



THE INFORMATION SYSTEMS COLLECTION

Daniel J. Power, *Editor*

Information Technology Security Fundamentals

Glen Sagers

Bryan Hosack



BUSINESS EXPERT PRESS

Information Technology Security Fundamentals

Information Technology Security Fundamentals

Glen Sagers, PhD
Illinois State University.

Bryan Hosack
Sr Analyst, BI, Reporting, and Analytics
Equity Trust



BUSINESS EXPERT PRESS

Information Technology Security Fundamentals

Copyright © Business Expert Press, LLC, 2016

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations, not to exceed 250 words, without the prior permission of the publisher.

First published in 2016 by
Business Expert Press, LLC
222 East 46th Street, New York, NY 10017
www.businessexpertpress.com

ISBN-13: 978-1-60649-916-0 (paperback)
ISBN-13: 978-1-60649-917-7 (e-book)

Business Expert Press Information Systems Collection

Collection ISSN: 2156-6577 (print)
Collection ISSN: 2156-6593 (electronic)

Cover and interior design by S4Carlisle Publishing Services
Private Ltd., Chennai, India

First edition: 2016

10 9 8 7 6 5 4 3 2 1

Printed in the United States of America.

Dedication

*To Sharon, our kids, and my
mother, for agreeing to a
grand adventure.*

—Glen Sagers

*First and foremost, anything I
do, create or strive for would
not happen without the
loving support of my family,
especially my wife Rebecca.
I would also like to thank
Glen who was willing to take
me along for not only this
ride, but many others over the
course of the years.*

—Bryan Hosack

Abstract

Information security is at the forefront of timely IT topics, due to the spectacular and well-publicized breaches of personal information stored by companies. To create a secure IT environment, many steps must be taken, but not all steps are created equal. There are technological measures that increase security, and some that do not do as well, but overall, the best defense is to create a culture of security in the organization. Such a culture makes each member ask themselves what security implications an action will have. The culture extends from someone at reception deciding to whether to admit a visitor to upper management determining whether a strategic alliance with another firm which links their corporate information systems.

The same principles that guide IT security in the enterprise guide smaller organizations and individuals. The individual techniques and tools may vary by size, but everyone with a computer needs to turn on a firewall, and have antivirus software. Personal information should be safeguarded by individuals, and by the firms entrusted with it. As organizations and people develop security plans, and put the technical pieces in place, a system can emerge that is greater than the sum of its parts. Improving computing security really means education, whether of oneself, one's employees, or one's family. Thinking "security first" may seem paranoid, but in today's world, experience shows that it reflects reality.

Keywords

Information Assurance, Computer Security, Personal Computing Security, Personally Identifiable Information (PII), Network Security, Encryption

Contents

<i>Preface</i>	<i>xiii</i>
Chapter 1: Security and Information Assurance	1
Information assurance and security in the enterprise	3
Interorganizational security	5
Physical asset protection.....	7
Looking ahead	9
Chapter 2: Operating System Security	11
What is the threat landscape?	12
How can a machine be attacked?	13
Patching.....	15
Hardening basics.....	15
Servers in the CIA model	16
Specifics for different operating systems	18
Open source operating systems	20
OSS security	22
Threat model for desktops: disgruntled or careless users ...	23
Rogue applications/malware.....	23
Remote access—intentional	24
Summary	25
Chapter 3: Data Security: Protecting Your Information	27
Cost of a breach	28
Internal versus external.....	28
DBMS security features	29
Types of database threats.....	30
Data quality aspects of information assurance	31
Master data management	32
Data security strategy	33
Summary	34
Chapter 4: Keeping the Electronic Highways Safe.....	35
Using virtual local area networks.....	36
Security concerns with convergence	37

- Virtual private networks, firewalls, and other “secure” networking practices..... 38
- Importance of using secure networks..... 39
- Types of VPNs..... 40
- VPNs for remote workers on unsecured WiFi networks 41
- Firewalls..... 42
- Death of the perimeter 44
- Other firewalls 44
- Other security tools..... 46
- Wireless security..... 46
- Summary 50
- Chapter 5: We Released What?!? (Application Security)..... 51
 - The need for a secure developer!..... 51
 - How are the applications using our data and networks?..... 53
 - Securing the environment, test data, and making the migration happen..... 53
 - Testing applications 54
 - Summary 55
- Chapter 6: Cracking the Code (Cryptography) 57
 - What is it?..... 57
 - Modern ciphers in layman’s terms..... 61
 - AES & SSL/TLS 61
 - How is encryption used to secure resources?..... 64
 - Where should encryption be used? 65
 - Cryptography is not a cure-all 67
 - Summary 68
- Chapter 7: Danger! Danger! Danger! (Penetration Testing)..... 69
 - Internal vs. external testing..... 70
 - How penetration testing is performed 71
 - Volunteer penetration testers..... 77
 - Summary 79

Chapter 8: Disaster Recovery	81
What is a “disaster”?	81
Securing against catastrophe	82
What to consider?	82
Making your DRP a reality	84
Summary	87
Chapter 9: Integrating Your Security Plan across the Enterprise	89
What does the policy contain?	90
To whom does it apply?	93
Developing a security policy	93
Summary	94
Chapter 10: Conclusion	95
Security trends & future concerns	95
SCADA security	95
Big Data	97
Cloud security	99
What is next?	99
Home and SOHO security	101
Backups	109
Personal security	112
Final thoughts	113
<i>Glossary</i>	<i>115</i>
<i>Appendix A</i>	<i>131</i>
<i>Endnotes</i>	<i>147</i>
<i>Index</i>	<i>153</i>

Preface

IT security is at the forefront of overall IT concerns today. Spectacular and well-publicized breaches of company databases, with subsequent theft of personal information, are all too common. Today's businesses need to develop a culture of security, starting from the top down. The costs of repairing the damage after a break-in are rising, and the costs in lost reputation and goodwill may exceed the direct costs. An organization with a secure culture can avoid many costly attacks, and also reap direct financial benefits. These benefits accrue because a company can confidently form partnerships and alliances with other organizations, knowing their systems are prepared for connection to outsiders.

This book is designed to teach the fundamentals of IT security management, and some of the underlying technology. While technology is not the primary focus of the book, effective management requires some knowledge of the tools of the trade. Security products evolve rapidly, but many fundamentals remain the same; the most modern firewalls still filter on the same basic levels, and add additional features. A familiarity with these fundamental technologies will enable understanding of newer tools as they are developed. As a manager, knowing the basics of the tools is sufficient.

The intended audiences for this book are Master's level students, particularly in MBA or executive MBA programs, and practicing managers who have gone through an MBA program. Unlike IT security students and line employees, they do not necessarily need details of each tool, but instead need to see how the various parts of a security scheme fit together. Especially in today's business environment, where a misstep by any employee can compromise sensitive information, a multilayered defense is critical. Implementing disparate security measures according to a comprehensive plan results in a system is greater than the sum of its parts, able to successfully ward off attacks.

A student or manager with basic computer skills should be able to understand the book, however some background in computer networking

would be advantageous. The level of knowledge required would be approximately that required to set up a home network, so well within the grasp of most computer users.

The book is organized by topics, but as with any categorization system, not everything fits neatly. That means there is some discussion of encryption before encryption is really described, and so on. The book can easily be read cover to cover, and enough information is given about novel topics to bring the reader up to speed and point to chapters where specifics are discussed. A reader can certainly skip around between chapters, but the authors recommend reading the first chapter as an overview before too much skipping. The final chapter also deserves special mention, as it's designed to help anyone in their personal security. The measures described there can be implemented by anyone with reasonable PC skills, especially with the help of many excellent online tutorials.

As you read through the book, we would recommend considering not only the examples given of computer attacks and breaches of data security, but also the many that unfortunately appear in the headlines daily. In doing so, try to analyze what happened behind the scenes of each news report. Further, ask yourself "Does it apply to me or my organization?" If so, what can be done to manage that risk? There are four main ways of dealing with risk; reduction, acceptance, transference, and avoidance. Each of these has advantages and disadvantages, a full discussion of which is outside the scope of this preface, but always remember that the goal of information assurance and security is to reduce risk to an acceptable level for an acceptable price. Eliminating a risk is almost never possible, and even if it were, the price would be too high. As you gain experience with security tools and methods, you will start to see patterns repeated in news accounts. Many computer crimes are committed using the same old techniques in use for a decade or more, because we as organizations and individuals do not seem to learn from others' mistakes. We hope this book can change some of that, and that managers and individuals alike will spend the time and money needed to be secure. The good news is that the expense and effort can be spread out by prioritizing concerns, fixing problems as time and money allow.

CHAPTER 1

Security and Information Assurance

People are concerned about data and information security threats. Both internal and external data breaches are a concern.¹ What is security? What is information assurance? How are they the same and how are they different? And perhaps, most importantly, why does it matter whether we call it information assurance or security? The last question is the easiest to answer, put simply, it does not matter much. Information assurance is an overarching construct that includes information security, network security, data security, and a few other “securities” thrown in. In other words, information assurance is the enterprise view of security, highlighting the fact that the reason for all security measures a firm takes is to ensure that vital company information remains secure.

A commonly used model in information assurance is known as the CIA model. CIA stands for confidentiality, integrity, and availability.² These three tenets cover (almost) all the needs of managers to assure the control of company information. Confidentiality entails making sure that only authorized users have access to information. Integrity, or more properly, data integrity, requires that data be accurate and trustworthy, and moreover, that any unauthorized alteration of the data, whether malicious or accidental, can be detected. Availability simply means that authorized users can access information at any time. There are many ways to accomplish the goals of CIA, which will be outlined in this book.

A concept related to information assurance is risk. Risks, and risk management, are part and parcel of information assurance. The goal of all information assurance is the management of risk associated with generating and storing information, whether on a computer, on paper, or in any other format. Bruce Schneier, a security guru, stated that “Security is both a feeling and a reality. And they are not the same.”³ Schneier notes that

true security is mathematical, calculated based on the probability of risk versus the effectiveness of countermeasures. But there is also a psychological component to security, whether our personal security or information security. For example, you may feel very much at risk of identity theft, but feel that your home is relatively invulnerable to burglary. However, these perceptions may not match your real risk of either event. If we misestimate the true risk we face, we will not take adequate precautions or implement proper countermeasures.

Security management focuses on managing and mitigating risk. The goal of information assurance is to correctly estimate the risk in order to get adequate security for a reasonable price. There is no such thing as perfect security, and the strength of a countermeasure should be chosen appropriately for the sensitivity of the asset. An e-commerce firm's database of product descriptions may not be especially confidential and may be protected by only long, complex passwords. Their customer information database, containing credit card information, is much more sensitive and may require both a long, complex password and a fingerprint to allow access.

Deciding how much risk your organization faces is a very difficult process, and classical risk analysis is of little help. Several factors contribute to the fact that classical risk analysis does not work. First, there is usually a many-to-many relationship between protection measures and the resources protected. For example, one firewall might protect your server and multiple desktops. That same server is likely protected not only by the firewall but also by antivirus software, an intrusion detection system, and other security measures. Thus, determining how much of the cost of protection can be attributed to one asset is difficult if not impossible. The other, and perhaps more daunting, challenge is that the likelihood of a certain type of event occurring is largely unknowable. Even knowing what types of attacks the organization faced last year does not predict what will happen in the next year. These and other factors make it nearly impossible to even pin down whether a given investment is "paying for itself" in terms of return on investment.

All is not lost, however. Instead of trying for hard numbers, a firm can be well served by prioritizing assets based on their criticality and sensitivity of the information contained on the systems. Security improvements can

then be prioritized, and in a given year, the most critical remaining assets can be protected, within the allowances of that year's budget. For example, as operating systems reach end-of-life, as recently occurred with Windows XP, and soon with Windows Server 2003,⁴ the threat of attacks against software that no longer receives fixes increases greatly, to say, nothing of simple failures of old equipment.⁵ Therefore, priority should be given to replacing these resources, then turning attention to the next-most critical assets.

Information assurance and security in the enterprise

All companies face variations on the same threats, regardless of their size or industry. Every firm faces both internal and external risks, as well as risks created by connections to other firms, whether suppliers, consultants, or partners. Firms also face physical security risks that impact their information technology (IT) systems.

Internal security has many components; however, one that cannot be overlooked is the concept of insider threat. Insider threat is simple enough conceptually; those on the inside of the organization can represent the biggest threat to its security. The problem is that these same individuals are also the biggest asset to the firm. This dichotomy makes it very difficult to police those who have the most knowledge and therefore could do the most harm. Perhaps the most dangerous are those individuals who manage IT and security; they know the most about the systems and ways around them. Recent events, including Edward Snowden and others delivering classified documents to various "leak" websites and media outlets, only serve to underscore the magnitude of the threat.⁶

What can be done to manage the insider threat? There are various small measures that can be taken. Discussing all of them is outside the scope of this chapter, or indeed, this book, but a list of a few is appropriate.⁷

1. Monitor logs. Log monitoring software looks for patterns indicating improper actions. Monitor logs of critical assets and actions of critical employees more closely.

2. Rotate job roles. Rotation makes it harder to carry out complex attacks.
3. Use separation of duties. Those who can make changes should not be able to approve those changes.
4. Organize data according to sensitivity. Grant access to sensitive data to only those who “need to know.”
5. Enforce least access. Give only the bare minimum access for employees to do their job, no more.

External threats to the organization may be myriad, but the majority are common to all organizations. The classical, or perhaps more accurately stereotypical, “hacker” is mostly a Hollywood construct. There are certainly antisocial introverts bent on wreaking havoc, defacing websites, and gaining “cred” with their peers, but they are likely not the most dangerous. While there may be a thrill in placing electronic graffiti, the real money is in money. Increasingly, criminals are the main enemy. Blackmail, theft, extortion, and similar crimes may be easier to accomplish in the virtual world than the physical, but the crimes themselves have not changed in thousands of years. Criminals and organized crime represent a real threat to today’s firms. Other threats include competitors, who may engage in industrial espionage, and even national espionage. Finally, malware such as viruses may not be directly aimed at your company, but there are many automated attacks looking for easy targets. In fact, 92 percent of breaches can be attributed to nine basic patterns, according to Verizon’s annual report⁸:

1. Point-of-sale intrusions
2. Web application attacks
3. Insider privilege misuse
4. Physical theft or loss of computing assets
5. Miscellaneous human errors such as e-mailing confidential information
6. Crimeware (such tools as bank information theft malware and so-called ransomware, which locks files unless a ransom is paid)

- 7 Card skimmers (which steal credit/debit card numbers as the card is swiped at a point-of-sale device)
8. Denial-of-service attacks
9. Cyberespionage

These threats run the gamut of ways that attackers get to confidential information. As can be seen, at least three of the nine are directly related to obtaining money, and several more likely lead to information that can be used to extort money from the victim.

Interorganizational security

Today's organizations engage in partnerships and supplier/client relationships with many other organizations. While this practice is nothing new, the last decade has changed those relationships in a very real way. Electronic data interchange (EDI), also known as business-to-business (B2B) or electronic order systems, and the related concepts of "just-in-time" ordering and delivery mean that automated machine-to-machine (M2M) transactions flow at an unprecedented rate. A large company in the 1990s might place thousands of orders a week with suppliers, and some automation was in place, but most orders were handled by a human at some point in the process. Whether a human faxed the order, or entered it into a computer system, a sanity check was in place. Today, many orders are simply placed and fulfilled automatically. If a factory's automated inventory system is tampered with, too few or too many key components for the company's flagship "Widget Y" will be delivered, stopping production or causing logistical errors when there is no place for the excess parts.

The dangers related to EDI and M2M communication do not stop with ordering systems. Many B2B systems share private data with partners, and firms must be able to trust that only the correct information flows between partners and that it is only seen by authorized parties in the other firm. Consider the healthcare industry. A doctor's office, a lab, a pharmacy, a hospital, and an insurance company may all have information about patient James S. His doctor has a comprehensive history of all visits, his own diagnoses, records of tests, and a list of prescriptions that he takes. The lab needs only certain information to

positively identify James when he comes in for a test, along with data indicating which tests to perform, but not information on previous diagnoses. The pharmacy needs to know what medications are prescribed, but does not need lab results or a history of all the drugs James has taken in the past. The hospital needs much the same information as the doctor, but many of the doctor's previous diagnoses are immaterial to the current illness; last year's flu does not impact a gallbladder problem this year. Last, the insurance must know what has been diagnosed, and what tests were performed and medications dispensed in order to pay the providers. The Health Insurance Portability and Accountability Act (HIPAA) mandates that only relevant information be shared among parties; even if a lab wanted historical data about a patient, they likely could not obtain it without the patient's written consent. If the information of James S. is disclosed to an unauthorized party, HIPAA provides for financial penalties against the discloser.⁹

Besides ensuring that only the right partner firm gets access to information, businesses need to be sure that within the partner organization, only authorized individuals have access to data. In our healthcare scenario, the doctor needs to be sure that the orders sent to the lab can be read by only lab techs in order to perform the tests, but that a receptionist, for example, would not be able to access a full history of all tests performed on a patient. This would avoid the scenario of a receptionist giving away James' medical history to a reporter when he decides to run for public office, or an insurer trying to deny claims based on a preexisting condition. Before entering into B2B relationships with other companies, a firm should exercise due diligence in ensuring that the partner's information assurance practices, policies, standards, and procedures are in line with their own and any regulatory requirements.

As with any confidential data, a firm must ensure that B2B data is passed securely between partners. Two basic modes of securing documents can be used; a firm could encrypt the documents before transmission, and the partner would decrypt them, or the communications pipeline could be secured from end-to-end. Both approaches have advantages and disadvantages, discussed in Chapters 4 and 6.

One other avenue of attack that is sometimes overlooked in security is making sure that outsiders employed by your firm are vetted. Whether hiring a consulting firm or a janitorial service, an organization must be

sure that adequate background checks are being performed on employees by the other organization.¹⁰ The depth of the background check required will vary; a janitorial service cleaning only public areas of the firm's buildings may be less of a security risk than one hired to clean private offices. Similarly, vendors should be vetted before being allowed into private areas; and unexpected visits from vendors (or worse, someone unknown wearing a vendor's shirt!), should be viewed with suspicion. Receptionists and others should be trained to make a phone call to confirm identity and purpose of unscheduled visits or unknown people. After all, it is quite easy for a visitor to take pictures of confidential documents with a camera phone.

Physical asset protection

IT assets take many forms. The information stored on a machine is often much more valuable than the computer itself, but that does not make the server cheap to replace. Further, physical access to the server may defeat many electronic controls; someone standing at a keyboard in the data center does not have to get around firewalls to get in. This should not be construed to mean that insiders are the only threat to physical assets. If a firm does not properly secure IT assets, others may be able to get access. Someone who steals an entire server, and then has unfettered access to it for days or weeks, could retrieve a great deal of information, to say nothing of the cost of a replacement server or downtime suffered as a result of the theft. The aforementioned impostor vendors may be able to remove physical documents or media, or simply plug a thumb drive into an unused PC and copy documents. A copier repairman using a laptop to "diagnose the machine" might, in actuality, be plugged into the network port used by the copier and may be reading traffic on the network or accessing shared files.

How can a firm avoid these nightmares? A firm can avoid these by physically protecting its assets. These protections include, but are not limited to¹¹:

- Lock the server room door. It seems simple, but a simple locked door will stop many unauthorized visitors. Locks can be mechanical or electronic.

- Surveillance cameras. They are cheap and effective as a deterrent, but footage must be recorded and reviewed.
- Train employees to not allow “tailgating.” Every person going into a secured location must individually sign in or use his or her swipe badge or other credentials. No exceptions can be made, and your IT security policy (refer Chapter 9) must contain penalties for violations of security protocols.
- Do not allow nonemployees into certain areas of the building, at least unescorted.
- Lock office doors automatically when not occupied. If the door is shut, it should be locked from the outside. This prevents unauthorized snooping or use of another’s workstation.
- Secure areas should not allow the use of removable media or recording devices, including phones and media players that could be misused that way. Further, Universal Serial Bus (USB) ports can be disabled on sensitive machines, either electronically or by simply filling the port with glue.
- Data centers should be located in the interior of the building, have proper (not water!) fire suppression, raised floors, and be away from overhead water or sewer lines.
- In highly secure facilities, such as data centers, guards may be appropriate to monitor entry.
- Alarm systems. Install fire, motion detection, burglar, glass break, and other sensors as needed.
- Fences and other physical barriers. Retail stores have large metal and concrete posts in front of the entry doors to stop vehicles from ramming through; does your facility not deserve similar levels of protection?
- Recovery and remote wipe software on mobile devices. They are easy theft targets, and can contain passwords, documents, and other valuable organizational information.

Ultimately, with all protection measures, remember that the goal is to ameliorate the risk to a sufficient degree for a cost that matches the sensitivity of the asset. If a company does not have a large data center, it would be ludicrous to hire a security guard to sit outside the server room door. It would not be unreasonable to install a \$1,000 electronic lock system to keep unauthorized personnel out, nor would it be unreasonable to expect IT personnel to take care of cleaning that room so that no janitor is ever allowed inside after hours. An alarm system is likely already part of a factory; adding fire and motion detection alarms in the server room is likely a small additional cost. You already train your employees in policies and procedures; why not include a module on physical security?

Looking ahead

In the following chapters, more details about many aspects of security are presented. The overarching theme throughout the book is to protect the assets, whether electronically, physically, or by training personnel. This strategy is known as defense-in-depth. This means setting up a combination of defenses such that an attacker must breach each in series in order to get access to the target. Defense-in-depth requires that each asset be protected by multiple measures. No matter which facet of information assurance we examine, the goal is always to present ways to ensure the CIA of the information and the systems that house your organization's most valuable assets.

CHAPTER 2

Operating System Security

The operating system (OS) of a computer, or of a tablet or phone, for that matter, is the basic software that transforms the machine from an collection of electronic components into a usable device. Currently, Windows, Mac OS, Linux, iOS, and Android are the most familiar operating systems. Most operating system software is written by a handful of companies, except for Linux, which is written by volunteers, although even there, various companies oversee much of the development.

Operating systems have undergone great changes over the last 30 years. In the early 1980s, PC operating systems were fundamentally designed for only one user. The concept of security was mostly ignored, and no version of Disk Operating System (DOS, used on most PCs in the 1980s and early 90s) had even rudimentary support for separate users, much less for passwords.¹² Anyone who could start up the computer could access all files on the hard drive, and while a few applications could set passwords for their own files, nothing was done at the system level. Fast forward to the first “modern” versions of Windows, such as Windows 95 and 98, and we see the concept of users and passwords, but fundamentally, all users could still see all other users’ files. With Windows NT and its successors, such as Windows 2000, XP, Vista, 7, 8, and 10, vast improvements have been made, allowing a user to lock off access to his or her files from all others. Phone and personal digital assistant software has undergone similar changes, to the point that on modern iOS, and to a lesser extent on Android devices, passwords or Personal Identification Numbers (PINs) protect the device, and one application cannot even access files created by another application.

So, why do we need OS security? Specifically, to protect against unauthorized users gaining access to files belonging to others. This fulfills the confidentiality requirement of the confidentiality, integrity, and availability (CIA) triad, as well as preventing some unauthorized changes,

giving a measure of data integrity. This applies to files on a single PC that might be shared with others, as well as to files on shared network drives that are used by many. Further, the operating system should protect against applications misbehaving and claiming privileges of other users, such as a regular user being able to use an application as an administrator. As a case in point, a famous antivirus software of the late 1990s had a bug that allowed a user to become an administrator. When a virus was found, the user was prompted to quarantine it, and could browse to the folder where they wanted to put the quarantined file. This file-browsing window had administrative permissions, allowing a savvy user or attacker to browse to and then open any application with system-level privileges. These types of issues have plagued various operating systems over the last three decades, but are slowly going away. As vendors find and fix bugs, and take a more proactive stance against security vulnerabilities, operating systems exploits have decreased.

Before discussing the changes to operating systems, a definition is in order. The term threat landscape refers to all possible security threats to a system. The threat landscape typically includes, but is not limited to, individuals who have or might gain access to a system, including hackers and insiders, software threats such as malware, and anything else that might allow unauthorized use of or damage to the system. In short, the threat landscape includes all the threats enumerated in the Verizon report mentioned in Chapter 1, and many other threats.¹³

What is the threat landscape?

Over the last five to ten years, the threat landscape has changed.¹⁴ A decade ago, the biggest threat to an information systems was likely a virus or Trojan horse that could attack the core operating system. Operating system vendors have steadily improved their systems, and today, traditional operating system viruses, while still circulating in the wild, have far less significance. Both the amount of damage done by a given virus and the number of attacks has decreased. While some spectacular virus outbreaks have been seen in the last few years, the strides made by Microsoft, Apple, and other software vendors to protect their operating

systems have made it difficult for malicious operators to use viruses as hacking tools. Antimalware software also shares some credit in preventing these threats.

Today, the threat landscape has changed to the point that applications are one of the most preferred targets for attackers. There are at least two reasons behind this phenomenon, perhaps the chief reason is that there are simply many times more applications than operating systems; tens of thousands versus a mere handful. This multiplicity of applications gives attackers many targets of opportunity. A second reason that the majority of attacks take place on applications is that most are written by small software development firms or individuals. Such developers may not have the training, time, or other resources to make sure their software is secure. The subtle nature of software bugs means that testing in dozens of different scenarios is needed to expose the bugs, something a small company may not be able to feasibly do. This stands as a lesson for companies doing in-house development of software for their own use: Do not trust it! In planning custom software, time and money should be allocated to design security in from the planning phases, and allow for the testing process.

How can a machine be attacked?

Even when an operating system may be directly attacked, not all attacks are equal. Whether a bug in software is exploited or a virus is run on a machine, there are two major classes of attacks. The first is a local exploit. In a local exploit, the attacker must be present at the physical machine. The second type of attack is a remote exploit, in which an attacker can attack the system from afar, over a network, such as the Internet. In both classes of attack, typically the attacker has the privileges of the logged-in user.

It may seem that remote exploits are more dangerous or potentially damaging than local exploits, and to some extent, this is true. However, remote exploits are comparatively rare by comparison to local exploits, and often, your systems are protected from them by firewalls. The local exploit, then, may actually be more dangerous. An authorized user,

meaning an employee, can utilize local exploits. Disgruntled or curious employees can find these holes in the system and probe further into areas where they are not allowed. The bug present in the antivirus software mentioned at the start of the chapter is a prime example of this.

This brings us to the definition of hacking, which is: “use of a system in excess of authorization.”¹⁵ Certainly, your users are authorized to log in to your computers; the system is there to help them do their job. When a worker goes beyond this authorized use, prying into other accounts, or getting system-level privileges, this can be considered hacking. A company likely has many more authorized employees using the machines daily than there are outside hackers attacking the system. If even a small fraction of those workers are less than satisfied with their job or pay, they may try to get past security measures. These factors combine to make the local exploit much more dangerous.

If good practices are followed, and all users are limited users, then a successful local or remote attacker will have these limited privileges. In other words, by compromising that account, they will be able to access the applications and files that user would have access to, but not other users or system-level privileges. After gaining access, then, most attackers will try to perform a privilege-escalation attack, that is, to use other bugs present in some systems to become the administrator of the computer. If successful, this can be devastating.

In September 2014, Home Depot was the target of such an attack, which led to the disclosure of 53 million customer e-mail addresses, and 56 million credit card accounts were also taken. In order to get the information, the attackers compromised the username and password of a vendor, giving access to the network, but only with that vendor’s privileges, which certainly would not have included customer information. The attacker then acquired elevated privileges by exploiting flaws in other devices on Home Depot’s corporate network. This elevated access was used to navigate further into the network and install custom software on the self-checkout kiosks in the United States and Canada.¹⁶ While Home Depot did not disclose exactly which flaw was used to gain administrative rights, most corporate networks, with their mix of old

and new operating systems and applications, provide multiple targets. This same attack pattern has been carried out in many cases, including the breach at Target stores the previous year.

Patching

If a modern operating system is so secure, why do we need to worry about securing it? The answer is simply that a modern OS is much more secure than previous versions, but not inviolable. Even the best software has bugs, and as these are found, vendors develop solutions to protect against the bug and release the solution as a patch. Patching is simply applying these vendor-supplied fixes to PCs and other devices. Microsoft, for example, releases many patches on the second Tuesday of every month, a day called “Patch Tuesday” in the IT world. When patches are released, a system administrator will apply the patches to a test system, and work with it for few days or even weeks to test that the patch works, and that it does not break something else. Once this requirement has been satisfied, the patches are rolled out to production systems. Patching is an ongoing first step in securing an operating system. In fact, as soon as an operating system is installed on a computer, the first step to hardening that operating system is to apply all the patches, often called service packs, to the computer. For a home user, this is taken care of fairly automatically, with Windows Update or the App Store on a Mac offering to install available updates. The process is more manual on a server, but must still be done.

Hardening basics

Once the operating system is installed and patched, other hardening steps need to be taken. These steps ensure that common vulnerabilities are resolved before the system is put into use.¹⁷

1. Choose secure configuration options
 - a. Choose good passwords
 - b. Get rid of all default passwords
 - c. If any unneeded users are set up, delete them

- d. Make sure all users are limited users, not administrators
- e. Change configuration defaults as needed
2. Install only absolutely needed software
 - a. A file server does not need web server software installed
 - b. No productivity software
3. Patch installed software
 - a. Patch the operating system, and where appropriate, set up automatic updates
 - b. Patch any installed software
4. Set up users and groups
 - a. Give least permissions needed to do their job
 - b. Create a good password policy for users to follow (see Chapter 10)
 - c. Deactivate, but do not delete, unneeded users as they leave company
5. Backup, Backup, Backup!
 - a. Policy governing what should be backed up, when, and how long it should be retained
 - b. Offsite backups are a must

This list is not all-encompassing, but covers most holes in a brand-new system. Not installing extra software is vital, as all software has bugs, and software which is not installed cannot be exploited. Backups are equally important, and when deciding on backup policies, legal and regulatory requirements for retention must be considered.

Servers in the CIA model

Most medium to large firms today have one or more servers in place to centralize their computing. Servers are simply machines that have enough processor power and memory to allow many people to simultaneously utilize their resources. In other words, servers represent a computing resource that contains large amounts of data, and often the most sensitive data. As such, servers represent a temptation few hackers can withstand. Since a server may contain the keys to the kingdom, protecting them by patching and hardening is vital. To return to the CIA model of information assurance, all three must be ensured for a server.

There are many facets to protecting confidentiality, and not all will be discussed here, but some of the basics are to patch the system, activate firewalls, and turn on encryption. Patching has been discussed previously. Firewalls are discussed in more detail in Chapter 4, but briefly, besides the firewalls that protect the whole network, a server should have its own firewall. Finally, encryption can be used to protect data at rest and in transit. Encryption is discussed in much more detail in Chapter 6. When encryption is used to protect the whole hard drive, confidentiality for anyone but authorized users is accomplished. When data in transit across the network is encrypted, it is protected against those who would “sniff” network traffic to listen in on transactions between machines. All these protections should be implemented on servers.

The integrity of data on a server must be unquestionable. This means that data cannot have been altered while stored, while being processed, or in transit; or that such alterations are detectable. Whether the changes were intentional due to attackers or simple errors does not matter, altered data cannot be trusted. Data transfers can be protected by checksums, and files can be protected both by checksums and journaling file systems. Journaling is also used extensively to protect database transactions. Briefly, checksums work by calculating a value based on every bit of the data. This number, unique to that particular data, is stored. At some future time, the file can be checked again, and if the numbers match, the data has not been changed. Journaling means that the system, whether a journaling file system or a database, keeps track of all changes made to stored data in a separate area. In the event of a system crash or power failure, the changes that were not saved are still available in the journal and can quickly be replayed to give a complete file.

The last facet of CIA for servers is availability. If a server is unavailable, productivity and profits suffer. Workers cannot accomplish tasks, and customers cannot make purchases. In a large e-commerce firm, server downtime may cost hundreds of thousands of dollars an hour. In these scenarios, it is worth almost any cost to ensure high availability. Like most aspects of computing, there are multiple ways to accomplish this goal, but most fall in the category of overprovisioning. Overprovisioning is simply having more computing resources available than needed at any given time.

In the simplest case, two identical servers are purchased and installed, and kept up-to-date on patches simultaneously. Further, data on the two systems is kept identical, and one system is configured to take over in the event failure of the first system. This is known as failover, which is a fairly simple way of ensuring constant availability. A second method, more expensive but often more reliable, employs several to hundreds of identical servers. These machines are arranged as a group called a cluster, and a separate machine determines which server in the cluster handles an incoming request. These clusters can handle millions of requests per minute, and most clustering systems scale well, meaning that if performance suffers, one or more new servers can be added to alleviate the overload. Either clustering or failover increases server availability, but at differing costs and degrees of reliability.

Specifics for different operating systems

Thus far, this chapter has handled generalities of patching software and operating systems. We turn now to specifics of different operating systems, and how they need to be managed for effective patching and security.

Windows Server is Microsoft's offering in the server space. The various editions of Windows Server can fulfill almost any role required in the modern enterprise, from web servers to domain controllers to database servers. High-availability configurations are also possible, with clustering and failover capabilities. Windows server has been touted as the easiest-to-learn server software, thanks to the familiar graphical user interface used on Windows workstations. This, however, has likely led to some poor implementations, since the ability to point and click does not necessarily equate to being a good system administrator. The skills needed to securely configure and administer any server are very different from those needed to compose and format a document in a word processor.

The general principles of setting up and hardening a server explained in the first half of this chapter apply to Windows Server. Specifically, all unneeded services and software should be removed, users and groups should be securely administered, backups configured, and a host-based

firewall (Windows Firewall) should be set up. Additionally, some Windows-specific steps remain after these general tasks are completed. First, in the case of domain server, proper Active Directory schema design is necessary. While discussion of how to achieve this is far beyond the scope of this book, it should be noted that making changes to an Active Directory database is more complex than designing it properly to start. Second, also in the case of a Windows domain controller, the use of Group Policy Objects (GPOs) is strongly recommended. GPOs allow an administrator to permit or deny access to almost any part of the computer hardware, software, or network. For example, a GPO can be configured to prevent a user from turning off antivirus software. Another GPO can be used to disable USB ports on sensitive workstations or servers, making it more difficult for a malicious user to steal data. GPOs are available to accomplish almost any control an administrator wishes to set, and can work in both online and offline scenarios, meaning that even laptops taken on the road away from the domain controller can still be controlled.

A final aspect of Windows patching revolves around Microsoft's Patch Tuesday. As previously mentioned, on the second Tuesday of each month, Microsoft releases many of their patches. Some emergency patches are released on different days, but most are released on that day. When system administrators get these patches, they will typically test them for a few days to weeks, or sometimes months in the case of servers, and then push them out to clients. This push operation, facilitated by software such as Windows Server Update Services, allows the administrator to force the update via a group policy, and push only the vetted, approved patches out. This allows for "automatic updates" without relying on Microsoft's like-named service, and without the likelihood of something breaking in case of a regression.

Apple has largely left the server market. The Mac OSX Server edition operating system remains available for download, but they no longer produce specific hardware for servers, meaning the software must run on machines designed for use as workstations, which are not usually as powerful as typical servers. If Mac OSX Server is employed in an organization, it is a robust operating system, but the core principles of

hardening apply. Specifics for Mac OSX largely consist of turning services on or off to allow it to communicate with Windows or other operating systems, as desired.

UNIX is a family of operating systems that is often used for servers, and occasionally for scientific or engineering workstations. Various products can use the UNIX trademark, including Mac OSX, the Berkeley Software Distribution (BSD) family of open source operating systems discussed in the next section, and products such as Oracle's (formerly Sun Microsystem's) Solaris, HP-UX from Hewlett Packard, and AIX, also from HP. These operating systems, unsurprisingly, share the same basic hardening steps as other computers. There are specific additional requirements for each of the members of the UNIX family, but they will not be discussed in detail here.

Open source operating systems

Another type of operating system exists besides the proprietary Windows, Mac OSX, and UNIX systems that most users are familiar with. These operating systems, collectively called Open Source Software (OSS) operating systems, include Linux and the BSD family, FreeBSD, NetBSD, and OpenBSD. All are UNIX-like, meaning they share the same file naming and organization conventions, similar ways of configuring, starting, and stopping services, and similar user interfaces. More uniquely, volunteers develop them all. Linux is a well-known example of this. Linux is developed by a core group of volunteers who are not directly remunerated for their time writing the software. These volunteers have various motives for participating, but the software produced represents more or less the collective will of the developers. The software produced in this fashion is generally of high quality, and more companies are taking note of OSS, and implementing it in the enterprise.¹⁸

A final interesting aspect of open source is that it is “free,” sometimes expressed as “free as in freedom, not free as in beer,” or “libre, not gratis.” The code is copyrighted, but the licensing terms allow everyone to use it on certain terms. Because of this, firms who make money on open source products often do so by selling service and support, rather

than the software product itself. In other words, while OSS is often freely downloadable (*gratis*), the real freedom comes from the way it can be used.

The *gratis* aspect of OSS has been a draw for many to download and use the products. After all, licensing fees for software are often a large cost to an organization. However, what does free really mean? For some OSS projects, such as the well developed and thoroughly tested Firefox, Apache, or Linux, the degree of help or technical support needed is minimal. In other cases, a fair degree of support is required. Two main avenues of support are available; commercial and from other users via the Internet. Commercial support, where available, may be a good choice for an enterprise user. For the Linux operating system, for example, support is available from a number of vendors who sell the product with support included. Dell, HP, and others all sell servers with Linux preinstalled, and support it to the same degree they support machines running Windows.

From a managerial standpoint, then, it likely matters little which operating system is purchased. As long as support personnel have some degree of experience with the chosen OS, any additional help can be obtained from the vendor's technical support lines. For products with no available commercial support, or if the firm does not wish to pay for them, abundant technical support is available on the Internet for almost every product. This support is obviously freely available, but only free in the sense that no additional monies must be paid out, not in the sense that employee time must not be used. In other words, OSS is not truly free (*gratis*) for an enterprise. But, for that matter, neither are proprietary operating systems. The choice of which to use depends largely on what personnel are best-trained on, but many firms are deciding to skip the licensing costs and move to OSS solutions, since the support costs are comparable. Initially, many firms start this transition on the server side, and some have moved to running Linux on the desktop.¹⁹

OSS security

Like any operating system, or any software, for that matter, open source operating systems require patching. There are several differences between patching Windows and Linux. First, each Linux vendor or distribution manager releases patches on a different schedule. Typically, once a bug is discovered, such as the “Heartbleed” bug discovered in 2014 in the OpenSSL encryption package, the project’s developers patch it. Once this “upstream” patch is released, each vendor takes the new code, compiles it with their specific tools, against the specific current version of their OS, and after testing, releases it to the general public. This process goes fairly quickly, especially for severe bugs. When a specific firm or user applies the patch depends on their own schedule; for an individual user, the update manager software checks every few days for new patches, and prompts the user to download and install them. In a firm, the network administrator likely performs the typical process of downloading the patch, installing it on a test system, and once vetted, releasing it to users, just as in a Windows environment.

A second major difference between Windows and OSS operating systems updates is that since a Linux distribution includes both the operating system and a collection of tools to go with it, the update manager logically checks for updates to all software installed with the system. On a typical user’s workstation, the software update manager would check for Linux kernel updates; updates for the web browser and plugins; and for music players, office suites, and even games if installed. All pending changes would be presented to the user, who would merely need to click “Install” to update the system, or this approval and installation process can be centralized.

To summarize, hardening servers, and keeping them patched and up-to-date is a formidable task. It is not, however, one that can be put off. The threat landscape has changed in the last decade, but attackers have done nothing but increase their attempts to compromise systems. In the final analysis, it really matters little whether the hacker got in through a flaw in the operating system or through a compromised spreadsheet downloaded by an employee. If the hacker was able to steal customer credit card numbers either way, the cost to repair the damage will be the same.

Threat model for desktops: disgruntled or careless users

The threat model for a desktop PC largely centers on users. The employee works at a given workstation day in and day out, and has large blocks of unsupervised time to plan and carry out an attack. While it may seem unsavory to think of our biggest asset as our greatest threat, experience has shown that many cyber-attacks, to say nothing of common crimes such as embezzlement, stem from insiders. As previously discussed, a disgruntled employee may actively seek chances to break into information systems, or careless employees could install applications that compromise corporate security. Careless users provide for about 25 percent of all incidents, and malicious users another 18 percent, according to Verizon.²⁰ We cannot overlook this chink in the corporate armor.

Rogue applications/malware

Before discussing how to prevent users from installing rogue software, some definitions are in order. Several different types of software pose a threat to corporate (or indeed personal) computer security. These threatening applications are called “malware.”

Malware is any type of malicious software. The category includes viruses, worms, Trojan horses a.k.a. “Trojans,” and spyware, as well as various blended threats that do not fit well into the above classifications. Viruses, like biological viruses, must have a host to operate. They work by attaching to a program on the computer, and when it runs, they spread themselves to other installed applications. Worms, on the other hand, can spread by themselves and often seek out server software to infect. Trojan horses, as the name implies, masquerade as innocent software, but in fact have a malicious payload, and spyware often has the similar property of masquerading as useful software, but secretly tracks a user’s web browsing or other activities. All malware shares the commonality that it attempts to cause damage to a system or steal information. Malware may be countered to some extent by antimalware software, but unfortunately, most of this software is only about 50 percent accurate in correctly identifying damaging software. To supplement the low detection rate, companies should

train their users not to open e-mail attachments they were not expecting, even from someone in the company. Employees should also be trained not to install software of any kind, and technical measures to enforce this policy, such as not giving administrative access, are required.

Another common type of malware is a Remote Access Trojan (RAT). This software acts innocently, as all Trojans do, but the payload allows a hacker remote access to the machine. The attacker takes over control of the user's system, and may then install other software to make the machine part of a so-called "botnet," which is simply a large collection of computers that act together when commanded to do so by the hacker. Botnets have been implicated in some of the largest attacks in computer history; since by having thousands of machines available, a determined attacker can overwhelm the defenses of any website or server.

Pirated software is another common security threat. Regardless of the legality or perceived morality of software theft, the simple fact is that much pirated software contains malware of one sort or another. The cost in damages from that malware may be high and is certainly too high to justify not paying for a licensed copy of the software. Beyond the malware issues, the legal issues of installing unlicensed software can include high costs in fines and other penalties. It is simply not worth the potential costs to install illegal software.

Remote access—intentional

Remote access provides workers with access to their work desktop from other locations, such as home or at a client's site. This type of access is intentional, and in no way related to RATs. Such access is often provided by Microsoft's Remote Desktop or Terminal Services, Citrix, or on other platforms by programs such as Virtual Network Computing (VNC). All fulfill the same basic need, allowing a user to log in to their PC from afar. While these programs serve a useful purpose, they also serve to broaden the threat landscape, by allowing another point of access to the company's IT resources. As such, their use should be carefully controlled.

As a first control step, the firm's firewall should block the ports used by these programs, and access should only be allowed after a user makes a Virtual Private Network (VPN) connection. Second, the group of users allowed access should be carefully controlled; users who do not work from home do not need a remote desktop. Finally, PCs should be monitored for unauthorized installations of any remote access software; even if software installation is blocked, it may be worth scanning for this specific type of software in case a user is able to install it. Remote access software is potentially very useful, but also potentially dangerous. The dangers should be mitigated by good management of the software and authorized users.

Summary

Operating systems are the foundation of our digital enterprise. Whether a computer, a tablet or phone, for that matter, it is the basic software that allows the machine to go from a collection of electronic components to a usable device. While great security strides have been made in OS security, hardening is still a vital part of the initial configuration of devices. Monitoring and managing who has access to the device or service to maintain CIA is one of the primary goals of a security team within an organization. Protecting against threats from the server side by applying patches or establishing policies for both use and physical access is key to creating a secure organization and threat management. Different operating systems have different base levels of security, but all can be made more secure. Following a checklist of steps each time an operating system needs hardening will aid administrators in consistent provisioning.

CHAPTER 3

Data Security: Protecting Your Information

Data is the foundation on which businesses are built. Businesses continually generate data on customers, employees, suppliers, partners, or systems via transactions, automatic logging, and aggregating historical data for analysis, rarely discarding anything along the way. In fact, the information a company or organization generates, analyzes, and processes is one of its most valuable resources. As technologies advance and we can evaluate more and more structured (i.e., data stored in databases) and unstructured (i.e., documents and streaming feeds) data, protecting it becomes even more critical.

The data environment consists of a Database Management System (DBMS) that resides on a server or multiple servers in our organization. The data from our business applications, webpages, and internal/external processes create, update, and delete records. The data is stored in a structure that allows us to identify where these records are and imparts the ability to quickly retrieve them for display and manipulation within our applications. This is our structured data environment. The most common transactional system is a relational database, though hierarchical and object-oriented databases, among others, exist in modern organizations. Structured Query Language (SQL, “Sequel”) is used to build, manipulate, and manage data in our relational DBMS. The data can be retrieved and processed using SQL in applications or preprocessed using stored procedures on our data server. The database administrator (DBA) is in charge, in whole or in part, of managing the DBMS, database structure, and data stored within the database.

Increasingly unstructured data is becoming as important as structured data to an organization. Unstructured data resides in our file storage, desktop machines, and log files for all manner of hardware or software services. This data can account for tens or hundreds of times the amount

of structured data stored by a firm. NoSQL (“No-Sequel”) and big data resources such as MongoDB and Hadoop are used to retrieve, organize, and churn data into valuable information.

In this chapter, we will focus on data security breaches and methods for securing against threats to our structured and unstructured data. Recovery of valuable data, in the face of both man-made and environmental catastrophes, is covered in a later chapter.

Cost of a breach

Historically, the cost and number of breaches have been difficult to determine, yet a consensus can be reached that the number of breaches that are visible to the public are increasing rapidly and the severity of damage of such breaches is substantial. In a 2007 study of the cost of a data breach, the Ponemon Institute found that the average cost of a data breach per record was \$198 across 35 companies from varying industries. Breach size was from 4,000 to over 125,000 records in 2006.²¹ By 2013, the report had expanded to include 277 companies across the globe, and cited the average cost of a record as between \$136 and \$199 in the United States.²² To avoid skewing the results, breaches over a 100,000 records were not used in the study! As that last statement sinks in, we see the global impact of data breaches is more concerning today than it was six years previously.

Internal versus external

As you consider data security in your organization, you will examine security from not only an external threat perspective but an internal loss perspective, as well. External threats will primarily be intentional threats in which vital organizational data is targeted as well as leaving mechanisms in place that give the ability to retrieve data in the future or over the long term. Internal threats can be either intentional or unintentional. Similar approaches can be applied to intentional threats, but care should also be given to unintentional threats. Unintentional threats can take many forms, including data loss due to negligence, application failure, and poor design.

DBMS security features

From an internal perspective, an entire suite of resources is required to protect and monitor your data. The key to effective data security is to have policies in place to govern the data environment. The chance of stopping every breach is a near impossibility, but being aware and monitoring your data for unusual activity will limit the impact of a breach. This in turn requires you to know and understand what is typical for your organization and have the necessary technical and process elements in place to monitor your data and the networks that transport it.

Using the features of the DBMS will allow you to restrict users, and monitor and log information. Limiting access to an “as-needed” basis and reviewing these privileges often is important. Using the password features and having a policy in place for changes is as important as for employees’ desktop machines. Other database tools provide additional features such as ensuring inputs to the database are “clean” or free from harmful syntax. Strategies should include policy specifics on:

- storing the data,
- accessing the data,
- managing the DBMS,
- monitoring the transfer of data,
- monitoring access to data environment, and
- consistently evaluating and reevaluating policies.

Putting actionable plans in place is necessary to minimize the impact of a breach and damage from lost data. Testing the plans and policies helps to secure data against both internal and external threats.

As with the servers discussed in Chapter 2, protecting the data residing on the DBMS means limiting access to the hardware. Many organizations use a combination of scan card logging, biometric devices, live video streams, and other means to secure the servers and physical resources. Just as we see in spy movies and cat burglar capers, providing ample protection to our organizations data and network servers should be part of every organization’s IT security plan.

Types of database threats

SQL injection can occur when we provide direct access to our data through web-based forms. Since our queries can directly interact with our DBMSs and data and these queries have a characteristic structure, they are exploitable. The structure of a relational database SQL query is such that we “SELECT ... [columns] ... FROM [tabular data objects] ... WHERE ... [conditions are evaluated];” (the semicolon terminates our query). Similarly, we INSERT, UPDATE, and DELETE, not just select data. Data objects such as tables are CREATE-d and DROP-ped. This structured language, coupled with the fact that many column names in databases are human-readable, such as “users”, “passwords,” or “card number” means that some degree of guessing can help an attacker to do SQL injection.

Sanitizing inputs is important and providing a layer of separation between receiving data from the outside world and inserting it into our data environment is a key. Sanitation is done by deleting characters that do not make sense in a specific context. This avoids a situation in which a developer puts the following code into a web page:

```
SELECT * FROM users WHERE name ='' + userName + '';
```

That particular code snippet simply pulls all the users from a database column called “users” where the “name” field matches the given “username”. If an attacker adds a bit of code to that web page, by filling in a form with “or '1'='1',” instead of a username it will give the following statement, which always is logically true:

```
SELECT * FROM users WHERE name ='' or '1'='1';
```

That statement, since it’s always logically true, pulls all the usernames from the “users” column, not just the one matching a specific user name. If the same statement said to DROP the column, all users would be erased.²³ While a well thought out, tested security and recovery plan (Chapter 8) should allow us to recover our user database after this attack, the impact to our organization in downtime and potential

data loss could be large. We should also be mindful of how the data is being transferred between our external customers and vendors, as well as between our internal, geographically distant locations. The need to store data offsite is necessary for recovery, but being aware of how the data gets there and planning for secure transport minimizes the risk.

Further security concerns arise as organizations evaluate moving to the cloud. Cloud computing is leveraging a host site that will provide server and storage resources as well as providing other computing resources “as a service,” such as Software as a Service (SaaS). A common example of SaaS is Google Docs. A cloud solution can be internal or external. Building and hosting an internal cloud environment has the benefits of increased control of the environment and flexibility to customize policies or procedures to fit the organizational needs, but it also has the drawbacks of increased start-up costs and the burden of yet one more thing to manage. Using an external vendor’s cloud solution also comes with benefits and drawbacks. Benefits include low start-up costs, minimal hardware management or maintenance, and as-needed/on-demand pricing. Yet using a hosted solution exposes organizations to some degree of risk since they have less control of the environment and the data stored there, and rely on service level agreements (SLAs) to ensure the necessary up-time.

Data quality aspects of information assurance

We must focus on another feature of data that potentially has negative effect on an organization – data quality. Information assurance goes hand in hand with security efforts. Reviewing data and understanding how deviations occur in data can improve the image of your company. If your company is using faulty data to contact former customers, interact with current customers, or recruit new customers, then these customers most likely will not trust your organization, if you cannot get their personal details right.

If a goal of security is to ensure trust, using bad data due to poor information assurance erodes that trust. These issues of data quality can be expensive to solve and can require substantial resources to review and reconcile throughout the organization. Even if we automate much of the

search for these records, decisions will need to be made about whether the different records represent the same entity. The process of accounting for data in planning and design becomes increasingly more important. Involving the data team early and often in the development process is key. Leveraging their knowledge and allowing for the allocation of the data team's time to these efforts is as important as managing the production systems. Improving data quality then allows the organization to build out analytics environments on which trustworthy decision making can occur. Customers see a view of themselves that is expected and reconcilable with their perceptions of how they engage with the organization.

Master data management

The days of having all data located in a single location on a single server are long gone. To further complicate data security, most organizations deploy a variety of technologies, platforms, and versions of applications, as well as legacy systems that are being nursed along. For the customer, this complicates issues of information assurance as we see multiple versions of the customer's data represented across the organization. For the decision makers in an organization, it becomes hard to trust the data in the system as representative enough of our business to be a successful base for strategic decisions.

The tools available in this space come from a variety of vendors. A whitepaper by Wolter and Haselden²⁴ provides a typical overview of what should be considered when evaluating master data management:

1. Identify the source of master data.
2. Identify the producers and consumers of the master data.
3. Collect and analyze the metadata for your master data.
4. Appoint data stewards.
5. Implement a data-governance program and data-governance council.
6. Develop the master data model.
7. Choose a toolset.
8. Design the infrastructure.
9. Generate and test the master data.

10. Modify the producing and consuming systems.
11. Implement the maintenance processes.

The outcome of a full analysis allows for the organization to move forward in a state of confidence, logging and addressing master record reconciliation in real time. Evaluating production data across multiple sites where mirroring is used is necessary to ensure data integrity across the enterprise and serves to meet the organizational goals of CIA.

Data security strategy

Any strategy, policy, or plan in an organization should include the following when managing data security across the enterprise:

1. *Always use the DBMS built-in security.* You paid for it, and you might as well use it! Built-in security in a DBMS includes features such as stored procedures, which allow workers to only execute specific preapproved queries; username, role, and password management; and encryption for fields, columns, and full tables.
2. *Review, review, review.* Review SQL as it goes into code, review workloads across the enterprise and review the quality of data in storage. Review recovery plans. Can you recover everything you have lost in the last hour, last day, or last week by tomorrow morning start of business?
3. *Maintain an active, top-down approach.* Data security and information assurance is EVERYONE'S concern not just the DBA's. Good IA starts from the top of the organization and should be driven down. If management makes data security a priority, then surprises from a breach or system failure can be minimized. Furthermore, we can improve the trust of our customers or suppliers, increase the capability of our decision makers, and leverage the most valuable resource our organization may possess.
4. *Expecting the unexpected.* You should have a plan for mitigating and publicly addressing breaches or loss. This is discussed further in Chapter 8, but the plan must be in place, practiced, and ready to execute the moment a breach is identified.

Summary

The key to data security is to manage what we can manage (control what is within our control), and have a plan in place to back up data and recover the data. A firm must practice that recovery, and monitor backup processes to ensure we are indeed protecting our data and ready for the unexpected. Furthermore, monitoring the consumption and movement of our data is increasingly important as insider security breaches appear to be on the rise. Unfortunately, the exfiltration of data may happen long after the original intrusion, as intruders bide their time before striking. Data security and IT security in general truly follows the adage of “time, effort, and money” to yield the promise of success.

CHAPTER 4

Keeping the Electronic Highways Safe

According to the FBI, “Internet-based social networking sites have created a revolution in social connectivity. However, con artists, criminals, and other dishonest actors are exploiting this capability for nefarious purposes.”²⁵ This same statement also applies to physical networks; the revolution in Internet connectivity is being exploited for nefarious purposes. Networks carry data from place to place. However, thinking of a network as a data carrier only tells half the story. We speak daily of networks, of which the Internet is the largest, and the data available on them. Without breaking stride, we move on to speak of cellular telephone networks and could rightfully speak of landline telephone networks, although they are rarely called such. In television, we speak of networks in terms of the content providers (the stations), but really, the network for television lies above us as electromagnetic waves in the air and below our feet in cables. What do all these networks have to do with each other? They carry data; but not just data, they carry information in the form of web pages, financial transactions, voice calls, video streams, and television programs. Exactly how this information is carried is really beside the point; it is enough to simply know that data is carried across a network.

So what is a network? At its simplest, a network consists of two or more machines (“hosts”) connected to each other, and exchanging some form of communication. The hosts, in a typical computer network, are individual PCs, whether clients or servers. The “hosts” in a cable TV network include the head end hardware that transmits the signal, the wires through which the signal is carried, and the televisions that display the signal. We need not concern ourselves particularly with TV, cable,

or cell phone network security. But, even a typical computer network no longer just carries static web pages or online shopping transactions. Thanks to efforts called convergence, Voice over Internet Protocol (VoIP) phone calls, video streams, and e-mails all ride over the same physical wires or fiber optic links.

Converged networks promise to simplify network administration since all communications run over the same physical lines. Unfortunately, both for security and performance, not all traffic is equal. Some traffic is more vulnerable to attacks, and other traffic must be delivered quickly and reliably. Without getting into too many details, it can safely be said that an e-mail can be delayed a few seconds, whereas voice calls must be delivered quickly in order to be understandable. There are various ways to accomplish this; and at least one offers a degree of added security.

Using virtual local area networks

There are two general methods of managing traffic flows on a network, traffic shaping or prioritization, and virtual local area networks (VLANs). Both have advantages and disadvantages from both a security and traffic management standpoint. Traffic shaping or prioritization, as the name implies, gives higher priority to certain types of traffic. These bits of data, typically voice or video, are sent most quickly to the destination, while other messages, such as parts of web pages or e-mails, must wait a few extra milliseconds. Traffic shaping has very few implications for security, other than perhaps someone tampering with the policies set by a network administrator. If a company relies on VoIP for phone calls, tampering could become problematic, but should be fairly quick to diagnose and fix.

VLANs work by setting up multiple virtual networks that carry data concurrently across one set of cables. Each VLAN is logically segmented from others, and someone on another VLAN cannot read data carried on one VLAN. This is a boon to security, as eavesdropping on voice calls on one VLAN is not possible from the data VLAN. VLANs, however, work against convergence to some degree; the individual network segments are isolated from each other, meaning some of the promised efficiencies may disappear. However, the loss of efficiency is of lower

concern than the advantages of protecting traffic from unauthorized listeners and keeping it flowing well. Both strategies, prioritizing or using VLANs, are good from a managerial standpoint.

Security concerns with convergence

When data, voice, and video are converged, some security concerns emerge. Data networks are often subject to “sniffing,” in which an unauthorized operator connects to the network and runs a program that receives all traffic passing over the network. This traffic, by virtue of the fact that most web browsing and e-mails are not encrypted, can be reassembled into human-readable messages, which can then be used by the interceptor. The same applies to voice and video; in some cases, the full conversation can be recorded. This loss of confidentiality may reveal business secrets, client data, plans, and the like. This behavior is of course illegal under wiretap laws, but various hackers have used it to their advantage.

Data networks which do not use encryption to scramble packets can also have those packets altered, resulting in the loss of data integrity, the second facet of our CIA triad. These altered packets may not be detected. As a simple example of data alteration, the program *ettercap* can inject arbitrary pictures into web pages, replacing the original images. This is sometimes done as a joke at hacker conferences, replacing the images on a web page with “You’ve been hacked” messages, especially for unaware users. Similar things can be done with words on a web page, which would be less noticeable. At least one program easily available online waits for a browser to visit a news site, and replaces the real story with spoofed stories. This, again, is mostly a prank, but could easily be made more malicious. A final threat to data networks is for someone to replace management messages with malicious management commands, to shut down parts of the network, or reroute various types of traffic. This could include the aforementioned attacks on traffic shaping, leading to loss of voice services on the converged network.

VoIP calls, as previously mentioned, may be vulnerable to eavesdropping. This is not the only possible issue with voice, however. Another threat is spam over Internet telephony (SPIT), in which an attacker places advertising messages into a voice conversation. Also, previously mentioned was the disruption of packet flows, resulting in garbled or dropped calls.

Before leaving the topic of VoIP, one program, in particular, deserves discussion. Skype is a widely used program to make free computer-to-computer calls, and call regular telephones also. While Skype is encrypted, and cannot readily be intercepted by outsiders, the keys for the encryption are held by the vendor, currently Microsoft. The keys in an encryption process, much like physical keys, allow their possessor to access the protected areas. That means, frankly, that the conversation can be eavesdropped upon by Microsoft.

Microsoft has not admitted anything at all about their use of the encryption keys nor have the original Skype developers or Microsoft ever discussed what algorithm (computer formula) is used for the encryption. It is clear, however, from one experiment that Microsoft can and sometimes does decrypt and monitor the call. In this experiment, a user created a brand-new domain and website, and sent a few links to pages on the site via Skype to a friend within an hour or so of setting up the site. Those specific links were visited within a few minutes by an IP address belonging to Microsoft.²⁶ Microsoft has also been very clear that it will cooperate with law enforcement to monitor traffic. In today's environment, where high monitoring by government and law enforcement is the norm, it seems clear that a firm concerned with privacy would want the most security possible. In choosing video conferencing services, a savvy manager would choose a product that is secure and can be counted upon to remain secure in the face of various eavesdroppers, which means the firm must control the encryption keys.

Virtual private networks, firewalls, and other “secure” networking practices

As the Internet has grown, it has become the preferred means of transferring data between companies and individuals. Firms use it to develop business deals, send e-mails to suppliers and customers, process payments, and sell products. All of these tasks are done daily across the Internet, often with little thought about the security of sending such messages, except perhaps when financial data are involved. While handy, the Internet is also a fairly open environment from a security standpoint, and much of

the data transmitted across it is unencrypted, or unsecured. The need for securing online conversations of any nature is not new, but has become more pressing over the last decade. There are two main methods of securing transmissions across any network: first, one can secure the network pathway across which the message travels; and second, one can secure the message itself. Both have advantages and disadvantages. Protecting the pathway via virtual private networks (VPNs) is discussed here, while encryption to protect individual messages is discussed in Chapter 6.

Importance of using secure networks

The first means of scrambling traffic to prevent interception is perhaps the most used. VPNs are typical of securing the channel across which data is transmitted. There are numerous ways of classifying VPNs, which will be discussed later, but all share some commonalities. All VPNs form a secure tunnel through an otherwise insecure network. This can be thought of as a pipe within a pipe. The outer pipe contains the stream of mixed traffic, while the inner pipe contains a “pure” stream. The outer pipe, in this case, is the Internet, and the inner pipe, the VPN tunnel.

The concept of using a VPN comes from the idea of using a leased line, that is, a telephone line, which was private to the firm. All the firm’s voice or data calls were routed across this particular wire. This provided relatively good security; even though the content was not scrambled, it could only be accessed by someone who could physically tap the line. This was a good solution to carry private information, but it came at a steep price, in dollars. As Internet speeds have gone up, and reliability has increased, many firms realized that the old leased private networks were very expensive, and as long as confidentiality could be maintained, it would be better to communicate over the Internet.

Types of VPNs

VPNs can be categorized in many ways, which are not mutually exclusive. The first way to classify them is which type of technology they use for encryption. There are two main standards, and a few others. First,

Secure Sockets Layer or Transport Layer Security (SSL/TLS, hereinafter simply SSL for simplicity) can be used to create a VPN. In fact, every SSL connection, including all secure web browsing (HTTPS connections) is a VPN, from one client to a server. SSL has the distinct advantage that it is present in all web browsers and requires little or no setup on the client side for the network administrator. All that must be done is to set up a secure server, and point the client to the server. SSL works best to encrypt web pages and e-mail. Therein lies the great disadvantage of SSL; software must be written specifically to work with it. Many types of communication, including mapping network drives for remote access to shared company files, do not work.

There is one major exception to SSL's lack of compatibility with arbitrary traffic, named OpenVPN. OpenVPN is an open-source SSL-based VPN. In order to implement it, an administrator must do some server side and some client side configuration. First, a VPN profile must be configured on the server, then client software must be installed on each device (computer, phone, tablet, etc.), and the VPN configuration copied to the device. The user then must click on the software and initiate the connection. They can then access network resources as if they were sitting at their desk on corporate premises. This type of VPN works well, but lacks some management features of the IPsec standard, discussed next.

IPsec, short for Internet Protocol Security, is a set of standards for encrypting IP packets. It works at the network layer of the Internet Protocol Suite, while SSL operates at the next higher layer. Do not worry about understanding the layers used in communication, what this means for the average user is that IPsec can wrap any traffic running across a network in a blanket of security. Whereas SSL requires that software be written to work with it, IPsec has no such shortcomings. The trade-off for this simplicity is that it is not especially simple to set up. In order to use an IPsec VPN, an administrator must create the VPN profiles, similar to the process for OpenVPN, then set up client software on the device. The advantage and complexity of IPsec over OpenVPN, however, is that there are many more configurable options, including more choices for methods of encryption (more algorithms) and more fine-grained

control over which algorithm can be used with any given device. An authentication server enforces these policies. From a managerial standpoint, IPsec is the best choice, but OpenVPN can be equally secure.

The other main way to categorize VPNs is based on what devices are connected. There are three types, host-to-host, host-to-site, and site-to-site VPNs. In a host-to-host VPN, a single device connects to another single device. This type of VPN is typified by SSL connections for secure web browsing, and is also sometimes used for server-to-server connections. Traffic between the two hosts is encrypted end-to-end.

Host-to-site VPNs, also called remote access or road warrior VPNs, are what most think of as VPNs. A host-to-site VPN enables a worker on the road, or telecommuting from home, to access corporate resources exactly as if sitting at their desk. In a host-to-site VPN, the traffic is encrypted from the client to a VPN concentrator, a.k.a. VPN server, which sits at the edge of the organization's network. This secures the traffic across the most dangerous portion of its journey, that is, across the Internet.

The final category in this group is site-to-site VPNs. In this scenario, a VPN concentrator at the edge of the firm's network connects to another concentrator at a remote site, such as a branch office. All traffic from all devices at the branch office goes through the concentrator and is then decrypted at the main office site. Any traffic from the main site to remote follows the reverse pattern, and is again encrypted as it passes across the most dangerous portion of its journey.

VPNs for remote workers on unsecured WiFi networks

There are many managerial reasons to insist that a firm implement VPNs everywhere possible and then train users to use them properly. One of the most important scenarios for using a road warrior VPN is very common today. WiFi hotspots are everywhere, from hotels to coffee shops, and are used by many for business and personal purposes. These hotspots are often unsecured, so that any customer may easily join. This leaves data passing across the network vulnerable to eavesdropping. The use of a VPN while connected to WiFi is an excellent security measure that should be taught to all employees until it becomes

habit. While SSL adds security to https:\\ web browsing sessions, it is not a cure-all. A certain type of attack, called a man-in-the-middle (MITM) attack, allows some sessions to be captured. Successful MITM attacks require the user to accept an invalid SSL certificate; however, experience has shown that when users are presented with an invalid certificate, they often just accept it because they “want to get online.”²⁷ The only real defenses against MITM attacks are to train users to never accept invalid certificates, and to always use a VPN when on WiFi.

Firewalls

Firewalls are the least useful and least effective security tool in the IT security arsenal. At the same time, they are the most used, the most overused, and most over-trusted devices in said toolkit. Finally, they are absolutely essential to securing a modern network. While this may seem to be hyperbole, all of the previous statements are true. Before discussing why firewalls are both required and overused, it is first necessary to understand how firewalls work.

A basic firewall functions by blocking Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports. TCP and UDP are two network protocols whose function is to make sure that one application on one machine can speak to a corresponding application on another machine. For example, standard web browsers contact a web server on TCP port 80. Both browser and server are programmed to use this port, as it is a standard for exchanging web pages.

TCP and UDP ports, then, represent a door or gateway through which one program talks to another. To extend this metaphor, if the door is closed, no one can communicate through it. Similarly, if a door is never opened, that is, if no program is set up to utilize that port, no communication takes place. This scenario occurs on every server on every network. When a server is configured, different applications are set up on it. Each of these programs opens one or more network ports. All other ports remain closed, as there is no software program listening on that port. In other words, if a server is set up with only web server software, and nothing else, it should have only port 80 open. All other ports remain closed.

Assuming our example server above is supposed to be accessible to the outside world, a firewall that protects the network must also have port 80 open. The firewall allows requests from outside to pass unhindered to the web server. These requests usually take the form of a web browser asking for a web server page, and the server then answers with the requested page. Some requests, though, are malformed, or malicious; in other words, they represent attacks on the server. A basic firewall cannot distinguish between the legitimate request packets and the attack packets. This means that it provides exactly zero protection to our simple server. Even if the firewall has all other ports blocked, those blocks are doing nothing. Those same ports that the firewall is blocking are not even open on our server! Even if no firewall was used, and someone tried to connect to any other port on our server, there would be no answer. This means that the simple firewall is of no use. It is much like answering the door at home. If you have no peephole or other way to see out, you really do not know whether you are opening the door to a friend, a salesperson, or a thief. The server simply answers to the incoming packets passed by the firewall, and gets both the good and bad.

Of course, this only covers the simplest type of firewall. There are more advanced firewalls, and tools such as intrusion detection or intrusion prevention systems (IDPS), which can distinguish some kinds of attacks from legitimate messages. These technologies will be discussed shortly, but some aspects of network design, which have implications on the effectiveness of firewalls, should first be discussed. In most modern networks, the main firewall is called the border firewall. It is placed at the edge of the organization's network, so that all incoming and outgoing traffic to the outside world is filtered through the device. Additional firewalls are often placed between departments, blocking access from insider threats. Finally, servers often have software firewalls, also called host-based firewalls, installed as additional protection for the valuable data they contain. The border firewall of a decade ago protected the entire perimeter of the network. In the intervening years, a phenomenon known as the "death of the perimeter" has made border firewalls less relevant than previously.

Death of the perimeter

The death of the perimeter describes a condition in which only the smallest networks have a single point of access. In a medium to large firm, there may be several independent connections from different Internet service providers (ISPs) so that if one fails, the firm still has Internet connectivity. Each of these connections represents a possible incursion point for an attacker, and all must be protected equally. VPN connections represent another possible incursion point. If improperly secured, or if their incoming traffic is simply trusted, rather than passing through the firewall, they represent an insider threat. Trusting incoming VPN traffic is the usual practice, for performance reasons, but it leaves some vulnerability.

A final breach of the old perimeter lies in wireless networks. WiFi is an exceptionally useful way to network the many mobile devices that all of us carry together. If misconfigured, however, it represents another attack surface. Again, as with VPNs, traffic that enters the network via WiFi is generally trusted. With that said, a properly configured wireless network is mostly immune to attacks. Wireless security will be discussed shortly, but it must be remembered that in many cases, not all the devices that connect to the wireless network are trustworthy. As bring your own device (BYOD) programs become more popular, it is imperative that managers and administrators ensure that those devices follow corporate security policy. As the perimeter becomes more open, firewalls become less effective at keeping intruders at bay.

Other firewalls

There are more complex firewalls that can control network traffic in more sophisticated ways than simply based on which port a packet is passing through. First, there are firewalls that look at streams of packets, rather than individual packets in isolation. Called stateful packet inspection (SPI), the process considers the state of each packet. Then, a simple rule set decides whether the packet is part of an established connection, and permits or denies that traffic. These rules also drive another common technology, Network Address Translation (NAT). NAT is used by

many small firms, so that they only need purchase one IP address from their ISP, which is then shared among all PCs in the firm. Most home Internet connections also use NAT. Like an SPI firewall, NAT tracks the state of packets, making a NAT device a very effective firewall for small firms and home users, so that they probably do not need any other firewall.

The final class of firewalls includes many different types, which vary in their exact filtering mechanisms and level of sophistication. These devices, generically called unified threat management (UTM) firewalls, include the previously discussed port-based filtering and SPI, and additionally can filter based on other criteria. Some can even filter at the application level, meaning that they can filter based on text, or a URL. Examples of this include the firewalls typically used at schools and libraries to keep children away from questionable websites.

This application-layer text filtering can also be used to sanitize incoming traffic, for example, an application firewall can clean out the special characters used to hack databases via SQL injection as discussed in Chapter 3. Other UTM firewalls can filter out many viruses and other types of malware. However, both of these types of filtering require large amounts of processor power, which equates to high costs, practically limiting their use mostly to larger firms.

The last thing firewalls cannot protect against, in most cases, is an internal user downloading threats. While some UTM firewalls can filter viruses, given that typical antivirus software is only about 50 percent effective, this leaves a fairly gaping hole. If a user downloads opens an infected spreadsheet, as happened to the security firm RSA,²⁸ then an attacker will have a way back into the system, bypassing any firewalls in place. If an internal user brings in a virus on a thumb drive, which then “calls home” to the virus’ author, firewalls are likewise circumvented. This scenario, sometimes called advanced persistent threat (APT) in the media, has been repeated many times.²⁹

So, where does that leave our firewall discussion? They are absolutely necessary, and horribly over-trusted. Firewalls are very good at keeping certain types of attacks at bay. There are many other types of threats that firewalls cannot handle, because they are simply the wrong tool for that particular job. Finally, firewalls only handle provable attack packets in

any scenario and ignore those things that are merely suspicious. In other words, firewalls must be implemented, but managers should take an extremely jaundiced view of anyone who glibly states “the firewall will block that.”

Other security tools

Related to firewalls are a device that can handle suspicious packets, not just provable attacks. These devices, called intrusion detection systems or intrusion protection systems (IDS/IPS/IDPS), work by filtering the long stream of packets that makes up an online conversation. IDPSs usually work at the network, transport, and application layers, and filter traffic-based on packet header information, as well as actual content. IDPSs can be very effective in alerting administrators to possible attacks, or even stopping the attack in progress. This is the difference between IDS and IPS systems; an IDS notices the attack and sends an alert to an administrator to manually intervene. The IPS, on the other hand, can take actions, such as rewriting firewall rules to block a certain IP address (the attacker’s IP), on the fly. While effective, IDPS systems are plagued by false positives, or identifying benign traffic as malicious. This gives rise to the “boy who cried wolf” phenomenon; and network administrators become complacent in checking alerts. Tuning the IDPS will result in fewer false positives, but they are currently a major problem for these systems. That does not mean they should not be installed, but realize that tuning, and possible job rotation to avoid boredom, should be part and parcel of the implementation.

Wireless security

Wireless devices and wireless networking are here to stay. The convenience of having access to the network anywhere, from almost any type of device, is so great that companies will almost certainly increase investment in newer iterations of wireless. Wireless technically refers to any technology that can communicate without wires; early radios in the 1920s were called “wireless sets,” and the Palm Pilots and laptops of the 1990s could send data over infrared light beams. However, in common use today, we mean the 802.11

family of standards, sometimes called WiFi (short for Wireless Fidelity) when we speak of wireless. A brief introduction to 802.11 networking precedes our discussion of the security of WiFi.

The 802.11 standards got their start in 1997, when the Institute of Electrical and Electronics Engineers (IEEE) finalized a standard that allowed network packets to be sent over radio waves, rather than over cable. The standard ran at 1 or 2 megabits per second (Mbps). The very popular 802.11b standard, which ran at 11 Mbps came in 1999, and was followed by 802.11g in 2003. 802.11g worked with the older technology, but ran much faster, at 54 Mbps. In 2009, 802.11n was approved, running on 2.4 and optionally 5 GHz, with speeds up to 600 Mbps. The 2.4 GHz devices can work with the older “b” and “g” devices, smoothing the upgrade path. The latest amendment to the standards is 802.11ac, approved in late 2013, which can run up to 1300 Mbps. More standards are in the works, and will of course bring faster speeds, less interference, and improved range, all of which are sought by organizations implementing wireless networks.

All of the 802.11 family currently share similar security concerns. The first concern is simply one of availability of the signal. When a signal is trapped within a wire or fiber optic cable, it is guided to its destination. As with any other radio signal, WiFi is a broadcast technology, meaning that it is sent, typically in all directions, to anyone in the area with a suitable antenna and tuner; e.g., the WiFi card built into the laptop. Just like anyone within a given geographical area can tune to a certain radio station, say 98.5 MHz on the FM dial, and hear it, anyone within range can and does receive the signal from the WiFi access point. In normal operation, the device simply ignores packets not addressed to it. Most wireless devices, however, can be set to so-called promiscuous mode, in which it accepts any and all signals received. If the signal is not encrypted, a device in promiscuous mode can eavesdrop on others in the area. In some cases, discussed below, even encrypted transmissions can be intercepted and decrypted.

WiFi has three main protocols for security. The first, Wired Equivalent Privacy (WEP) came with the original 802.11 standard in 1999. In 2001, a method of attack was published that took advantage of a weakness in the way encryption was performed, and subsequently, automated

tools were developed to crack WEP. The result is that today WEP can be cracked in a matter of seconds by even a non-skilled attacker and can be done reliably from almost a mile away with specialized antennas.³⁰ Because of this, WEP is thoroughly deprecated and should never be used. Sadly, as of late 2014, 16 percent of networks still use this standard, but that number has decreased greatly from a high of 45 percent in 2010.³¹ The use of the WEP protocol was implicated in the 2007 theft of 45 million debit and credit cards from T.J. Maxx, at that time the largest breach of credit cards in history.³² The credit card industry banned the use of devices implementing WEP after 2010, but that step was already long overdue. To summarize: Do not use WEP!

To address the shortcomings of the WEP, the 802.11 Alliance released two new standards for encryption. The first, WiFi Protected Access (WPA), became available in 2003. This standard was always intended as a stopgap measure.³³ WPA was designed so that devices without suitable processing power, such as older WEP-only access points, could be upgraded via a software update. In 2004, the much more robust WPA2, also known as 802.11i, was released. Subsequently, parts of WPA have been cracked, and it was officially deprecated in the 802.11 standard in 2012. Unfortunately, despite the intentions of the industry to use WPA as a stopgap, it is still included as an encryption option in modern production access points, and so it is still sometimes used. As of late 2014, approximately 10 percent of networks still use WPA.³⁴ As a manager, it would be wise to replace or reconfigure any and all devices still using WPA, or in other words: Do not use WPA either!

WPA2 is the current gold standard in wireless security. It has two modes of operation, pre-shared key mode (PSK, a.k.a. personal mode) and enterprise mode. Enterprise mode is the stronger of the two, but requires more hardware and software configuration. As the name implies, it is most often used in enterprise settings, not by home or small office users. In this mode, a user must authenticate via the same username and password they use for all other corporate resources. Assuming that the password chosen is strong, WPA2 networks in enterprise mode may be considered nearly inviolable as of mid-2015.

Personal mode is designed for smaller networks that do not need the complexity of a separate authentication server. Odds are good that you, the reader, use this mode on your home network. To implement this mode, the access point (or wireless router) is configured with a passphrase. This passphrase is then entered into every device that needs to connect to the wireless network. The shared passphrase is used only for initial authentication, that is, to prove the device is authorized to connect. After connecting, the access point shares a separate, unique cryptographic key with the device that is then used to encrypt all traffic. Since that key is unique, even an eavesdropping attacker cannot intercept the conversation. WPA2 PSK can generally be considered almost as strong as WPA2 enterprise, subject to the constraints below.

Both WPA and WPA2 are vulnerable to a few attacks. The first, weak passwords, can affect both PSK and enterprise modes, with PSK being more vulnerable. If the user chooses a weak password or passphrase, and the authentication device does not limit the number of attempts before locking out an account, a brute-force attack can be tried. In brute force, an attacker simply uses a tool to try all possible password combinations until a match is found. Authentication servers used for enterprise mode can be set to limit both the rate and the number of password attempts. No such protections exist for PSK mode, so a short password can be easily cracked. It is generally considered that “short” means a passphrase of less than 20 characters for WPA or WPA2; the upper limit for passphrases is 63 characters. A passphrase may contain any characters, including spaces, an easy-to-remember sentence is a good passphrase, as long as it is not a famous quote. For example, “Every dog gets fleas starting in June in Wisconsin!” is reasonably memorable, easy to type, and long enough to be strong.

A more serious vulnerability for WPA and WPA2 in PSK mode is called WPS PIN recovery. Some routers, especially home and small business devices, implement a technology called WiFi Protected Setup (WPS). WPS promised to eliminate weak passphrases by automatically generating a strong key when a button on the router is pressed or an eight-digit pin is entered. Unfortunately, there is a flaw in the feature that allows an attacker to recover the WPA/WPA2 passphrase in a matter

of a few hours. The PIN is typically written on a label on the outside of the router, and cannot be changed if compromised, meaning that once compromised, the device is essentially forever vulnerable. The only solution is to turn off WPS; however, some routers do not allow this. Even worse, some, like a device owned by one of the authors, claim to allow it, but the checkbox that controls the feature actually does nothing. Luckily for the security of the enterprise, few business access points implement the WPS feature.

Where does this leave WiFi? If secured properly, and users are trained properly, it is an extremely useful tool. Proper security means using WPA2, with either good corporate passwords or long (more than 20 character) passphrases. Proper training consists of teaching users to always use a VPN when connected to public WiFi hotspots.

Summary

Keeping the information highways safe in an organization is a daunting task. There is no single tool that can magically fix network security. VPNs and other types of encryption can protect most of the traffic on the network by encrypting it. Firewalls protect the network from some types of intrusions, but often are trusted to keep everything out, including malware and internally initiated actions. Protecting data is relatively easy compared to voice or video traffic. VoIP traffic has historically been easy to intercept, but this is changing in today's threat landscape, as standards which can encrypt voice calls emerge and are adopted. More and more VoIP products can place encrypted calls from device to device, but none can protect a call to a landline or cell phone. When selecting video or voice conferencing tools, managers need to pay attention to security concerns. WiFi networks are more vulnerable than typical wired networks, but using the right security standards and enabling a VPN when running over a public WiFi hotspot can still protect the traffic.

CHAPTER 5

We Released What?!? (Application Security)

Application security is a major component of maintaining a high-security profile for the organization. Whether you are releasing a new mobile app or purchasing an HR solution, security needs to be considered in all phases of the project and at all levels of the organization.³⁵ Complicating the security in this space is the diversity of technology platforms, delivery methods, and availability expectations as well as accounting for the needs of stakeholders within and outside of our organization. This chapter will introduce the idea of a secure developer and provide strategies for thinking about application development security in your organization.

The need for a secure developer!

Our primary goal in writing this book is to stress the importance of integrating security into all aspects of the IT effort. It is insufficient for security be considered solely as a separate, stand-alone endeavor, monitoring the state of your organization. By including representation from the security team, a manager can bring context and understanding of security from across the enterprise to the development team. This knowledge allows the development team to situate their application in the context of the enterprise and enables them to create the application within the framework of the organization's security practices.

An organization may choose to train analysts and developers in organizational security practices. The training creates a new role of secure analyst or secure developer that can bring their knowledge to the analysis, design, and development process. Ultimately, the goal would be to have an entire organization that has a deep understanding of security across the enterprise, but that may not be realistic in the early stages. Security

training provided over time will build that deep understanding. A good initial goal would be to relieve some burden on the security team by providing targeted training to key resources that work on key projects. As these resources go out into the organizational space their knowledge will transfer. Continuing the deeper training over time and adding to the pool of secure developers in your organization will ultimately drive a security culture in your organization.

Integrating security into the development process adds a layer of complexity that will impact the time to delivery. The trade-off is the uncertainty of the cost of a breach. The cost of lengthening development is measurable, but evaluating the cost in the face of breach that may never occur is a tough sell. The push to incur this extra cost will need to come from above, and buy-in from program or project managers at the ground level will be absolutely necessary for this strategy to work. The security culture will repay in terms of minimizing the potential for organization-destroying breaches. Thus, it should be recognized that concrete costs will be incurred but these costs may be negligible in light of the potential damage that a breach may cause.

So how do we build this secure culture? Mapping our organizational technical architecture, application delivery, internal and external stakeholder access, and understanding where our assets reside (data, files, etc.) are key to ensuring security. The organization needs a framework to guide analysts, developers, and business partners about security procedures. The challenge is in creating something that is concise and easy to consume for the layman with security resources who have a deeper expanded framework to aid their evaluation across the enterprise.

Creating a framework based on guidelines, such as those provided by the International Organization for Standardization (ISO), or National Institute for Standards and Technology (NIST), etc., provides a starting point for developing a systematic approach to security that is present in all levels of the organization. But the organization must not stop there maintaining and assuring that the framework is followed and materials are evaluated.

How are the applications using our data and networks?

Secure development is important for in-house development, purchased solutions, or customizing existing applications in either category. As many IT professionals will attest, the user is able to discover surprising and novel ways to use our applications. To best evaluate the security of applications or potential purchase will require evaluation from a broad range of organizational users and support staff. The trade-off is always a cost/time issue, which is a common theme in this book. As organizational leaders, you will have to make a decision as to whether is it worth spending known dollars on staff evaluating and testing applications and technologies versus the unknown cost of a breach.

As organizations are increasingly more reliant on external, off the shelf applications, our view of security needs to incorporate the ability to evaluate before, during, and after purchase. The need is substantially higher in this case since the organization does not have a foundational understanding of how the application was built and how it may align with the organization's security protocols.

Securing the environment, test data, and making the migration happen

All IT research, test, and development environments in the organization need to be built, managed, and migrated following the security framework that incorporates all aspects of development. The challenges arise in keeping diverse environments secure across varying degrees of external exposure and use. A security framework needs to be able to enable access to internal and external resources while providing test data that is secure and protected.

Research is often key for developing or evaluating new technologies in the context of your business. This research environment is also typically the most exposed to external threats and the most difficult to populate with test data. Given the level of exposure and the potential of security threats as well as the possibility of a technology failure, the research environment should be essentially be stand-alone and partitioned off from other organizational resources.

Development environments require the balance of simulating the reality of the production setting and also allowing for the process of building the applications. So, a key for our plan is to incorporate sufficient resources and anonymous test data to allow our developers to produce functional applications that meet our requirements and the needs of relevant stakeholders.

In a production setting security is of the utmost importance. Depending on the industry, regulations and standards may apply, complicating the protection of data and resources. It is important to understand the range of security issues that impact development of systems facing external customers or suppliers and internal systems that are delivered to the appropriate employees. In production we close the circle of involving security in the early stages of planning, analysis, and development. Drawing the requirements out of our production environment and addressing them early in development, while not guaranteed to make implementations smooth and seamless, should help cut down on issues both during and after implementation.

Testing applications

Once the application is built, the application development process is not complete. It is not sufficient for the application to simply work; we need to assess the quality of the application within the context of the enterprise. Further complicating quality assurance is the multitude of environments in which the application is required to function. As applications are tested, the effect of the application on the overall security to the organization should be reviewed. Typical questions to ask are:

- Is the application for internal or external use?
- Who has access to the application?
- What data is consumed or altered by this application?

The key to making sure all goes well in implementation is testing aspects of security throughout application development. As development progresses, tests need to account for increasingly more real-world-like scenarios. An application working securely in a test environment does not

necessarily mean “we’re good to go!” Exploring facets of an application with stakeholders and receiving quality evaluation from testers outside of the project are good strategies for mitigating the impact of application flaws on security. Many organizations add a separate stand-alone testing environment that simulates as closely as possible production workloads and environment configurations.

Tools and methodologies are available to better incorporate testing into development. Test-driven development (TDD) is an approach to start with the simplest test for a problem, based on stakeholder requirements, writing code that passes that test, refactor or clean up the code, then proceed to the next, more complex test. The cycle continues until a functional solution is provided. As approaches, such as TDD, become more prevalent, nonfunctional requirements, such as security, are being considered in the methodology.³⁶

As our workforce and the world shift to a more mobile focus, an organization is forced to monitor not only applications on the web but also apps or sites that are mobile, on devices that connect in coffee shops, home offices, and airports around the globe. Mobile technologies introduce additional device-specific security challenges in both development and technical architecture.

Other challenges include testing patches and fixes to operating systems and database management systems as we need to update and maintain our current infrastructure. The development and testing environments become vital for evaluating the impact of these fixes. Researching migration to new versions of an OS or DBMS takes place in our research environment. The tangible benefit of such environments is not necessarily evident in the bottom line, rather the benefit is seen in avoidance of costly breaches, smoother transitions between versions, and better understanding of our production environment.

Summary

The goal of establishing CIA in the enterprise is important in all aspects of IT, including development and implementation. Creating an atmosphere to encourage and support “secure” developers should be considered a priority in any organization. As applications are built, security should

be considered in all phases. Methodologies and tools are progressing to a state that facilitates the integration of security into development. The final takeaway from this chapter is providing developers with resources to securely develop applications. Strategies include building a secure development workforce and providing adequate test environments. Security should always be a proactive with consideration of breaches across environments planned and evaluated on a regular basis. Security protocols should be communicated throughout the organization with advocates at all levels promoting these vital efforts.

CHAPTER 6

Cracking the Code (Cryptography)

This chapter is likely the most technical chapter of the book, but it is absolutely necessary for a manager to understand how encryption works, at least at basic levels. Knowing how encryption works will ensure that the manager understands which security problems can be solved by encrypting data. Typically, for example, we think of cryptography as ensuring only the confidentiality of information, but as we will see, it can help fulfill at least two of the three tenets of the CIA triad.

From a managerial standpoint, not only does encryption protect our data, but it's also required for compliance with many current regulations. PCI DSS (Payment Card Industry Data Security Standard), for example, requires that certain types of encryption be used on credit card processing machines and on stored data. Fines can be levied against firms who accept credit cards as payment, but do not comply with PCI DSS.

What is it?

Cryptography, from the Greek “kryptós,” meaning hidden or secret and “graphein,” meaning writing, is the science that designs and implements ways of keeping communications secret. Usually, it involves *encryption*, turning *plaintext* that is human-readable into *ciphertext*, which is the scrambled version. *Decryption* is the reverse process. A brief note is in order here on cryptography versus secret codes.

Secret codes differ from cryptography in that they use a plaintext word or phrase to mean a different word or phrase. For example, “The hen has escaped from the coop” could mean “The prisoner escaped” (very transparent!), or “We will attack at dawn,” or “Meet me at the post box on the corner of Main and Vine at 5:00 on the 17th.” Codes, then

seek to hide meaning in the open, while cryptography scrambles the message to hide the meaning. Codes have very limited applicability in computing, so we will discuss only cryptography in this chapter.

We have already discussed many products that use cryptography, including the encryption used to secure hard drives and to keep communications via VPNs secure. Cryptographic recipes or algorithms, properly implemented, can keep information completely secure, but unfortunately, proper implementation is difficult, and the history of cryptography is littered with failures.

Early cryptography involved such methods as the Greek *scytale*, a stick of a known diameter around which a paper tape was wound in a spiral. The text was then written in the usual fashion, one word or letter at a time along the length of the stick, then the paper was unwound and the long paper tape sent to a recipient who had a stick of the same diameter.³⁷ The shortfall of this method is obvious; it's fairly trivial to find the right diameter with very little trial and error. Another classical Greek method, cited by Herodotus, was to shave a slave's head, tattoo the message into the skin, and wait for the hair to regrow to cover the message.³⁸ This technically would be called *steganography*, for "hidden writing," but the net effect was the same.³⁹

In slightly more recent times, the basics of modern cryptography were born. Julius Caesar used a *substitution cipher*, in which a letter was substituted for another letter in the plaintext. Most readers probably played with a similar cipher in grade school. Caesar apparently used a three-character shift to communicate with his generals. An example of a three-character right shift is shown below.

Plaintext: Attack the city at dawn
Ciphertext: Dwwdfn wkh flwb dw gdzq

This is obviously easy to decode, as many a third-grader has found out to their dismay when a note was intercepted by a teacher, but again, this was an early cipher.

A similar method can be used with a key that is kept secret by both sides, where one letter is substituted according to its position in the alphabet with another letter in the key, where the key contains all letters of the alphabet, merely in random positions. An example of this is shown below.

Plaintext alphabet:

a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t	u	v	w	x	y	z
h	f	s	k	e	g	j	m	x	q	y	b	d	z	r	a	o	c	v	w	t	u	i	l	p	n

Mixed alphabet (key): hfskegjmxybdzraocvwtuilp

Plaintext message: Attack the city at dawn

Ciphertext message: hwwhsy wme sxwp hw khiz

Often, spacing is removed in substitution ciphers to mask the length of words or sentences, resulting in ciphertext of “hwwhsywmesxwphwkfiz.” This is still fairly easy to decode, with even a rudimentary knowledge of the frequency of letters in English text; it can be determined that “w” and “h” in the ciphertext probably correspond to one of “e,” “t,” or “a,” the three most common letters in English text, and indeed, they do.

The other common type of basic cipher is the transposition cipher. These are often seen in the comics section of newspapers, as puzzles, or elementary school word scrambles. In a transposition cipher, letters remain the same, but are moved to a different order. The order is determined by a key kept secret by both communicating parties. A common example of transposition is columnar transposition: the message is written out in rows, then read back column by column, in a pre-determined order, as shown below. In the example, the length of the first word was chosen as the number of columns, and five letters added to the end as padding. The columns are arbitrarily set as 6-3-4-2-5-1.

Plaintext message: Attack the city at dawn

Key:

6	3	4	2	5	1
A	t	t	a	c	k
t	h	e	c	i	t
y	a	t	d	a	w
n	o	u	x	r	z

Ciphertext message: atyn thao tetu acdx ciar ktwz

Working backwards, the decryption process would be to write the groups of four bottom to top, in the six groups, and then read across. The spacing between the groups in the ciphertext is optional, and can be arbitrary, making decoding a little harder.

Other ciphers followed these simple ciphers. Some were based on multiple substitutions or transpositions, while others were based on combinations of transposition and substitution. This latter method forms the basis of modern cryptography. From the late 16th to early 20th centuries, various mechanical devices were used to automate the process, culminating in the German Enigma machine used in World War II. This same time period saw the start of computerized cryptography and *cryptanalysis*, that is, codebreaking, typified by the efforts at Bletchley Park in the United Kingdom. The Colossus, the world's first electronic computer, was successful in breaking the Nazi's secret codes. The work of Turing and others during the war was recently popularized in the movie "The Imitation Game."⁴⁰ While not completely factual, the movie is a good synopsis of the work. The same processing power that enabled better code breaking in the war also allowed the development of much more complex ciphers since then.

Modern ciphers in layman's terms

A modern cipher uses many thousands of transposition and substitution operations to completely mask the original text. A few common ciphers will be briefly outlined below, as representatives of various classes or methods of encryption. A modern cipher provides confidentiality to resources, insures the integrity of the message, and additionally may provide authentication of the identity of the sender.

Two main types of encryption are used in modern ciphers, symmetric ciphers and asymmetric ciphers, also known as public key cryptography. In symmetric ciphers, both sender and receiver use copies of the same key. Much like if you wanted to allow a friend access to your house or car, you would simply copy your key, and give it to them in person. This works well, unless the person is not local, in which case, you would need to securely deliver the key to them, perhaps via the mail or a courier. For a house key, this might be acceptable, but what if the key were to a safe or vault? Who could you trust? This is where public key cryptography becomes valuable.

Public key cryptography involves creating pairs of keys, which are mathematically related in such a way that something encrypted by one of the pair can only be decrypted by the other member of the pair. The keys are known as public and private keys, and as the names imply, the public key can be distributed at will to anyone. They are often stored on public key servers, or posted on websites. However distributed, a public key can be used for anyone in the world to encrypt a message to the key's owner. The owner, as the sole possessor of the corresponding private key, can decrypt the message. This owner then responds to the original sender, securing the message with the original sender's public key, and signing it with his own private key. This signature attests to the ownership of the key and therefore the identity of the sender. This solves the key distribution problem nicely and allows for *nonrepudiation*, that is, the sender cannot deny having sent the message.

AES & SSL/TLS

Two very common ciphers in use today are SSL/TLS and AES. SSL was previously explained in Chapter 4. AES, the Advanced Encryption

Standard, is a US Federal standard for encryption; it is currently considered the gold standard in encryption. SSL is a public key cryptography method while AES is symmetric.

SSL, as a public key encryption method, requires the server operator to generate a public-private keypair. These keys are then installed to specific locations on the web server and used to secure e-commerce transactions and email. Before the installation of the keys, though, another operation needs to take place, that of signing. When a new key pair is generated, the server operator sends the public key to a certificate authority (CA), such as Verisign or Thawte. The CA checks the credentials of the server operator, ensuring that the individual or business really is who they claim to be, then signs the public key, which is then called a digital certificate. The CA plays a role analogous to a notary public, attesting to the identity of the operator or business.

When a customer wishes to purchase an item from the website, their browser requests the digital certificate from the server. The server sends this certificate to the customer's PC. The customer's browser then checks the certificate's signature against a list built into the browser, which contains the credentials of several hundred certificate authorities. Assuming the signature of the CA on the certificate is validated, the current date is within the valid date encoded in the certificate, and that the name of the firm on the certificate matches the domain name the browser is visiting, the certificate is accepted. The browser then displays a lock icon, and the customer can be assured that they really are visiting the e-commerce site they meant to, and not an impostor's site. The browser then uses the public key contained in the server certificate to encrypt a temporary session key that is then transmitted to the website. This symmetric key is used to protect all subsequent communication between the customer and website, until such time as the customer disconnects or closes the browser.

It is worth mentioning that even though the communications are secured, once the customer's information reaches the server, it is decrypted. If the server operator decides to store credit card records unsecured, they are vulnerable to anyone breaking into that storage, something that the world has seen too much of in recent years. Similarly, the CA validates the credentials of the business to ensure it's really the business it claims to

be; they do not validate business practices. In other words, if “Bob’s Fly-By-Night Electronic Megastore” really is a business, with a correct address, business licenses, and such documents, then the certificate will likely be signed. Whether “Bob’s” really ships the 97-inch TVs they advertise for \$150 is beyond the scope of the CA’s job.

One final issue with security certificates lies in how users handle invalid certificates. Often, as with software licenses and prompts by malware software, we have been trained to simply click through error messages. However, when a user clicks through the warning for an invalid SSL certificate, they open themselves up to two possible attacks. First, as just mentioned, the site they are visiting could be an impostor’s site, set up to steal credentials or credit card information. The other attack is the MITM attack. In this scenario, an attacker sits between the victim and the website to be visited, often via a fake public WiFi hotspot. They connect normally to the intended website, then present the victim with an invalid certificate. If the victim accepts this invalid certificate, their session with the MITM attacker is encrypted, then the attacker decrypts their information, including things like passwords and credit card information. Then, the attacker simply sends the information to the intended website and passes all replies from the site to the victim. The victim only has the first warning (the invalid SSL certificate) that something is wrong. The only way to avoid this type of attack is educating users to never click through an invalid certificate warning. This also means that self-signed certificates (generated by the company, but not signed by a recognized CA) should not be used in an organization. They are sometimes used for testing, but should never be employed in a production environment, as it simply trains employees to accept invalid certificates blindly.

AES, as mentioned, is a symmetric key encryption method. If public key encryption solves the key distribution problem so well, why use symmetric keys at all? The answer is simply a question of resources. Public key encryption uses many more CPU cycles than symmetric key encryption. Authenticating the server must be done via public keys, but after that, the bulk of the communication can be done with the much “cheaper” symmetric key, a small file which has been exchanged securely, having been encrypted by the server’s public key.

Key size is a critical issue in any cryptographic system. Generally speaking, the longer the key, the more secure the encryption, and the longer it takes to break that encryption. However, one cannot simply say that a 1024-bit key is better than a 128-bit key, without some additional qualifiers. First, required key lengths depend on the type of algorithm. For example, an AES key of 128 bits is (at the time of this writing) the minimum required length for security, a property shared with many other symmetric algorithms.

Asymmetric keys, on the other hand, are typically much longer; 1024 bits would be a minimum for SSL (although in the wake of the POODLE and Logjam attacks on SSL in early 2015, 2048 bits seems advisable). Comparing key lengths and strengths really involves comparing apples to apples; for a given algorithm, and largely among algorithms of the same type, longer is better. Key length has received new scrutiny in the wake of so many break-ins to steal customer information, and the NSA domestic spying scandals. Irrespective of the algorithm, many experts now recommend generating longer keys; 256 bits or longer for symmetric ciphers and 2048 or more for public key algorithms.⁴¹ These longer keys do require somewhat more CPU horsepower to decrypt, but as processor power increases, this is acceptable. Increasing the size of keys used is generally simply a software setting, and can often be adjusted on the fly, or at worst, it might require a restart of the server. Google, for example, adjusted their asymmetric key size from 1024 to 2048 bits in mid-2013, well before the 2015 attacks on SSL.⁴²

How is encryption used to secure resources?

Now that you know the basic mechanics of encryption, what can it do for you? The obvious answer is to make your data confidential, protected from eavesdropping attackers, but the benefits do not stop there. Encrypted data may be protected both while at rest and during transmission to another party, and the protection also guards against accidental damage to the data, ensuring integrity. When data is encrypted, the encryption algorithm creates a checksum⁴³ that can be used to ensure the data has not been modified. Second, at least in the case of public key algorithms, the sender can be

authenticated. This allows both parties to be sure that they are communicating with the right person. This is an obvious benefit when the parties may be on opposite sides of the country or globe.

Where should encryption be used?

Where should encryption be used? Everywhere and always! Recent data breaches and governmental spying, not restricted to the United States, should be adequate evidence of the need for encryption. Three basic locations to encrypt will be discussed, but they are not all-encompassing. First, data in transit should be encrypted. This means using a VPN as discussed in Chapter 4. The second place to encrypt data is during storage; this may be subdivided into three types. First, full-disk encryption encrypts every bit of data on the computer's hard drive. Second, the built-in encryption features of databases can protect the valuable information contained therein. Finally, individual files can be encrypted.

Encrypting a file system is part of hardening a device, as discussed in Chapter 2. Mechanisms to accomplish this are built into almost every modern operating system and device. Before discussing when, where, and how they work, it is imperative to note when they *do not* work. A device or disk that is powered on, and unlocked with a password, so that it can be used normally, has already been decrypted. If a user is logged on to the device, the storage is unlocked, at least for that user's files. A device that is not powered on is fully protected. The specifics of what is unlocked, and at what stage of the boot or power-on process is beyond the scope of this book, but beware that a powered-on phone, for example, has no more protection than that provided by the unlock screen; e.g., a four-digit code. So, why use full-disk encryption?—To protect against data recovery following a theft or loss of the disk.

- In the case of a USB stick or drive, if it's encrypted, it cannot be read by someone finding it and plugging it in, without the proper password. Encrypting these disks is trivial, and should be done before issuing them to users.
- In the case of a hard drive installed in a server, if the server is stolen, it would likely be powered down in the process, and when booted, would need a passcode entered.

- For laptops, often a theft target, if the laptop is powered down, it cannot be powered on without the passcode. To enforce power-off behavior, the power properties can be adjusted to disallow hibernation or sleep mode, and force power-off when the lid is closed. Again, this should be done before the machine is issued to the user.
- In the case of backup drives or tapes, especially those stored off-site, as backups should be, encryption is a must. This protects the data against theft if the package is lost or stolen.
- In the case of a phone, it cannot be read if it was turned off when found; if it was on, the only protection provided is the screen lock code. Any cards already inserted in the device can be read on that device if the screen can be unlocked, but removing the card will render it encrypted again, and it cannot be read by an attacker.

Most modern DBMSs have protections built in that protect the contents of the database from unauthorized access. These protections include not only usernames and passwords for authorized users, but also encryption for individual database fields or columns. As with disk encryption, if a database is in use, the contents may be decrypted already, meaning that the sum total security available is simply a function of the strength of the usernames and passwords chosen for authorized users. The other downside of some types of database encryption is that they may impact performance. Despite these drawbacks, the use of encryption of critical database fields, such as credit card numbers, should be strongly considered. Given the high-profile thefts of hundreds of millions of credit card numbers, it is obvious that attackers are intent on extracting information from databases, as they contain the crown jewels of the firm.

Storing files in encrypted form is nominally somewhat more secure than full-disk encryption. Unlike protecting a full drive, where a signed-in user decrypts all files on the disk by logging in, single-file encryption only removes protection from the file when it's accessed. All other files not currently in use are still scrambled for confidentiality. This makes it more secure than full-disk encryption on a running system, and equally secure on a powered-off system. However, single-file encryption is

fraught with one simple problem: ensuring that the user re-encrypts the file after using it. Most single-file encryption systems require the user to explicitly decrypt the file, then open the unprotected version using the default program. The user then saves the file and encrypts the new version. This workflow tends to be disruptive. That said, there are cryptographic systems which automate the process, and tools like Microsoft Office can encrypt and decrypt their own documents on the fly, but usually not by default. Where it can be used, through automation and by training users, single-file encryption is a valuable weapon in the arsenal against hackers, who find they cannot open stolen files.

Cryptography obviously fixes confidentiality issues. An encrypted file cannot be read by an attacker. Encrypted data is also typically electronically “signed” in such a way that alterations can be detected. This gives integrity to the file or data packet. Cryptography also provides authentication of the sender. Finally, while availability is not directly influenced by encryption, there are indirect effects on availability through integrity; if the file is known not to be changed, it can confidently be used.

Cryptography is not a cure-all

As Bruce Schneier has emphasized, “Cryptography is not magic security dust that you can sprinkle over [your] software and make it secure.”⁴⁴ And, as another security pundit said, “If you think cryptography will solve your problem, either you don't understand cryptography, or you don't understand your problem.”⁴⁵ In other words, there are certainly problems that cannot be solved by encryption. Among these are the handling of encrypted file systems; if the user is logged in, all data is decrypted. Encryption, on its own, without proper key management, can lead to disclosure of data. Further, if a firm does not escrow users' keys, but allows self-management, when an employee leaves the firm, they could simply refuse to decrypt the data, meaning the firm has lost access to it.

Summary

While not a magic solution, cryptography is an absolute must in today's computing environment. The types of attacks that can be prevented by simply having the data in unreadable form justify any inconvenience associated with encryption.⁴⁶ The added data integrity and authentication benefits are icing on the cake. Much like firewalls, knowing which problems can be solved by encryption and which need another solution is vital. Users must also be educated in how to use encryption, and the dangers of things like accepting invalid SSL certificates.

CHAPTER 7

Danger! Danger! Danger! (Penetration Testing)

How do you stop attackers from getting into your network? Of course, there are many ways to protect your assets; we have already discussed a number of them. But, the security provided by a firewall, WiFi encryption, or any other protection, including passwords, is only as good as the implementation. Given the complexities of setting up security solutions properly, penetration testing is required. Penetration testing, also called pentesting in the security world, involves using the same tools an attacker would use to get into the network illicitly to find the holes in a system.

Professional penetration testers, sometimes called *white-hat hackers*, are professionals who try to breach defenses before *black-hat hackers*, the attackers, can find the weaknesses.⁴⁷ Penetration-testing personnel are trained in the same tools and techniques that their less scrupulous counterparts use, including taking an organized, systematic approach to their work. A properly conducted penetration test, with well-documented results, will allow the firm to fix vulnerabilities before they can be exploited by adversaries. Properly executed, a penetration test can ensure compliance with regulations such as HIPAA, PCI DSS (Payment Card Industry Data Security Standard), or Sarbanes-Oxley. In fact, requirement 11 of PCI DSS actually requires regular testing of security systems and processes.⁴⁸ The reports resulting from such a test can be used to demonstrate to auditors that proper measures have been taken to protect IT resources. A network or system fixed following the recommendations of a good test can save the firm millions in legal costs and restitution that would result from a major breach.

Internal vs. external testing

Penetration testing may be performed by internal or external teams. Each approach has merits, and each has disadvantages. Internal penetration testers are retained by the firm, and work day in and day out on the same systems. External penetration testers, on the other hand, may work for a week on one system, then move to the next firm for a new engagement. This puts internal testers in a better place to be familiar with the intricacies of the firm's defenses, and also may lead to shortsightedness, or overlooking obvious flaws. An external tester may be able take a naive view that allows them to discover problems that internal employees consider solved.

Another danger of any penetration testing, internal or external, is that you often hand the keys of the kingdom to the penetration tester in order for them to do their job. This may include IP addresses of all devices on your network, network names, and usernames. It is not necessary nor advisable to give out passwords, in most cases. This information can obviously be used by those who are less than ethical to steal information from the company, plant *backdoors*⁴⁹ on systems, and wreak other havoc. Some suggest that external firms pose a greater danger; but realistically, either a penetration-testing firm or the company which hires internal testers has the same chance of hiring unethical employees.

Who can be trusted to perform penetration testing? While some standards exist for certification of individual pentesters, the decision comes down to thorough background investigation coupled with careful monitoring. Another common discussion among security professionals is whether ex-black-hat hackers, especially those who have served jail time, should be hired to perform penetration tests. They are obviously skilled enough to have penetrated the defenses of at least one system. The big question is whether they can be trusted. Many security professionals share the opinion that no convicted hacker can ever be trusted again. This same logic pervades most human resource departments; rarely would someone with a criminal background be hired. This logic is sound, but there may be exceptions. The particular skills needed for certain penetration tests may be most readily found among those who have done time for their talents. This does not apply to routine pentesting; in

most cases the advice of not hiring convicts should be followed. With that said, there are a few high-profile ex-hackers who have made good names for themselves in the business by founding consulting firms.

When selecting an external firm to perform tests, the choice really comes down to reputation. Whether an external penetration-testing firm decides to hire ex-hackers or not, the firm itself should be thoroughly vetted. Consulting with business partners and other companies can help direct your choice. There are several large and many smaller firms with established reputations, and certainly other rising stars.

Once a firm is selected, contracts must be signed. Like any agreement between firms, the paperwork should be verified by the legal department. Especially important in penetration-testing contracts is a description of what will happen with the results of the test. Even if penetration testers are not given a great deal of information to start their tests, their findings will show the vulnerabilities in the network. These results have enough information to make the job of any would-be attacker much easier. Therefore, distribution of the reports should be restricted to only those who have a need to know. Sanitized summaries or executive reports can be prepared if wider distribution is needed, they should contain only brief descriptions of the problems found, without identifying system information or anything that might be confidential.

How penetration testing is performed

Penetration-testing methodology closely mirrors real attacks by hackers. It is instructive to walk through the steps of a real test to show how attackers might get into IT systems. The steps that a hacker would not perform are the first, third, and last steps.

The basic steps to be performed, as outlined by the Penetration Testing Execution Standard, are the following:⁵⁰

1. Pre-engagement interactions
2. Intelligence gathering
3. Threat modeling
4. Vulnerability analysis

5. Exploitation
6. Post-exploitation
7. Reporting

Pre-engagement interactions are the interactions between the firm and the penetration testers. In these meetings and documents, the key elements to be addressed are permission to test and the scope of the engagement, including some of the following questions:

- What are the (specific) goals of test? It is not sufficient just to say “We want you to test our network to make sure it’s secure.” It is much more feasible to say “Please test all hardware and software firewalls at our borders and on servers A, B, C, & D.”
- Is this test required for specific regulatory compliance?
- Will the test be performed during or after business hours? Pentesting during business hours has greater potential to interrupt productivity.
- How many IP addresses will be tested?
- Once a machine is penetrated, how should the testing team proceed? Should they attempt to elevate privileges? Should they crack passwords on that machine? Should the team attempt to exfiltrate information or shut down services?
- Will physical penetration testing (attempting to break into doors, windows, and buildings) be part of the test?
- Is social engineering (attempting to trick people into doing things against their own best interest, such as revealing passwords) part of the testing?

One of the main things that a firm should understand in terms of scope is that the penetration testers, especially those external to the firm, must manage scope carefully. Scope creep is the death knell of any pen-testing project, and possibly even the pentesting business. It is very easy for a firm to request more work as new vulnerabilities are discovered during a test, but the firm must realize that this additional work will be billed by the pentesters. As with any business, pentesters sell their time and expertise, and changes to the original plans must be compensated.

Intelligence gathering is performed by legitimate pentesters and hackers alike. Both will likely use many of the same sources, and the findings of the pentester will inform the firm of their weaknesses in leaking information. The first steps in information gathering consist of finding simple information such as names, addresses, and ownership records for the company and its premises. Following this, tax information, physical security measures in public view, and all network information that is public can be gathered. In this first step, the focus is on public information. While it may seem that public information cannot be harmful, consider that these same sorts of activities are performed by burglars while “casing the joint.”⁵¹

Other sources of intelligence can include company directories, job openings, and charity affiliations. The first, which may, but should not be, posted on the company website, can be used for social engineering. The same goes for charity affiliations. Job postings give a wealth of information about the technologies in use in the firm; a firm asking for a database administrator with 15 years of Oracle experience almost certainly has a large investment in that software, meaning a hacker now can focus on attacks that target Oracle’s products specifically. RFP or RFQ (Request for Proposal/Quote) documents expose similar information. One final source of open information, that is increasingly difficult to control, is the information leaked by individuals on social networks. Finding the names of everyone who works in a firm may be difficult, but once one is found, they are likely to be friends with their coworkers online. Reading their posts and profiles may be a treasure trove of information about company internal politics and plans.

What can a firm do to avoid this information leakage? In some cases, not much. Much of the public information is mandated by law, but specifics about the firm can be omitted from IP address registries. These are easy to sanitize, usually with just an email to the domain registrar. Removing company directories and organization charts from the web is also relatively easy. Controlling social network information is much more difficult, but implementing a policy about what may and may not be posted is vital to the firm. Such policies must have consequences, and a firm must be willing to enforce those policies.⁵²

Threat modeling is performed only by ethical penetration testers and involves analyzing each asset in an information system and the corresponding threats it faces from attackers. Each asset obviously faces different dangers, which vary according to the value and sensitivity of the asset as well as the skill of the attacker. Unfortunately, this is the most difficult part of the pentest process. Knowing which targets an attacker might most desire is probably the easiest portion; if there is money or personal information involved, the system is a target. Other highly vulnerable systems include those with business plans or research and development documents. The organization's remaining systems are much more difficult to assign value or danger levels to.

Asset sensitivity is also difficult to rate for two reasons. First, sensitivity depends on what exploits are available for the attacker to use. These change every week as new exploits are discovered and existing vulnerabilities are patched. Second, as mentioned in Chapter 2 in terms of risk analysis, is the many-to-many relationship that exists between assets and protections. This relationship means that it's very hard to decide how sensitive or vulnerable a certain asset is. A given system may be protected by a border firewall, an intrusion detection system, a host-based firewall, antivirus software, and other security measures. It's almost impossible to determine which of these reduces the vulnerability of the system the most. Likewise, most of these security measures protect more than one computer, and it's difficult to determine which shields that machine the best.

One final avenue that may be explored while modeling threats is to look at relevant news of compromises to similar organizations. These reports may be very difficult to obtain, as companies may not make them public. Worse, what is openly released only represents a small fraction of the overall incidents that a firm has faced. It is worth seeking these reports, but care must be taken in terms of what may be believed. Ultimately, it must be said that while threat modeling can and should be performed, finding credible results can be extraordinarily difficult.

Vulnerability analysis is a nice term used by pentesters to cover the processes used by hackers to find the holes in your system. Details of the process vary with the component being tested or attacked, but the basic steps are the same for all. First, active and passive scanning of the network

or system takes place. Passive scanning involves looking at things like metadata in publicly available office documents (such as authorship, when it was created/saved, and so on), and monitoring traffic to and from the target. Active scanning, perhaps the most “famous” type, as seen in movies, involves using tools to probe defenses. The results of active scanning are probably the most valuable, but it also is the most dangerous for the attacker, as it leaves traces. Typically, an active scan will examine the IP address associated with the target, looking for open ports, which represent potential ways into the system. The result of an active scan is a list or map of the target system, which allows an attacker to prioritize where to focus their efforts in the exploitation phase.⁵³

Exploitation is the first part of the actual attack. In exploitation, an attacker employs the prioritized list of targets generated during vulnerability analysis. Each target might have multiple potential vulnerabilities or attack vectors. For example, a single machine might be acting as a web server, a file server, and a database server. Each of these applications may have multiple possible vulnerabilities. The attacker must use deeper probes of the server to see whether the particular installed software versions on that server actually are vulnerable. If the installed software does not respond to the exploit, then the attacker simply crosses it off the list, and moves to the next potential vulnerability. From a penetration-testing standpoint, that server or service has passed the test; from the standpoint of an attacker, it means they have been thwarted. However, if the application responds to the attack by allowing unwanted access, then the hacker has reached the first goal. Once access is gained, the adversary has a number of possible steps they could take next.

Of note during the exploitation phase for pentesting is the fact that it is the first part of the process that represents a real danger to your systems. Exploitation, by definition, usually consists of attacking software that is broken in some way. How this software responds to the break-in attempt varies; but many exploits are able to crash at least the particular software being targeted, and perhaps even the whole system. In a pentesting course taught by one of the authors, students are frequently frustrated when they mistype a command in lab, and as a result the target system crashes, rather than allowing them in. At that point, they must manually reboot the system before they can even try again; it simply freezes.

Depending on how deep a pentester (or attacker) probes, it is well within the realm of possibility that they will crash a server. Whether this is intentional depends largely on the scope defined in the pre-engagement meetings, but even with a very limited scope, accidents do occur. For this reason, the contracts signed with the penetration-testing firm will contain a hold-harmless agreement. This protects the pentester, as long as they stick exactly to the plan. Any deviations from the plan, then, to examine things found in the course of testing, would require written authorization from the organization being tested.

During the postexploitation phase, after gaining access, a hacker will always take one of several steps; the behavior of the pentester depends on the rules for this particular engagement. An attacker will attempt to elevate privileges, exfiltrate data, or install backdoors into the system. They may also attempt to directly damage the software on the server such that it must be formatted and reinstalled before it's usable. Attackers also seek to cover their tracks, so that their presence is undetectable. The best attackers do a very good job of this; even after the fact, it may be almost impossible to find out how they got in. The pentester will often attempt privilege escalation or data exfiltration, but will generally stop short of doing intentional damage or installing backdoors on the system. The ultimate goal of most attackers today is to steal data. Once the hacker has met this goal, they may leave the system intact, merely cleaning up evidence they were ever there, or damage the system to brag or prevent its use by the organization. This is end of the hacker's activities.

A legitimate pentester goes one step further and generates organized reports to present to the client. These reports will detail all vulnerabilities found, by system or in some other logical fashion. Once the reports are received, the organization must act on them. There is rarely enough time or money to fix everything at once. The firm, aided by the penetration-testing team, must prioritize the repairs by the value of the asset and severity of the threat. Much like initially assessing threat to your IT systems, this can be difficult, but the firm must take the steps to fix the most important threats to their security. Within budget, the firm must fix everything found by the tests.

Volunteer penetration testers

“Volunteer” pentesters can be categorized as hackers or opportunists, or alternatively as good Samaritans. Volunteers come to light when their activities, which may not be considered legal, uncover and report a vulnerability in the firm. These reports may or may not be anonymous. Once received, the firm must decide what to do with the report, assuming it’s credible. Two examples may help to illustrate both how this happens and the issues surrounding such an event. Both involved students of one of the authors.

The first incident occurred in a local restaurant. Like many, they offer free WiFi to their customers. More than once, the student had been there for lunch, and not been able to get WiFi access. After a discussion one day in a networking course, the student surmised that the problem was likely that the access point was configured to only hand out a few IP addresses, insufficient for the lunch rush. Rather than ask the manager, the student used another bit of information easily found online; that is, the default username and password for that particular brand of wireless router. Using this, the student logged into the router, and changed the setting to allow the router to hand out more IP addresses.

Shortly thereafter, the student informed the instructor what they had found and the actions taken, then asked for advice about whether to tell the owners about the issue. Realizing that the student’s actions probably constituted a felony (under hacking statutes, any unauthorized access to a computer system is criminal), the instructor advised the student to proceed cautiously; ultimately the student decided to approach the manager. While the student did not share exactly how the findings were presented, the upshot was that the manager asked the student to fix the problem, and also asked the student to fix the issue of the default username and password. In exchange, the manager gave the student a few free lunches.

In the second case, another student found vulnerabilities in the websites of two local businesses (both, coincidentally, within a block of the first establishment). This student had been doing some amateur probing of the websites of eateries from which they had ordered, ostensibly to ensure their credit card information was safe. In one case, the student

found that the shopping cart was vulnerable to SQL injection, and in the other, the website was vulnerable to cross-site scripting (XSS) attacks which could result in leakage of customer's personal data.

This student also approached the author, asking what he should do in each case. In the SQL case, the responsible party would be a firm that wrote the e-commerce shopping cart app. After some research, the student found that the firm had tended to be litigious in response to unsolicited vulnerability reports. The student was faced with the following choices:

- Inform the authors of the software of the bug anonymously, in which case, nothing would likely be fixed.
- Inform the authors of the software in a fashion that could be traced to the student, in which case, legal action might be taken against the student.
- Inform the business owner that the software they were using was flawed. If anonymous, likely no action would be taken, as the owner probably would not understand the issue anyway. If credible reports were given to the owner, the means by which they were obtained was questionable.
- Go public, anonymously or otherwise, and publish the bug on one of several bug-tracking databases. In this case, a number of hackers would likely see the bug, develop exploits for it, and use the exploits to steal personal data of others.
- Do nothing. Perhaps the bug would remain hidden, perhaps not.

Each option given has some ethical and some legal implications. After some discussion, it was determined that the best course of action for the student in this case was likely to go public, perhaps anonymously, with the bug. While it might result in some information theft, the authors of the software would be forced, by the court of public opinion, to fix the problem.

In the second incident with this same student, the XSS vulnerability, the firm responsible was a local web development firm. The choices were largely the same as for the SQL injection, but the firm had not

shown any propensity for lawsuits. In this case, the student decided to approach the firm privately, show them the issue, and, if needed, offer assistance in repairing it. As this book was going to press, no further details were available as to how either firm had responded to the student's actions.

These three cases show how vulnerabilities can be found by such "volunteers," as well as the range of what a firm can do in response to such notifications. Whether the manner of finding was accidental, and therefore legal, or due to some quasi-legal or outright illegal activities really does not matter much. If someone comes forward with a potential vulnerability, the firm should take the matter seriously and investigate whether it represents a real attack vector for their firm. It is this author's opinion, not shared by all, that the firm should not pursue legal action against someone who points out problems with their systems, as they have likely done the firm a favor. This does not apply, of course, to someone who takes advantage of an exploit, and does harm to the firm, merely to those who point out issues.

Summary

Whether it's called ethical hacking or pentesting, having someone who can be trusted break into your network for testing is exceptionally valuable. By finding the holes in IT systems, software and hardware configurations can be validated and break-ins prevented, rather than merely stopped. While good penetration testing does not come cheap, it represents an investment in security that can pay large dividends. The payout is only realized if the firm actually fixes the issues found and reported. Even if a "volunteer" provides a report of vulnerability, the firm should take it seriously and fix the issue. Whether they decide to pursue legal action is secondary to securing their systems.

CHAPTER 8

Disaster Recovery

The focus of disaster recovery should be answering the following question: "When something bad happens, how fast can your business be back up, online, and functioning at levels acceptable to our business and any governing bodies?" The outcome or answer to this question should lead to a testable, actionable disaster recovery plan (DRP). The goal of this chapter is to provide an overview of what types of disasters should be considered in a DRP and serve as a launch point for creating a plan or revising a current plan to meet your organization security goals.

What is a “disaster”?

A disaster can be an act of nature or man-made occurrence that shuts down your organization in part or in whole. A disaster might be a hurricane taking out your East Coast data center, a blackout in the west that has shut down all of your operations in California, or malicious hackers who have orchestrated a successful denial of service attack on your organization’s website. Disasters can occur due to failure of resources such as a server or corrupted data transactions. A good disaster plan will account for these as well as other items specific to your business, such as supplier shortages or employee strikes. While most disasters go beyond affecting just IT, IT plays an ever increasing part in most organizations. From an IT perspective you should focus on addressing hardware, connectivity, data and application availability.

Depending on your organization’s industry, you may have government or industry guidelines that must be considered in your DRP. Standards that may need consideration to determine if you are in compliance in a particular industry include NIST documents, Federal Information Security Management Act (FISMA), PCI DSS, Family Educational Rights and Privacy Act (FERPA), or HIPAA.⁵⁴ Many

organizations provide frameworks or guidelines to meet these standards such as the NIST Risk Management Framework (RMF) and ISACA's COBIT (Control Objectives for Information and Related Technology) framework for information security.⁵⁵

Securing against catastrophe

Regardless of where you fit in an organization, IT or otherwise, protecting against disaster is important for the longevity of all businesses. It is not a question of whether a disaster will strike your organization; the question is when it will strike. Proactively considering disaster recovery is vital to be ready when disaster strikes. The first step is determining who will be in charge of the DRP as well as identifying key stakeholders who should be involved in the development and maintenance of the DRP.

The challenge when building this team is getting buy-in at all levels for spending money on something you hope you never have to use! It is a challenge to estimate return on investment for security breaches or disaster-driven losses. The best estimates are driven by calculating costs incurred if a breach or loss occurs. Scenarios presenting the high, medium, or low estimates on an asset-by-asset basis provide a gauge for what could possibly happen.

What to consider?

Once the team is assembled, several key items should be considered as the team begins to formulate the DRP. Some items to address or achieve common understanding on include the following:

- a) **Determining disaster threats, affected assets, and impact.** Questions to ask include: What is the size of the breach? Who is affected? Is only the firm a victim, or customers, employees, or business partners? A breach may affect our "almost" employees, does our applicant system protect the information of job applicants as they apply? Calculations about cost should include costs for hardware replacement and labor to run the data recovery jobs, and cost of legal actions if needed.

- b) **Putting time into training and planning.** Since return on investment can be difficult to calculate, putting time into training and planning can be difficult to justify to decision makers. Leveraging other high-profile breaches as well as providing clear data based on the likely outcome of potential scenarios can help you make the case. Training and planning are like R&D in that the payoff comes over the long run and the companies that are most successful have these strategic efforts in place.⁵⁶
- c) **Approaches to recovery.** Most database management systems include tools to log transactions and database changes and manage checkpoints to rollback/rollforward when data is corrupted or lost. Furthermore, entire suites of bolt-on tools or applications can help to further manage data and file backup and recovery. The key for a DRP is to have the tools in place and monitor them. Making sure critical data and files are backed up offsite is integral to successful recovery. Resources such as virtual data centers or site mirroring can be leveraged to successfully recover data and application resources in the event the production environment goes down.
- d) **Backup strategies.** Backups are covered in detail in Chapter 10, as they are essential for home users and enterprises alike. Rather than rehashing all details of backups, a few differences for businesses will be highlighted here. As with home users, backup strategy revolves around what should be backed up, and how often. For businesses, data is the most critical thing to back up, but it may make sense to back up programs in image-based backups, especially for critical servers. Short recovery time is vital for servers, and image-based backup allows the fastest restoration. In terms of frequency, businesses will need to back up more often, both because they create more data and because that data is more mission-critical. Backing up every few hours is probably within reason for many businesses, and no less often than once a day would make sense. The last question is what medium to use. For a business, tape drives are a likely solution. Tape drives are fairly slow, and may seem like old technology, but have two advantages, large capacity and easy offsite storage, allowing for more effective disaster recovery.

- e) **Evaluation of successful recovery exercises.** Metrics for success should be specified in a DRP. Metrics are useful for evaluating both the practice sessions and actual recovery. Metrics to consider include: How fast threats are identified, how quickly the threat can be neutralized, and how fast your critical systems can be back up and running. Certainly, many other metrics could be defined for your particular business or industry. Beyond metrics, qualitative measures include identifying issues that were encountered that could be improved upon.

Forming the team and assessing the expectations and requirements for the DRP is the necessary first step. Once the team is formed and requirements have been considered, the next phase is making it happen.

Making your DRP a reality

Once the key issues that are important to your organization are defined, the team can begin to map out the DRP. Templates for DRPs abound and can be a good place to start.⁵⁷ Evaluate many options and find one that fits your business. The key aspects include formulating the plan, implementing or executing the plan, testing the plan, and regularly revisiting or updating the plan.

1. **Plan for everything and the kitchen sink.** The DRP should account for a variety of scenarios. A temporary outage due to an accident that took out a data trunkline, and a weeklong blizzard that has shut down the entire East Coast require different approaches. The DRP should account for all conceivable disasters. What contingencies need to be in place? How can we leverage resources in the cloud versus stand-alone datacenters?
2. **Implement.** Depending on your organization, you may have much or all of your plan implemented. This section will present things to consider from the viewpoint of an organization that does not have a DRP in place or the DRP is out of date. The list below provides an overview of the major components that should be considered as

part of DRP in no particular order as each is equally important in its own right; a complete checklist can be found in Appendix A. Key questions to ask when reviewing your DRP are provided.

i. *Physical Resources*

In the event of a disaster, can all resources be secured both physically and from unauthorized access (i.e., use of passwords, locked and environmentally controlled rooms, limited remote access, etc.)? An inventory of resources and the corresponding application dependencies should be included in the DRP. What systems are critical in the running of the business? What temporary failover resources can be leveraged to get key systems online while physical issues are addressed? Are hot or cold sites available? Are offsite backups stored at a location that will be accessible if a local natural disaster occurs?

ii. *Data and Applications*

Data: The data should be replicated, mirrored, or backed up in at least one offsite location. What transactions were potentially lost during the crisis?

Applications: Applications should be inventoried and a list of vendors generated for third-party software the organization utilizes. If custom applications are running on the server, where does the backup code or executables reside? If it is a vendor supplied solution, what backup procedures should be followed? What recovery services or special cases should be considered? Are software keys and media backed up to a safe location where they can be accessed for reinstallation?

iii. *Networks and Connectivity*

Both internal and external networks need to be considered during DRP implementation. In the face of a disaster, what internal routing is necessary for our core systems to function? Are our outside facing resources visible? How do you overcome or deal with outages outside our control in regions where the firm conducts business?

iv. *Workstation Resources*

Finally, it is necessary to consider the resources that are the foundation of your organization: the files on your employees' workstations and shared drives. These documents, presentations, spreadsheets, and other personal productivity resources are what ultimately drive your business. While standing up critical networks, servers, applications, and supporting data are key in the short-term recovery, what mechanisms are in place to back up and recover the unstructured data captured in business process documentation or even employee emails? Using a content management system and providing backup of all support systems is important for the long-term health of an organization in the face of a disaster.

Ultimately, the implementation process should allow for the identification and prioritization of all critical systems in the case of a disaster and full organizational shutdown. The process will initially be expensive but yields the benefit of a better understanding of how your organization's IT functions. The exercise may also provide cost savings from identifying noncritical overlap or opportunities for business process reengineering!

3. **Practice.** The adage "practice makes perfect" is applicable in disaster recovery. Testing and evaluating the DRP in real-world scenarios is important to prepare our employees and provide key information to decision makers. Some key items to consider are the following:⁵⁸
 1. Define specific exercise objectives upfront.
 2. Include business stakeholders.
 3. Rotate staff responsibilities to allow cross-functionality.
 4. Develop specific risk scenarios for your exercises.
 5. Run joint exercises with business continuity (BC) teams.
 6. Vary exercise types from technical tests to walk-throughs.
 7. Make sure to test all IT infrastructure concurrently at least once per year.
 8. Identify members for the core disaster recovery response team.

9. Learn from your mistakes.
10. Report results to stakeholders.

This process then creates a feedback loop for the entire team, which through cross-training and involvement of both technical and business stakeholders helps foster a security-focused enterprise.

- (4) **Revisit, revise, and report.** The last two items in the list above should stand on their own. Having a DRP and testing it is great, but if you do not share what is learned from the testing process, then the risk of your DRP failing in a time of need is great. A review period should be built into the DRP where each aspect is assessed in light of current security trends and concerns. Addressing aspects such as how your organization has changed since the last review, including new technologies or processes, as well as how the external threat landscape has changed is imperative. Steps should be taken to revise the plan to meet these changing criteria and incorporating lessons learned from testing. Finally, reporting out the changes and results of DRP testing closes the loop and provides an opportunity for feedback from stakeholders outside the disaster recovery team.

Depending on the size of your organization the DRP might be a large document or might be a summary document that points to unit level DRPs. Testing might also be done on a unit-by-unit basis, but testing a full organizational shutdown might indicate issues the sum of the parts will not. Whatever the approach selected, someone should understand the big picture and keep an eye on results of DRP testing and revisions.

Summary

Having a DRP in place is vital in today's organizational landscape. The potential of malicious attacks or some aspect of your business being affected by a natural disaster is a reality that will eventually occur, and proactively anticipating the potential response might be the difference between hundreds of thousands or millions of dollars of loss to your business. The primary key to a successful DRP is having a team in place to develop and

keep it current. The DRP should have evaluation built in to help determine the success of the recovery. Finally, practicing and testing the plan are key, evaluating the outcomes and incorporating the feedback into plan revisions as needed. The key is to lock it down and limit access while continually reviewing policies and procedures, evaluating new technologies that will allow for increased security, but beware of the bleeding edge!

CHAPTER 9

Integrating Your Security Plan across the Enterprise

Every organization needs a written security policy. Every organization has a stance on information assurance, but if it's not recorded and formalized, it might as well not exist. The organization will eventually encounter situations that require taking action about a security matter. If no written policy exists, many actions they would wish to take become legally impossible. As a simple example, consider the case of an employee who wastes work time surfing the Internet for sports scores. This is not illegal (illegal actions have clear legal penalties, giving built-in consequences and procedures for the organization to pursue), but most firms would frown on this waste of company time and resources. The employee is warned, but continues with the behavior. At that point, the firm wishes to fire the employee. When they attempt to do so, their legal counsel warns them that this is not possible, as they do not have a policy that prohibits personal Internet use on company time. This situation could be avoided with a relatively simple Internet acceptable use policy.

The policy actually typically consists of a series of policy documents that work with each other. One document is an overarching policy that contains the basic statements of the firm's position on information security. A set of subsidiary policy documents seeks to govern specific aspects of security, such as the aforementioned acceptable use policy. Care should be taken in developing these documents to ensure that they support the main security policy document, and do not contradict it in any way.

Who should develop the firm's security policy? Certainly it must be the domain of IT, meaning the IT security staff should develop the policy, and then present it to management and employees, right? Truly, this is likely to be a failed effort. The likely outcome of such a course of action is

for IT to write a policy that is not only unacceptable to other employees, meaning it will not be followed, but also the policy will probably focus too much on technical issues and ignore behavioral and legal issues. Best practice dictates that security policies be developed with top management support.⁵⁹ This means the policy will be based on business value, and what value specific assets have to the firm. Further, the policy needs to be a team effort. The team must be interdepartmental. As mentioned, an IT-centric team will focus mainly on IT matters; if HR were tasked with development, it would reflect their biases and worldview. The minimum requirements for a team would include top management, IT management, perhaps IT line employees, HR, and legal. Other units may be added if deemed necessary.

What does the policy contain?

A good security policy, at least the main, umbrella document, does not seek to cover every eventuality that may arise in the history of the firm. Besides being patently impossible, it would render the document extraordinarily unwieldy. Instead, the top-level document should state the aims of the policy, to whom it applies, and the importance of adherence, along with the basic consequences of nonadherence.⁶⁰ These basic tenets then guide more specific documents, often called standards documents or procedure documents. Standards documents contain details pertaining to common situations that may occur in security matters. The security policy document guides individuals on why and how they should follow the policy, but not exactly what they should do.

A good security policy, being concise, may only be a few sentences to a few paragraphs long. Some example statements from existing security policies may aid in understanding what should be in the policy. Two such examples appear below. Both are excerpts, and should not be construed as full policy documents; however, the first example approximates the minimum length of a policy.

Example 1

Purpose:

The purpose of this policy is to establish a University-wide framework for the protection of State University's information technology resources, computers, networking systems, and data. This framework aids the University in meeting its obligations with regard to information security and privacy.

Scope:

This policy requires controlling the security of and access to the University's IT systems and information, and allows the University to develop procedures to secure and audit their IT resources.

Compliance:

Policy violations shall be reported in accordance with University procedures. Types of violations include, but are not limited to, breach of confidentiality, breach of privacy of student or faculty records, theft of intellectual property, vandalism, or damage of IT resources or computer systems. The University will impose sanctions for policy violations, in accordance with the extent of the violation. Sanctions may include, but are not limited to, limitation of computer or network access, disciplinary actions including loss of student or faculty status, financial restitution, or legal actions.⁶¹

Example 2

Introduction:

Information is a key resource of our organization, without which operations would cease. Our information includes: all data; computer systems; and hardcopy documents. Whatever forms the information takes, it must always be protected.

Our security objective is to protect our organization from security problems that might interrupt our operations or damage our reputation.

(Continued)

Security problems may include breaches of our own or client confidentiality, integrity of our information, and loss of availability. This wide definition of security includes all threats to our information.

Scope:

All permanent and temporary staff, past and present, of the organization are bound by this policy to protect our information infrastructure. All staff will act in a responsible, professional and security-aware way that conforms to this policy.

The president of the organization has final responsibility for ensuring that all the organization's and third party information over which the organization has stewardship is adequately protected. Day-to-day duties in overseeing information security will be delegated to specific individuals.

Policy Awareness:

This policy will be made available to all staff upon hiring. The policy or sections thereof will be updated as needed to adapt to changing needs. These changes will be made public when they occur. All staff are expected to be familiar with such changes, and comply with all sections of the policy. Any staff seeking clarification of any part of this policy should discuss the issue with a member of the security policy development group.

Applicability and Enforcement:

This policy applies to all members of the organization who use its information resources. This policy forms part of the conditions of employment of all staff, at all levels.

Failure to comply with the information security policy could damage both the reputation of the organization and the impair its ability to achieve its aims. Failure to comply will result in sanctions, including, but not limited to disciplinary action, dismissal, or legal actions.⁶²

To whom does it apply?

An information security policy absolutely must apply to everyone in the firm, from the CEO on down to temporary employees. This requirement is needed because anyone in the organization can represent a potential security breach. When a breach occurs, it must be possible to discipline anyone in the organization.

Enforcement of security policy is not the domain of IT. This has some implications for IT security, who sometimes see themselves as “police,” or “military.” However, neither metaphor really works for IT security. First, IT security’s job is preventative. Neither military nor police really works in a preventative manner; police are concerned with catching miscreants, and military can use fatal force to punish. Security cannot even punish—enforcement falls to HR. IT security seeing themselves as military or police merely creates a negative view of users.⁶³ Although IT security is tasked with frontline defense and will likely discover violations of policy, their job is simply to pass this information to management and eventually to HR.

Developing a security policy

A cross-functional team should develop security policies, but policies are not developed in a vacuum. When writing or revising policy, authors should base their work on the best practices of others.⁶⁴ Many books have been written on developing security policies, the team should reference them and avoid reinventing the wheel. Among the resources that should be considered are the various standardized IT governance frameworks, which include guidelines for security planning and implementation. These frameworks include COBIT, the COSO internal control integrated framework, ITIL, and the ISO/IEC 27000 family, and others for specific industries. While these frameworks are written about planning and implementation, policy developers can work backward from implementation to the policy that will guide it.⁶⁵

When writing the policy, the team needs to focus on clear language and set a proper tone. Security policies should focus on the importance of information resources to the organization. Tone should be positive, emphasizing how the policy protects information resources, and spelling

out expectations and punishments. Do not center on punitive actions, but rather the benefits to the company when policy is followed. The policy should be concise, but explain rationale when appropriate. Likewise, define any terms that may confuse the reader, after all, you want employees to understand the policy.

Policies are not developed in a vacuum, nor in a walled castle, and simply tossed over to the rest of the firm. IT security policies need buy-in from the rest of the firm. This needs to start at the top. Upper management is subject to the policy and needs to publicly present the advantages of following policy. HR, who typically handles all training in the firm, can then educate users in both the importance of the policy and specifics of how the various elements should be followed.⁶⁶

Summary

It is no understatement to say that policy is the most important tool in the security manager's toolbox. Policy drives all IT decisions. The policy guides IT security (remember, they did not author the policy in isolation) in implementation guidance. In turn, this guidance informs those writing mandatory standards and guidelines for best practice. The standards, guidelines, checklists and baselines developed for specific systems limit the discretion of implementers. This ensures that all systems are installed and configured in accordance with policy. At any time, the firm can use the oversight provided by the policy to guide their implementation, as shown in Figure 1.

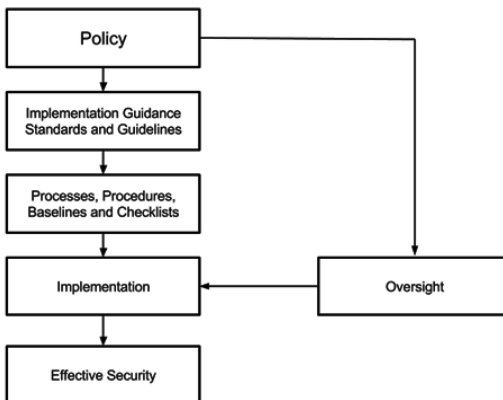


Figure 1 Security Policy Provides Oversight⁶⁷

CHAPTER 10

Conclusion

Throughout the book, we have made the effort to highlight the importance of security in all aspects of IT (and sometimes outside of IT) for your organization. The goal has been to provide guidelines (or at least food for thought) to help decision making. The next time you are in a meeting, planning to purchase the next big piece of hardware or software that will take your organization to the next level, you will be able to consider these concepts. In this chapter, we will highlight some current IT trends and security issues that might be encountered as well as leave you with some final thoughts on security and information assurance.

Security trends & future concerns

The one constant in the IT field is change. The principles defined in this book should provide a foundation that adapts to future changes. In this section, several current hot topics in security as well as some thoughts looking ahead into the near future are presented.

SCADA security

SCADA stands for Supervisory Control and Data Acquisition. SCADA networks are the systems that control industrial processes and monitor distribution systems. The nation's electrical grid, municipal water and sewer treatment systems, oil pipelines, and refineries all use SCADA to test water quality, make appropriate adjustments to the process, or open and close valves remotely. Factories use SCADA to control machinery, control flows of parts and chemicals, and watch for failures. Business buildings use SCADA for HVAC (heating, ventilation, and air conditioning) and mechanical controls. Even life support systems in hospitals

can be considered SCADA systems. SCADA systems require protection in their own right, but should also be protected because compromise of a SCADA system can give an attacker a foothold into corporate assets.

SCADA systems are IT systems and suffer from many of the same vulnerabilities, but have some unique features and flaws. First, SCADA systems control the physical world far more than a typical piece of software. The system gets inputs from sensors, makes a decision based on that data, and then controls some actuator to take an action. Because of this, SCADA systems are more critical than a typical IT system. In most cases, human life or safety is affected by failure. More protection is needed for SCADA than a typical server. Second, SCADA systems are built from custom hardware and software with a much longer life cycle. Organizational IT is usually refreshed on a 3- to 5-year cycle, but SCADA systems often have 15- to 20-year lifespans. The original manufacturer of the system may not even be in business by the time the installation is retired. This means that patching SCADA systems may not be possible. Even when patches are readily available, the process may be difficult. The geographic distribution of the controllers and monitors, coupled with the criticality of the system, means that the process must be carefully planned to take place on days when downtime is acceptable.

There are many specific measures that should be taken to assure that SCADA systems are not compromised. Excerpts from a short list developed by the US Department of Energy appear below, and many more be found in NIST guide 800-82.⁶⁸

1. Identify all connections to SCADA networks, including old modem pools and cell phone connections, then disconnect unnecessary connections.
2. Do not rely on proprietary protocols to protect SCADA installations.
3. Implement the patches and security features provided by vendors.
4. Physically protect SCADA assets, a challenging job given the geographic distribution of the parts of the system.
5. Focus on a risk management process to assess risk to SCADA, and manage changes to the system.

As computers control more of our smart homes, factories, and refineries, SCADA security will become even more important. Hopefully vendors will focus on standards-based hardware and software, and implement better security in future products. In the meantime, managers and system administrators will need to protect existing systems through application of outside controls.

Big Data

Organizations are increasingly storing and using large amounts of data. The need for incorporating security into big data retention and analysis has also grown. In light of this, we must consider the impact of losing these strategic resources. Each presents a separate set of issues with regard to security and information assurance. Data warehouses, “Big Data”, and business intelligence are an interrelated set of systems⁶⁹ that require evaluation within the CIA framework.

- Data warehousing

A data warehouse is a large data repository and may come in many forms. It will include data extracted from the source that is stored and summarized in a form that can be consumed by applications to support strategic analysis. The source data may include transactional records, business process results, and/or external data. The data warehouse has grown from a corporate advantage in the 90s into a requirement for most organizations. Regardless of the approach, the data warehouse serves as the corporate memory and archive. A well-designed data warehouse is a repository for data to be extracted, mined and analyzed to assess performance, mined for discovery, and searchable for key insights.

Security in a data warehouse is vital, especially for personally identifiable information that might be stored within it. In the event of a breach, the firm may have to pay damages and for credit monitoring for clients whose records were compromised. The impact of security on a data warehouse can also be based on the potential loss of this organizational memory. Additionally, ensuring that the stored data accurately reflects the events that are stored is vital to organizational memory. While not critical

for short-term recovery, the integral nature of the data warehouse in most organizations will require that it is available as the business recovers from disaster.

- **Big Data**

The current hot topic for many organizations is the concept of big data and the tools that support its collection and analysis. While many opportunities exist for businesses to successfully leverage the ever-increasing amount of structured and unstructured data in an organization, Big Data is still in its infancy and we need to consider biases in the tools as well as ethical issues around the use of Big Data.⁷⁰ These include protecting access beyond hiding solutions behind firewalls and synthesizing how all of these solutions will work together, but most importantly deciding who has access!

- **Business intelligence**

In the context of the organization, the business intelligence applications report out results based on data stored in the data warehouse or Big Data environments. The key for business intelligence tools is less about exposure of critical data, though competitors may find value in the strategic presentation of your data, but rather trustworthiness for decision making. As the tools become more embedded in decision making at all levels of your organization, the biggest threat faced could be the quality and integrity of the data on which decisions are being based. Like many concepts in this book, continual review and process sharing are necessary to address this threat. By providing transparency in the decision process, including the tools and data utilized, will make an organization stronger, but as with most issues in IT, this might be harder than it seems.

Data and how it is used is rapidly changing in every industry. The key is to consider the form and impact of the data as it moves from one environment to another. Is there information lost in the transition? What context is gained or lost? Understanding the life and flow of data in an organization will provide the integrity and assurance that decisions made using these technologies are based on a solid foundation.

Cloud security

Cloud services are a new buzzword, and many of us have one or more cloud storage services, music or movie streaming, and online email accounts. For a firm, cloud services expand to include cloud web hosting, all types of software as a service, hosted email, and many more. All of these deserve consideration in our overall security posture. For example, if we store files in an online service such as OneDrive or Dropbox, how secure is that data? Is it encrypted? How strong are the passwords used to access it?

For individuals and enterprise alike, the first step in cloud security is to make sure a good password is chosen, and two-factor authentication used if offered by the site. Then, the next question, especially for firms, is how data is encrypted. Is it encrypted before being sent to the cloud service? Is it encrypted in transport? What about in storage? Encryption in all of these cases would be the best scenario. The last aspect of encryption concerns who owns the encryption keys. For many online services, the keys are controlled by the service provider. This means that they can decrypt all information stored there, although their terms of service may say they will not. Further, it means that if their servers are compromised and keys seized by attackers, all files will be readable by the attackers. Last, it means that the firm will likely turn over decrypted data in the face of subpoenas or other legal demands. On the other hand, if the firm controls the encryption keys, only the firm can decrypt the information. In making the decision about whether to trust cloud providers, it is worth considering that almost every online service has been compromised over the years, and personal data exposed.

What is next?

Several trends in security are evident looking into the near future. Security touches all aspects of the organization from your organization's foundational architecture to your customers through mobile or other application services. Some key trends in each space to be considered include the following:

- **Physical platform.** The future will see more reliance on built-in security in hardware; e.g., within the switch or

router, next generation firewalls and intrusion detection, and iterative improvements in threat identification tools.

- **Mobile delivery.** Malware detection may or may not improve, but hopefully companies will expend similar efforts on mobile app security as they do on other software. In the face of so much malware, greater reliance will need to be placed on device hardening and enterprise policy application to mobile devices, whether company-owned or BYOD.
- **Applications.** Security analytics and identity verification software use will only increase. We will also see security-aware software that masks data in real-time, and security services that fit within the service-oriented architecture.
- **Data security** will be the likely focus of upcoming development. Given the extent of recent data breaches, the cost of fixing the issues will become less than the cost of cleaning up after. Data protection will include mobile data security, data eradication at system decommissioning, and hopefully at user request (which is already possible in the European Union), and greater use of encryption of Big Data and data warehouses as well as in the data center as a whole.
- **Cloud computing**, although a buzzword, will likely persist. The flexibility of virtual servers that can scale as new capacity is needed, and offloading some of the costs and expertise of IT to others will allow the firm to focus on core competencies. As cloud computing grows, the need for security in the cloud will also increase.
- **Privacy and policy management.** As more data is stored about each of us, the IT industry will hopefully evolve to allow personalization and control of the privacy of information, again, something already possible in the European Union. This will include privacy of data in embedded devices, such as smart appliances, light controllers, and other smart home devices. Hopefully too, organizations will be able to develop strategies for managing exposure of their information in social media and other technology venues.

The biggest change coming down the pipeline is not even necessarily on this list. Understanding the entire scope of the organization's IT use and exposure from the infrastructure to customer or business partner facing solutions is imperative. Individuals on the security team should be held accountable for monitoring trends in governing security technologies. Equally important is identifying the overlap between technologies to streamline the security planning and policy development process. In light of the rapid pace of change in IT and organizations, the best a security and information assurance team can hope for is a proactive approach to the security effort by all individuals in the organization.

Home and SOHO security

Home and SOHO (small office/home office) security really is not so different than enterprise security. The biggest changes lie in budget and scale, but the same assets need to be protected, in much the same way. Home users need a firewall, anti-malware software, and passwords. Security patches must be applied, software upgraded, and accounts and permissions managed. Personal data is just as valuable to you as customer data is to your firm.

How can you ensure your home computers are protected and your personal information safe? It takes time, but luckily, it's well within the skill level of the average office worker. We will start our tour of security at the point where your Internet connection comes into your home. Most home users have a cable, DSL, or fiber-optic connection to the Internet. Most also have a dynamic IP address assigned from their ISP, with NAT enabled to allow all machines in the home to share one IP address. These two facts represent the first line of defense for a personal connection. With dynamic IP addressing, it becomes much harder for an attacker to try to attack a specific person, as the address changes occasionally. More importantly, NAT provides an extremely efficient firewall. The NAT router, usually built into the cable or DSL modem, and/or into the wireless router, allows all outgoing connections and blocks all incoming connections by default. This behavior, which mimics an SPI firewall, means that no outsider can get into the network directly from the Internet. The best

feature of NAT on the modem or router is that it's automatically configured; the end-user does not need to do any configuration to turn it on. To test to be sure that the firewall is truly closed off, a very useful tool is ShieldsUP! from Gibson Research Corporation.⁷¹ This performs a legitimate port scan (a type of penetration test) against your network; for most home users, no ports should appear open. If they do, and you have not explicitly opened that port, further investigation is needed.

The next device that needs investigation is the wireless access point or router. In the past, this has represented a large security hole for the average user, as they shipped in open configurations, and the setup routines were fairly cryptic. Today, most routers have security turned on out of the box, but the default passphrase may not be. There are two passwords that need to be configured. The first controls login to the router's administrative web pages. This is typically accessed by connecting to the router, either wired or wirelessly, and entering the IP address 192.168.0.1 or 192.168.1.1 into the browser's address bar. The router then prompts for an administrative password. If the default has not been changed, the user is typically "admin" or "administrator," and the password is blank, "admin," "pass," or "password." The router's manual or manufacturer's website will contain information about the default.⁷² Change this password to something long and complex, and record it. The space provided in the front of the router's manual actually is a good spot for this; if someone has physical access to the manual, they can probably do more damage than simply stealing the password.

The second password that needs to be changed on a wireless router is the passphrase used to connect to WiFi. Often, the passwords used by default are printed on the bottom of the router, but are not very secure. Some are derived from the address of the router and others are simply too short. Once you have logged in to the wireless router, look for a tab across the top or a link on the left side of the page that says something like "wireless settings" or "basic wireless settings." Then, make sure WPA2 is chosen as the security method in the dropdown box, and set a long (at least 20 characters), complex passphrase. The passphrase can be up to 63 characters long. Various websites exist to generate long, random passwords of this length,⁷³ but honestly, typing the random characters

"\${XB'Be6MgW?%vLFb)Gk8&x%q*H[*Aa\$KO{7_(\$5B/+eVs[]493#O6B2ocvbVA-&}" into the on-screen keyboard of a phone is not much fun. Since length trumps complexity in passwords and passphrases, a better strategy is to stick with a plain-English phrase. For example, if you had owned a 1982 Chevy Chevelle at age 16, you might type "My first car was a lime-green 1982 Chevy Chevelle in 1993!!!" This gives 60 easy-to-type characters, and even if someone knew about that particularly sad chapter in your past, they would be very unlikely to guess that you had used that exact phrasing. Record this password in the user manual also, and then save your changes.

A last issue in WiFi security may be harder to address. The WPS protocol promised to allow one-button secure connections to wireless routers. Unfortunately, as previously discussed, the implementation was very flawed, and today a WPS-enabled router can be attacked within a matter of minutes. The successful attacker then retrieves the WPA2 passphrase, allowing them to connect to the router. WPS is usually turned off with a simple checkbox on the same page as other basic WiFi settings. Uncheck the box, and save the settings. Then, the hard part is determining whether the setting actually worked. Unfortunately, about half of the routers on the market by 2014 allowed turning WPS off, but the change had no effect. To check, open the Windows wireless network settings tool, and click on the name of your wireless network. If Windows still shows a box saying "You can also connect by pushing the button on the router," the change was not successful.⁷⁴ At this point, your choices are to:

1. Check for updated firmware from the router's vendor (this should be updated anyway, occasionally);
2. Buy a new router, after some research to determine whether the new model has the same issue; and
3. Live with the vulnerability, realizing that an attacker (your neighbor's kid!) could break in within hours, thanks to online help videos.

One last WiFi issue that represents a lesser security flaw is the network name. Many people set the network name, formally known as the ESSID (Extended Service Set ID), to be their name or some other personal information. This makes it easy for an attacker to determine which of the 20 or so networks visible in a neighborhood belongs to the victim. Again, this is not a huge risk, but a better course of action would be to choose a more generic name.

After properly securing the router, the remaining devices on most home networks are PCs, tablets, phones, game consoles, and printers, along with set-top boxes for satellite or streaming devices such as Apple TV. Several of these devices have very few security settings that can be configured by the end-user. Printers, satellite receivers, game consoles, and streaming devices typically only need updates applied for best security. Streaming devices, satellite boxes, and game consoles usually prompt the user to apply upgrades; this should be done as soon as possible. Printer software may require manual updates, but the driver software installed on PCs often prompts the user and handles the changes automatically. Beyond updates, however, little configuration is needed for security.

Tablets and phones do not have many security options that really affect the network's security, but the devices themselves are vulnerable to attacks. The first thing to do on a phone or tablet is to make sure a passcode is set. Even if it's only "semi-secret," in other words, other family members know it, it will protect the device if it's lost or stolen. Many of us have a great deal of personal information stored on mobile devices. A simple 4-digit PIN code allows for 10,000 combinations, and as long as 1234 or 0000, birthdays, or similar simple patterns are not used, it provides security against casual snooping. Remember that all current devices support more than 4-digit codes, for improved security, a longer password or even a longer PIN should be used.

Once a PIN has been set, anti-malware software should be installed, and location software configured. One recent study by IBM found that half of companies developing mobile apps devote no budget to security development, making anti-malware software vital for mobile platforms.⁷⁵

Both iOS and Android devices have ways to locate lost devices. For Apple, it's through iCloud; for Android, it's via the Google Play store or the Amazon Manage Your Device service. Many other applications are available to perform these tasks, including a notable open source app called Prey.⁷⁶ Take a few minutes to familiarize yourself with the process for locating a lost device; it can save valuable minutes if stolen, and make finding it when lost around the house a lot easier too.

Setting up full-device encryption on your phone or tablet is easy and stops most snooping, including, in some cases, law enforcement. This sounds technically the most challenging, but simply searching the Internet for “Your Phone Model’ encrypt device” will give a number of tutorials, the total time required is less than about 20 minutes.

PCs have the most settings for security. Most are software settings; PC hardware does not usually have security flaws. Further, most are application software settings; the operating system only has a few settings. The first thing to consider is whether the device is a desktop or laptop. Laptops, like tablets and phones, tend to go missing more often. For this reason, locator software should be installed on laptops, and could be installed on desktops also. Prey or commercial software such as LoJack for Laptops should be installed now in case the device goes missing. Prey is available for Windows, Mac, and Linux systems, as well as phones and tablets.

One major setting for all computers in a home or small office is to have individual user accounts with good passwords or passphrases. Even if a home user feels no need to hide anything from other family members, having separate accounts is good practice for organization of files and bookmarks. Even better, user accounts can be set up as limited accounts rather than administrator accounts, preventing unauthorized software installation. Setting up users is simple on either Mac or Windows, under System Preferences or Control Panel, respectively. Once a new user is set up, a good password should be set. Modern versions of both Windows and Mac OSX allow for passphrases with spaces, and long passphrases can be used. These are easier to remember, but a little more cumbersome to type. If shorter passwords are desired, there is a fairly easy way to generate a memorable password.

First, think of a sentence that is around 10 words long. “My dog gets fleas every year in June and scratches all summer” is 12 words. Then, take the first letter of each word, “mdgfevijisas.” Next, add mixed case, with the result of “mDgfeviJasaS.” If desired, a number, symbol, or both can be added, leaving us a 14-character password “mDgfeviJasaS3!” that is still fairly memorable, because it’s based on a phrase. This strategy should not be construed to mean that dictionary words can be modified the same way and be secure. The password “antidisestablishmentarianism61!” is still subject to many types of hybrid dictionary attacks.

Software on PCs, tablets, and phones needs to be kept updated. Luckily, most systems today provide automatic update capability. Unlike enterprise systems, where each update should be tested, it’s usually safe to enable automatic updates on personal devices. Even in the event something breaks due to an update, it’s an annoyance, but does not generally impact the bottom line. When prompted to install updates, this should be done as soon as practical. Operating systems are not the only software that needs updating, but again, many applications automatically check for updates on a weekly or monthly basis.

Now is a good time to make sure that the firewall on the PC or Mac is enabled. Having a host-based firewall on the target machine is simply another layer an attacker would need to breach. This principle, described earlier in the book as defense-in-depth, just means that as many obstacles as practical should be thrown up to prevent unauthorized access. The host-based firewall is configured on Windows by default, but must be enabled on Macs. As new applications want access to the network, the user is prompted to allow the application. Any unknown applications should be viewed with extreme suspicion, as they may represent malware “phoning home” with private user information.

Anti-malware software is a must in the home environment. If the subscription that came with purchase of a new PC has expired, it should be renewed. A number of free programs are available; some free only for non-commercial home use, others for home and small office. One recent study showed that most free antivirus software is equally effective as the paid offerings, although both detect only around 50% of viruses.⁷⁷ Microsoft, for example, provides Windows Security Essentials, and searching for “free

antivirus” will return a number of good free programs, although due diligence should be exercised to make sure the program selected has a good reputation. Under no circumstances should software be downloaded due to a popup ad. Such popups often say, “A virus has been detected. Click here to download our antivirus software.” The software downloaded from such a popup likely contains malware of its own, and may simply disable rival (real) antivirus software.

Filtering software may be installed to prevent home or small business users from accessing forbidden sites. This could be used to prevent children from accessing inappropriate material, or to prevent workers from wasting work hours checking sports scores. A full discussion of the options available is beyond the scope of this work, but solutions exist that are sized to meet the demands and budgets of homes and small offices.

One final area to address in security for home users is passwords. The password for the PC was previously discussed, but everyone has multiple website accounts. To protect them, the user must create good passwords. This is especially important for accounts where money is involved. Online banks should obviously have strong passwords, but do not forget e-commerce sites. Many shopping sites save your credit card information for easier checkout. Each site should have a unique, strong password. The same applies to email accounts, which can be used to reset passwords for other sites.

For most of us, the problem with passwords today is the sheer number. Trying to remember a password for the home PC and work PC is bad enough, let alone the various email accounts, social networking websites, and so on. So, what do most users do? They recycle the same password for multiple sites, or choose passwords based on names, dictionary words, or birthdays. All of these can be easily guessed. The best strategy for remembering so many passwords is simply not to do so. In other words, remember the few that must be used often, and the rest can be safely stored.

One way to (relatively) securely store passwords is low-tech. Passwords can be printed out on a slip of paper and put in a wallet. The logic behind this is that if a wallet is lost or stolen, you would cancel the credit cards it contained, and could change passwords at the same time. This

will lead to additional work in the event of loss or theft, and you should never keep the passwords for various credit card or bank accounts in the same location (wallet) as the cards themselves.

Another, perhaps better, low-tech strategy is to print the passwords and put them in a safe location, such as a home lockbox or safe or locked filing cabinet. A further copy could be stored in a safe-deposit box if desired. Much like the WiFi passwords recorded in the manual, if an outside adversary is close enough to get access to the password list, they can simply steal the systems. On the other hand, a password list in a small business may be an invitation for insider abuse. If printed, such lists should be kept in a safe location. Finally, the file that contained the passwords should be deleted; a copy on the computer where others could find it negates any value of keeping the paper copy safe!

A higher-tech and better way to store passwords is a password manager. Password managers are simply small databases that encrypt usernames and passwords for websites, credit card numbers, and short notes. The encrypted passwords are secured with a master password. A long, complex password can be chosen for this master, and the user must remember only that password. Password managers can also generate strong random passwords, so the user does not have to do it each time. All modern web browsers have password managers built in. Mobile apps are another popular way to store passwords, the advantage being that they can be carried anywhere. Some of these mobile apps have browser extensions or plug-ins that allow them to automatically fill in password fields. A final option is a web-based solution. This can be accessed from anywhere the user has an Internet connection. Password managers are really the only sane way to manage the huge volume of passwords each of us must remember.

Another way to secure some websites is to use two-factor authentication. Many banking websites have enabled two-factor authentication for added security. Two main methods are used for the consumer market after two-factor is enabled. First, the user may be sent a text message containing a one-time password or PIN. The other way uses a one-time password generator in the form of an app stored on a mobile phone or tablet. To log in, the usual username and password is entered, then the

user is prompted for the one-time password. This one-time password can only be obtained with something the user has, namely, the mobile device. Even social networking sites and email servers have enabled two-factor authentication. Several regularly updated lists of sites that support two-factor can be found online, and it should be enabled wherever possible, as it exponentially increases the degree of security.⁷⁸

Education is key to providing information assurance in small organizations. Teaching users, whether children, spouses, or employees, to use good passwords, avoid phishing emails, and not download software is vital. For children, this can be done as soon as they are allowed on the computer, with the depth of education matched to the child's comprehension level. Children, especially, should be taught to be suspicious of emails or social networking posts that seem too good to be true. Children, like employees, should also be taught never to accept invalid SSL certificates.

Backups

There are two kinds of people in the world, those who have never lost data, and those who make backups. Since almost all of us have lost data at one time or another, why have not you set up backups for your computer? Many of us would simply say that it is too difficult, and unfortunately, a good backup system takes some time. There are several good strategies, and some good software to use, taking much of the pain out of the process.

Backup strategies relate to what should be backed up, and how often. There are two main options for each. In terms of what should be backed up, the options are to do image-based backup, which copies the whole hard drive, or to back up only user-generated files. Image-based backup copies are large (the size of the full drive, at least initially), and slow to create. However, when restoration must happen, they are the fastest way to get back up and running. The other option, simply backing up the important files and settings, is much quicker and smaller, but takes longer to get back running at restore time. In this case, system files and directories are not backed up, the programs are simply re-installed, then the data restored from the backup. For home users, full-drive backups typically are too large to be manageable, and file-based backups make the most sense.

Both image-based and file-based backups have two further options, to backup each file every time or only files that have changed. The best strategy is probably a hybrid, to back up all files the first time, then only changes. Every week or month, the sum of the changes can be consolidated to a new full backup, and then changes kept for the next week or month. This should be done automatically by the software chosen; it's not something most would like to do manually.

The main strategy in how often backups occur is to decide how long to keep backups. How many generations of backups should be kept? For a home user, the answer is probably to keep the backups essentially indefinitely, especially of things like pictures and video clips. For a business, the answer is a legal matter. Financial records, client records, and so on all have different retention requirements. A firm should clarify these matters before developing a strategy. In terms of how often to back up, the answer really comes down to how much data you can afford to lose. Given the volume of data a home user generates, daily backups are probably fine. For a small business, hourly or every few hours might be a better choice.

The next choice in backups is what medium to use. For a home user, the choice mostly comes down to optical media or USB devices, with optical media decreasing in popularity as fewer machines come with a DVD or Blu-Ray writer. Even if optical drives are available, there are three issues with their use. First, in a time of multi-terabyte hard drives, the 4 to 8 gigabytes of a DVD or 25 gigabytes of a Blu-Ray recordable is fairly insignificant, meaning multiple disks will be needed for each backup. Second, it's difficult to automate the backup process with removable disks, especially if disks need to be swapped out partway through. Finally, although optical media was claimed to have a long life, on the order of 50+ years, studies and this author's experience have shown that even in near-ideal storage conditions, with no scratches on the disk, many disks have failed in around 5 to 10 years and are no longer readable. This leaves the primary option of a USB drive for most home users. Sizes and cost of hard drives, or even high-capacity thumb drives, are reasonable, and the process is as simple as plugging in the drive and setting up the backup software. The main decision left is simply choosing which files and folders to save copies of.

For a business, tape drives also deserve consideration. Tape drives are fairly slow and seem like old technology, but have two big advantages. First, the capacity of a tape can be fairly easily matched to the volume of data to back up. Second, tapes are removable storage and can be sent to offsite storage locations. As discussed next, this can greatly improve the reliability of backups.

Another great backup strategy is to do offsite backups. Offsite backups add reliability and safety to the process. If a computer is stolen, or a fire or flood occurs, there is another copy of valuable data kept somewhere else. For a home user, this can mean buying two drives, backing up to both, then give one to a family member who lives elsewhere, or stash a copy in your office desk. Then, every week, month, or some other period, swap that drive for the one at home that is getting backups daily. For a business, the same strategy of using two drives, or a tape drive, can be followed. The business would likely ship the drive or tape to a secure offsite facility, such as one of the numerous document archiving services. Such drives or tapes should be encrypted when they are created for extra security.

The other principal offsite backup method is cloud-based backup. Multiple companies provide online backup services; however, cloud-based sync, such as DropBox, cannot be considered backup because when a file is deleted in one location it is deleted in all. Instead, services such as iDrive, Carbonite, and Mozy, along with many others, provide services that copy files to their servers and store them for a monthly fee. If disaster strikes, the files can simply be downloaded to a new computer. For relatively small quantities of data, say, a few to 20 gigabytes, these services work well. If a user has hundreds of gigabytes of photos and files, though, the speed of their Internet connection becomes a factor in both backup and recovery. When choosing a cloud backup provider, be sure to check their reputation, the cost of their services, and whether they offer encryption for security of their files. If encryption is offered (and most do), who controls the keys? Like all cloud-based services, encryption keys can be managed by the backup provider or by the client. If the backup provider controls the keys, they probably would need to release

them for legal requests, such as subpoenas, but it also means that they can recover the keys in case they are lost. If the user creates and keeps the keys, they have full control over who can access the files, but if they lose the keys, there is no recourse; the backup is irretrievably lost. Both service-controlled and client-controlled keys are viable options, but make sure you understand your needs and the legal and security implications when choosing.

All of the above can seem a bit daunting. However, each step can be tackled separately, and no one step, except perhaps backups, should take even an hour. A checklist for home security can be found in the Appendix A.

Personal security

Beyond just those in this last section, many of the security guidelines found in this book can be applied to your personal life as well. Since we live in a digital age with our mobile phones or online purchases being logged, monitored, and marketed to, you should be aware of the implications to your personal as well as professional life.

The data that your organization collects is the same data that other organizations collect about you or individuals near to you. As you interact with online sites, mobile apps, or even make a purchase in the grocery store, consider how your data is being used. Similarly, how much of my communications am I exposing with my home WiFi setup or using my Bluetooth personal hotspot at a hockey game? What implications are there for working at home, in a hotel, or at the airport? Answering these questions are important and ultimately we do not expect you to change your habits but rather want you to explicitly consider how much of your data and information are you exposing.

In light of all this, what is your personal disaster recovery plan? In many ways, considering our personal enterprise in the same way we approach our organizational enterprise is sound practice. The challenge, much the same as in your organization, is getting the buy-in of key stakeholders (i.e., family members, roommates, etc.) to help minimize

exposure by evaluating and choosing to visit or conduct business with only trusted sites, turning off Bluetooth access, or changing passwords regularly, for example. By openly discussing security and planning for disaster, you may personally minimize “organizational” threats and losses.

Final thoughts

Moving forward you should be constantly aware of the ever-shifting security landscape. If you take only two things away from this book, we would challenge you to have:

1. Heightened awareness—Be aware of your organization’s security strengths and weaknesses in all aspects of the infrastructure, data, and application tiers.
2. Embrace change—Revisit your security and recovery plans incorporating new internal and external knowledge to address deficiencies.

By embedding these concepts within the organization and your approach to the entire IT enterprise, you will build a culture of secure IT professionals. Having a team that thinks “secure first” will enable your organization to navigate future threats and breaches. Does this make your organization invulnerable to a major breach or catastrophe? Certainly not, but it does enable to respond quickly to minimize the damage from the event. Furthermore, it might just make you a little less desirable target to external malicious attacks or even internal incidents.

The only promise that we can make is that security will continue to be at the forefront of IT challenges for many years if not decades to come. Additionally, the security solution that you put into place yesterday may not be applicable in a month, week, or even today! But a well thought out security and disaster recovery plan can carry your organization into the future.

Glossary

Algorithm: A recipe for doing something on a computer. Algorithms may be logical, or mathematical; for example, “If the temperature exceeds 73 degrees, turn on the air conditioner” or “Add two numbers together and print the results.” Algorithms power every decision a computer makes.

Auditing: Just like financial auditing, IT auditing looks for indications of misbehavior. IT security audits look specifically for evidence of break-ins by an attacker. Together with Authentication and Authorization, these make up the AAA of information assurance.

Authentication: An entity proving that they are who they claim to be. Together with Authorization and Auditing, these make up the AAA of information assurance.

Authorization: The permissions granted to a user once successfully logged in. Together with Authentication and Auditing, these make up the AAA of information assurance.

Availability: See CIA Triad.

B2B: Business-to-business communications. These are electronic transactions between businesses, which facilitate automatic ordering or financial exchanges. See also EDI and M2M.

Backdoor: A software program installed by an attacker which allows them to take control of a system, even if the original avenue of access is closed off by the defender. It is usually installed by a hacker after successful intrusion into a system.

Breach: Successful unauthorized access to a system. A breach, also called a security breach, means that an adversary was able to get past system defenses. It does not mean that anything was stolen, nor that the attacker left backdoors.

Certificate: See Security Certificate.

Certificate Authority: See Security Certificate.

Checksum: A checksum is a “sum” used to “check” data integrity. The sum is not truly just a sum; much more complicated mathematical formulas (algorithms) are used to calculate the value. The value is much shorter than the original data, typically on the order of 20 to 50 characters, and can be easily stored or transmitted along with the original data. To be useful, a checksum must have two properties. First, the value must be repeatable. A certain piece of data must always give the same number. In this way, if the value calculated before transmission or storage of data matches that calculated at a later time, it can be confidently said that the data have not been altered. The second property is uniqueness. The value calculated for a given file must be unique to that file. If applying the formula to any other file could produce the same checksum, one could not confidently say the original file was unaltered. Several checksum algorithms have failed this “collision” test in the last decade, specifically the MD5 and SHA-1 algorithms. These algorithms are deprecated and should no longer be used.

CIA Triad: Confidentiality, Integrity, and Availability. These three factors are the basis of information assurance. Confidentiality means that data is protected against unauthorized access. Integrity, also known as data integrity, means that data cannot be altered or at least that any potential alterations can be detected. Availability means that information is always available, at least to authorized users. A security countermeasure should fulfill at least one of the three elements of the CIA triad, and preferably two or more elements.

Cloud: The cloud has been defined in many contradictory ways. From a security standpoint, however, the cloud presents some new challenges. Chief among these are ensuring confidentiality and availability. Availability is a particular challenge, because Internet connections can be slow or unreliable, or the cloud provider could have reliability issues. Availability challenges can be counteracted by SLAs with cloud providers and ISPs. Confidentiality may be an issue. First, if the provider does not encrypt data, unauthorized access will result in disclosure. If encryption is used, a firm must decide who should have control of encryption keys. If the provider controls encryption keys, they will be able to read the data stored in their cloud service. The data may be protected during transfer to and from the cloud provider, and even while stored, but they will be able to read it and can turn over the keys to law enforcement. If the firm controls the encryption keys, the firm is the only party who can access the data; this is the only viable option in the authors' opinion.

Cold Site: In disaster recovery, a cold site is simply rented space to which company can move servers and client computers in case of disaster. The site has only electrical and HVAC provisions, no Internet connection or computers in place. Contrast with Hot Site.

Confidentiality: See CIA Triad

Convergence: Convergence, or converged networks, is a way of building a network such that voice, video, and data traffic all run over the same network connection. Convergence promises more efficient network management, but can also cause some security vulnerabilities. These vulnerabilities can be partially alleviated by moving certain types of traffic to separate physical or virtual networks, but this reduces the efficiency gains. See VLAN.

Cryptography: Cryptography is the science of scrambling data in such a way that it can be read only by authorized users. The process of scrambling, so-called "plaintext," human or machine-readable data into "ciphertext," the version that cannot be read is called encryption. Decryption is the reverse process, making data readable again. The algorithm used to perform the encryption is called a cipher; there are two main types of ciphers, symmetric key and public key or asymmetric ciphers. A key is much like a physical key, allowing the data to be locked and unlocked, or encrypted and decrypted. In cryptography, the key is simply a long (several hundred) set of numbers.

Symmetric ciphers are much like a common key and lock, in order to decrypt the data; a user must have a copy of the key. Symmetric ciphers are much cheaper

computationally to create, and so they are used whenever possible. They suffer from one major flaw, known as key distribution. Just like a house key, if you want someone to have a copy of the key, you must give it to him or her personally. This is challenging in a geographically dispersed environment like the Internet.

Asymmetric ciphers, on the other hand, create a key pair with a public and private portion. These are mathematically related in such a way that something encrypted with one key can only be decrypted with the corresponding half of the key pair. The user keeps the private key and can distribute the public key to anyone. This distribution takes place in digital certificates in SSL/TLS. When a customer encrypts their credit card information with the firm's public key as distributed to the customer's browser, the data cannot be stolen as it's transmitted across the Internet. When a user cryptographically signs a file, it can be positively said that this user was the source of the file, and further, that the file has not been altered in transit, or checking the signature would fail. Cryptography is the backbone of confidentiality in information assurance, and is also very useful in data integrity. See also Security Certificate.

Data Exfiltration: Sending proprietary data outside of permitted areas, typically outside the firm. This may be intentional or accidental; in either case the firm could face consequences. Intentional exfiltration usually involves someone stealing data for sale to another party. Accidental exfiltration could be as simple as someone mistakenly attaching the wrong file to an email, or quoting a patient's social security number in an email. Data exfiltration prevention solutions exist which monitor outgoing traffic on the firm's network looking for files that are marked electronically as "Do Not Distribute," or patterns such as credit card numbers, social security numbers, and so on. Exfiltration software or hardware may be required for compliance with various industry certifications.

DBMS: A database management system. This is the software, such as Oracle, MySQL, DB/2, or even Access that allows creation of the database and manages queries of the data. DBMS software also incorporates security features to allow users and passwords to be created. Based on the authorizations of the logged-in user, the DBMS may allow only certain operations, allowing read operations, but not updates, for example, or only allow the user to read certain fields in the database, such as a username, but not the password.

Decryption: See Cryptography

Disaster Recovery Plan: Disaster recovery plans are the means by which a firm deals with the unexpected. They may cover any type of disaster, including physical or electronic break-ins, fire, flood, or other man-made or natural disasters. IT and IT security are only a small part of the overall plan, and the events that trigger disaster recovery may not be IT in nature. IT will have a vital role to play in bringing the business back to productivity. Also known as business continuity plans.

EDI: Electronic data interchange allows firms to exchange data with each other without human intervention. EDI powers financial transactions, orders, and

even automated technical support. EDI also enables just-in-time ordering. See also B2B and M2M.

Encryption: See Cryptography.

Exploit: Exploit can be a verb or a noun. As a verb, an exploit is the act of taking advantage of vulnerabilities in a firm's defenses. As a noun, an exploit is the piece of software used to take advantage of the vulnerability. Local exploits are those in which the attacker must be physically in the same location as the machine to be attacked. Remote exploits allow the attack over a network connection. Remote exploits are often thought more dangerous, but local exploits are more common. Further, those who are local may be insiders who know the locations of valuable data and can access it quickly and quietly, if only the local exploit gives them privileges to do so.

External Threats: External threats are those outside the firm. Classical and modern hackers, competitors, and malware all form part of the external threat environment. See also Internal Threat.

Firewall: Firewalls block traffic into (ingress) or out of (egress) a network. Classical firewalls worked on the basis of TCP and UDP port numbers (see Port), while more modern firewalls look at many different layers of network traffic, including filtering based on text in a web page, content of emails, and so on. These modern firewalls, called UTM firewalls, can filter based on these multiple layers and also perform malware scanning. Data exfiltration prevention software, network filtering software, and IDPS systems can all be considered types of firewalls. See also Data Exfiltration, IDPS, SPI Firewall, and NAT.

Full-Disk Encryption: As the name implies, this protects all files on a hard drive. All modern operating systems, with the notable exception of home editions of Windows, including for phones and tablets, allow drives to be encrypted. This type of encryption unlocks all files on the drive once the user has entered the proper password, before the system even starts up. At that point, as long as the user is logged in, all the files remain accessible. Once the system is shut down, the data is again encrypted. The true security afforded by full-disk encryption hinges on the strength of the password chosen. Laptops and portable devices are especially good candidates for full-disk encryption, due to the likelihood of loss or theft.

GPO: Group Policy Object. Used in Windows domain networks (Active Directory) to control what an authorized user may do. GPOs exist to prevent almost any actions on a computer, and they can prevent such actions even when the user is disconnected from the corporate network.

Hacker: This term has been redefined over the past few decades. Originally, a hacker was someone to be respected, a programmer who could write software that could adapt to any situation, or was a hardware wizard. However, in common use today, the term is used to mean a cracker, someone who breaks into computer systems or networks. Two main categories of hackers exist, black hat

and white hat hackers. Black hat hackers are those who break into systems for fun and profit; white hat hackers are those who work as penetration testers to find vulnerabilities in defenses before they can be exploited. Black hat hackers used to ply their trade for reputation among peers; today, cybercriminals are more common, looking for money or trade secrets.

Hacking: Use of a computer system without authorization or in excess of authorization. Historically, hacking was defined as the acts of a hacker, as noted in the history of hackers above, and was a good thing!

Hardening: Protecting a host. Usually done as the first step once a new computer is purchased or installed. Some typical steps in hardening are turning off unneeded services, uninstalling unwanted software, deleting default accounts, setting strong passwords, patching installed software, and configuring backup on the system. See also Patching.

Hot Site: A facility where a firm can open up shop immediately as part of their disaster recovery. A hot site, unlike a cold site, contains computers and Internet connections all the time; all that is left for the firm is to copy data to the systems. Contrast with Cold Site.

Information Assurance: Commonly called computer security, but more of an overarching term that includes the concept that the information, not just the systems that contain it, is what needs protection.

Integrity: See CIA Triad.

Internal Threats: Internal threats originate inside the firm. As such, they are mostly composed of employees and their actions. Internal threats can be the most damaging, as the actors have the most intimate knowledge of the firm's assets. See also External Threat.

IDPS: Intrusion detection/prevention system. Both systems track suspicious network traffic, either based on known attack signatures or anomalies from normal. IDS systems detect the attack and notify a network administrator, by phone, email, or text. IPS systems stop the attack in progress, usually by rewriting firewall rules.

Journaling: Journaling is used for file systems and databases. In journaling, the system stores data about each write operation before committing the action. If the system crashes during the process of the write, the journal will be replayed, meaning the data can be restored.

Least Permissions: Also known as least access, or the principle of least permissions. Least permissions means giving employees the minimum access needed to do their job.

Linux: Linux is one of several open source operating systems that power servers and desktops. Linux is a Unix work-alike, originally designed by a Finnish computer science student named Linus Torvalds. Today, Linux is written and maintained by a community of volunteers and paid contributors. Linux technically refers to the "kernel," the core of the operating system, but in common use

more often refers to a Linux “distribution,” which consists of the kernel and all of the tools and user applications that run on a modern computer. Common distributions are Ubuntu, Debian, Red Hat, and Fedora, but literally hundreds of other distributions exist. See also Open Source Software.

M2M: Machine-to-machine communication. Unlike B2B or EDI, M2M usually implies automatic notifications from one machine to another. Examples include a server sending status reports to the vendor and various equipment tracking devices sending periodic updates to a central inventory database. These communications are often simply logged, and a human sees them only if an exceptional situation occurs. See also B2B and EDI.

Malware: Malware is any type of malicious software. There are many types of malware, such as viruses, worms, trojans, spyware, botnets, and ransomware. The methods by which they cause damage varies, as does the type of damage caused, but all are trying to damage a system or information. True damage to the system, as in damaging the hardware of the system, is exceptionally rare, Internet myths notwithstanding. System damage caused by malware typically involves deleting files so the system cannot boot. Anti-malware software detects malware either by looking at specific bits of the actual malicious file or by watching for anomalies—deviations from normal behavior by programs. However, even the best anti-malware software is only about 50% effective.

MITM: Man-in-the-middle attack. This occurs when an attacker can intercept the communications between a client and server. The attacker, being in the middle, can intercept all communications between the two parties. Common MITM attacks involve communication with secure sites. In most cases, the attacker is prompted to accept an improper security certificate, and does so. At that point, the data from the victim to the attacker is encrypted, transmitted to the attacker, and then decrypted by the attacker. The network traffic is read, then re-encrypted to be passed on to the site the victim originally intended to visit. The best defense against MITM is training users *never* to accept an invalid security certificate.

NAT: Network Address Translation. This technology is used by many organizations, but is especially prevalent in small office or home environments. NAT is implemented to allow multiple computers inside the organization to share one IP address. NAT has a side benefit from a security standpoint. NAT operates in much the same way as stateful packet inspection (SPI) by keeping track of each outgoing connection and allowing incoming traffic only if it’s related to an already-permitted outgoing connection. Thus, it acts as a very effective ingress firewall. NAT does pose a few problems for running a webserver on the inside of the organization, but by enabling port forwarding on the NAT device, incoming traffic will be allowed only on that port. Port forwarding should be used with caution, as it does open gateways to your network from the outside world, but works well when needed. See also SPI.

Open Source: Open source software is software that is licensed under an open source license. These licenses all have the commonality that the source code to the software is available to be changed. This freedom allows software developers to customize the program to their exact needs. Further, much open source software is free to download and use. Many organizations have made good use of open source software, because it is freely available and free to change. Some well-known open source software packages are the Linux and FreeBSD operating systems, Mozilla Firefox, the web browser, Apache, a web server powering many millions of websites, and OpenOffice/LibreOffice, office suites like Microsoft office. Communities of software developers who collaborate via the Internet write the software and give away their products under open source licenses. In the last decade, more and more firms have made financial and code contributions to many open source projects.

The exact terms of open source licenses vary; some allow reuse of the code in commercial products under almost any conditions, some require that the original developer be credited, while one of the most popular, the GPL, states that if any product is released that contains code licensed under the GPL, any changes made to the original code must also be released. This does not prevent a company from selling a product based on open source code, it merely means that their customizations and changes must be given back to the community. Often, companies who sell open source software make much of their money from service and support of the software. If a corporation wishes to incorporate open source code into their products, they must be certain that they comply with the license terms.

Operating System (OS): The software that allows a computer system go from a collection of electronic components to a functioning device; in other words, to boot up. The OS provides a way for the user to interact with the computer, such as a graphical user interface, touchscreen, mouse, and keyboard capabilities. The OS also provides a uniform way for programmers to do common system tasks, such as play music, print a document, or open a file. Operating systems have historically had many security vulnerabilities, but most operating system vendors have spent a great deal of money and man hours over the last decade to eradicate bugs. Today, applications typically have more vulnerabilities than the operating system.

Patching (Patch Tuesday): Patching is the act of applying a small change to software. The software's author develops the patch after a bug is discovered. The patch changes just those parts of the software that have security (or performance) issues, without having to reinstall everything. A patch should first be tested by a system administrator to ensure that it does not break anything else, and that it actually fixes the problem it was designed to. After running for a few days or weeks on a test system, the administrator installs the patch on production systems. Patches can be pushed out to client machines, automating the process. Patch Tuesday is the second Tuesday of each month in North America,

and is the unofficial name for the day that Microsoft regularly releases security patches. Sometimes Microsoft releases updates on other dates, usually for severe vulnerabilities. Other software vendors follow their own schedule, but because of Microsoft's dominance, the day has earned its own term. See also Zero Day.

Passphrase/Password/PIN: Passphrases, passwords, and PINs are related ways of authenticating a user to an IT system. Most systems have built-in password capabilities, including mobile phones, tablets, and computers, as well as web applications. The difference between the three lies mostly in the length and complexity. PINs are typically numeric and 4 to 10 digits long. Passwords may be much longer, anywhere between 5 and about 20 characters, but do not have spaces. Passphrases are phrases, including spaces and possibly punctuation, and may be 10 to 30 or more characters. With this length and complexity comes increased resistance to cracking, particularly brute-force attacks.

All three are subject to several kinds of attacks; the simplest is a dictionary attack. An attacker simply tries every word in the dictionary (which may have several million words, in many languages). These can be tried at rates of several hundred thousand a second. If the user has selected a dictionary word, even a long word, it will quickly be found. Hybrid dictionary attacks are effective against passwords based on words, but with the addition of a character or two; for example, "password1" or "65Mustang." These are only slightly slower than plain dictionary attacks. The last and slowest type of attack is brute force. In this case, the attacker simply tries every possible combination from "a" to "ZZZZZZZZ" or even "Zz99zZ99!" This runs much slower than dictionary attacks, and as of this writing, takes years or even centuries for a password longer than about 10 characters. A final type of attack is called rainbow tables, in which the hacker has pre-computed password hashes for all passwords that could be brute-forced. Since they are pre-computed, the attack runs very quickly, but comes at a great cost in disk space.

All of this combines to mean two things in terms of choosing passwords, length and complexity. Adding even one character to a password or passphrase increases the time required for a brute-force attack by a factor of about 80. Making the rules more complex, for example, requiring mIxED CaSe or punctuation characters, only increases the time by a factor of 2 to 4. Thus, length trumps complexity. This makes passphrases very strong, as long as a common phrase is not used. At the time of this writing, a non-dictionary password or passphrase of less than about 10 to 12 characters can be easily cracked with brute force or rainbow tables, if the attacker has physical access to the machine. Passwords or phrases longer than about 12 to 15 characters are still safe. The oft-repeated idea that 8-character passwords, with MixED CasE and special characters, are safe does not hold anymore. Corporate security policies should be updated to reflect this reality, or corporations should move to two-factor authentication. See also Two-Factor Authentication.

PCI DSS: Payment Card Industry Data Security Standard. This standard was developed by the payment card industry, specifically Visa, MasterCard, American Express, Discover, and JCB. It does not include private label cards. The standard has 12 major requirements for compliance that are designed to protect cardholder data.

1. Install and maintain a firewall configuration to protect cardholder data.
2. Do not use vendor-supplied defaults for system passwords and other security parameters.
3. Protect stored cardholder data.
4. Encrypt transmission of cardholder data across open, public networks.
5. Use and regularly update antivirus software on all systems commonly affected by malware.
6. Develop and maintain secure systems and applications.
7. Restrict access to cardholder data by business need-to-know.
8. Assign a unique ID to each person with computer access.
9. Restrict physical access to cardholder data.
10. Track and monitor all access to network resources and cardholder data.
11. Regularly test security systems and processes.
12. Maintain a policy that addresses information security.

The standards allow the credit card company to prevent a merchant from accepting credit cards from customers, but in practice, this is rarely done. More often, the credit card issuer levies a stiff fine against the merchant found not to be in compliance. Criticisms leveled against PCI DSS include the size of the fines, and the fact that despite only officially having 12 requirements, each is divided into dozens of sub-requirements. According to some industry experts, this means true compliance is essentially impossible. Finally, compliance with any industry standard does not really equate to true security.

Penetration Testing: Penetration testing, pentesting for short, is the act of trying to break into a system in order to protect it. A firm hires a pentester to test their systems for potential vulnerabilities in the same way an attacker would. Pentesters then work within the scope of the firm's requirements in terms of which systems to test and how deeply they should probe. The results of a penetration test inform the firm of vulnerabilities they should patch. The firm should patch all vulnerabilities found, but of course, budgetary constraints may prevent fixing all. In that case, the list should be prioritized, and the most serious vulnerabilities should be fixed.

Physical Security: Physical security of the firm's premises may or may not fall to the same individuals as those who maintain system and network security, but the two groups must coordinate efforts. Breaches in physical security can lead to

the theft or loss of information assets. Contrariwise, in today's world of electronic door locks, breaches in IT security can lead to easy physical break-ins. Physical security of the firm's electronic and information assets should be taken at least as seriously as technological measures like firewalls.

Policy: Policy drives the implementation of technological security solutions in a firm, and provides oversight to the process. Security policies should be written clearly to emphasize the benefits to the firm if policy is followed, but must have penalties built in for noncompliance. Ideally, security policies consist of a set of documents that have a top-level document stating the firm's aims in information security and other documents that spell out specifics. One such sub-policy document would be an Acceptable Use Policy (AUP), which governs how an employee uses the electronic resources, including the Internet. This document is signed by new hires upon joining the firm. A strong security policy, accepted by all from top management down, gives vision to the firm.

Port: In networking and security, a port is a "door" on a computer system that is opened by a program to allow other machines to communicate with it. A server program opens a TCP (Transport Control Protocol) or UDP (User Datagram Protocol) port when it runs. For example, a web server running on a computer opens port 80, and web browsers are programmed to automatically talk to port 80 on that computer. The port allows communication between the two machines. Blocking a port with a firewall prevents that communication from occurring, which of course is not very useful if the machine is supposed to be a web server! If, on the other hand, no one should be communicating with that machine, firewalling the port is the right action. Finally, blocking a port at the firewall level has very little effect if a port is not opened on a machine. In other words, if that particular machine is not running a web server, port 80 will not be open, and blocking port 80 on a firewall has little practical effect. *But:* Firewall ports should be closed unless explicitly needed. Why? In case there is a vulnerability in a piece of software that allows a port to be opened in an unauthorized fashion. This has happened in the past, and will happen again. So, the best policy is to close all ports by default, and only open ports on an as-needed basis. See also Port Scanning.

Port Scanning: Also simply known as scanning, this consists of using a piece of software to find open ports on machines connected to the network. This is the fourth step in an attack or penetration test. Both attackers and penetration testers carry out port scans. The results of a port scan on a single machine show what ports are open on a remote machine and therefore what applications are running and ready to talk to authorized or unauthorized users. An attacker would use this information to develop a plan of attack on that machine. A penetration tester, on the other hand, would use the same information to develop the opposite, a plan to secure the services offered by that machine. See also Port.

Privilege Escalation: Getting the permissions of an administrator, who has unfettered access to the system. Privilege escalation attacks are a very valuable tool for an attacker, and are commonly the first thing sought once access to a system has been obtained. Given that insiders have regular user-level access already, privilege escalation can be especially devastating when an insider goes rogue.

Remote Access: Remote access may be intentional or unintentional. Intentional remote access is often used to allow workers to telecommute. In this context, remote access usually means giving the user a remote desktop. The user simply logs in to the remote machine and can work as if on corporate premises, with the same access privileges. This presents some security risks, namely, if the user's credentials are compromised, the attacker gains all those privileges. Other types of remote access are also possible; for example, a corporate Intranet or various databases may be configured to allow remote workers to connect to them. The primary way to protect such assets is to allow access only via either a VPN connection or over a web connection secured by SSL/TLS.

Unintentional remote access, on the other hand, is nothing but a security risk. This occurs when a server is misconfigured, or a firewall allows access to machines that should not be exposed to the outside world. Preventing unintentional remote access is best achieved by regular penetration testing.

Security Certificate: Also called a “server certificate,” “certificate,” or “cert.” A security certificate is used in SSL/TLS for authentication and encryption. The certificate is derived from the public part of a public/private key pair. Typically, the server administrator generates the key pair, entering some basic information such as the name and contact information of the firm. The administrator then sends the public half of the pair to a certificate authority (CA) to be signed. The certificate authority plays the role of a notary public. After checking the credentials of the firm asking for a signature, the CA electronically signs the public key; at this point, the signed key is called a certificate. The certificate is sent back to the firm, and the administrator installs it onto a specific server. The private half of the key is also copied to that server, in a different location where only the server software has access to it.

When a user wants to browse a secure web page, they enter the `https://` URL into their browser's address bar. The browser then connects to the desired server, and requests the server's certificate. The user's browser then checks three key parts of the certificate. First, did a recognized CA sign the certificate? Second, does the name on the certificate (as entered by the administrator) match the domain name of the website the browser connected to? Third, is today's date within the valid range of the security certificate (typically 1–3 years)? If the certificate passes all these tests, the browser considers the site authenticated, allows the connection, and loads the website. If any test fails, the browser presents an error message telling the user that the connection is untrusted. The user can, but should not, continue to this unauthenticated website. The main reason a user should not continue to the untrusted website is that the website may be compromised, or another

website may be impersonating the desired website. Impersonation can be as simple as a user wanting to visit <https://www.mygreatbank.com>, but instead typing <https://www.mygreatbanks.com>, which is run by a rogue operator capitalizing on typographical errors. The attacker designs her website to look like the real “My Great Bank” website, including a place to log in. When the user logs in, that username and password are then captured by the attacker, and then used to access the user’s account on the real bank’s site. This attack has happened with a number of websites, leading companies to buy misspellings of their own domains, such as pespi.com, or paypal.com. After the domain is purchased, it’s simply redirected to point to the company’s website.

Finally, the public key contained in the key pair can be used to encrypt data to and from the server. To do this, a symmetric encryption key is generated, which is then exchanged between the user’s browser and the server. One way to protect this symmetric key as it’s transmitted to the server is to encrypt it with the server’s public key. Once the symmetric key is securely exchanged, it’s used to encrypt the user’s account data. See also Cryptography.

Sniffing: Also known as traffic sniffing or network sniffing. System administrators legitimately use a sniffer to look at communications between machines. An administrator might investigate traffic between machines to debug a connection that is not working, for example. Attackers also look at conversations between client and server, but with the goal of discovering information such as passwords. If a server’s communications are encrypted, the attacker will not be able to read the information.

Social Engineering: Tricking a user into doing something against their own best interests. For example, an attacker might call the administrative assistant of a manager and impersonate a member of the IT staff. The attacker could claim that there was a problem with the manager’s email account, and ask the assistant for the manager’s password, so that it can be “reset.”

Social engineering is also used in “phishing” emails. In these emails, the attacker impersonates someone else. For example, a recent phishing email received by one author claimed to be from eBay, claiming that there was a problem with the author’s account. The author was asked to log in to correct the problem. The email looked mostly legitimate, even including graphics from eBay. However, the link that the victim should click on did not lead to eBay, but rather to a site called “<http://www.ebay-myfunnydomain.com>.” This was a dead giveaway.

The only real way to decrease the threat from social engineering is user education. Such education is increasingly important. As the number of exploitable vulnerabilities in operating systems and other software decreases, the number of attacks using social engineering appears to be on the upswing.

SPI Firewall: Stateful packet inspection firewalls work by analyzing the state of a packet. The rules used are outlined below and are very simple, but SPI firewalls are extremely effective. SPI filtering is the most used mechanism today.

The rules used by SPI are:

1. Is the packet trying to open a connection? (Only about 1% of all packets):
 - a. If the packet is coming from inside the organization, and outbound, typically the packet is allowed. Exceptions to deny certain types of traffic can be based on specific rules.
 - b. If the packet is coming from outside the organization, and inbound, it's typically denied. Again, exceptions can be made, for example, to allow access to an internal webserver.
2. If the packet is not trying to open a connection (99% of all packets):
 - a. If the packet is part of a previously allowed connection, allow the packet. This allows a user to browse to an outside website, and the reply to be routed to the user.
 - b. If the packet is *not* part of a previously allowed connection, deny the packet.

The simplicity of SPI rules allows packets to be processed very quickly, but still gives great effectiveness, as all packets either do or do not try to open a connection, and those that do not simply need to be looked up in a small table called the tracking table. See also NAT.

SQL: Also known as “Sequel.” Technically, the pronunciation S-Q-L is probably more correct, as that is how the official ISO documents spell it, but both pronunciations are in common usage. SQL stands for Structured Query Language and is the default language used to query databases. The security implications of SQL are centered around an attack called “SQL injection,” in which the attacker makes a statement that is always logically true, and then adds another query to the end of the statement. An excellent example is shown in a cartoon at <http://xkcd.com/327>. SQL injection is a common attack on websites and can be prevented by sanitizing database inputs, such that things like special characters are not directly passed on to the database software. Those inside the organization can also perform similar attacks; prevention consists of properly managing database permissions and using stored procedures to allow users to access only specific parts of the data and database.

Tailgating: Allowing more than one individual at a time through a physical access control point. For example, if two employees are entering the server room or data center, each should scan their own access badge and enter their own PIN. If one opens the door and holds it for the other, this is tailgating. The problem with tailgating is that there are visitors for whom there is no record. These individuals could cause damage to the system or steal equipment or information, without evidence they were ever in the room. Educating users and punishing violators is the best solution to tailgating. Security cameras positioned to watch entrances are also an effective deterrent.

Threat Landscape: The threat landscape is the sum of all threats to an organization's IT resources. Common elements in the threat environment are hackers, internal threats, malware, theft, and the risk of natural disasters.

Two-Factor Authentication: Two-factor authentication means using two of the three classes of ways that an individual can prove their identity. The three classes are as follows:

1. Something the user *knows*, like a password, passphrase, or PIN
2. Something the user *has*, like an ID badge or access card
3. Something the user *is*, meaning biometric methods like fingerprint readers

When two factors are used, the security of the authentication increases exponentially. The most common two-factor authentication schemes involve swiping a card, then entering a PIN, like in an ATM or debit card transaction. Other common methods include a fingerprint reader coupled with a password, or a password coupled with a physical token that shows a one-time password on a small screen.

VLAN: Virtual local area network. VLANs logically segregate a single physical network into several networks. Setting up VLANs can reduce some security concerns. For example, if traffic from accounting and engineering run on separate VLANs, an engineer could not snoop on the accounting transactions. If voice traffic runs on a separate VLAN, not only are many eavesdroppers thwarted, but also call quality can be improved. VLANs have had some security issues in the distant past, where traffic from one VLAN could be forced to jump to another, but these vulnerabilities are unlikely to be encountered in today's production networks.

VPN: A Virtual private network allows traffic to be protected inside a "pipe" or "tunnel" as it passes across the Internet. In the past, organizations leased telephone (or data) lines from the phone company. These lines were mostly protected from eavesdropping, because no one else had physical access to them. A VPN's name comes from the fact that it gave the same level of privacy without having a physical leased line. However, cryptographic VPNs, the most common type today, give even better privacy than the old leased lines. A VPN should be the first line of defense against unauthorized users gaining access to company resources. Best practice dictates that remote offices and telecommuting employees make a VPN connection before sending data to headquarters.

Vulnerability: A software flaw that allows unauthorized actions. Some software bugs simply cause the application to misbehave or even crash. Unless the crash results in a way for an attacker to gain access to the system, they do not represent vulnerabilities. When vulnerabilities are found, the developer of the software writes a patch to fix the problem. See also Zero Day and Patching.

WiFi Security: WiFi (802.11) security is a complex topic. The basics are that there is only one secure protocol for WiFi today, WPA2. WEP, the original security protocol, can be cracked within a few seconds to minutes, allowing an attacker

to recover the encryption key. At that point, all other data encrypted by the WEP key can be read by the attacker. WPA, an interim protocol, is partially broken. WPA2 (and WPA) run in two modes, pre-shared key (PSK) and enterprise mode. PSK mode is the type used in home and small office environments, where all users authenticate to the wireless router with one passphrase. Each is then given a unique key that encrypts all transmissions to that router. In enterprise mode, each user authenticates using their organizational credentials. In either mode, the security of WPA2 hinges on the strength of the password or passphrase. In PSK mode, up to 63 characters are allowed; at least 20 are required to consider a passphrase secure. For enterprise mode, corporate password policy should be followed.

UNIX: Also less properly spelled “Unix.” UNIX is a family of operating systems that is mostly used for servers. Mac OS is based on a version of UNIX. UNIX has a reputation of stability and security, but objectively, it may not be any more or less secure than Windows.

XSS: Cross-site scripting. XSS vulnerabilities affect many websites and allow attackers to steal private data from visitors. There are several types of XSS vulnerabilities, but all depend on the idea that everything coming from one website has the same level of trust and permissions, but something coming from another website must have separate permissions granted. In cross-site scripting, attackers place malicious information into the content being delivered from a website they have compromised. Once the content arrives, the browser assumes it’s coming from the proper website and interprets the code based on the permissions granted to that site, meaning the malicious code also runs.

Zero Day: Can be applied to zero day vulnerabilities or exploits. Zero day simply means that a vulnerability exists which has not yet been patched by the vendor. In the case of an exploit, it means the attack program takes advantage of such vulnerability. Zero day has nothing to do with how long the vulnerability has been known; some zero day vulnerabilities are years old and will never be patched, because new versions of the software have been released. See also Patching.

APPENDIX A

Checklists

These checklists are designed to help make certain that all basic steps have been taken for each security measure. They do not give all details about how to take certain steps, nor should a checklist be considered exhaustive; many more steps can be taken after completing the checklist. In fact, these checklists, and similar checklists from various standards bodies and industry advisory groups, can be used to develop your own checklists for your particular organization. Customized checklists ensure that you implement best practices in a way that fits your organizational security goals while utilizing the protection technologies your firm has invested in. A set of checklists needs to be an integral part of your IT policy, procedures, and practices.

Chapter 1: Security and Information Assurance

	Yes	No	N/A	Remarks
The CIA (Confidentiality, Integrity, Availability) model stands as the base of security. Any security measure implemented needs to address at least one of these areas.				
Our perception of risk (or security) may not match mathematical realities, leading to over- or underspending on countermeasures.				
Goal of information assurance is not perfect security, but adequate security for the sensitivity of the asset, for a reasonable price.				

	Yes	No	N/A	Remarks
Since hard numbers are hard to pin down in IT security, firms must prioritize their expenditures based on criticality and sensitivity of assets. End-of-life software is one example of a critical asset to replace as soon as budget allows.				
Be sure to address insider threat in your organization by monitoring critical assets' logs, rotating job roles, and using separation of duties.				
Enforce the principle of least access for all information.				
When deciding whether to enter into B2B, EDI, or M2M relationships with other firms, audit their IT security, and encrypt data transferring between the entities.				
Physical asset protection is a first line of defense against both internal and external attacks.				
Lock the server room door.				
Use surveillance cameras in sensitive areas as a deterrent.				
Do not allow tailgating.				
Do not allow visitors in sensitive areas.				
Lock office doors automatically when unoccupied.				
Secure areas should not allow removable media or recording devices.				
Locate data centers according to industry best practices, away from water lines, with proper fire suppression.				
Use alarm systems, and if appropriate, guards for secure areas.				
Erect fences and physical barriers around sensitive premises.				
Install recovery and remote wipe software on mobile devices.				

Chapter 2: OS and Application Security & Hardening

	Yes	No	N/A	Remarks
Install only required software on each device. Servers typically do not need productivity software installed, for instance, and software not installed cannot be exploited.				
Remove unneeded bloatware that came pre-installed with the computer.				
Disable unneeded services (and uninstall associated unneeded software) on each device. A database server likely does not need a file-sharing service.				
Check and set default configuration options. For example, a web server might allow access to certain files that would represent security vulnerabilities; changing a configuration option could block this.				
Use GPOs in Windows environments. These can control almost anything.				
Configure backup software. See the checklists for Chapters 8 and 10 for backup guidelines.				
Patch operating system vulnerabilities and application vulnerabilities.				
Vet all patches before installation in an enterprise setting.				
Use automatic installation services to install vetted patches.				
In SOHO environment, use automatic updates to ensure software stays up-to-date.				
Manage users securely.				
Get rid of unneeded default users.				
Change passwords for default users that remain.				

	Yes	No	N/A	Remarks
Set password policy enforcement on domain controllers, matching organizational password policies.				
Set up individual accounts for all users of the computer. Do not use group accounts for any purpose; instead, assign roles to those needing common access. Group accounts do not allow proper auditing.				
Do not use administrator accounts for everyday use. Normal users should not have administrator privileges, and administrators should only run programs as administrator when required.				
Malware				
Install anti-malware software on all devices, including mobile devices. Even though only approximately 50% effective against new threats, it's the first line of defense. The idea that certain operating systems do not need anti-malware protection is simply false.				
Install file alteration monitoring software on servers.				
Scan email for malware in at least two locations, at the mail server and at the client.				
A number of good free anti-malware options exist for home and SOHO use.				
Use special-purpose anti-malware that runs on demand to supplement the basic malware program or to clean infections.				

Chapter 3: Data Security

	Yes	No	N/A	Remarks
Internal threats apply to both structured and unstructured data. Prevent unstructured data exfiltration (such as business plans and technical documents) with data exfiltration prevention software at network border.				
Limit physical access to hardware in data centers.				
Use DBMS security features.				
Encrypt critical data fields such as credit card numbers.				
Implement passwords to control access to databases, whether central corporate passwords or separate for specific databases.				
Employ stored procedures so that users need not be granted general access to the database to get report data.				
Implement a separate SQL code review process to guard against unintentional data loss.				
Protect against SQL injection.				
Use stored procedures for both local and web-based queries.				
Sanitize database inputs by filtering database control characters from queries.				
Review data periodically to ensure data quality.				
Manage master data to provide record consistency.				
Include data backup and recovery as part of your overall disaster recovery plan. See Chapters 8 and 10 checklists.				

Chapter 4: Network Security

	Yes	No	N/A	Remarks
Use VLANs or prioritization to ensure timely delivery of VoIP packets.				
Consider encryption of VoIP within the organization using Secure SIP or similar protocols. Do not trust protocols such as Skype where the firm does not control encryption keys.				
Implement VPNs, either by using SSL/TLS to protect web traffic to certain servers, or by using IPSec or SSL-based VPNs.				
Protect remote workers with host-to-site (remote access) VPNs.				
Protect remote offices with site-to-site VPNs.				
Protect critical individual servers (for example, database servers that replicate to each other) with host-to-host VPNs to encrypt their vulnerable data transfers.				
Train users not to accept invalid security certificates. Do not use them, even for internal-only sites, as it undermines training.				
Implement firewalls at network borders and on individual machines.				
Test your firewall configuration with penetration testing.				
Use specific types of firewalls for specific purposes; e.g., an application-layer firewall to filter SQL injection traffic.				
Install and configure IDPS systems, and train them to reduce false positives.				
For SOHO users, a NAT firewall may be all that is needed.				

	Yes	No	N/A	Remarks
WiFi				
Only allow WPA2 encryption. For enterprise users, ensure passwords comply with corporate policy and limit login attempts to avoid brute force attacks. For SOHO preshared key mode, use passphrases of at least 20 characters.				
Change the default login password for WiFi access points or routers.				
For individuals, set a good generic network name (ESSID/SSID).				
Turn off WPS on SOHO routers. Then, check to be sure it's really off.				
Do not bother with useless WiFi security measures like MAC address filtering or hiding the SSID.				

Chapter 5: Secure Application Development

	Yes	No	N/A	Remarks
Consider security at each stage of the development process.				
Use existing frameworks for secure application and system development.				
Ensure that data in test and development environments is stored safely. Data used in testing should be cleaned of sensitive information before use in tests.				
Have developers trained in secure coding processes. Very few colleges and universities teach secure coding, meaning developers will probably need additional training.				
Monitor mobile app development to guard against data leakage.				
Thoroughly test in-house applications at each stage of development for security, including having professional penetration tests against them.				

Chapter 6: Cryptography

	Yes	No	N/A	Remarks
Realize that encryption cannot solve every security problem. Further, even the best encryption relies on password or passphrase strength.				
Take advantage of both the encryption and authentication features of asymmetric encryption to validate message origins. This applies to users who can be issued individual keys for authentication via smart cards or email, and to machine-to-machine traffic.				
Use long keys (2048 bit for public keys and 256 bit for symmetric keys, at the time of this writing) to ensure data cannot be decrypted.				
Use full-disk encryption, especially for mobile devices and critical servers.				
Encrypt USB drives wherever possible.				

Chapter 7: Penetration Testing

	Yes	No	N/A	Remarks
Avoid information leakage (about IT systems) in public sources, such as job postings, RFP and RFQ documents, and Internet registrars.				
Engage in-house or outside professional penetration testers to probe for network and application vulnerabilities.				
Internal testers may be more familiar with systems and know where to probe.				
External testers may see things that internal testers miss because they are used to systems.				
Carefully enumerate systems to be tested in order to give pentesters a proper scope.				
Use test results to:				
Demonstrate compliance with regulations such as HIPAA or PCI DSS.				
Fix issues found, within budget.				
Sanitize test reports before distribution, as they may contain details that would help a would-be hacker.				
Create corporate policy to determine how to handle vulnerability reports from outside “volunteer penetration testers.”				

Chapter 8: Disaster Recovery

	Yes	No	N/A	Remarks
Establish a core disaster recovery team, with representation from management, IT, legal, HR, and operations divisions.				
Develop a charter that includes plan goals and definitions.				
Keep change logs or revisions of plans.				
Evaluate and practice plan through tabletop exercises as well as dry-run or even live tests of the plan.				
Develop a team tasked with classifying risks and analyzing threats to IT systems.				
Develop strategies for many potential disasters, including scenarios of effects on the organization.				
Define disaster recovery phases and procedures.				
Consider hot vs. cold sites. Where is data stored now? How quickly can it be moved to new site?				
Are backups stored in an accessible site if a large natural disaster occurs?				
Are software media and keys stored in accessible location?				
Document the plan, disseminate it to all stakeholders, and maintain the disaster recovery plan as conditions change inside or outside the firm.				

Chapter 9: IT Security Policy

	Yes	No	N/A	Remarks
Every firm has policies that govern actions of employees. IT policy fits within overall corporate policy, and IT security policy is a further subset of all IT policy.				
IT security policies must be written. A written policy enables enforcement if needed. The policy also needs to apply to everyone, from the CEO down.				
An effective security policy must contain penalties for noncompliance. If the policy has no teeth, it will not be followed.				
Policy should be developed by a multidepartment team. If any one department develops policy in a vacuum, other viewpoints will be overlooked.				
Policies should be developed with reference to standards documents for the organization's particular industry.				
Tone of the document should be positive, focusing on what the policy protects and what it enables.				
Security policies should be concise, a few sentences to a few pages. Standards and procedures documents flow from the policy document, and contain more detail.				
Visible upper management support for security policies will set a tone of security within the organization.				
Policy helps develop implementation guidance, processes, and procedures and provides oversight to implementation.				

Chapter 10: Conclusion

	Yes	No	N/A	Remarks
SCADA systems will likely become a bigger security issue in the near future, as more systems are connected to the Internet.				
Patch SCADA systems as soon as patches are available.				
Evaluate vulnerabilities of SCADA systems already in place.				
NIST publication 800-82 has very good guidelines for SCADA protection.				
“Big data” security challenges include physical security of the data warehouse and security of the data.				
As records are stored longer, the chance of previous customers being affected grows. Can we delete old records?				
Scrub PII (personally identifiable information) from records that are used for data mining purposes.				
Encrypt PII if it must be retained.				
Consider ethical implications of data retention and mining.				
Business intelligence tools are only as good as the data on which they operate. Transparency about data quality and procedures and tools will enable better decisions.				
	Yes	No	N/A	Remarks
Cloud security revolves around three factors.				
Access controls—Is access to the data protected by passwords? If so, security is only as good as the passwords, so strong passwords are a must.				
Is the data encrypted during storage?				
Is the data encrypted during transmission?				

	Yes	No	N/A	Remarks
Does the firm control the encryption keys?				
Does the cloud provider control the encryption keys?				
Is the uptime guarantee sufficient for firm needs?				
SOHO Security				
Personal information				
Shred sensitive documents before recycling or disposal.				
Destroy storage devices or wipe them securely before recycling, and wipe securely before selling.				
Do not post personal details on social networking sites, and do not post detailed travel plans with dates and locations.				
Be selective about what personal information you give out for loyalty or rewards programs. Use disposable emails and/or phone numbers, and never give more than absolutely needed.				
Network protections (see also Chapter 4 checklist)				
Use ShieldsUP! or another online port scanning tool to check your firewall from the outside.				
If ports are open, determine why. Check the port forwarding section of your modem or router setup page.				
Use a VPN when accessing secure websites from public locations such as coffee shops or hotels. Inexpensive options abound, but check reviews for the service.				
WiFi protections				
Set the device password, and record it in the user manual, then file the user manual away safely.				

	Yes	No	N/A	Remarks
Set the device to use WPA2 as the WiFi encryption method, and set a long (more than 20 character) passphrase. Choose a phrase that is easy to remember and easy to type, but not something like a famous quote or movie tagline. Record this in the manual also.				
Turn off WPS on devices where it's present. Then, ensure that it was actually switched off; if the device does not disable it properly, decide whether it's a risk worth living with.				
Change the network name to something nondescript, not a family name or address.				
Do not bother with MAC address filtering or SSID hiding, which are ineffectual.				
Portable devices				
Use a good PIN or password, even if it's known to other family members; it protects against access if lost or stolen.				
Use more than 4-digit codes where practical.				
Use anti-malware software on tablets.				
Learn how to use the "lost device" feature for your tablets and phones. It can speed recovery, and is great for a device misplaced in the house, too. For laptops, consider locating software such as Prey.				
Set up full-device encryption for phones, tablets, and notebook computers.				
Passwords				
Do not use dictionary words, common phrases, or personal information like birthdates in passwords.				

	Yes	No	N/A	Remarks
Choose a long (10~12 words) phrase, then take the first character of each word. ChaNGe CaSE of letters, and add or substitute numbers or symbols. This gives a relatively memorable, strong password.				
Use a unique password for each website or account.				
Use two-factor authentication (such as a one-time code sent via text message) on sites that support it.				
Use a password manager to deal with the many passwords that each of us must manage.				
Backups				
Make backups automatic. If backups are manual, they will not happen. Either pick a program that has built-in scheduling, or use a scheduler such as Windows Task Scheduler to run the program periodically.				
Decide what to back up. For most home users, this will be only their own files and folders, not programs.				
Decide how often to back up. Backups should be conducted often enough to preserve (almost) all changes, without being done so often as to be burdensome.				
Decide how long to keep backups. Most individuals will probably keep files indefinitely.				
Decide what medium to use. USB drives are usually the best choice, or cloud-based backup.				
If physical media such as a hard drive is used, make two copies and store one offsite, with a family member, in a safe deposit box, or at the office.				
Encrypt backups for added security.				

	Yes	No	N/A	Remarks
Other home protections				
Set up individual accounts for each user.				
Use a host-based firewall on all machines.				
Install anti-malware software on each device. Several good free options exist for home and some small business use.				
Use filtering software to prevent users from accessing specific sites or types of sites.				
Educate users to never accept invalid security certificates on websites.				
Train users never to respond to “too good to be true” offers via email, popup ads, or instant messages. They are probably phishing attempts designed to steal passwords or other personal information, or install malware.				

Endnotes

Chapter 1

1. <http://www2.trustwave.com/rs/trustwave/images/2014TrustwaveSecurityPressuresReport.pdf>
2. The origins of the CIA triad are lost to history. Some of the ideas behind the triad were certainly known in ancient times, military leaders of antiquity ensured confidentiality of messages, and verification of who sent the message was equally important. Being able to create authentic-sounding, conflicting orders that could be delivered to one's enemy would have been a great coup. In more modern usage, the CIA triad forms the basis of many information assurance standards, including those from the ISO. <http://www.techrepublic.com/blog/it-security/the-cia-triad/>
3. http://www.schneier.com/essays/archives/2008/01/the_psychology_of_sc.html
4. <http://www.microsoft.com/en-us/windows/enterprise/end-of-support.aspx>
5. <http://blogs.microsoft.com/cybertrust/2013/08/15/the-risk-of-running-windows-xp-after-support-ends-april-2014/>
6. <http://www.nbcnews.com/feature/edward-snowden-interview/edward-snowden-timeline-n114871>
7. <http://www.networkworld.com/article/2280365/lan-wan/13-best-practices-for-preventing-and-detecting-insider-threats.html> and [https://www.us-cert.gov/sites/default/files/publications/Combating the Insider Threat_0.pdf](https://www.us-cert.gov/sites/default/files/publications/Combating%20the%20Insider%20Threat_0.pdf)
8. <http://www.verizonenterprise.com/DBIR/2014/>
9. <http://www.ama-assn.org/ama/pub/physician-resources/solutions-managing-your-practice/coding-billing-insurance/hipaahealth-insurance-portability-accountability-act/hipaa-violations-enforcement.page?>
10. <http://www.azfamily.com/news/Cleaning-crew-at-state-buildings-arrested-in-ID-theft-raid-263111321.html>
11. <http://www.techrepublic.com/blog/10-things/10-physical-security-measures-every-organization-should-take/> and Boyle, R.J. & Panko, R.R. (2012) "Corporate Computer Security," 3rd ed, Pearson, Upper Saddle River, NJ, USA.

Chapter 2

12. <http://theinvisiblethings.blogspot.com/2010/08/ms-dos-security-model.html>
13. [https://en.wikipedia.org/wiki/Threat_\(computer\)](https://en.wikipedia.org/wiki/Threat_(computer))
14. <http://www.verizonenterprise.com/DBIR/2014/>
15. Boyle, R.J. & Panko, R.R. (2012) “Corporate Computer Security,” 3rd ed.
16. <http://www.eweek.com/security/home-depot-breach-expands-privilege-escalation-flaw-to-blame.html>
17. Boyle, R.J. & Panko, R.R. (2012) “Corporate Computer Security,” 3rd ed.
18. <http://www.eweek.com/enterprise-apps/enterprise-linux-adoption-increasing-steadily-study>
19. <http://www.techrepublic.com/article/linux-on-the-desktop-isnt-dead/>
20. <http://www.verizonenterprise.com/DBIR/2014/>

Chapter 3

21. Ponemon Institute (2007) “2007 Annual Study: U.S. Cost of a Data Breach: Understanding Financial Impact, Customer Turnover, and Preventative Solutions.” http://eval.symantec.com/mktginfo/enterprise/other_resources/b-cost_of_data_breach_ponemon-institute_2007.pdf
22. Ponemon Institute (2013) “2013 Cost of Data Breach Study: Global Analysis.” <http://www.ponemon.org/library/2013-cost-of-data-breach-global-analysis?s=global+analysis>
23. Munroe, R. “Exploits of a Mom.” www.xkcd.com/327
24. Wolter, R. and Haselden, K. (November 2006) “The What, Why, and How of Master Data Management.” Microsoft Developer Network, <https://msdn.microsoft.com/en-us/library/bb190163.aspx>

Chapter 4

25. <http://www.fbi.gov/about-us/investigate/counterintelligence/internet-social-networking-risks>
26. <http://arstechnica.com/security/2013/05/think-your-skype-messages-get-end-to-end-encryption-think-again/>
27. http://www.slate.com/articles/technology/future_tense/2015/02/ssl_warnings_users_ignore_them_can_we_fix_that.html
28. http://bits.blogs.nytimes.com/2011/04/02/the-rsa-hack-how-they-did-it/?_r=0
29. <http://www.verizonenterprise.com/DBIR/2014/>

30. <http://www.cs.berkeley.edu/~daw/papers/wireless-cacm.pdf>
31. <https://wagle.net/stats>
32. <http://www.informationweek.com/tj-maxx-data-theft-likely-due-to-wireless-wardriving/d/d-id/1054964?>
33. https://en.wikipedia.org/wiki/Wired_Equivalent_Privacy
34. <https://wagle.net/stats>

Chapter 5

35. <http://www.veracode.com/solutions/by-need/secure-development>
36. https://www.owasp.org/index.php/OWASP_Secure_TDD_Project

Chapter 6

37. <https://en.wikipedia.org/wiki/Scytale>
38. <https://en.wikipedia.org/wiki/Histiaeus>
39. <https://en.wikipedia.org/wiki/Steganography>
40. <http://www.imdb.com/title/tt2084970/>
41. <https://blog.cloudflare.com/why-are-some-keys-small/>
42. <https://nakedsecurity.sophos.com/2013/05/27/anatomy-of-a-change-google-announces-it-will-double-its-ssl-key-sizes/>
43. Checksums were discussed briefly in Chapter 2, but essentially a checksum is just what the name implies. A number is calculated (a “sum”) for a certain file, by performing a mathematical operation on every bit of a file. That number is then saved, and later used to “check” the data by calculating the same sum again. If the two number match, it may be safely said the files are the same. See the Glossary for a more detailed explanation.
44. Schneier, B. (2004) “Secrets and Lies: Digital Security in a Networked World,” 2nd ed., p. xxii, Wiley, Indianapolis, IN, USA.
45. This statement has been attributed to Roger Needham, Butler Lampson, and Bruce Schneier, but none of those have ever admitted to saying it!
46. These inconveniences can be lightened by automating the encryption and decryption process. Lowering the perceived degree of difficulty is essential to getting user buy-in.

Chapter 7

47. The terms white-hat and black-hat hacker come from old western movies. The hero almost always wore a white hat; the villain, a black hat. The

same terminology applies to hackers, as well as so-called gray-hat hackers, who may switch sides depending on who is paying the most.

48. PCI DSS Requirements and Security Assessment Procedures, Version 2.
49. Backdoors are a type of malware that is placed on a system to allow a hacker easy access after they have penetrated a system once. By placing a backdoor, they can login at their leisure, without having to go through defenses again.
50. http://www.pentest-standard.org/index.php/Main_Page
51. Oriyano, S.-P. (2014) "Hacker Techniques, Tools, and Incident Handling," 2nd ed. Jones & Bartlett, Burlington, MA, USA.
52. *Ibid.*
53. http://www.offensive-security.com/metasploit-unleashed/Port_Scanning

Chapter 8

54. The following links provide information on several current standards:
 - FISMA: <http://www.dhs.gov/federal-information-security-management-act-fisma>
 - PCI DSS: <https://www.pcisecuritystandards.org/>
 - HIPAA: <http://www.hhs.gov/ocr/privacy/>
 - FERPA: <http://www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html>
55. The following links provide information on several frameworks to help organizations address compliance with standards:
 - NIST RMF: <http://src.nist.gov/groups/SMA/fisma/framework.html>
 - IASCA COBIT 5: <http://www.isaca.org/cobit/pages/default.aspx>

Many vendors and consulting services also provide security compliance support and guidance.
56. <http://www.datacenterknowledge.com/archives/2011/04/26/disaster-recovery-plans-practice-makes-perfect/>
57. Here are some example templates:
 - <http://publib.boulder.ibm.com/iserics/v5r2/ic2924/info/rzaj1/rzaj1disastr.htm>
 - http://blogs.technet.com/b/mspfe/archive/2012/03/08/a_2d00_microsoft_2d00_word_2d00_document_2d00_template_2d00_for_2d00_disaster_2d00_recovery_2d00_planning.aspx
 - http://cdn.tgtmedia.com/searchSMBstorage/downloads/SearchSMBStorage_business_continuity_plan_template.docx

The first two are for larger organizations and the last link for small businesses. The number of tools and services also are plentiful on the web and vary from

free to fee-based services (e.g., http://www.redanvil.net/dr_tool/ or <http://www.drj.com/resources/sample-plans.html>).

58. Dines, R. (January 2012) “How to Improve Disaster Recovery Preparedness,” CIO. Accessed April 10, 2015: <http://www.cio.com/article/2400373/disaster-recovery/how-to-improve-disaster-recovery-preparedness.html>

Chapter 9

59. McConnell, K.D. “How to Develop Good Security Policies and Tips on Assessment and Enforcement,” SANS Security Essentials, GSEC Practical Assignment, Version 1.3 and Microsoft Security Risk Management Guide.
60. <http://www.isaca.org/Journal/Past-Issues/2005/Volume-6/Pages/JOnline-Creating-and-Enforcing-an-Effective-Information-Security-Policy1.aspx>
61. Adapted from <http://policy.illinoisstate.edu/technology/9-8.shtml>
62. Adapted from <http://www.ccr.gov.uk/datasets/policystatement.shtml>
63. Boyle, R.J. & Panko, R.R. (2012) “Corporate Computer Security,” 3rd ed.
64. <https://www.sans.org/reading-room/whitepapers/policyissues/preparation-guide-information-security-policies-503>
65. <https://www.cert.org/historical/governance/implementation-guide.cfm?>
66. Boyle, R.J. & Panko, R.R. (2012) “Corporate Computer Security,” 3rd ed.
67. Adapted from Boyle, R.J. & Panko, R.R. (2012) “Corporate Computer Security,” 3rd ed.

Chapter 10

68. <http://energy.gov/oe/downloads/21-steps-improve-cyber-security-scada-networks> and <http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf>
69. <http://cacm.acm.org/magazines/2011/8/114953-an-overview-of-business-intelligence-technology/fulltext>
70. <https://hbr.org/2013/04/the-hidden-biases-in-big-data/>
71. <https://www.grc.com/x/ne.dll?bh0bkyd2>
72. A comprehensive database of router default passwords can be found at <http://www.routerpasswords.com/> or by simply searching for “your model number” and “default password.”
73. <https://www.grc.com/passwords.htm>
74. Another handy tool to check whether WPS has been disabled is WigleWiFi, available for Android devices on the Google Play store.
75. <http://www.networkworld.com/article/2899128/mobile-apps/ibm-mobile-app-security-stinks.html>

76. <https://preyproject.com/>
77. http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf
78. Among other sites, <https://twofactorauth.org/>, <http://lifelifehacker.com/twofactorauth-lists-all-the-sites-with-two-factor-authentication-1547219713>, <http://lifelifehacker.com/5938565/heres-everywhere-you-should-enable-two-factor-authentication-right-now> and <http://evanhahn.com/2fa/> have lists of sites enabling two-factor authentication.

Index

- Active Directory
 - database, 19
 - schema design, 19
- Advanced Encryption Standard (AES), 61–64
- Advanced persistent threat (APT), 45
- Algorithm, 38, 40, 41, 58, 64
- Android, 11, 105
- Antimalware software, 13, 23
- Apache, 21
- Apple, 12, 19–20, 104, 105
- Application security, 51–56
 - environment, 53–54
 - migration, 53–54
 - secure developer, need for, 51–52
 - test data, 53–54
 - testing applications, 54–55
 - using data and networks, 53
- Asset sensitivity, 74
- Asymmetric ciphers, encryption in, 61
- Authentication, 49, 63, 67
 - two-factor, 99, 108, 109
- Authorization, 14, 76
- Availability. *See* CIA model

- Backdoors, 70, 76, 150n49
- Backups, 16, 83, 109–112
- Berkeley Software Distribution (BSD), 20
- Big data, 28, 97–98
- Black-hat hackers, 69, 149–150n47
- Blu-Ray writer, 110
- Border firewall, 43, 74
- Breach cost, 28
- Bring your own device (BYOD)
 - programs, 44, 100
- Business intelligence, 98
- Business-to-business (B2B). *See* Electronic data interchange (EDI)

- Caesar, Julius, 58
- Carbonite, 111
- Catastrophe, securing against, 82
- Certificate. *See* Security certificate
- Certificate authority (CA), 62, 63
- Checksum, 17, 64
- CIA (confidentiality, integrity, and availability) triad, 1, 9, 33, 37, 55, 57, 97, 147n2
 - servers in, 16–18
- Ciphertext, 57–60. *See also specific ciphers*
- Citrix, 24
- Cloud-based backup, 111
- Cloud computing, 31, 100
- Cloud security, 99
- Colossus, The, 60
- Computer security, 23. *See also* Information assurance
- Confidentiality. *See* CIA model
- Control Objectives for Information and Related Technology (COBIT), 82, 93
- Converged networks, 36
- Convergence, security and, 37–38
- COSO internal control integrated framework, 93
- Cross-site scripting (XSS), 78
- Cryptanalysis, 60
- Cryptography, 57–67. *See also* Encryption
 - definition of, 57–60
 - public key, 61
 - symmetric key, 61–64

- Database administrator (DBA), 27
- Database Management System (DBMS), 27, 55, 66
 - built-in security, 33
 - security features, 29

- Database threats, types of, 30–31
- Data exfiltration, 76
- Data privacy, 100
- Data security, 27–34, 100
 - breach cost, 28
 - database threats, types of, 30–31
 - DBMS security features, 29
 - information assurance, data quality
 - aspects of, 31–32
 - internal vs. external, 28
 - master data management, 32–33
 - strategy, 33
- Data transfer, 17
- Data warehousing, security in, 97–98
- “Death of the perimeter,” 44
- Decryption. *See* Cryptography
- Defense-in-depth strategy, 9
- Desktops, threat model for, 23
- Development environment security, 53–54
- Disaster, definition of, 81–82
- Disaster recovery plan (DRP), 81–88
 - considerations for, 82–84
 - implementation of, 84–86
 - practice, 86–87
 - reporting, 87
 - revising, 87
 - revisiting, 87
 - securing against catastrophe, 82
- Disk Operating System (DOS), 11
- DropBox, 99, 111
- DVD, 110

- Electronic data interchange (EDI), 5, 6
- Electronic order systems, 5
- Encryption, 17, 57. *See also*
 - Cryptography
 - full-disk, 65
 - need for, 65–67
 - public key, 62–63. *See also*
 - Asymmetric ciphers
 - to secure resources, 64–65
 - single-file, 66–67
 - symmetric key, 63–64. *See also*
 - Symmetric ciphers
- Ettercap, 37
- Exploitation
 - local, 13
 - in penetration testing, 75–76
 - remote, 13
- External threats
 - to data security, 28
 - to organization, 4

- Family Educational Rights and Privacy Act (FERPA), 81
- Federal Information Security Management Act (FISMA), 81
- Firewalls, 2, 17, 19, 38–39, 42–46, 74, 106. *See also specific firewalls*
- FreeBSD, 20
- Full-disk encryption, 65

- Google Docs, 31
- Group Policy Objects (GPOs), 19

- Hackers, 69, 70, 71, 74, 77, 149–150*n*50
- Hacking, definition of, 14
- Hadoop, 28
- Hardening, operating system
 - security, 15–16
- Health Insurance Portability and Accountability Act (HIPAA), 6, 69, 81
- Heartbleed bug, 22
- Home Depot, 14–15
- Home security, 101–109

- IDrive, 111
- Imitation Game, The, 60
- Information assurance, security and, 1–9
 - data quality aspects of, 31–32
 - in enterprise, 3–5
 - physical asset protection, 7–9

- Insider threat, management of, 3–4, 44
- Integrity. *See* CIA model
- Intelligence gathering, in penetration testing, 73
- Internal security, 3
- Internal threats to data security, 28
- Internet Protocol Security, 40
- Internet service providers (ISPs), 44, 45, 101
- Interorganizational security, 5–7
- Intrusion detection/prevention system (IDPS), 43, 46
- iOS, 11, 105
- ISO/IEC 27000 family, 93
- ITIL, 93

- Journaling, 17

- Least permission/access, 4, 16
- Linux, 11, 20, 21, 22, 105
- Local exploit, 13

- Machine-to-machine (M2M) transactions, 5
- Mac OS, 11
- Mac OSX Server, 19–20
- Malware, 4, 12, 23–24, 100
- Man-in-the-middle (MITM) attack, 42
- Master data management, 32–33
- Microsoft, 12, 15, 18, 19, 24, 38, 67, 106–107
- MongoDB, 28
- Mozy, 111

- NetBSD, 20
- Network(s). *See also* WiFi networks
 - converged, 36
 - definition of, 35–36
 - disaster and, 85
 - importance of using secure, 39
 - managing traffic flows on, 36–37
 - security, 37–39, 136–137
 - virtual local area, 36–37
 - virtual private, 25, 38–41
 - WiFi. *See* WiFi networks
- Network Address Translation (NAT), 44–45, 101–102
- NoSQL, 28

- OneDrive, 99
- OpenBSD, 20
- Open source operating systems, 20–21
- Open Source Software (OSS), 20, –21 security, 22
- OpenSSL encryption package, 22
- OpenVPN, 40–41
- Operating system (OS) security
 - hardening, 15–16
 - open source, 20–21
 - patching, 15
 - specifics for, 18–20
 - threat landscape, 12–13

- Passive scanning, 75
- Passphrase, 49–50, 102–103, 105
- Password, 49–50, 102–103, 105
- Patching, 15–19, 22, 96
- Patch Tuesday, 15, 19
- Payment Card Industry Data Security Standard (PCI DSS), 57, 69, 81
- Penetration testing, 69–79
 - internal vs. external, 70–71
 - performance of, 71–76
 - volunteer testers, 77–79
- Penetration Testing Execution Standard, 71–72
- Personal computing security, 37
- Personal Identification Numbers (PINs), 11, 49–50, 104, 108
- Personally Identifiable Information (PII), 97
- Personal security, 112–113
- Physical asset protection, 7–9
- Physical resources, disaster and, 85
- Pirated software, 24
- Plaintext, 57

- Policy, 100. *See also specific policies*
- Port, 42
- Pre-engagement interactions, in
 - penetration testing, 72
- Privilege escalation, 14, 76
- Production setting security, 54
- Public key cryptography. *See*
 - Asymmetric ciphers
- Remote access, 24–25
 - VPNs. *See* Virtual private networks (VPN), host-to-site
- Remote Access Trojan (RAT), 24
- Remote Desktop, 24
- Remote exploit, 13
- Risk Management Framework (RMF), 82
- Road warrior VPNs. *See* Virtual private networks (VPN), host-to-site
- Rogue software, 23–24
- Sarbanes-Oxley, 69
- Schneier, Bruce, 1–2, 67
- Secure developer, need for, 51–52
- Secure Sockets Layer (SSL), 40, 61–64
- Security
 - application, 51–56
 - certificate, 63
 - concerns with convergence, 37–38
 - data, 27–34
 - information assurance and, 3–5
 - internal security, 3
 - interorganizational, 5–7
 - management, 2
 - network, 37–39, 136–137
 - Open Source Software, 22
 - operating system, 11–25
 - policy. *See* Security policy
 - tools, 46
 - trends in, 99–101
 - wireless, 46–50
- Security policy, 89–94
 - development of, 93–94
 - examples of, 91–92
- Servers, in CIA model, 16–18
- Skype, 38
- Sniffing, 37
- Snowden, Edward, 3
- Social engineering, 72, 73
- Software as a Service (SaaS), 31
- SOHO (small office/home office)
 - security, 101–109
- Spam over Internet telephony (SPIT), 37
- Stateful packet inspection (SPI)
 - firewalls, 44
- Steganography, 58
- Structured Query Language (SQL), 27
- Substitution cipher, 58
- Supervisory Control and Data Acquisition (SCADA), 95–97
- Symmetric ciphers, encryption in, 61
- Tailgating, 8
- Tape drives, for backups, 111
- Terminal Services, 24
- Test-driven development (TDD), 55
- Testing applications, 54–55
- Thawte, 62
- Threat landscape, 12–13
- Threat modeling
 - for desktop, 23
 - in penetration testing, 74
- Traffic shaping or prioritization
 - method, 36
- Transmission Control Protocol (TCP), 42
- Transport Layer Security (TLS), 40, 61–64
- Transposition cipher, 59
- Trojan horses (trojans), 12, 23, 24
- Two-factor authentication, 99, 109
- Unified threat management (UTM)
 - firewalls, 45
- Universal Serial Bus (USB), 8
 - for backups, 110
- UNIX, 20
- User Datagram Protocol (UDP), 42

- Verisign, 62
- Video conferencing services, 38
- Virtual local area networks (VLANs), 36–37
- Virtual Network Computing (VNC), 24
- Virtual private network (VPN), 25, 38–39
 - host-to-host, 41
 - host-to-site, 41
 - importance of, 39
 - for remote workers on unsecured WiFi networks, 41–42
 - site-to-site, 41
 - types of, 39–41
- Viruses, 23
- Voice over Internet Protocol (VoIP), 36, 37
- Volunteers, for penetration testing, 77–79
- Vulnerability analysis, in penetration testing, 74–75

- White-hat hackers, 69, 149–150ⁿ⁴⁷
- WiFi networks. *See also* Networks
 - VPNs for remote workers on unsecured, 41–42
 - wireless security, 46–50
- WiFi Protected Access (WPA), 48
- WiFi Protected Setup (WPS), 49, 103
- Windows 10, 11
- Windows, 11, 21, 22
- Windows 2000, 11
- Windows 7, 11
- Windows 8, 11
- Windows 95, 11
- Windows 98, 11
- Windows NT, 11
- Windows Security Essentials, 106–107
- Windows Server, 18
- Windows Server 2003, 3
- Windows Server Update Services, 19
- Windows XP, 3, 11
- Wired Equivalent Privacy (WEP), 47–48
- Wireless access point, 102
- Wireless router, 102
- Wireless security, 46–50
- Workstation resources, disaster and, 86
- Worms, 23

THE BUSINESS EXPERT PRESS DIGITAL LIBRARIES

EBOOKS FOR BUSINESS STUDENTS

Curriculum-oriented, born-digital books for advanced business students, written by academic thought leaders who translate real-world business experience into course readings and reference materials for students expecting to tackle management and leadership challenges during their professional careers.

POLICIES BUILT BY LIBRARIANS

- *Unlimited simultaneous usage*
- *Unrestricted downloading and printing*
- *Perpetual access for a one-time fee*
- *No platform or maintenance fees*
- *Free MARC records*
- *No license to execute*

The Digital Libraries are a comprehensive, cost-effective way to deliver practical treatments of important business issues to every student and faculty member.

For further information, a
free trial, or to order, contact:

sales@businessexpertpress.com

www.businessexpertpress.com/librarians



BUSINESS EXPERT PRESS

Information Technology Security Fundamentals

Glen Sagers • Bryan Hosack

Information security is at the forefront of timely IT topics, due to the spectacular and well-publicized breaches of personal information stored by companies. To create a secure IT environment, many steps must be taken, but not all steps are created equal. There are technological measures that increase security, and some that do not, but overall, the best defense is to create a culture of security in the organization.

The same principles that guide IT security in the enterprise guide smaller organizations and individuals. The individual techniques and tools may vary by size, but everyone with a computer needs to turn on a firewall and have antivirus software. Personal information should be safeguarded by individuals and by the firms entrusted with it. As organizations and people develop security plans and put the technical pieces in place, a system can emerge that is greater than the sum of its parts.

Glen Sagers is an associate professor at Illinois State University, teaching networking and security courses. He received his PhD from Florida State University and has published articles about the processes used to create open source software, and wireless security. Most recently, he contributed a chapter on threats to wireless privacy to the book, *Privacy in the Digital Age, 21st Century Challenges to the Fourth Amendment*.

Bryan Hosack currently works as a senior analyst in business intelligence, reporting and analytics in the financial industry. He has taught, worked and consulted in a variety of IT areas across a variety of industries. He received his PhD from Florida State University.

THE INFORMATION SYSTEMS COLLECTION

Daniel J. Power, Editor

ISBN 978-1-60649-916-0

