, phishing is a fraud that exists because of the difficulty of authentication—verifying that an entity is who it claims to be. It can be hard to identify the real sender of an email. It can be hard to tell whether a web site really belongs to a given bank. Banks and other institutions that conduct business online have the same problem in reverse. They can find it difficult to identify their customers when someone shows up at their web site to log in. As with spam, the ability to perform phishing attacks is facilitated by the global, largely anonymous nature of the internet. In January 2006, more than six billion emails were recorded as part of 15,000 different phishing scams.

Criminals use phishing attacks because they work. In a test of people's ability to distinguish real email from fake, only 6% got all the answers right, and only half of real emails were recognized as being real. Even so, many companies that do business online have not yet adopted some simple measures that would help protect their customers. Phishing attacks use fake web sites to harvest the personal information of victims, so companies that do business online should advise their customers to never click a hyperlink in an email. Companies should also never send their customers links in an email. Customers should be told that whenever they want to visit the company online, they should use a bookmarked web address, and that web address should ideally be delivered using traditional postal mail. (This advice is intended for those companies that have ongoing relationships with their customers, and who send them occasional alerts.) Rather than take these measures, many companies have instead made things more difficult for their customers by registering new web addresses, using confusing web addresses, and using certain technologies in their web pages that make it easier for fraudsters to camouflage their actions.

To be fair, some companies have sought to address the problem of phishing by implementing a new breed of authentication technologies. In theory, these products help the customers identify when they are at a real web site rather than a fake. In practice, they don't seem to work. For example, in a 2007 study, one of the market-leading products in this space was shown to be ineffective 92% of the time.

As we depend on email more and more, its security weaknesses become ever more apparent.

Hostile Code

Viruses continue to plague our computers. The first viruses were created in the early 1980s. Early viruses were handcrafted, and their creators had some degree of skill. Virus creation became much easier with the introduction of powerful virus-creation toolkits. This has led to a dramatic upswing in the number of viruses. This problem with viruses is not unique to any one vendor of computer software. Viruses affect a wide variety of systems, from mobile phones to mainframes.

For the last twenty years, the majority of anti-virus (AV) products have relied on explicit knowledge about every virus that exists in the world. That knowledge is codified within a *signature*. When a piece of AV software can match the bits in a file to a signature in its library, it blocks or deletes what it presumes to be a virus. This approach is effective close to 100% of the time when the AV software has a signature for the particular virus that happens to be attacking the computer. When it doesn't, this approach doesn't help. The value of an AV product therefore hinges on two things: the AV product vendor must identify new viruses and create signatures for them, and those signatures must reach the end user's computer as quickly as possible. Most AV products are updated daily or weekly with new signatures, but this is a never-ending race between the virus writers and the AV product vendors. Even if you run AV software, your computer might become infected by a virus before a signature is installed. The dramatic changes in virus creation over the past quarter century contrast with the rather tepid evolution of AV products.

Commercial AV products have typically been signature-based. Vendors have periodically brought products to market that use heuristics, such as analyzing behavior, to try to identify viruses. The idea is to remove the dependency on signatures by learning how viruses tend to act. But this technology can struggle with distinguishing between hostile and benign actions, and it can have an error rate of 50% or more. We certainly have fewer problems with computer viruses due to the degree of protection that AV software can provide. But we have only treated the symptoms. Viruses continue to be created at a very high rate. We haven't solved the problem with existing technology, and millions of people continue to be affected. With no cure in sight, it seems that viruses will be with us for some time.

Specialists refer to self-propagating network viruses as worms. On November 2, 1988, Robert Morris, Jr., a student at Cornell University, released the first internet worm. Morris claimed that his intention was not to create damage, but to attempt to determine the size of the internet at the time. It had a bug that caused it to infect machines too quickly.

The Morris Worm, as it became known, pre-dated a raft of damaging internet worms that took root on the internet and within enterprise networks from 2001 onward. There was no fundamental difference between the methodology or techniques used by those modern incarnations of worms and the original Morris Worm. (The Morris Worm targeted the most popular operating systems on the internet, just as subsequent worms have done.) A decade passed between the Morris Worm and those later incarnations.

Viruses, worms, adware, and other hostile code are now lumped together under the generic term *malware*, meaning software that no one wants around. We have gained more knowledge of malware, and the defensive technologies we can employ have become more robust. But modernity is little consolation if we continue to fall victim to the same problems.

Security Breaches

In mid-2006, the *New York Times* and the Associated Press revealed that a laptop containing the personal information of 26.5 million U.S. veterans had been stolen. This is about 9% of the U.S. population. The 26.5 million individuals who were affected were all living veterans who had been discharged since 1976. When the data breach was announced, much uproar occurred in the press and among veterans. The question most often asked was, how could this happen? The reality was that many other organizations of all sorts and sizes have suffered similar breaches in their information security. The organizations affected by these security breaches range from government departments to nonprofit organizations and multinational corporations. Only some states require companies to publicly disclose breaches. Reports are most prominent (or at least most visible) in the English-speaking world, so we are most able to discuss breaches that affect Americans.

TJX is an example of a company that announced a breach. TJX owns well-known brands in the U.S. such as T.J. Maxx and Marshalls, and it has retail stores in Canada and Europe. TJX announced on January 17, 2007 that its computer systems had been hacked. The personal data that was compromised included customer information related to purchases and returns, and it contained credit and debit card numbers. The number of credit and debit card numbers compromised by the attackers is unknown, but estimates (and opinions) range from about 45 million to as many as 200 million cards. According to a TJX press release, TJX believes that its systems were intruded upon from as early as July 2005 until January 2007. Eighteen months was enough time for the attackers to thoroughly ransack the TJX computer network.

Some of the data that was stolen from TJX was used to commit crimes. Police in Florida arrested six people suspected of a fraud scheme that used the stolen credit card data. Unfortunately for TJX, one of the victims was Massachusetts Attorney General Martha Coakley, whose information was used to fraudulently purchase a Dell computer. That probably contributed to the early momentum of the investigation.

Over half of all Americans have been sent notices that their personal data may have been compromised in one of the many breaches that have been disclosed. This number seems low given the vast number of databases containing personal information, the rates of reported laptop theft, and how personal information is bought, sold, and traded. One effect of these "breach notices" is that the sorry state of information security has become more visible, and people want to know why things are so bad.

Chapter 4 is devoted to breaches, so we won't dwell on that topic here. Suffice it to say that security breaches can cause real pain to individuals whose personal data has been compromised, and one of the major causes of concern with such incidents is the threat of identity theft.

Identity and the Theft of Identity

Imitation is the sincerest form of flattery, but no one is flattered by having their good name and credit used for fraud. Such frauds include emptying your bank account, applying for credit, or getting medical care in your name. Personally identifying information such as your full name, national identity number, bank account details, and so on are valuable precisely because they can be used by someone else to impersonate you.

The desire to commit fraud is an important part of the rapidly growing and widely misunderstood crime known as identify theft. Before we can discuss it, we need to describe identification, authentication, and authorization. These three concepts are often confused. *Identification* concerns the labels we provide for things. Much like *The New School of Information Security* identifies a book, "John Wilson" identifies a person. We use other identifiers to identify people, such as "Dad." Dad is not a *unique* identifier, but most people are pretty sure whom they mean when they say it. A bank with eight customers named John Wilson needs to be able to differentiate between them. Anyone can claim to be John Wilson, so how can we tell if he really is? The answer lies in *authentication* to figure out which John Wilson is *authorized* to take money from account number 1234.

You may plan to have coffee with John, and he might tell you that he is tall, bald, and is wearing a green shirt today. Those are authenticators. They help you recognize John at the coffee shop. But if you're a bank, you want to make sure that John is authorized to withdraw money, so you might check his signature, password, or PIN. Identification and authorization are tricky. Too many organizations believe that anyone who knows your social security number (SSN) is you.

The same information about us is stored repeatedly, by different organizations and in different places. Tremendous duplication occurs, and many organizations continue to design processes that depend on these little pieces of data. The problem is that many of these identifying fragments were never designed for the ways in which they are being used. The SSN was not designed to be secret, and yet it is widely believed to be secret and often is treated as such. The result is that SSNs are used as both an identifier and an authenticator. We are told it is important not to hand out our SSN willy-nilly, but at the same time, everyone demands it.

If something is valuable, it should be protected, and we should give our personal information to only trustworthy organizations that really need it. Unfortunately, most organizations seem to think that they are trustworthy and that they must have our personal information. Landlords, utility and insurance companies, employers, hospitals, governments, and many others all profess to be completely trustworthy. It's likely that these organizations, storing the most personal information imaginable, will authorize hundreds of thousands of other completely "trustworthy" people at a variety of organizations to see it, increasing the possibility that it will become compromised.

Why do these approaches persist? The idea that we have a "core identity" that is truly "us" seems to be both strong and pervasive, as does people's desire to build on it. These drivers seem to be deep-seated, despite the practical problems. The willingness to build identity systems without testing our ideas mirrors and reinforces a willingness to build security systems on faith. The deep-seated desire to make identity-driven systems work is not only emotional, but also economic: the use of SSNs to identify us is inexpensive to the people designing the systems. Other systems might cost more to deploy, might be harder to use, or might be more intrusive on the surface.

One outgrowth of such faith is the fastest-growing crime in America today, identity theft. This term calls to mind the deep sense of violation that many of its victims feel, because we often believe that our identity is our "good name" and one of the most important things about us.

To get a credit card in the U.S., all you need is a date of birth and an SSN that match a record in a database. Criminals who obtain credit take on as much debt as they can and then disappear. The loan is reported to credit bureaus and collection agencies. Collection agencies attempt to track down the person identified, thinking that he is the person responsible for the debt, and a Kafka-esque nightmare ensues.

Credit fraud is not the only goal of identity fraudsters. They can obtain medical care under false names, leading to a risk that medical records will be unfortunately intertwined. They can obtain driver's licenses and passports under false names, leading to repeated arrests of innocent individuals. As more and more systems are based on the notion of identity, the value of identity fraud will grow. Some states have proposed "identity theft passports" to help victims of identity fraud. However, the more we tighten the security of identity systems, the less willing authorities will be to believe they can be compromised and defrauded. This will increase the value of compromising these systems and make victims' lives more difficult.

Addressing identity theft will likely involve some investment in technology, and perhaps more importantly, an understanding of the motivations of the various participants that make it such a problem. One of the themes of this book is using economic analysis to increase our understanding of systems and using that understanding to reach better outcomes. Looking at identity theft allows us to see that all the players behave rationally. That rational behavior imposes costs on everyone who touches the financial system.

Should We Just Start Over?

Describing the many failings of information security could easily take an entire book. We have described only some of the most visible problems. Given the nature of these issues, perhaps we should consider the radical step of rebuilding our information technologies from the ground up to address security problems more effectively.

The challenge is that building complex systems such as global computer networks and enterprise software is *hard*. There are valid comparisons to the traditional engineering disciplines in this respect. Consider the first bridge built across the Tacoma Narrows in Washington state. It swayed violently in light winds and ultimately collapsed because of a subtle design flaw. The space shuttle is an obvious example of a complex system within which minor problems have resulted in catastrophic outcomes. At the time this book was written, the Internet Archive project had 85 billion web objects in its database, taking up 1.5 million gigabytes of storage. During the 1990s, such statistics helped people understand or just be awed by the size of the internet, but the internet *is* undoubtedly one of the largest engineering projects ever undertaken. Replacing it would be challenging.

Even if we "just" tried to recreate the most popular pieces of computer software in a highly secure manner, how likely is it that no mistakes would creep in? It seems likely that errors in specification, design, and implementation would occur, all leading to security problems, just as with other software development projects. Those problems would be magnified by the scale of an effort to replace all the important internet software. So, after enormous expense, a new set of problems would probably exist, and there is no reason to expect any fewer than we have today, or that they would be any easier to deal with.

Much of the usefulness of the internet comes from its open-platform nature that allows new ideas to be developed and incubated. The ability of people to invent the world wide web, instant messaging, and internet telephony stems in part from limited (if any) restrictions on who can do what. Imagine if Internet Service Providers (ISPs) were required by law to collect and keep copies of passports from their customers, or if an official "internet certification board" had to approve new software. The rate at which individuals came online and at which new products were brought to market would be substantially slower. The internet's success depends to a large degree on an open philosophy, which in turn requires accepting a certain amount of insecurity.

In recognizing this reality—that security threats and vulnerabilities will always exist—the question becomes, how efficient and effective can we make our response to those security challenges? If we are not making good decisions today, why not? Creating balanced solutions requires that we understand the true nature of problems. We need good information with which to make the right decisions.

The Need for a New School

Criminals and thugs seek to take advantage of the increasingly electronic nature of our lives. Some crimes occur in the physical world, and others take place purely in the realm of computers. These problems can contribute to distrust of the internet as a medium for commerce and interaction. Problems such as data breaches and identity theft portend doom, but the mere fact of their existence raises important questions. Perhaps our approach to information security is flawed. If it is, a dollar spent on information security is unlikely to be spent well.

We wrote this book not because we are pessimists, but to help coalesce and accelerate the rise of a New School of Information Security. That New School is focused on putting our ideas and beliefs through tests designed to draw out their flaws and limitations. By testing our ideas, we can learn to do better than simply following our superstitions and ingrained beliefs. Such testing allows us to improve on the status quo. The New School is concerned with analyzing on what basis we make security decisions today and with seeking data to support rational decision-making. The New School also believes we can make better decisions by learning from other sciences, such as economics. If there were a single information security community, we could say that parts of the New School have been percolating through it for a while. We hope to help organize, add context to, and extend these ideas into a coherent whole.

Some might say that we are already doing enough, that our current approaches and existing levels of investment are sufficient. If we were to implement new approaches—new training, new technologies, and new processes—would their cost be justifiable? Our answer is that investing in new ways of thinking is inexpensive.

Some security practitioners are beginning to question the received wisdom of their profession. In parallel, the way in which businesses view their information security needs is changing. Organizations want to know how to protect themselves in this new world, but they also want to ensure that they are making security decisions that are both effective and fiscally responsible. A skeptical, pragmatic, and forward-thinking outlook is emerging and will become a new consensus. That consensus is the New School of Information Security.

A psychologist friend likes to say that there are three ways to deal with any problem: you can change it, you can accept it, or you can go nuts. This book is offered in the hopes that we can effectively change some things, accept others, and fail to go nuts.

Index