



Managing Risk and Information Security

Protect to Enable

—

Second Edition

—

Malcolm W Harkins

Apress
open

www.dbooks.org

Managing Risk and Information Security

Protect to Enable

Second Edition



Malcolm W. Harkins

Apress
open

Managing Risk and Information Security: Protect to Enable

Malcolm W. Harkins
Folsom, California, USA

ISBN-13 (pbk): 978-1-4842-1456-5
DOI 10.1007/978-1-4842-1455-8

ISBN-13 (electronic): 978-1-4842-1455-8

Library of Congress Control Number: 2016949414

Copyright © 2016 by Malcolm W. Harkins

ApressOpen Rights: You have the right to copy, use and distribute this Work in its entirety, electronically without modification, for non-commercial purposes only. However, you have the additional right to use or alter any source code in this Work for any commercial or non-commercial purpose which must be accompanied by the licenses in (2) and (3) below to distribute the source code for instances of greater than 5 lines of code. Licenses (1), (2) and (3) below and the intervening text must be provided in any use of the text of the Work and fully describes the license granted herein to the Work.

(1) License for Distribution of the Work: This Work is copyrighted by Malcolm Harkins, all rights reserved. Use of this Work other than as provided for in this license is prohibited. By exercising any of the rights herein, you are accepting the terms of this license. You have the non-exclusive right to copy, use and distribute this English language Work in its entirety, electronically without modification except for those modifications necessary for formatting on specific devices, for all non-commercial purposes, in all media and formats known now or hereafter. While the advice and information in this Work are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

If your distribution is solely Apress source code or uses Apress source code intact, the following licenses (2) and (3) must accompany the source code. If your use is an adaptation of the source code provided by Apress in this Work, then you must use only license (3).

(2) License for Direct Reproduction of Apress Source Code: This source code, from Intel® Trusted Execution Technology for Server Platforms, ISBN 978-1-4302-6148-3 is copyrighted by Apress Media, LLC, all rights reserved. Any direct reproduction of this Apress source code is permitted but must contain this license. The following license must be provided for any use of the source code from this product of greater than 5 lines wherein the code is adapted or altered from its original Apress form. This Apress code is presented AS IS and Apress makes no claims to, representations or warranties as to the function, usability, accuracy or usefulness of this code.

(3) License for Distribution of Adaptation of Apress Source Code: Portions of the source code provided are used or adapted from Intel® Trusted Execution Technology for Server Platforms, ISBN 978-1-4302-6148-3 copyright Apress Media LLC. Any use or reuse of this Apress source code must contain this License. This Apress code is made available at Apress.com/9781484214565 as is and Apress makes no claims to, representations or warranties as to the function, usability, accuracy or usefulness of this code.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark. The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Cover image designed by Freepik.

Managing Director: Welmoed Spahr
Lead Editor: Robert Hutchinson
Development Editor: James Markham
Editorial Board: Steve Anglin, Pramila Balen, Aaron Black, Louise Corrigan, Jonathan Gennick, Robert Hutchinson, Celestin Suresh John, Nikhil Karkal, James Markham, Susan McDermott, Matthew Moodie, Natalie Pao, Gwenan Spearing
Coordinating Editor: Melissa Maldonado
Copy Editor: Mary Behr
Compositor: SPi Global
Indexer: SPi Global
Artist: SPi Global

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springer.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales-eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code/.

Printed on acid-free paper

About ApressOpen

What Is ApressOpen?

- ApressOpen is an open access book program that publishes high-quality technical and business information.
- ApressOpen eBooks are available for global, free, noncommercial use.
- ApressOpen eBooks are available in PDF, ePub, and Mobi formats.
- The user friendly ApressOpen free eBook license is presented on the copyright page of this book.

This book is dedicated to my family.

Contents at a Glance

Foreword	xv
Praise for the second edition of Managing Risk and Information Security	xvii
About the Author	xxi
Acknowledgments	xxiii
Preface	xxv
■ Chapter 1: Introduction	1
■ Chapter 2: The Misperception of Risk	17
■ Chapter 3: Governance and Internal Partnerships: How to Sense, Interpret, and Act on Risk	31
■ Chapter 4: External Partnerships: The Power of Sharing Information	49
■ Chapter 5: People Are the Perimeter	65
■ Chapter 6: Emerging Threats and Vulnerabilities: Reality and Rhetoric	81
■ Chapter 7: A New Security Architecture to Improve Business Agility	99
■ Chapter 8: Looking to the Future: Emerging Security Capabilities	117

■ CONTENTS AT A GLANCE

- **Chapter 9: Corporate Social Responsibility: The Ethics of Managing Information Risk** 129
- **Chapter 10: The 21st Century CISO** 139
- **Chapter 11: Performance Coaching** 155
- **Appendix A: References** 171
- Index** 181

Contents

Foreword	xv
Praise for the second edition of Managing Risk and Information Security	xvii
About the Author	xxi
Acknowledgments	xxiii
Preface	xxv
■ Chapter 1: Introduction	1
Protect to Enable®	5
Building Trust.....	8
Keeping the Company Legal: The Regulatory Flood	8
The Rapid Proliferation of Information, Devices, and Things	12
The Changing Threat Landscape	13
A New Approach to Managing Risk	16
■ Chapter 2: The Misperception of Risk	17
The Subjectivity of Risk Perception.....	18
How Employees Misperceive Risk.....	18
The Lure of the Shiny Bauble.....	20
How Security Professionals Misperceive Risk	20
Security and Privacy	22
How Decision Makers Misperceive Risk	23

How to Mitigate the Misperception of Risk	24
Uncovering New Perspectives During Risk Assessments.....	25
Communication Is Essential	26
Building Credibility	28
■ Chapter 3: Governance and Internal Partnerships: How to Sense, Interpret, and Act on Risk	31
Information Risk Governance	32
Finding the Right Governance Structure	34
Building Internal Partnerships.....	37
Legal	38
Human Resources	42
Finance	43
Corporate Risk Management	44
Privacy	45
Corporate Security.....	45
Business Group Managers.....	46
Conclusion.....	47
■ Chapter 4: External Partnerships: The Power of Sharing Information.....	49
The Value of External Partnerships	51
External Partnerships: Types and Tiers.....	52
1:1 Partnerships	55
Communities.....	57
Community Characteristics	57
Community Goals.....	59
Sharing Information about Threats and Vulnerabilities.....	59
Sharing Best Practices and Benchmarking	60

Influencing Regulations and Standards.....	62
Corporate Citizenship	63
Conclusion.....	63
■ Chapter 5: People Are the Perimeter	65
The Shifting Perimeter	65
Compliance or Commitment?.....	66
Examining the Risks.....	68
Adjusting Behavior	69
A Model for Improving Security Awareness	71
Broadening the Awareness Model.....	74
The Security Benefits of Personal Use	74
Roundabouts and Stop Signs	75
The Technology Professional.....	77
Insider Threats.....	78
Deter	79
Detect	79
Discipline	80
Finding the Balance.....	80
■ Chapter 6: Emerging Threats and Vulnerabilities: Reality and Rhetoric	81
Structured Methods for Identifying Threat Trends.....	82
The Product Life Cycle Model	83
Understanding Threat Agents	88
Playing War Games.....	90
Trends That Span the Threat Landscape	91
Trust Is an Attack Surface.....	91
Barriers to Entry Are Crumbling.....	92

The Rise of Edge Case Insecurity	92
The Enemy Knows the System	93
Key Threat Activity Areas.....	94
The Industry of Malware.....	94
The Web Expands to the Internet of Things.....	94
Smartphones.....	96
Web Applications	97
Conclusion.....	97
■ Chapter 7: A New Security Architecture to Improve Business Agility.....	99
The 9 Box of Controls, Business Trends, and Architecture Requirements	101
9 Box of Controls	101
IT Consumerization	102
New Business Needs.....	103
Cloud Computing	104
Changing Threat Landscape	104
Privacy and Regulatory Requirements.....	105
New Architecture.....	105
Trust Calculation.....	106
Security Zones.....	109
Balanced Controls.....	113
Users, Data, and the Internet of Things: The New Perimeters	115
Conclusion.....	116
■ Chapter 8: Looking to the Future: Emerging Security Capabilities.....	117
Internet of Things	120
Consistent User Experience Across Devices	121

Cloud Computing	122
Big Data Analytics	122
Artificial Intelligence	122
Business Benefits and Risks	123
New Security Capabilities.....	123
Baseline Security.....	124
Context-Aware Security.....	126
Conclusion.....	127
■ Chapter 9: Corporate Social Responsibility: The Ethics of Managing Information Risk	129
The Expanding Scope of Corporate Social Responsibility	130
The Evolution of Technology and Its Impact	132
Maintaining Society’s Trust	134
The Ethics of Managing Information Risk	135
Conclusion.....	137
■ Chapter 10: The 21st Century CISO	139
Chief Trust Officer.....	139
The Z-Shaped Individual.....	141
Foundational Skills.....	142
Becoming a Storyteller.....	143
Fear Is Junk Food.....	144
Accentuating the Positive	145
Demonstrating the Reality of Risk.....	146
The CISO’s Sixth Sense	147
Taking Action at the Speed of Trust	148
The CISO as a Leader	148
Learning from Other Business Leaders	149

■ CONTENTS

Voicing Our Values	150
Discussing Information Risk at Board Level	151
Conclusion.....	153
■ Chapter 11: Performance Coaching.....	155
How to Use the Tables	156
Independence and Initiative	157
Efficiency and Effectiveness.....	158
Commitment.....	160
Professionalism	161
Discipline	161
Teamwork.....	162
Problem-Solving.....	163
Communication.....	164
Goal-Setting.....	168
Conclusion.....	169
■ Appendix A: References.....	171
Index.....	181

Foreword

Security and first-person shooter video games have one obvious thing in common: if you're not continuously moving, you're dead. In this second edition of *Managing Risk and Information Security*, Malcolm Harkins helps us move our thinking into areas of risk that have become more prominent over the last several years.

Because there is so much new content in this edition, I will focus on a topic that has risen to greater prominence since the first edition: people are the perimeter. When we reflect on what has changed in recent years, with an eye to the vulnerabilities that result in real-world compromises, a pattern emerges: virtually all the major breaches that we have seen involve manipulation of people. When nearly everyone has heard of phishing, we have to ask ourselves: why is it still such an effective tool?

The obvious theory is that we haven't managed people risk as well as we should. Perhaps we have been standing still and need to learn how to dodge and experiment with the way we drive better people-security outcomes. Unfortunately, the path is not 100% clear. Unlike technology, the field of influencing human behavior in security is remarkably complicated and supported by limited research.

Malcolm provides us with a great foundation and framework to build our "security engagement" functions. I like to use the word "engagement" because it speaks to how the security organization relates to the workforce in a manner that isn't simply bounded by the more traditional term "training and awareness." Engagement encompasses anything that shifts the desired behavior outcome in the direction we want it to go. I have seen remarkable shifts in measured behavior from the use of non-traditional tools such as security gamification and simulation.

The way Malcolm differentiates between "compliance" and "commitment" is key. *Managing Risk and Information Security* is an ever-evolving classic in the field of security management.

—Patrick Heim
Head of Trust & Security, Dropbox

Praise for the second edition of *Managing Risk and Information Security*

We assign Malcolm's book to our Carnegie Mellon CISO-Executive Program students on their first day of class. It is relevant, pragmatic, and solution oriented. Our adversaries are changing their practices and so must we. Malcolm's book is a terrific tool for the modern-day info sec leader who wants to shift from security as a restriction to security as a business enabler.

—Andy Wasser
Associate Dean, CMU Heinz College

Malcolm is a top-notch executive, security leader, and innovator, with a keen ability to convey thought-provoking and valuable insights. His latest effort demonstrates remarkable foresight into the skills necessary to excel as a security leader today and tomorrow.

—Clayton J. Pummill
Executive Director, Security Advisor Alliance

*I could go on and on about what I liked specifically—there was much, including the discussion about governance models and social responsibility—but here is the net: this is the first time I've seen someone be able to speak to security specifics while also raising the conversation to a much higher level. It begins to take on an Alvin Toffler feel from his astounding book, *The Third Wave*. Malcolm's thoughts are philosophically sweeping while at the same time imminently practical.*

—Todd Ruback, Esq., CIPP-US/E, CIPT
Chief Privacy & Security Officer & V.P. Legal Affairs, Ghostery

Malcolm Harkins is a foremost expert at managing risk and information security. In this latest book, he further expands his Protect to Enable philosophy and does so in a way that offers practical and actionable initiatives that any risk manager or CISO can implement to protect their enterprise while enabling business growth. A must-read for CISOs and their teams!

—Tim Rahschulte, Ph.D.
Chief Learning Officer & Content Officer, Evanta

Malcolm Harkins is a visionary thought leader on cyber security and risk management. Managing Risk and Information Security is a must read. Malcolm helps readers immediately take the information and apply it to their own organizations. You will find that this book cuts through the fog and provides a clear picture of where and what to focus on to effectively manage cyber business risk.

—Phil Ferraro
Global CISO and Cyber Security Consultant

The CISO is more than just a technology expert; she must be savvy about leadership, influence, and change across complex organizations; someone who sees her mission not to just drive implementation of a large system, but to foster sustainable culture change at every level. As an organizational psychologist, I recognize Harkins' keen eye for group dynamics and leadership tactics that enable CISOs to enhance enterprise security. He puts his finger on the habits, assumptions, and decision processes typical of many employees and teams, as they unknowingly increase security risk, and for that alone this book is a gem. It should be required reading for aspiring CISOs and for anyone who has a role in the recruitment and hiring of CISOs.

—Marc Sokol, PhD
Executive Editor, People + Strategy

Malcolm Harkins' take on information security and risk is a refreshing change from the increasingly frequent alarm bells raised in the press with regard to the "brave new world" where technology is presented as an ever-escalating conflict between our seemingly insatiable appetite for connectivity, cool applications, and customized information, on the one hand, and a desire to control who has our information and how they may use it, on the other. Harkins instead offers a cool, clear-eyed perspective where managing information and risk are placed in a wider context. His prescriptions and frameworks are recipes for well-managed organizations in the broadest sense. They allow us to embrace our new-found

technological abilities without fear because we have defined their purpose capaciously enough to be a positive good, to be of service to all a company's stakeholders. That is, once we set a truly human course, technology serves rather than threatens us. Organization purpose, when defined in this way, is an expression of our values and is empowered by that fuel. Harkins' book is a practical as well as purposeful guide to a values-driven implementation of information technology.

—Mary C. Gentile, PhD

Author of *Giving Voice To Values: How To Speak Your Mind When You Know What's Right* (Yale University Press)

*In today's rapidly evolving security landscape, security professionals are navigating a complex set of dynamics across the enterprise. In *Managing Risk and Information Security*, Malcolm Harkins draws on his rich security experience to present a connected view of where companies should be focused. He puts forth a valuable perspective, as organizations around the world look to create a necessary balance of protection and innovation, which ultimately enables business success.*

—Bret Arsenault

Corporate Vice President and CISO, Microsoft Corporation

Malcolm generously shares through personal experiences and story telling the formula for a successful 21st century CISO. It is one part multi-disciplinary leader and one part trusted advisor to the business, combined with behavioral models required for balanced risk decision making. A must-read for all new CISOs. Malcolm lives his beliefs.

—Nasrin Rezai

GE Corporate Security & Compliance Officer

In the second edition of his book, Malcolm seamlessly articulates the future horizon of cyber security and the critical role that the CISO and security professionals will need to fulfill in order to defend both the company and consumers they serve. The guidance he provides into the skills, leadership, and approach required for successfully navigating the emerging challenges of securing a digital economy is invaluable. Regardless of your current role, this is a must-read for everyone who has accepted this great responsibility and privilege.

—Steven Young

CISO, Kellogg Company

While other security officers are looking to the traditional or the latest “cool” product, Harkins goes against the tide and asks the questions that need addressing. His forward-thinking mindset and Protect to Enable approach inspire others to innovate and go beyond the mainstream. If you cannot bring Harkins to your company for mentoring, this book will at least spark thought and will change how your engineers view security within the business.

—Charles Lebo
Vice President and CISO, Kindred Healthcare

Malcolm’s vast experience makes him one of the most credible security leaders on the international stage and serves as the perfect platform for this book. Rational, compelling, and authoritative writing is far too rare in the world of risk and information security, but Malcolm completely nails it in Managing Risk and Information Security with invaluable advice and recommendations for anyone planning a future in the security world. His extensive experience in business before becoming a CISO is one of the missing ingredients in many security executives’ professional toolbox, which is why this is such an important book. Make sure to keep a highlighter and notepad handy because there are a lot of nuggets in here you’ll want to remember on your journey to becoming a better security professional.

—Mark Weatherford
Chief Cybersecurity Strategist at vArmour and
former Deputy Under Secretary for Cybersecurity
at the US Department of Homeland Security

I’ve had the privilege of working with many talented CISOs over the years and Malcolm is one of the best. His logical, methodical approach to solving the most complex cybersecurity problems is reflected in his lucid style. An enlightened approach to understanding risk that unites all stakeholders and a systemic intelligence-based approach to security infrastructure are the only ways to reduce the threat to manageable levels. This is our best path forward if we are ever to realize the vast potential of the innovative digital world we are creating. In Managing Risk and Information Security, Malcolm shines a light on that path in a comprehensive yet very readable way.

—Art Coviello
Former CEO and Executive Chairman, RSA

About the Author



Malcolm Harkins is the Chief Security and Trust Officer (CSTO) at Cylance Inc. In this role, he reports to the CEO and is responsible for enabling business growth through trusted infrastructure, systems, and business processes. He has direct organizational responsibility for information technology, information risk, and security, as well as security and privacy policy. Malcolm is also responsible for peer outreach activities to drive improvement across the world in the understanding of cyber risks and best practices to manage and mitigate those risks.

Previously, Malcolm was Vice President and Chief Security and Privacy Officer (CSPO) at Intel Corporation. In that role, Malcolm was responsible for managing the risk, controls, privacy, security, and other related compliance activities for all of Intel's information assets, products, and services.

Before becoming Intel's first CSPO, he was the Chief Information Security Officer (CISO)

reporting into the Chief Information Officer. Malcolm also held roles in finance, procurement, and various business operations. He has managed IT benchmarking and Sarbanes-Oxley-compliance initiatives. Harkins acted as the profit and loss manager for the Flash Product Group at Intel; was the general manager of Enterprise Capabilities, responsible for the delivery and support of Intel's Finance and HR systems; and worked in an Intel business venture focusing on e-commerce hosting.

Malcolm previously taught at the CIO Institute at the UCLA Anderson School of Management and was an adjunct faculty member at Susquehanna University in 2009. In 2010, he received the RSA Conference Excellence in the Field of Security Practices Award. He was recognized by Computerworld as one of the Premier 100 Information Technology Leaders for 2012. (ISC)² recognized Malcolm in 2012 with the Information Security Leadership Award. In September 2013, Malcolm was recognized as one of the Top 10 Breakaway Leaders at the Global CISO Executive Summit. In November 2015, he received the Security Advisor Alliance Excellence in Innovation Award. He is a Fellow with the Institute for Critical Infrastructure Technology, a non-partisan think-tank that provides cybersecurity briefings and expert testimony to the U.S. Congress and federal agencies. Malcolm is a sought-after speaker for industry events. He has authored many white

■ ABOUT THE AUTHOR

papers and in December 2012 published his first book, *Managing Risk and Information Security*. He also was a contributing author to *Introduction to IT Privacy*, published in 2014 by the International Association of Privacy Professionals.

Malcolm received his bachelor's degree in economics from the University of California at Irvine and an MBA in finance and accounting from the University of California at Davis.

Acknowledgments

I received valuable feedback from many readers of the first edition of this book. That feedback helped me to expand the book with additional insights, clarifications, and updated examples. It also encouraged me to add two more chapters to the second edition: one on corporate social responsibility, and the other on performance coaching.

Special thanks to Mike Faden: without his help this book would not have happened.

As I noted in the first edition, many people during my journey at Intel helped me learn and grow. A number of them published material that is still referenced in this second edition.

Other experts who have helped me come from a variety of different peer groups. They include members of the Bay Area CSO Council, the Executive Security Action Forum, the members and staff of CEB and its Information Risk Leadership Council, participants in the Evanta CISO Executive Summits and the CISO coalition, as well as the Security Advisor Alliance.

Finally, I wish to thank Stuart McClure for giving me the opportunity to join Cylance.

Preface

*If you don't believe in the messenger, you won't believe the message.
You can't believe in the messenger if you don't know what the messenger
believes.
You can't be the messenger until you're clear about what you believe.*

—James Kouzes and Barry Posner,
in *The Leadership Challenge*

A great deal has transpired since the first edition of this book was published in January 2013, both in the world of information risk and in my personal life and career. To briefly cover the latter, in January 2013, I was named Intel's Chief Security and Privacy Officer. My broad role was one of the first of its kind in corporate America: I was charged with managing and mitigating risk for Intel's products and services worldwide, in addition to Intel's internal IT environment. In June 2015, I left Intel to become CISO at Cylance Inc., and in May 2016, I was named Cylance's Chief Security and Trust Officer.

These career changes occurred during an extraordinary period of escalating information risk, as evidenced by an almost continuous stream of major hacks and breaches, and a corresponding rise in society's awareness of risk. Some key examples:

- May 2013: Edward Snowden flies to Hong Kong after leaving his job at an NSA facility in Hawaii. The following month, he reveals thousands of classified NSA documents. The disclosures, including previously unknown government surveillance programs, continue to cause worldwide repercussions today.
- December 2013: The blog Krebs On Security reports a massive data breach at Target. The company confirms the breach the next day. Within months, Target's CIO and CEO both resign amid the fallout.
- May 2014: A U.S. grand jury indicts five Chinese military officers on charges of hacking American companies and stealing trade secrets.
- November 2014: Employees at Sony Pictures arrive at work to discover their network has been hacked. Attackers steal and then erase data on thousands of systems, forcing studio employees to revert to using fax machines and pen and paper. The attackers then dump huge batches of confidential business and personal information online.

- March 2015: Google’s Project Zero hacking team demonstrates the ability to exploit a fundamental flaw in DDR3 SDRAM to perform privilege escalation attacks on systems containing the chips. Some mitigation approaches are available, other than replacing the DDR3 memory in millions of systems worldwide.
- June 2015: The US Office of Personnel Management announces a data breach targeting the personal data of up to 4 million people. The attack, which includes security clearance-related information, is one of the largest-ever breaches of government data. By July, the estimated number of stolen records increases to 21.5 million.
- February 2016: The Hollywood Presbyterian Medical Center in Los Angeles says it has paid a bitcoin ransom to attackers who held its systems hostage, encrypting data and blocking access by hospital staff. Some believe the healthcare industry is the next major target for cyber criminals.

Given this escalating cycle of risk, and the potential catastrophic societal implications of today’s attacks, we must all be ready to be held accountable. This may require a large mental shift for those used to simply assigning responsibility and blame for a breach to the people who traditionally perform post-attack cleanup: corporate IT departments, internal information security teams, and investigations and computer forensics groups. Everyone, from corporate executives to security practitioners, shares responsibility for security and privacy. We must all step back and contemplate our own personal responsibilities, not only to the organizations we work for and the customers we serve, but also to society as a whole.

The challenge we sometimes face is how to characterize that responsibility. Is our responsibility to limit liability for our organizations? Or is it a duty of care to the people whose information we store? What values are we using when we make decisions about cyber risk, and what bias do those values create in our decisions? Are we forward-looking enough, or will the decisions we make to fix our problems today create other problems in the future? As Benjamin Franklin once said, “All human situations have their inconveniences. We feel those of the present but neither see nor feel those of the future; and hence we often make troublesome changes without amendment, and frequently for the worse.”

As security and privacy professionals, a key part of our role is to ensure the right dialogue and debate occurs. We need to ask “high-contrast” questions that sharply define the implications of the choices our organizations make. We need to make sure that the opportunities are as clearly defined as the obligations to mitigate risk, so that our organizations make the right decisions. And we need to take equal responsibility for the outcomes of those choices, as opposed to abdicating that responsibility solely to the business. Once the choice is made, we must transition out of the debate about what is right and focus on taking the right actions—on making tomorrow better than today.

We can think of this as doing what’s right. We can think of it as protecting our customers and partners and keeping our markets healthy for everyone. No matter what motivates us, thoughtfully building systems to support a culture of genuine responsibility for privacy and security is not only good corporate responsibility; it is also good for

business. For computing to continue to improve the world we live in rather than endanger it, it needs to be trustworthy. And for that trust to be deliverable, we need to ensure the data we enter into our computers is both secure and private. As an organization, we demonstrate and build trust through our approach to solving these cyber-risk challenges.

In the preface of the first edition, I said “*Managing Risk and Information Security* is a journey, but there is no finish line. Our approach to managing information risk must continue to evolve as rapidly as the pace of business and technology change. My hope is that people will read this book and begin their own journey.”

I still firmly believe what I said then. But I also believe that, as General George Marshall once said, “The only way human beings can win a war is to prevent it.” We are at war against adversaries who wish to harm the users of technology. But there is also a battle among those responsible for protecting security and privacy. On one side are organizations that would like to continue on the current path because they profit from the insecurity of computing, or that approach the duty of care with a bias towards limiting liability rather than protecting their customers. On the other side are those who believe that our role is to generate trust. We do that by protecting to enable people and businesses. It’s a hard road; I know, because I experience it every day. But we shouldn’t back away from something just because it is hard. We need to plant our feet and stand firm. The only question is where we plant our feet.

CHAPTER 1



Introduction

There are two primary choices in life: to accept conditions as they exist, or accept the responsibility for changing them.

—Denis Waitley

In January 2002, I was hired to run a new Intel internal program called Security and Business Continuity. The program had been created following the major security events of the previous year (9/11 and the Code Red/Nimda viruses) and it focused primarily on the availability risks at that time. I had no background in technical security, but I had been at Intel for nearly 10 years in a variety of business-related positions, mostly in finance. As I learned about information risk during the first few months, it became apparent to me that the world was starting to change rapidly and that a “perfect storm” of risk was beginning to brew. In June 2002, I put together a diagram (Figure 1-1) to explain the risks to my manager, Intel’s CIO, and anyone who would listen to me. The diagram has been updated slightly since then to more explicitly highlight the geo-political forces that are a key part of the threat, vulnerability, and regulatory risk landscape.

The Perfect Storm of Risk




Figure 1-1. *The perfect storm of information risk*

Today, it is clear that my view of the world was essentially accurate. Security breaches and intrusions are reported almost daily at organizations of all sizes, legal and regulatory issues related to technology use continue to grow, and geo-politics have surged to the forefront of some of these discussions in a post-Snowden era. Cyber attacks and data breaches are now considered the biggest threats to business continuity, according to a recent survey (Business Continuity Institute 2016).

But the key question that I asked in the first edition of this book is still valid. Is information security really effective? Given the rapid evolution of new technologies and uses, does the information security group even need to exist?

Obviously, this is a somewhat rhetorical question. I cannot imagine that any sizeable organization would operate well without an information security function. But the real issue is whether the information security group should continue to exist as it does today, with its traditional mission and vision. It is clear from the prevalence of breaches and compromises that we have not kept up with the threats, and we appear to be slipping farther behind as the world grows more volatile, uncertain, and ambiguous. It is no wonder that we have fallen behind: as the world of technology expands exponentially, so do the technology-related threats and vulnerabilities, yet our ability to manage those security and privacy risks has progressed only at a linear rate. As a result, there is a widening gap between the risks and the controls. In fact, many organizations have essentially given up actively trying to prevent compromises and have defaulted to reliance on after-the-fact detection and response tools.

As information risk and security professionals, we should be asking ourselves pointed questions if we wish to remain valuable and relevant to our organizations. Why do we exist? What should our role be? How are new consumer and Internet of Things (IoT) technologies shaping what we do, and can we shape the world of these new technologies and usage models? How is the evolving threat landscape shaping us, and can we shape the threat landscape? Given the bewildering pace at which technology changes and new threats appear, how do we focus and prioritize our workload? What skills do we need?

Traditionally, information security groups in businesses and other organizations have taken a relatively narrow view of security risks, which resulted in a correspondingly narrow charter. We focused on specific types of threats, such as malware. To combat these threats, we applied technical security controls. In an attempt to protect against attacks and stop them reaching business applications and employees' PCs, we fortified the network perimeter using firewalls and intrusion detection software. To prevent unauthorized entry to data centers, we installed physical access control systems. Overall, our thinking revolved around how to lock down information assets to minimize security risks, and how to reactively detect and respond to risks as they presented themselves.

Today, however, I believe that this narrow scope not only fails to reflect the full range of technology-related risk to the business; it is detrimental to the business overall. Because this limited view misses many of the risks that affect the organization, it leaves areas of risk unmitigated and therefore leaves the organization vulnerable in those areas. It also makes us vulnerable to missing the interplay between risks and controls: by implementing controls to mitigate one risk, we may actually create a different risk. And by focusing primarily on detection and response, we are not preventing harm; we are just trying to limit the damage.

As I'll explain in this book, we need to shift our primary focus to adopt a broader view of risk that reflects the pervasiveness of technology today. Organizations still need traditional security controls, but they are only part of the picture.

There are several reasons for this. All stem from the reality that technology plays an essential role in most business activities and in people's daily lives.

Technology has become the central nervous system of a business, supporting the flow of information that drives each business process from product development to sales. In addition, as I'll discuss throughout this book, almost every company is becoming a supplier of technology in some form, as technology becomes a vital element of most products, services, and infrastructure from cars and household appliances to the power grid.

The role of technology in peoples' personal lives has expanded dramatically, too, and the boundaries between business and personal use of technology are blurring. Marketers want to use social media to reach more consumers. Employees want to use their personal smartphones to access corporate e-mail.

Meanwhile, the regulatory environment is expanding rapidly, affecting the way that information systems must manage personal, financial, and other information in order to comply—and introducing a whole new area of IT-related business risks.

Threats are also evolving quickly, as attackers develop more sophisticated techniques, often targeted at individuals, which can penetrate or bypass controls such as network firewalls, traditional antivirus solutions, and outdated access control mechanisms such as passwords.

In combination, these factors create a set of interdependent risks to a business’s information and technology, from its internal information systems to the products and services provided to its customers, as shown in Figure 1-2.

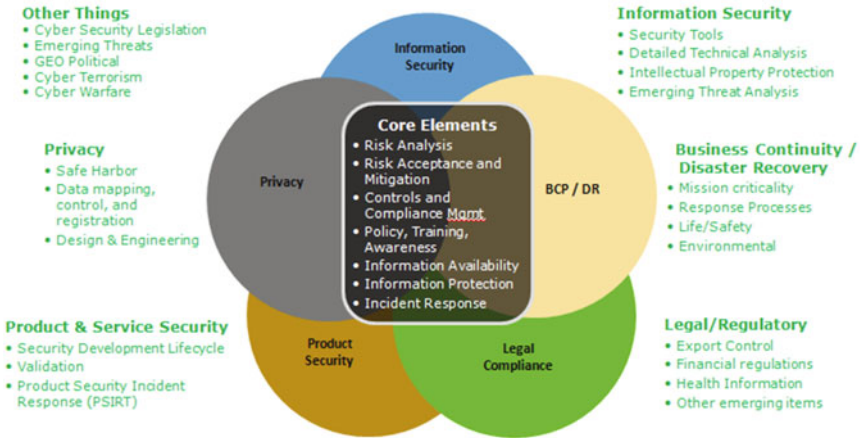


Figure 1-2. Managing the interdependent set of technology-related risks

Traditional security or other control type thinkers would respond to this situation by saying “no” to any technology that introduces new risks. Or perhaps they would allow a new technology but try to heavily restrict it to a narrow segment of the employee population. An example of this over the past few years was the view at some companies that marketers should not engage consumers with social media on the company’s web site because this meant accumulating personal information that increased the risk of noncompliance with privacy regulations. Another example was that some companies didn’t allow employees to use personal devices because they were less secure than managed business PCs.

The reality is that because IT is now integrated into everything that an organization does, security groups cannot simply focus on locking down information assets to minimize risk. Restricting the use of information can constrain or even disable the organization, hindering its ability to act and slowing its response to changing market conditions. A narrow focus on minimizing risk therefore introduces a larger danger: it can threaten a business’s ability to compete in an increasingly fast-moving environment.

THE CHALLENGES OF RISING SECURITY COSTS AND SKILLS SHORTAGES

Growing recognition of the importance of security and privacy, triggered largely by highly publicized breaches, has led to sharply increasing security spending and an accompanying skills shortage. If the current trajectory continues, Gartner Inc. predicts that by 2017 the typical IT organization will spend up to 30 percent of its budget on risk, security, and compliance, and will allocate 10 percent of its people to these security functions. That is triple the levels of 2011 (Gartner 2015b). At the same time, skill shortages may worsen; more than a third of security managers surveyed in 2015 reported significant obstacles in implementing security projects due to inadequate staffing (Morgan 2015). One question is how much of the projected cost increase is due to under-investment in the past, and how much is due to the fact that organizations have invested in technologies that do not adequately reduce risk. To break the cycle, as I'll explain in Chapter 7, we need a new security model and tools that create a demonstrable decrease in the risk curve, with a greater focus on effective prevention and machine learning to reduce cost and manual effort.

Protect to Enable®

To understand how the role of information security needs to change, we need to re-examine our purpose. We need to *Start with Why*, as author Simon Sinek argues convincingly in his book of the same name (Portfolio, 2009). Why does the information security group exist?

As I considered this question back in 2010, and discussed it with other members of the risk and security team that I led at Intel, I realized that we needed to redefine our mission. Like the IT organization as a whole, we exist to enable the business, to help deliver IT capabilities that provide competitive differentiation. Rather than focusing primarily on locking down assets, the mission of the information risk and security group must shift to enabling the business while applying a reasonable level of protection. To put it another way, we provide the protection that enables information to flow through the organization, our partners, and our customers. We also provide the protection for the technology that our organizations create to provide new experiences and opportunities for our customers.

The core competencies of information security groups—such as risk analysis, business continuity, incident response, and security controls—remain equally relevant as the scope of information-related risk expands to new areas, such as technology-enabled products and services, as well as privacy and financial regulations. But rather than saying “no” to new initiatives, we need to figure out how to say “yes” and think creatively about how to manage the risk.

During my time at Intel, the security group's mission evolved toward this goal as we helped define solutions to a variety of technology challenges. For example, my team recognized as early as 2002 that implementing wireless networks within Intel's offices could help make the workforce more productive and increase their job satisfaction by letting them more easily connect using their laptops from meeting rooms, cafeterias, and other locations. At the time, many businesses avoided installing wireless networks within their facilities because of the risk of eavesdropping or because of the cost. We learned pretty quickly that when we restricted wireless LAN deployments or charged departments additional fees to connect, we actually generated more risks. This was because the departments would buy their own access points and operate them in an insecure fashion. We recognized that the benefits of installing wireless LANs across the company outweighed the risks, and we mitigated those risks using security controls such as device authentication and transport encryption. By 2004, that approach had enabled ubiquitous wireless and mobile computing that propelled productivity and actually reduced risks.

A more recent example that many organizations have experienced: for years, Intel didn't allow employees to use personal smartphones for business, due to concerns about privacy and other risks such as data theft. However, we experienced growing demand from employees soon after the launch of the iPhone 3 in 2009. We realized that letting them use these consumer devices to access e-mail and other corporate systems would help boost employee satisfaction and productivity.

By working closely with legal and human resources (HR) groups, we defined security controls and usage policies that enabled us to begin allowing access to corporate e-mail and calendars from employee-owned smartphones in early 2010. The initiative was highly successful, with a massive uptake by employees, overwhelmingly positive feedback, and proven productivity benefits (Evered and Rub 2010, Miller and Varga 2011). The success of the initiative led to its selection for an in-depth Ivey Business School case study (Compeau et al. 2013).

The transformation within the information security group was reflected in changes to our mission statement and top priorities over the years. In 2003, the internal mission statement reflected the traditional focus and scope of information security organizations: the overarching goal was to protect information assets and minimize business disruption.

By 2010 it was clear to me that we needed to simplify our purpose and also broaden the scope. So in 2011, I changed our mission to Protect to Enable to express the idea that our primary goal was to find ways to enable the business while providing the protection necessary to reduce the risk to an acceptable level.

For a few years after this, I thought of information risk and security as a balancing act. I felt that we needed to try to find the right balance between providing open access to technology and information to enable the business and locking down assets. Providing open access allows greater business agility. The business can move more quickly with fewer restrictions. Employees can work more freely, and the faster flow of information allows the company to grow and transform.

But as my responsibilities grew to encompass security and privacy not only for internal systems but also for all aspects of products and services, I realized that a balancing act was the wrong analogy. We should not start from a position of making trade-offs between risks and enablement, or between security and privacy. So I began using a different model that I now feel more accurately represents the challenges of managing information risk: we should take on the harder task of optimizing what is really a multivariate equation of risk dynamics and business objectives in order to create solutions that are "tuned to target," as shown in Figure 1-3.

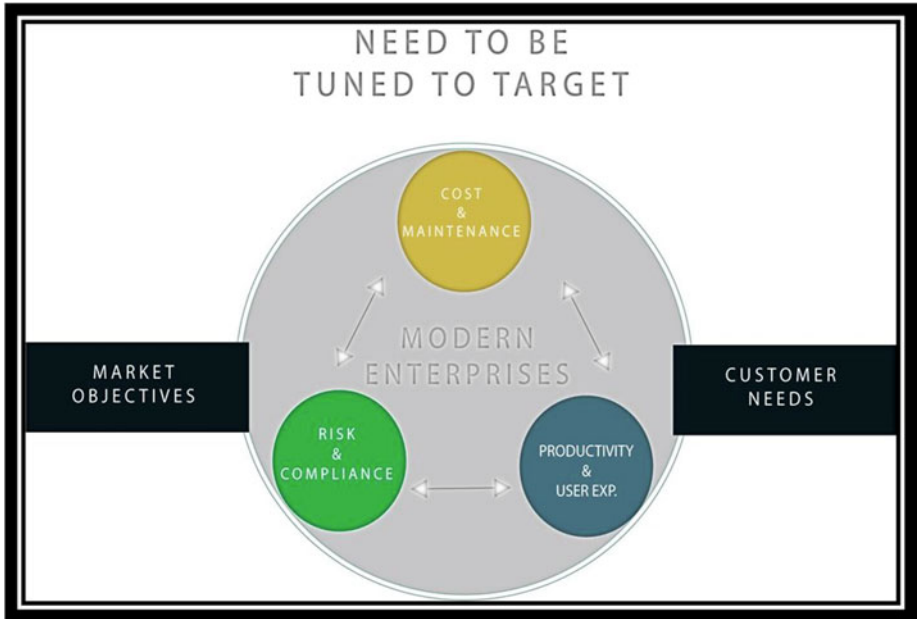


Figure 1-3. *Tuned to target: optimizing the equation to meet business objectives and customer needs*

For each problem and solution, we try to optimize or “tune” five primary variables:

- **Risk and Compliance:** Meeting security, privacy and compliance requirements, based on the organization’s risk tolerance and security and privacy principles.
- **Cost and Maintenance:** The total cost of controls, factoring in deployment and maintenance costs.
- **Productivity and User Experience:** The extent to which controls hinder business velocity by making it harder for users to do their jobs. I call this *control friction*. In addition, if we make it difficult or time-consuming for users to follow security policies or use security tools, they’ll ignore them, thus creating more risks. (See the discussion of the 9 Box of Controls in Chapter 7).
- **Market Objectives:** The company’s goals, such as increased market share.
- **Customer Needs:** Our customer’s privacy and security needs, as well as their overall experience.

Ultimately there may be cases where we cannot fully optimize each item and we need to make trade-offs, but that doesn’t mean we shouldn’t try.

I hope that this model may help information security groups at other organizations think about how these priorities relate to their own businesses. The optimization points for each variable and objective will depend on factors such as the organization's overall culture, technical acumen, and appetite for risk.

Building Trust

I believe that if computing is to continue to improve the world we live in, rather than endanger it, it must be trustworthy. Unfortunately, as I describe in Chapter 9, the privacy and security breaches that have hit the headlines in recent years have weakened the public's trust in technology, according to the Edelman Trust Barometer, a widely used indicator. The rapid implementation of new technologies emerged as a new factor in depressing trust overall. "By a two-to-one margin, respondents in all nations feel the new developments in business are going too fast and there is not adequate testing," the study concluded (Edelman 2015).

To rebuild trust in technology, we need to ensure the data we enter into our systems is both secure and private. At Cylance, we strive to cultivate a work environment where security, privacy, and trust are an integral part of the evolving culture of the company and foundational to the design, development, and delivery of our products and services.

To analyze the context that led to my approach to the risk and security mission, and helped to shape top priorities, I'll explore some of the key changes in the landscape: the rapidly expanding regulatory environment, the emergence of new devices and technologies, and the changing threat landscape.

Keeping the Company Legal: The Regulatory Flood

Until the early 2000s, I didn't see regulatory compliance as a top priority for information security. That's simply because there weren't many regulations that impacted IT, at least in the United States. There were a few exceptions that affected a subset of companies, including Intel, such as controls on certain high-tech exports. And in European countries, there were already regulations that sought to protect personal information. But in general, IT groups didn't have to dedicate much of their time, or budget, to regulatory compliance.

The change in the last decade has been extraordinary. We have seen a flood of new regulations implemented at local, national, and international levels. They affect the storage and protection of information across the entire business, from the use of personal information for HR and marketing purposes, to financial data, to the discovery of almost any type of document or electronic communication in response to lawsuits. And with growing concerns about cyberwarfare, cyberterrorism, and hacktivism, several countries are evaluating additional cybersecurity legislation in an attempt to protect critical infrastructure and make industries more accountable for strengthening security controls.

In most cases, these regulations do not aim to specifically define IT capabilities; however, because information is stored electronically, there are huge implications for IT. The controls defined in the regulations ultimately must be implemented in the organization's systems. These systems include more than just technology: they consist of

people, procedures, devices, and applications. The business risk includes a significant IT-related component, but we must take a holistic view of risk management. Noncompliance can damage a company's brand image, profitability, and stock price—not just through resulting legal problems, but through bad publicity.

Let's take a brief look at some of the key areas and regulations that are having the biggest impact.

Privacy: Protecting Personal Information

For many US companies, the wake-up call was the California data security breach notification law (State Bill 1386), which became effective in 2003. A key aspect of this law requires companies that store personal information to notify the owner of the information in the event of a known or suspected security breach. Businesses could reduce their exposure, as well as the risk to individuals, by encrypting personal data.

After this, other states quickly followed suit, implementing regulations that generally follow the basic tenets of California's original law: companies must promptly disclose a data breach to customers, usually in writing.

In addition, federal laws, such as the Health Insurance Portability and Accountability Act (HIPAA), have addressed specific categories of personal information. Further regulations have been added in other countries, too, such as the updated data-protection privacy laws implemented in Europe (European Commission 2011, 2012).

The implications of these local and national regulations extend beyond geographical boundaries. As companies do more business online, they're increasingly likely to acquire and store information about customers from other countries, and find that they also need to comply with regulations around the world. Those regulations may change, with implications for businesses in multiple countries. In late 2015, for example, Europe's highest court struck down the so-called "safe harbor" agreement that had allowed companies to move information about consumers between the European Union and the United States. The replacement EU-US Privacy Shield, agreed after three months of negotiations, aimed to address European privacy concerns with written guarantees that US intelligence agencies would not have indiscriminate access to Europeans' personal data stored in the US (Scott 2016).

The issue can become even more complex when businesses outsource application development or HR functions to providers located in yet another country. Now, software developers in India may be building and operating the systems that collect information about Europeans for US companies, making it even more difficult for businesses to navigate compliance with all relevant privacy regulations.

Personalization vs. Privacy

Privacy concerns are set to become even more important over time, as businesses increasingly seek to create online experiences tailored to the needs of individual users. The more a business knows about each individual, the more it can personalize services and offer targeted advertising based on income and preferences.

Many users also like personalized services. If a web site “remembers” them, they don’t need to enter the same information each time they visit the site, and they’re more likely to see content and offers relevant to their needs. In fact, companies may be at a disadvantage if they don’t personalize services because users may prefer a web site from a competitor that offers a more streamlined experience.

However, there’s an inevitable conflict between personalization and privacy. The personalization trend is fueling the growth of an industry focused on collecting, analyzing, and reselling information about individuals. This industry existed long before the Web; personal information has been used in mass-mailing campaigns for decades. However, the Web is both increasing demand for this information while providing new ways to collect it. Companies now have opportunities to collect information from multiple online sources, correlate and analyze this information, and then sell it to others. And of course, consumers’ fears that information will be lost or misused have increased accordingly.

For businesses, however, offering personalized services also can increase compliance concerns. As companies store more personal information, they are responsible for safeguarding that information and are liable for any loss or compromise. In many parts of the world, companies are also required to explain why they are collecting personal data, how they are protecting it, and how long they will keep it.

We can expect continuing tension due to conflicting desires for personalization and privacy—and more regulation as a result. Governments clearly believe that businesses cannot be relied upon to regulate themselves, so they will continue to add regulations designed to protect the privacy of individuals. Meanwhile, businesses will seek new ways to collect more information so that they can further personalize services. Developing compliance strategies and guidelines becomes even more pressing.

Financial Regulations

Financial regulation surfaced as a top priority in the United States with the Sarbanes-Oxley Act (SOX), which emerged from the public outrage over corporate and financial accounting scandals at companies such as Enron and WorldCom. These scandals cost investors billions of dollars and damaged public confidence. To help avoid similar catastrophes in the future, SOX imposed financial tracking requirements designed to ensure that a company’s financial reporting is accurate and that there hasn’t been fraud or manipulation. Once enacted, SOX required publicly held companies to meet specific financial reporting requirements by the end of 2004.

Although the Sarbanes-Oxley Act doesn’t mandate specific technology controls, it has major implications for IT. Ensuring financial integrity requires controls to be implemented within everyday financial processes. In practice, this means they must be enforced within the IT applications and infrastructure that support those processes. Purchases above specific thresholds may require approval from the finance group; the underlying applications have to support this workflow, and to be sure the applications function correctly, businesses need to establish the integrity of the underlying computer infrastructure. Compliance with financial regulations therefore creates a series of IT requirements, from making sure that applications provide the right functionality to implementing access controls and updating software.

E-Discovery

Regulations governing the discovery of information for litigation purposes officially extended their reach into the electronic realm in 2006. That's when the US Supreme Court's amendments to the Federal Rules of Civil Procedure explicitly created the requirement for e-discovery—the requirement to archive and retrieve electronic records such as e-mail and instant messages.

This created an immediate need not just to archive information, but to automate its retrieval. This is because records must be produced in a timely way, and manual retrieval would take too long and be prohibitively expensive. The business risks of noncompliance are considerable: unlike many countries, US practice allows for potentially massive information disclosure obligations in litigation. Companies that fail to meet e-discovery requirements may experience repercussions that include legal sanctions. The implications are correspondingly onerous. Lawsuits may draw on information that is several years old, so businesses must have the capability to quickly search and access archived information as well as current data. E-discovery is further complicated by the growth of cloud computing models such as software as a service (SaaS). As organizations outsource more business processes and data to cloud service suppliers, they need to ensure that their suppliers comply with their e-discovery needs.

Expanding Scope of Regulation

The regulatory universe continues to expand, with the likelihood of more regulations that explicitly address IT, as new technologies emerge and governments try to control its use and inevitable misuse. In the US, lawmakers have proposed legislation to increase the security and privacy of connected cars, following a widely publicized demonstration in which researchers hacked into a Jeep and took over its controls. The Food and Drug Administration (FDA) has published cybersecurity guidelines describing requirements for manufacturers of Internet-connected medical devices (FDA 2016).

The attempts by various governments to gain access to technology for the purposes of combating terrorism have generated considerable impact and controversy. In China, a new anti-terrorism law requires that technology companies hand over technical information and help with decryption when the police or state security agents demand it for investigating or preventing terrorist cases (Buckley 2015). In the US, even greater controversy was generated by the US Government's attempts to force Apple Computer to create “back doors” that make it easier to access information on iPhones used by terrorists or criminals. In India, after terrorists used unsecured Wi-Fi access points to communicate information about their attacks, the government created a legal requirement that any access point must be secured (Government of India Department of Telecommunications 2009).

In other countries, businesses that operate unsecured Wi-Fi access points (a common way to provide Internet access for visitors) may find themselves facing other legal problems. For example, unscrupulous individuals may tap into the network to access web sites for purposes such as illegally downloading music or pornography. Access appears to originate from the company hosting the access point, which may then find itself on the receiving end of correspondence or raids from the music industry or government agencies.

The Rapid Proliferation of Information, Devices, and Things

The computing environment is growing as rapidly as the regulatory environment. The sheer volume of information is exploding, and it is being stored across a rapidly growing array of devices. The Internet of Things will drive yet another exponential increase: Gartner, Inc. estimates that during 2016, 5.5 million new “things” will be connected every day, and Cisco expects 50 billion connected devices by 2020. In the not too distant future, almost any device with a power supply may have an IP address and will be capable of communicating—and being attacked—over the Internet.

Recent headlines have highlighted the growing threat activity focused on IoT, as I’ll discuss further in Chapter 7. Researchers hacked into a Jeep via its Internet-connected entertainment system and remotely controlled the vehicle’s functions (Greenberg 2015); other researchers showed that thousands of medical devices in hospitals are vulnerable to attack.

At the same time, the boundaries between work and personal technology have in some cases completely dissolved. Whether businesses officially allow it or not, employees are increasingly using their personal devices for work by sending e-mails from and storing information on their personal smartphones and computers. Furthermore, people may forward e-mail from business accounts to personal accounts created on external systems, without considering that when they signed up for the personal account, they agreed to a license that allows the external provider to scrutinize their e-mails.

The use of personal technology such as smartphones can considerably enhance business productivity because employees can now communicate from anywhere at any time. However, this also creates a more complex, fragmented environment with more potential points of attack. Information is now exposed on millions of new devices and disparate external networks, many of which do not have the same type of security controls as corporate PCs, and all of which are outside corporate network firewalls. Not surprisingly, mobile malware has become a major industry, and is still growing: one survey found more than 1,200 known families of Android malware in 2014, more than double the number found the previous year (Millman 2015).

The boundaries between work and personal lives are dissolving in other ways, too. Employees store more information on the Internet—on business and consumer social media sites, for example—than ever before. These sites are powerful tools for communicating with audiences outside the corporate firewall.

However, just as there’s an industry gathering and analyzing personal information for marketing purposes, information on the Web can be used for competitive intelligence or for less legitimate purposes. Users store snippets of information in multiple places on the Web. Although each of these snippets may not provide much information, when pieced together they can provide new intelligence not just about the individual, but also about the organizations to which the person belongs. Each item is like a single pixel in a digital picture. Alone, it doesn’t convey much information; but step back, aggregating information from a wider range of sources, and those pixels combine to form a portrait. In the same way, pieces of information strewn across a variety of unrelated web sites—the name of a department, workmates, pet names that might be used as passwords—can be linked together to create a picture of an individual and used for malicious purposes.

The Changing Threat Landscape

The threat landscape is evolving rapidly, with an increase in highly organized and well-funded groups capable of executing sustained attacks to achieve long-term goals, including cyberespionage, cyberterrorism, and cyberwarfare. These attackers, generally known as *advanced persistent threats* (APTs), were originally thought to focus mainly on governments but more recently have also been shown to target private-sector organizations, with the goal of stealing intellectual property or simply causing damage. APTs include nation-state organizations, “hactivist” groups attempting to publicize or further their cause, and organized crime. Hacktivists who said they were targeting oppressive regimes claimed responsibility for an attack that disabled more than 30,000 computers at the world’s biggest oil producer, Saudi Aramco. The FBI blamed North Korea for a crippling attack on Sony Pictures (Schmidt et al. 2015). In 2014, the US Justice Department indicted five Chinese military hackers for stealing trade secrets and other information from US companies in the nuclear power, metals, and solar industries (Department of Justice 2014); in 2016, the US charged seven hackers linked to the Iranian government with hacking US banks and dam operations (Nakashima and Zapotosky 2016).

The steady rise of organized cybercrime online is entirely logical. As the exchange of money and information has moved online, organized crime has followed, focusing on theft of valuable assets such as intellectual property. This has spawned a mature malware industry that increasingly resembles the legitimate software industry, complete with a broad set of services, guarantees, and price competition among suppliers. *Ransomware*, which encrypts a victim’s data until a ransom is paid, is a recent trend.

Stealthy Malware

This evolving set of threat agents is using new, more sophisticated tools and methods to mount attacks. Once upon a time, attackers were amateurish and often driven by personal motives such as the prestige of bringing down a big company’s network. Accordingly, the arrival of malware on a user’s machine was easy to detect: the malware announced itself with icons or messages, and the system often became unusable.

Now the trend is toward malware that is stealthy and uses sophisticated techniques to avoid detection. Attackers plant malware that lies undetected over a long period while it captures information. Another common technique is to quietly spread malware by injecting malicious code into an unsuspecting company’s web site; users who visit the site then unknowingly download the code onto their systems.

Accompanying this is a shift from spam mass e-mails to carefully crafted *spearphishing* attacks aimed at individuals or specific groups. These typically use social engineering techniques, such as providing enough contextual or personal information in an e-mail to tempt people to download malware or click on a link to an infected web site created specifically for that purpose. Though more expensive to mount, spearphishing attacks can be enormously profitable to cybercriminals; an analysis by a supplier of anti-phishing solutions found that they were the primary initial attack method used by APTs in 2015; 22% of attacks were motivated by financial fraud or other crimes (PhishLabs 2016). We can expect these stealthy and targeted attacks to continue, with new methods emerging as necessary to circumvent defenses.

Nine Irrefutable Laws of Information Risk

Over the years, I've identified a number of "laws" that encapsulate some of the lessons I've learned, and that seem to remain true despite the continually changing environment. I call these the Nine Irrefutable Laws of Information Risk (with acknowledgements to Culp (2000), Venables (2008), Lindstrom (2008), and other sources):

- **Law #1: *Information wants to be free.*** People want to talk, post, and share information—and they increase risk by doing so. Some examples:

A senior executive at a major technology company updated his profile on a business social networking site. In doing so, he inadvertently pre-announced a shift in his employer's strategy—a mistake that was promptly and gleefully picked up by the press.

An employee found a novel way to fix a piece of equipment more quickly and, to help others across the company, decided to videotape the procedure. Because video files are so large, it didn't make sense to e-mail the video, so the employee posted it online. Unfortunately, by doing so, he exposed confidential information.

At one time or another, many people have experienced this disconcerting event: when composing a message, the e-mail software helpfully autofills the address field, but it selects the wrong name from the address book. You hit Send without realizing the error, thus dispatching a company-confidential message to someone outside the organization.

It's worth noting that that this rule is not new. Information has always wanted to be free: think of the World War II slogan "loose lips sink ships." People communicate, and sometimes they share more information than they should. It's just the methods that have changed, and the fact that, with the Internet, a carelessly mentioned detail is instantly available to anyone across the globe.

- **Law #2: *Code wants to be wrong.*** We will never have 100 percent error-free software. In fact, the more widely used the software, the more malicious individuals will hunt for vulnerabilities in the code. They have found and exploited errors in the world's most widely used web sites, productivity applications, and enterprise business software.
- **Law #3: *Services want to be on.*** On any computer, some background processes always need to be running, and these can be exploited by attackers. These could even be security software processes used for everyday activities like keeping systems up-to-date with software patches or monitoring for malware.

- **Law #4: *Users want to click.*** People naturally tend to click when they see links, buttons, or prompts. Malware creators know this, and they take advantage of it. In fact, the entire phishing industry is based on the assumption that users will click on enticing e-mails, web sites, or pop-up ads, triggering the download of malicious code to their systems. The evolution of highly targeted attacks such as spearphishing has taken this to a new level, as when e-mails purporting to be letters discussing legal action from a circuit court were sent to senior executives at a number of companies.
- **Law #5: *Even a security feature can be used for harm.*** Security tools can be exploited by attackers, just like other software. This means that laws 2, 3, and 4 are true for security capabilities, too. Networking equipment supplier Juniper Networks discovered that its firewall software contained “unauthorized code” that surreptitiously decrypted virtual private network traffic (Goodin 2015). Security researchers have uncovered vulnerabilities that can be exploited by attackers in products from well-known security suppliers, including Kaspersky Labs and FireEye (Ashford 2015).
- **Law #6: *The efficacy of a control deteriorates with time.*** Once put in place, security controls tend to remain static, but the environment in which they operate is dynamic. Organizations tend to “set and forget”: to install security controls and then fail to update them with security patches or to properly maintain access lists. As attackers find new ways to circumvent or compromise the controls, their effectiveness progressively degrades. As Rob Joyce, who heads the National Security Agency’s elite hacking unit, put it, an organization with static defenses will drift to the back of the herd, where it is easily picked off by a predator (see Chapter 6).
- **Law #7: *Code needs to execute.*** All software, good or bad, needs to execute in order to perform its intended function. Malware is created with malicious intent, but until it executes, it is dormant and can do no harm. Exploits can therefore be intercepted and stopped by security tools that inspect code before execution, identify good from bad, and prevent bad code from executing.
- **Law #8: *Controls create friction.*** Security controls can slow users and business processes by impacting system performance or forcing them to use cumbersome processes. High-friction controls therefore impose a “drag coefficient” on business velocity. Users react to a high degree of control friction by circumventing the controls whenever possible; as a result, the controls can actually introduce new risks as business users go around IT to get their jobs done. Control friction is an important consideration when designing security architectures (see the discussion on the 9 Box of Controls in Chapter 7)

- *Law #9: As our digital opportunities grow, so does our obligation to do the right thing.* As technology becomes embedded into the fabric of our lives, exploits that take advantage of technology vulnerabilities may increasingly impact the well-being of almost everyone in society. So it is particularly important that we apply the right ethical values to shape the way we design, develop, and implement these technologies. As I explain in Chapter 9, security and privacy should now be considered a corporate social responsibility.

A New Approach to Managing Risk

Given the ever-broadening role of technology and the resulting information-related business risk, we need a new approach to information security built on the concept of protecting to enable. This approach should

- *Incorporate privacy and regulatory compliance by design, taking a holistic view of information risk.* Also, because all companies are moving toward using technology not only for internal operations but also in products and services, the information security organization must work closely with other business groups to understand and manage risk.
- *Recognize that people and information, not the enterprise network boundary, are the security perimeter.* Information is no longer restricted to tightly managed systems within data centers; it now also resides outside the firewall, on users' personal devices, and on the Internet. Managing risk therefore requires a range of new tools, including user awareness and effective security controls for personal devices.
- *Be dynamic and flexible enough to quickly adapt to new technologies and threats.* A static security model will inevitably be overtaken by the dynamic nature of threats. We need security architectures that can rapidly learn and adapt to new devices and evolving threats, with a high degree of automation.

Above all, we need to accomplish a shift in thinking, adjusting our primary focus to enabling the business, and then thinking creatively about how we can do so while managing the risk. Our roles will only increase in importance as technology becomes even more prevalent. Our ability to protect information security and privacy will be essential to building the trust that enables our organizations to take advantage of new digital opportunities.

CHAPTER 2



The Misperception of Risk

The moment we want to believe something, we suddenly see all the arguments for it, and become blind to the arguments against it.

—George Bernard Shaw

One hundred years ago, the “unsinkable” *Titanic* foundered after striking an iceberg off the coast of Newfoundland. More than 1,500 people died in what became one of the deadliest maritime accidents ever. Several factors contributed to this massive death toll, but perhaps the most critical was that there simply weren’t enough lifeboats. The ship carried 2,224 people, but fewer than half of them could squeeze into the boats.

As we know, passengers who didn’t get a spot in one of those lifeboats quickly died in the freezing waters of the North Atlantic. What’s less well known is that the *Titanic*’s supply of lifeboats was in full compliance with the British marine regulations in force at time. The law required the ship to carry 16 lifeboats; the *Titanic* actually had 20 lifeboats.

The ship’s owners did a good job of providing enough boats to address the regulatory risk of noncompliance. Unfortunately, meeting regulatory requirements did little to prevent the tragic loss of life.

This is a case of *misperception of risk*. The owners focused on mitigating the regulatory risk, apparently blind to the much larger risk of disaster. They framed the lifeboat issue as a compliance item that needed to be addressed so that the ship could start carrying passengers and generating revenue. One could argue that if they had stepped back and considered the potential consequences for the customers rather than the company’s short-term priorities, history might have unfolded differently. Reports suggest that the *Titanic* had enough capacity to easily add enough lifeboats for everyone on board, had the owners chosen to do so.

What does this example have to do with managing information risk? We encounter misperceptions every day within the realm of enterprise risk and security. Every organization has a greater responsibility than simply complying with regulations. We have to think about whom is ultimately at risk: the company or the customer? Furthermore, as I’ll show in this chapter, everyone in the organization has their own priorities and their own subjective view of risk. Unless we mitigate these misperceptions, they can have disastrous consequences. As a result, I believe that the misperception of risk is the most significant vulnerability facing enterprises today.

The Subjectivity of Risk Perception

As security professionals, we tend to think about objective ways to estimate risk—to assess the likelihood and extent of harm that can occur due to specific threats and vulnerabilities.

But in reality, the way people perceive risk has a strong subjective component. Economic and psychological factors greatly affect how each of us perceives the likelihood and potential impact of harm from specific actions or situations. Within an organization, each individual's perception of risk varies depending on his or her job role, goals, background, and peer group. This means managers, security professionals, and end users all may have a different view of the risk associated with a specific technology or action.

Misperceiving risk has serious consequences because our actions are shaped by our perception of risk. An employee may think that posting personal and work-related information on a social media site is relatively harmless. However, hackers might use this publicly available information in phishing e-mails to gain access to enterprise systems via the employee's computer, ultimately resulting in detrimental security breaches.

End users are not the only members of the organization who can misperceive risk. Everyone is capable of misperceiving risk, including risk and security professionals. As I'll explain later in this chapter, misperceptions occur at the group level as well as the individual level. Members of a group may share the same bias in their perception of risk and benefit.

The decisions that result from these misperceptions can weaken the entire organization's security posture. If an organization underestimates a risk, it will underspend on controls to mitigate that risk, increasing the likelihood and potential impact of major problems such as data breaches. On the other hand, if the organization overestimates a risk, it will allocate a disproportionately large share of its security resources to the risk, leaving other parts of the risk landscape underprotected.

In this chapter, I'll discuss how and why different people within an organization misperceive risk, whether they are acting as information technology users, security professionals, or managerial decision makers. To explore these misperceptions, I've drawn on research across the broader field of risk psychology, notably *The Psychology of Risk*, a book by Professor Dame Glynis Breakwell, Vice Chancellor of the University of Bath (Cambridge University Press 2007). I'll examine how these ideas about risk perception apply to information risk and security. I'll explain some of the consequences of those misperceptions, and I'll discuss some of the ways an organization can address them.

How Employees Misperceive Risk

Research shows that if we like an activity, we tend to judge its benefits to be high and its risk to be low (Slovic 2010). Conversely, if we dislike the activity, we judge it as low-benefit and high-risk. Because of this, the perception of risk by individuals and groups within an organization tends to be biased by their preferences, roles, and objectives. Everyone is trying to achieve their individual or group goals within the organization, so they tend to see activities and technologies that support those goals as beneficial, and therefore they tend to underestimate the risk.

So if employees like social media, their attraction to the technology skews their perception of benefit and risk. Because they judge the benefit to be high and the risk to be low, they feel comfortable posting information such as their job title, location, and even the projects they're working on. They may even allow sites to capture their location, using the global positioning system in their cell phone, and display the location in real time.

Unfortunately, these employees may not think about how a malicious individual could use the information. Today, as we've seen, an individual's use of technology can harm not only the individual but the entire organization. Attackers exploit publicly available personal information to craft spearphishing e-mails that are particularly convincing because they appear to demonstrate a relationship with the recipient, making the employee more likely to click on a link that downloads malware to the system. From there, the attack spreads to the rest of the corporate network. In addition, information posted by individuals is now routinely aggregated, analyzed to identify patterns, and sold, often to a company's competitors.

The risk and security team may also misperceive the risk of social media, but in the opposite direction: they overestimate the risk and underestimate the benefits. They may not like social media because it creates vulnerabilities, and their perception then drives them to focus on minimizing the risk by trying to block the use of the technology.

Other psychological factors also come into play in shaping end users' risk perception. People in general tend to believe they are personally less likely than others to experience negative events and more likely to experience positive events, leading to a sense of personal invulnerability (Breakwell 2007). In addition, users also are more likely to behave in risky ways if their colleagues do so. "It's conformity: being seen to be doing what everybody else is doing," Breakwell says (pers. comm.). Many social media sites encourage this conformist tendency; if all your friends are using a social media site, you're likely to join the site too because it enables you to see what they are doing and share information with them more easily.

The likelihood that individuals will behave in ways risky to the organization also increases when their individual interests don't align with the company's. This divergence is most likely when employees are discontented, resentful, demoralized, or simply don't trust IT or the broader organization.

In economic theory, the problem resulting from this lack of alignment is known as a *moral hazard*: a situation in which someone behaves differently from the way they would if they were fully exposed to the risk. A useful moral hazard analogy is renting a car with full insurance coverage. People are likely to be less careful with the rental car than they would be with their own car if they're not responsible for the consequences. The attitude is "if it's not mine, it doesn't matter."

In the realm of enterprise IT, moral hazards may be a bigger concern than many appreciate. A Cisco survey (2011a) found that 61 percent of employees felt they were not responsible for protecting information and devices, believing instead that their IT groups or IT service providers were accountable. Ominously, 70 percent of these surveyed employees said they frequently ignored IT policies.

One indicator of the extent of moral hazard within an organization may be how employees treat company-provided laptops. Higher-than-average loss or damage rates might suggest employees don't care about the laptops and may be an indication they don't care about other corporate assets either. As I'll discuss in Chapter 5, I believe allowing reasonable personal use of laptops can help reduce the risk of moral hazard because it aligns personal interests with those of the organization.

More broadly, organizations can address the moral hazard issue by taking steps to align the goals and concerns of everyone involved: end users, information security professionals, and executives. This returns us to the theme of the book: as information security professionals, our mission is to Protect to Enable. This mission aligns our security goals with those of the business. It helps maintain the perception of shared values. Research suggests that people with whom we share values are deemed more trustworthy (Breakwell 2007, 143). If employees trust us, they are more likely to believe our warnings and act on our recommendations.

The Lure of the Shiny Bauble

One further point to remember is that everyone in the organization, regardless of the job role, is an end user. Therefore, we can all fall prey to the same tendencies. Our attraction to new consumer technologies can also cause us to ignore the risks. I call this magpie-like attraction the *lure of the shiny bauble*; mesmerized by the appeal of gleaming new technologies, we downplay or even fail to notice the risks lurking in the shadows.

How Security Professionals Misperceive Risk

While end users tend to underestimate the risks of a desirable activity or technology, security professionals sometimes display the opposite tendency. We focus obsessively on the information risk associated with a specific threat or vulnerability. In doing so, we completely miss bigger risks.

This phenomenon is known as *target fixation*, a term originally coined to describe a situation in which fighter-bomber pilots focus so intently on a target during a strafing or bombing run that they fail to notice the bigger risk to themselves and crash into the target as a result (Colgan 2010, 44). As information security professionals, we can develop a similar fixation. We focus so intently on one risk that our awareness of larger hazards is diminished. This target fixation can also occur in other groups with “control” functions within the organization, such as internal audit, legal compliance, and corporate risk management.

Here is an example from my own experience at Intel. Several years ago, we discovered that malware had been introduced onto our network from an employee’s personal computer. We became so focused on this source of danger that we eliminated all personal devices from our network. We further fueled our target fixation by labelling these devices “non-Intel managed systems (NIMS),” a term that reflected the frustration over our lack of control. I vowed we would never again allow network access from devices that we didn’t fully control.

However, by becoming fixated on a single threat, we may have created some larger risks and additional costs. For example, we needed to issue contract employees with corporate PCs, each of which allowed broader access to the Intel environment. If we had instead focused on how we could provide limited access to the environment from “untrusted” devices, we might have managed the risk with lower total cost and obtained a head start in developing a key aspect of a more flexible security strategy, as I’ll describe in Chapter 7.

It's worth noting that security professionals can also suffer from a problem that's almost the opposite of target fixation: *alert fatigue*. At many organizations, security groups experience a constant deluge of thousands of alerts emanating from security tools across the enterprise. With so much noise, it's easy to become overwhelmed and miss important threats.

As security professionals, we also may misperceive risk due to the tendency to “set and forget” security controls. This common security loophole is described in the sixth Irrefutable Law of Information Security in Chapter 1, which states that the efficacy of a control deteriorates with time. Once in place, controls tend to remain static, but the threats they are intended to mitigate continue to evolve and change, sometimes in very dynamic ways. Controls that are initially very effective can become inadequate over time. Ultimately, an adverse event may occur and may even have disastrous consequences.

Think about the history of major oil tanker spills. For years, regulations allowed tankers to be built with a single hull, instead of a double (inner and outer) hull to provide additional protection in the event of a leak. Meanwhile, tankers grew steadily larger because bigger ships could transport oil more efficiently than smaller ones. It wasn't until the *Exxon Valdez* ran aground, puncturing its hull and creating a giant oil leak that contaminated huge stretches of Alaska's coast, that authorities were spurred to create new regulations requiring double hulls in oil tankers (EPA 2011).

Within enterprise IT, a typical “set and forget” error is the failure to keep controls up-to-date, particularly if the controls are designed to mitigate a relatively low risk. A case in point: *distributed denial-of-service (DDoS)* threats were a big concern more than a decade ago, due to widely publicized attacks by worms such as Code Red, Nimda, and SQL Slammer. These attacks disabled corporate web sites or flooded internal networks by overloading them with requests. To mitigate the availability risk, many organizations invested in defenses against DDoS attacks.

Over time, however, DDoS attacks became less frequent, and organizations were assailed by newer threats. With limited resources, information security groups focused on mitigating these new threats rather than continuing to build defenses against DDoS attacks. At the same time, though, businesses were increasing their online presence. Web sites evolved from being used primarily for advertising and displaying static corporate information to managing business-critical data and applications. Some organizations began conducting all their business online. Even traditional brick-and-mortar businesses moved customer support, order management, and other critical business processes onto the Web. The larger online presence multiplied the potential impact of a successful attack. As a result, when DDoS attacks from a variety of groups resurfaced in the past few years, they created even greater disruption to business operations as well as damage to corporate brands.

Another example: over the past few years, many organizations have become much more diligent about scrubbing data from the hard drives of old computers before disposing of them or reselling them. But they failed to follow similar precautions for other business devices that have evolved to include hard drives.

Nearly every digital copier contains a drive storing an image of each document copied, scanned, or e-mailed by the machine. When CBS News reporters visited a company that specialized in reselling used copiers, they found businesses and agencies had discarded machines containing lists of wanted sex offenders, drug raid targets, pay stubs with Social Security numbers, and check images. One copier's hard drive even contained 300 pages of individual medical records, including a cancer diagnosis, which is a potential breach of federal privacy law (Keteyian 2010).

Security and Privacy

As I explained earlier in the book, security professionals, and the broader security industry, can sometimes be tone-deaf when it comes to privacy concerns. In their zeal to collect data for security purposes, they may create risks that the data could be used in a way that may violate people's privacy, or at least their expectations of privacy.

The challenge of balancing privacy and security concerns in the enterprise bears many similarities to the broader issue of balancing security and privacy in society, an area that has been extensively explored by privacy legal expert Daniel J. Solove. As he explains in the book *Nothing to Hide: The False Tradeoff between Privacy and Security* (Solove 2011), the debate between security and privacy has often been incorrectly framed to imply that we must choose between one value and the other. "Security and privacy often clash, but there need not be a zero-sum game," he writes. "There is a way to reconcile privacy and security: by placing security programs under oversight, limiting future uses of personal data, and ensuring that programs are carried out in a balanced and controlled manner."

Solove's conclusion is equally applicable to information security. Many in the security profession think that security equals privacy. That is not the case. We need good security to achieve privacy, but the two are not synonymous. Some security industry solutions conduct broad-based bulk data collection, monitoring the activity of users and their machines, and siphoning the data to the cloud. That data is then used to build profiles and, combined with other information, to enable the solution to scan for potentially anomalous activity. Considered in isolation, some machine data has few, if any, privacy implications. However, the collection of thousands of pieces of information about what the machine is doing, when and how, while someone is using it and even when not, creates a detailed digital profile of an individual and his or her behavior. That profile is collected, stored in perpetuity and analyzed. As our lives become more digitized, the richness of that profile will grow and evolve. We need to step back and ask ourselves whether this is really necessary for our protection.

As I've discussed elsewhere in this book, I believe that security and privacy programs should be managed together as elements of an overall enterprise information risk management strategy. Security and privacy are like the two halves of a zipper: when meshed together, they create a strong bond, protecting the enterprise and the individual against risk. Managing them as isolated silos is more likely to result in dangerous misperceptions of risk.

MISMATCHING CONTROLS TO THREATS

Businesses sometimes devote considerable time and resources to implement security controls that are completely irrelevant to the threats the companies are trying to mitigate. These mismatches reveal a lack of understanding of the security technology and the threat. The controls may further add to the risk by providing a false sense of security. In reality, deploying the wrong control is like carrying a lightning rod to protect oneself from getting wet in a storm.

Typical mismatches include

- Using firewalls to prevent data theft from applications that are allowed to operate through the firewall
- Using standard antivirus tools that are effective only against previously identified threats, to protect against zero-day attacks
- Using controls at the operating-system level to detect application-layer attacks

This mismatch does not mean that these controls are worthless. It simply means that if our goal is to deal with a specific threat, we must understand both the attacks and the controls well enough to identify which controls are applicable, and where it is necessary to add other controls. For example, if a firewall cannot prevent attacks against an application, we might deploy an additional control behind the firewall.

How Decision Makers Misperceive Risk

Managers make decisions based on information from technical specialists and other experts. Therefore, the decisions that managers make are only as good as the information they receive. Decision makers can misperceive risk when their decisions are based on biased or incomplete information.

Bias can influence these decisions every day. If people are trying to sell a particular proposal or point of view to their manager, what are they likely to do? They tend to select data supporting their arguments and often ignore data contradicting those arguments.

The danger of misperception is particularly acute when decision makers rely on a narrow range of sources with similar viewpoints. Without obtaining a diversity of viewpoints, managers don't get a full picture of the risk. Like-minded individuals tend to agree with each other, as you might expect. When a group is composed solely of people with similar backgrounds and viewpoints, it may be particularly prone to *group polarization* (Breakwell 2007, 99) and the group's decision may be more extreme than the mean of their individual views. This problem may be especially acute when the people involved share the same mental model of the world, as is likely to be the case when the group consists only of specialists from the same organization.

An even broader concern is how a focus on business goals can drive people to make unethical decisions. When these decisions are made by managers at the organizational level rather than at the individual level, the impact is compounded by the potential for widespread disaster.

After the *Challenger* space shuttle exploded in 1986, extensive post-crash analysis revealed the tragedy was caused because an O-ring on one of the shuttle's booster rockets failed to seal due to the low ambient temperature at launch time.

However, it subsequently emerged that engineers had warned of the potential danger before the launch. Engineers from NASA contractor Morton Thiokol recommended the shuttle not be launched at low temperatures after analyzing data that indicated a link between low temperatures and O-ring problems. After NASA responded negatively to the engineers' recommendation, Morton Thiokol's general manager reportedly decided to treat the question of whether to launch as a "management decision." Against the objections of their own engineers, Morton Thiokol's managers then recommended NASA go ahead and launch, and NASA quickly accepted this recommendation (Bazerman and Tenbrunsel 2011, 13–16).

For Morton Thiokol's managers, the desire to meet the business goal of pleasing the company's customer, NASA, apparently caused the ethical dimensions of the problem to fade from consideration—with terrible consequences.

According to Tenbrunsel, this *ethical fading* is not uncommon. The way a decision is framed can limit our perspective. If the decision is framed purely in terms of meeting business goals, ethical considerations may fade from view. In fact, we may become blind to the fact that we are confronting an ethical problem at all (Joffe-Walt and Spiegel 2012).

Another infamous ethical lapse involved the Ford Pinto, whose gas tank exploded in a number of rear-end collisions, resulting in fatalities. As Bazerman and Tenbrunsel describe (2011, 69–71), Ford discovered the dangers in preproduction testing. However, facing intense business competition, the company decided to go ahead with manufacturing anyway. The decision was based on a cost-benefit analysis. Ford apparently considered the choice as a business decision rather than an ethical decision and determined it would be cheaper to pay off lawsuits than make the repair. The impact of dehumanizing this risk decision was disastrous.

In the past, many information technology risk decisions have often been considered only in terms of their potential business impact. As information technology is integrated into more and more products, decisions about information risk will increasingly affect the lives of millions of people, making it essential to consider the ethical as well as the business dimensions of information risks. It becomes even more important that we, as CISOs, keep ethical considerations to the forefront. What is the potential impact of a security breach when a car's sensors and control systems can be accessed via the Internet? Or when medical life-support equipment can be remotely controlled using wireless links?

How to Mitigate the Misperception of Risk

It should be apparent by now that the tendency to misperceive risk is universal. We need to find ways to help compensate for this misperception, given that it is our job to manage risk. As security professionals and managers, how can we mitigate the misperception of risk?

We can start by ensuring that we include a diversity of viewpoints when making risk management decisions. Whenever possible, we should involve a broad cross-section of individuals representing groups across the organization. This diversity helps compensate for individual biases.

However, assembling the right mix of people is only the first step in building a more complete picture of risk. As information security and risk professionals, we need to ensure that the discussion brings up new perspectives and views. We must ask penetrating questions designed to bring alternative viewpoints to the surface. I think of these as *high-contrast questions* because the process is analogous to adjusting the contrast or colors of a photograph to highlight key elements of possible interest. This questioning counteracts the inevitable bias due to target fixation. We can also help counter target fixation by simply recognizing it exists, and then consciously trying to see the problem from someone else's viewpoint.

In addition, we need to continually seek out the minority report, the view that is contrary to perceived wisdom. If the majority is telling us to turn right, are we missing something important that we'd find out by turning left? In a striking example, Israel's Directorate of Military Intelligence considered this viewpoint so important that it created a devil's advocate office as an institutional safeguard against group-think. The office's job was to criticize analysis coming from the Directorate's other divisions and write papers countering the analysis. In order to explore alternative assumptions and worst-case scenarios, it examined possible radical security developments scenarios, including those that the defense establishment considered unlikely. Notably, the office was staffed by experienced, highly regarded people known for their creative thinking, and its reports went directly to all major decision-makers (Kuperwasser 2007).

Uncovering New Perspectives During Risk Assessments

Risk assessment models can be valuable tools for helping to evaluate risks and to prioritize security resources. But all models have limitations. If we base our decisions solely on the results generated by a model, we may miss important risks.

Many organizations use a risk assessment model based on a standard methodology. The model scores each risk using the following formula:

$$\text{Impact of Asset Loss} \times \text{Probability of Threat} \times \text{Vulnerability Exposure} = \text{Total Risk Points}$$

For each risk, we assign a rating to each of the three contributing factors in the formula. To illustrate, I'll use a scale of 1 to 5. A high-value asset, such as a microprocessor design, might warrant a rating of 5.

We then multiply the three ratings to obtain the total risk points. In this example, the maximum possible risk score is therefore 53, or 125.

A simple approach to risk management, using the output of the model, would be to divide the security budget among the highest-scoring risks.

The model is valuable because it provides a consistent method for helping compare and prioritize a broad spectrum of risks. However, allocating resources based only on the overall risk score can miss potentially disastrous “black swan” events that have very low probability but extremely high impact (Taleb 2007). Because the formula simply multiplies three ratings to obtain the overall score, black swans tend not to score as highly as lower-impact events with higher probability.

To counteract this problem, we can examine the information in the model in more detail, from different perspectives. We can create a list of the 20 most valuable assets and consider whether they need additional controls. In the same way, we can examine the top threats and vulnerability areas.

The point is that any model used to calculate risk should be used as a framework to drive a dialogue about all the variables and options, rather than as a tool that generates the answers to our problems. By discussing the issues from a variety of perspectives, we may identify important concerns we’d miss if we simply look at the overall risk scores.

Before I moved into the information security field, I worked in finance. In our finance group, we found the same principle held true when conducting ROI (return on investment) analysis. Our ROI model generated forecasts. However, it was by discussing the model’s assumptions that we determined whether or not the model’s predicted financial returns were reasonable.

Another method for prioritizing information systems risk management is to examine systems from the perspective of critical business processes and to consider the impact of a loss of confidentiality, integrity, or availability.

An application that prints shipping labels may initially appear to be low priority because it is small, inexpensive, and doesn’t contain confidential data; it simply takes the information it needs from a customer information system on the network. However, if it’s unavailable because the network is experiencing problems, the impact is huge because the company cannot ship products.

The potential impact to a business process of losing confidentiality, integrity, or availability may also vary depending on the stage of the business cycle. Consider a payroll system. Information confidentiality and integrity are always important, but availability is exceptionally critical on payday.

Communication Is Essential

Communication is an essential part of any strategy to mitigate the misperception of risk. To alter the way people behave, we need to change their perception of risk. To effect that change, we must communicate with them.

Changing perceptions is difficult. We may need to address long-held preconceptions about what is risky and what is not. Once people form an initial estimate of risk, they can be remarkably resistant to adjusting their perception, even when given new information (Breakwell 2007, 59).

In addition, each person may have a different perception of risk. To communicate effectively, we may need to understand an individual's viewpoint and then tailor our communication accordingly. Consider the example of taking laptops to countries with a high risk of information theft (see sidebar). People who are extremely concerned may need a patient, thorough explanation of the risks and benefits of taking their laptop versus leaving it in the office. A less fearful individual may just need a quick reassurance and a few basic facts.

Although changing risk perceptions can be challenging, we don't have any choice but to try. Employees will use social media whether we like it or not. When they do, they may not only put themselves at risk; they could be putting the company at risk too, if they are not careful.

Communication can reduce the issue of misperception due to asymmetry of information. This asymmetry is created when security professionals know about risks but don't share the information with end users within their organization. When two parties differ in their knowledge of a threat or vulnerability, their perception of risk is likely to differ also. In other words, it is difficult for users to care about a hazard if they don't even know it exists.

To succeed in changing users' perceptions, we must communicate in ways that engage them, using language they understand rather than technical jargon. In my roles as a security professional, I have always tried to employ entertaining, interactive video tools to help engage users and teach them how to spot dangers such as phishing web sites. As I'll explain further in Chapter 5, I have found these methods have been highly effective in changing users' awareness and perceptions, and ultimately in shaping their behavior.

Patiently explaining to users the consequences of their actions can also help shape their perception of risk. In some countries, pirating software is so commonplace that it is almost an accepted part of the culture. This poses a problem for many multinational companies. Employees in these countries may not even believe copying software is wrong, let alone view it as an illegal act. It can be useful to describe the potential consequences of copyright infringement for the individual and for the organization. We can explain to employees that a decision to pirate software can expose the company to software license compliance risks. The consequences may be even more far-reaching if the copied software is then incorporated into the company's technology-based products or services. If a product is discovered to include stolen software, the company may be unable to ship it to customers, which means a significant loss of revenue. Of course, employees may experience personal consequences too: if they copy software, they run a high risk of losing their jobs.

Organizations as a whole may also be blind to risks, or may simply choose to ignore them. One way to overcome this misperception is to patiently build up a list of examples showing how other organizations ignored similar risks and experienced adverse consequences as a result, according to Breakwell, the University of Bath psychologist (pers. comm. 2012). The more examples in the list, the harder they are to ignore.

"Organizations stick their heads in the sand, ostrich-like," she says. "But if you have a database of examples illustrating where things have gone wrong elsewhere, it becomes harder and harder to find enough sand to stick your head in."

CHALLENGING PRECONCEPTIONS: TAKING LAPTOPS TO HIGH-RISK COUNTRIES

It may be necessary to challenge perceived wisdom in order to expose a clear picture of the real risks, and consequently make the right decision.

Some companies react to the higher rates of intellectual property theft in certain countries by barring employees from taking their corporate laptops on business trips to those countries. In some cases, the companies issue employees with a new “clean” system from which all corporate data has been purged.

The goal is to prevent situations in which information theft might occur, such as when an employee leaves a laptop containing corporate data unattended in a hotel room. A malicious individual could then get physical access to the system and copy the data or implant software that will surreptitiously steal information over time.

But does preventing employees from taking their familiar laptops really solve the problem? Let’s suppose we issue employees with a new, data-free laptop. To do their jobs, they’ll still need to use this system to log into their corporate e-mail and other applications—providing an opportunity for hackers to intercept the network traffic.

Furthermore, if attackers really want to target an individual, they have ways to do it without gaining physical access to the system. With a spearphishing attack, they can induce the individual to click on a malicious link that remotely downloads malware.

Preventing employees from taking their laptops and information also deprives the organization of the key business benefits of using a full-featured portable computing device; employees will likely be less productive as a result. So when assessing the risks of traveling with mobile devices, an organization needs to think through the tradeoff between risk and benefit, including the cost of providing what they believe to be a “clean” system and the impact on the user.

Building Credibility

Ultimately, our ability to influence people’s risk perception depends on our credibility. We need to build trusted relationships with executives and specialists across the organization to ensure our security concerns are seriously considered rather than seen as fear-mongering or target fixation.

Trust is built in drips and lost in buckets; it is hard to create and easy to destroy. If we create a security scare about a threat that turns out to be irrelevant or overblown, we may be seen as just another source of misperception. If business groups think we are providing unreliable and exaggerated information, will they trust us to provide their security?

We can establish credibility by demonstrating consistency, striving for objectivity, and showing that we can accurately predict the real security issues affecting the organization, and then communicate them in an effective and timely way. As I'll describe in Chapter 10, we need to communicate security issues more frequently at C-suite level; to do so, we need to be able to clearly explain security issues in terms of enterprise risk.

Credibility is also built on the competence that comes from understanding the business and technology as well as possessing core security skills. As the scope and importance of information security continue to expand, creating this credibility provides an opportunity to step into a more valuable, high-profile role within the organization.

CHAPTER 3



Governance and Internal Partnerships: How to Sense, Interpret, and Act on Risk

If we are together, nothing is impossible. If we are divided, all will fail.

—Winston Churchill

To reduce cost, our company's human resources group wants to move all HR-related processes to a SaaS provider, a cloud-based business that's less than five years old. At first glance, this might seem a low-risk decision. There's a clear business case, and outsourcing HR systems doesn't seem to create risks to corporate information assets such as intellectual property. Most businesses regard HR systems as commodity applications, so they might select the supplier who can deliver the required functionality at the lowest cost.

But there's more to consider. Employees' personal information will be transferred to the outsourcer, potentially creating new privacy concerns. And imagine the impact if thousands of our employees don't get paid because the supplier experiences system problems on payday and lacks adequate disaster recovery capabilities.

Clearly, the HR group owns the HR business processes. However, outsourcing these applications and processes can introduce risks for the entire business. The systems that support HR processes can create information risks. Outsourcing also involves procurement. The business needs a clear overview of all the factors, including the risks, in order to make the best decision. To provide this view, the HR, procurement, and information risk and security groups need to work together.

A typical organization makes many decisions that require this kind of internal partnership to manage the risk. A product group wants to outsource development work to bring a product to market more quickly. A marketing team wants to engage a developer for a new social media initiative.

Similar considerations also apply to internal technology transitions such as OS and application upgrades. Each new technology introduces new capabilities and risks. Sometimes, the technology also includes features or options designed to help reduce risk. By carefully analyzing the risk and security implications, including privacy and

e-discovery considerations, we can help manage the risk of the transition, and we can often capitalize on the new features to improve the risk picture overall.

For example, when Intel IT was considering whether to migrate to Microsoft Windows 7, the information security team partnered with other groups in a broad evaluation of the OS. We identified several features that could improve security compared with previous versions of Microsoft Windows, and these security capabilities were an important factor in the decision to deploy Microsoft Windows 7 across Intel's enterprise environment (Fong, Kohlenberg, and Philips 2010).

The ability to make these decisions with an accurate view of risk depends on having the right organizational structure in place. Because each organization is different, there's no single, standard risk management structure that applies to all organizations. But at any organization, building an effective risk management structure involves considering two key areas, which I'll discuss in this chapter:

- **Clearly defined information risk governance:** Governance defines who makes decisions, who can block them, and who is allowed to provide input.
- **Strong partnerships and multi-stakeholder collaboration:** Collaboration between the information risk and security team and other internal groups is critical in forming an accurate view of risk and managing risk overall. Some partnerships are formally defined as part of the risk governance structure; others are informal relationships. These formal and informal relationships are so important that I'll dedicate a large part of this chapter to them.

Information Risk Governance

Governance is about establishing a structure that enables the organization to effectively sense, interpret, and act on risk. Traditionally, information risk governance has been considered as a component of IT governance. The IT-centric view is encapsulated in a definition from the Massachusetts Institute of Technology Center for Information Systems Research (MIT CISR):

"... A framework for decision rights and accountability to encourage desirable behavior in the use of IT. Governance identifies who will make key IT decisions and how will they be held accountable."

But as every company becomes to some extent a technology company, we need to broaden this definition to include the information risk associated with technology-based products and services. Perhaps a better definition for this broader view is "Governance identifies who will make key *information risk* decisions and how will they be held accountable."

Information risk governance focuses on enabling the business while protecting the confidentiality, integrity, and availability of information, whether it is corporate data or personal information about employees or customers. It requires the involvement of the entire organization. To achieve effective information risk governance, the information risk and security team must work closely with other groups.

A company's primary areas of information risk are closely intertwined, underlining the need for an effective governance structure that embraces all of these areas. For example, a hacker might compromise the IT systems used by the company's product developers, and then use those systems as a way to introduce malware into the company's technology-based products.

Think about how easily security researchers were able to hack into Jeeps and other vehicles over the past couple of years, demonstrating their ability to remotely take control of the car with potentially life-threatening consequences. Clearly, security may not have adequately considered such a scenario when the car's product groups designed those features. Yet any big company, including automakers, typically has large teams of people dedicated to managing information risk. It seems that in the case of the automakers, the companies perhaps lacked an effective structure for managing information risk wherever it occurs, whether that is in the company's products and services or within back-office IT systems.

To some people, the word governance may imply unnecessary bureaucracy, or perhaps even a dictatorial approach. MIT CISR notes that "good governance is enabling and reduces bureaucracy and dysfunctional politics by formalizing organizational learning and thus avoiding the trap of making the same mistakes over and over again."

Research at MIT CISR shows that the more businesses leverage the structure, tools, and techniques of governance, the greater the potential benefits. In fact, MIT CISR's work suggests that firms with effective IT governance enjoy profits that average at least 20 percent higher than their competitors (MIT CISR 2012).

However, leveraging governance doesn't imply slavishly following rules and procedures. A few years ago, I encountered an IT professional who was regarded by some people, including himself, as one of the best managers in IT. He rigorously based his project decisions on the prescribed practices and procedures, and gathered the correct metrics for reporting progress. Yet the projects he was responsible for generally turned out to be large, expensive failures. His obsession with correct procedures often impeded, rather than facilitated, the projects he was working on.

To use an analogy, if you gave the same recipe to a top chef and an average cook, would you expect them to produce exactly the same result? Probably not. Expert chefs don't simply follow the rules; they continually make adjustments using their senses and experience to achieve the best results. The temperature of a cooking surface is not exactly uniform, so a chef may move the pots until they're simmering just right. Fresh ingredients vary from day to day; the experienced chef is alert to the differences and tweaks the recipe and seasonings accordingly.

Like the procedure-obsessed IT project manager, we may scrupulously adhere to the rules but fail to achieve the desired outcome.

This is one reason that partnerships with other groups are so critical. They provide channels for dialogue, helping us sense changing business priorities so that we mitigate risk based on those priorities rather than our preconceptions.

Without a governance structure that facilitates this dialogue, organizations may take too rigid an approach when applying controls to manage and mitigate risks. For example, some security groups try to ban the business use of social media due to the risks, but attempting to stop the use of external social media web sites is counterproductive and, in any case, impossible. At Intel, we found it was more effective to embrace social media and shape the way that employees use it, as I'll describe in Chapter 5. This approach, developed in partnership with other internal groups, enabled the organization to enjoy the benefits of social media while managing the risk.

Finding the Right Governance Structure

No single governance structure will fit all companies (see Table 3-1 and the sidebar “IT Governance Archetypes”). Furthermore, organizations may shift between different risk governance models over time. When most organizations’ information assets were primarily managed in centralized IT systems, it was natural for information risk to be a centralized function managed within the IT group. But now, information-related risks are much more distributed. To drive corporate revenue, many companies are developing technology-based products and services more or less independently from the central IT organization. At the same time, business groups are shifting to cloud-based applications that store corporate and customer information at external cloud providers.

Table 3-1. *IT Governance Archetypes. Source: Weill and Ross 2000*

Style	Who has decision or input rights
Business Monarchy	A group of business or individual executives (CxOs). Includes committees of senior business executives (may include CIO).
IT Monarchy	IT executives.
Feudal	Business unit leaders, key process owners, or their delegates.
Federal	C-level executives and business groups; may also include IT executives. Equivalent of central and state governments working together.
IT Duopoly	IT executives and one other group (for example, CxO or business unit leaders).
Anarchy	Each individual user.

IT GOVERNANCE ARCHETYPES

When considering the right risk governance structure for your organization, it may be entertaining to think about how your organization compares with the deliberately provocative governance archetypes, ranging from a feudal structure to anarchy, identified by MIT CISR in the influential book *IT Governance* (Weill and Ross 2000, 59).

In practice, organizations may shift between different risk governance models over time—from an IT-centric monarchy during the mainframe era, toward a feudal model or business monarchy as distributed systems emerged, swinging back to a federal model as they recognized there’s a role for centralized IT, then shifting again towards a business monarchy with the focus on technology-based products and cloud computing.

Today, many organizations may find that it makes sense to establish a hybrid governance model that balances centralized and decentralized risk management functions. At the same time, the need for a single, broad view of all information-related risks is driving organizations to create an executive role with overall responsibility for information risk. The executive often has the title of Chief Information Risk Officer (CIRO) or sometimes the Chief Security and Privacy Officer (CSPO). The executive’s broad responsibilities encompass the roles of Chief Information Security Officer (CISO)/Chief Security Officer (CSO) and Chief Privacy Officer (CPO).

To consider how this model works, let’s first think about all the interrelated risks that an enterprise needs to manage. Figure 3-1 shows each primary area and the core elements that are common to all of them. The CSPO’s role is to manage this “Rubik’s Cube of risk.”

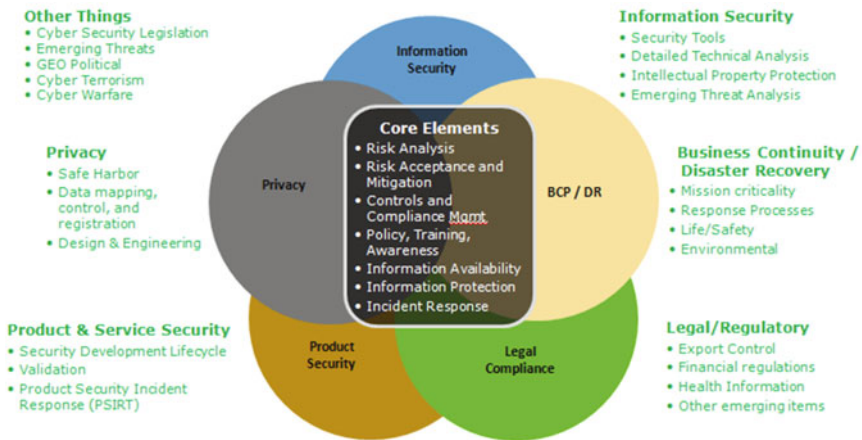


Figure 3-1. Security and privacy: the primary areas of information risk, and the core elements of information risk management that apply to each area

Now consider the governance model, the organization’s framework for managing those risks, shown in Figure 3-2. It consists of four main areas:

- **Oversight:** This area focuses on making informed risk decisions and reviewing risks. It includes committees and review boards that set strategic direction, and review key risk areas such as ethics, compliance, and corporate investigations.
- **Monitoring:** Monitor (sense) risk through external and internal sources. External sources include industry research and analysis. Internal sources include internal partners who inform us of new business risks or legal requirements. These internal sources also include our own security technology sensors.

- **Engagement:** Participate in industry workgroups and in partnerships and dialogues with trusted peer organizations. These external engagements provide a valuable risk-sensing function and help influence key security initiatives. I'll discuss external partnerships in more detail in Chapter 4.
- **Operations:** Day-to-day risk management activities and processes, including risk assessments, incident response, and exercises such as war games.

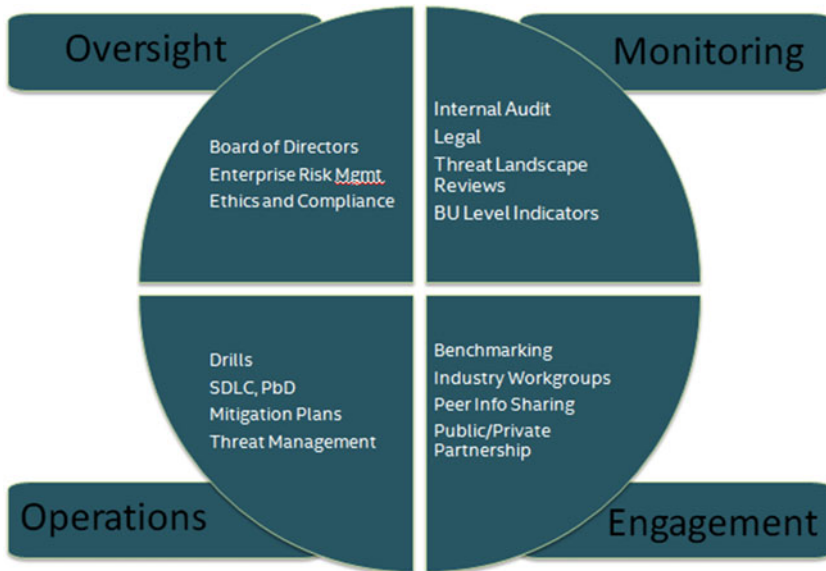


Figure 3-2. A corporate information risk governance model

Typically, the corporate governance model should achieve a balance of centralization and decentralization. At most large companies, risk is decentralized: at any one time, our companies are planning or managing many technology-related initiatives and events across practically every part of the business. Therefore, we need decentralized risk management processes; too much centralization can mean losing the ability to sense threats and respond in an agile way. But at the same time, we need a broad centralized view of the dynamic risk landscape and the ability to set organization-wide policies in areas such as security, ethics, and privacy. So the model must allow a centralized view and ownership of key risk functions, along with the ability for decentralized execution.

The CSPO and the information risk and security team are involved in all four quadrants of the model. The CSPO tends to be more focused on oversight and engagement, while the team's members naturally tend to be more involved with monitoring and day-to-day operations.

For most functions, the CSPO and team work with other parts of the organization, either taking primary responsibility or operating in a participatory role. In the Oversight quadrant, for example, the CSPO may sit on the ethics committee and participate in business unit risk management reviews. In monitoring, the CSPO's team may have primary responsibility for threat landscape reviews and threat indicators, but take more of a participatory role in internal audits and assessing business unit risks. In operations, the team may own responsibility for the security development lifecycle and privacy by design, while participating in change control. It should be apparent that all of these functions require collaboration with other groups within the organization.

Building Internal Partnerships

By providing vehicles for dialogue and decision-making, internal partnerships and multi-stakeholder collaborations enable information security teams to become more agile and responsive to business needs. The number of potential partnerships has grown as the scope of information risk has broadened to include a range of privacy and regulatory concerns as well as traditional security threats.

In mature and proactive organizations, the information risk and security team partners with many internal groups for a variety of functions, including risk management decisions, incident response, and monitoring. These groups include legal, finance, human resources, physical security, and business groups.

Partnerships may include formal structures such as standing committees as well as a large number of informal and ad hoc relationships. These are created and maintained through everyday communication with people in other groups. We might initially contact a business group to understand the potential impact of an emerging area of legislation. The business group identifies risks and opportunities that we hadn't even considered. Our initial request thus sparks a dialogue about requirements and controls, and ultimately evolves into a partnership that helps us monitor risks and mitigate them. We also gain business acumen, which helps us play a more valuable role within the organization.

In my roles running risk and security, partnerships and multi-stakeholder collaboration have been critical to my success in understanding the broader risk picture, helping the organization sense, interpret, and act on risk. Through these relationships, other groups can act as additional eyes and ears for the information security group, such as security threats and compliance concerns. For example, the HR legal group might alert us to an employment-related regulation that creates new compliance concerns. Information about risks flows in the other direction, too: we may alert our partner to new threats that we've encountered. As we leverage other groups to look out for our interests, they can also use us to look out for their interests. We also work with partners to interpret this shared information through analysis and decide how to act in response.

Internal partnerships may focus on just one of the areas shown in Figure 3-2, or they may intersect multiple areas. For example, we partner with HR for incident response (operations) and to learn about new employment laws (monitoring). Multiple partnerships may also be required within each focus area: with the growing number of regulatory requirements, partnerships with internal groups such as HR, legal, corporate security, and internal auditing become increasingly important and valuable in the area of operational investigations.

Because no two organizations are identical, each organization may require a different set of internal partnerships, depending on its structure and business needs. Every partnership should be created with a clear purpose. The organization should also clearly define who is involved and who makes the decisions. To determine the partnerships your information security group needs, as well as their structure and purpose, it may be useful to ask the following questions:

- Who do we need to partner with and why? To put it another way, who do I interact with every day, and why do I interact with them?
- What benefits do I receive from that interaction, and what benefits does my partner receive?

In the remainder of this chapter, I'll discuss some examples of important partnerships, describing how we can use them and the value they provide. I'll start by examining partnerships with fellow travelers who have complementary roles in managing business risk and liability: legal, finance, human resources, corporate security, and corporate risk management groups. Then, I'll examine partnerships with business group managers.

Legal

Legal groups are among the information security group's most important partners because of the many areas where their roles intersect with ours. They own the responsibility for legal compliance and legal review. They interpret laws, analyzing the implications and relaying the relevant information to the rest of the organization. Key partnership areas include privacy, litigation, intellectual property, contracts, and compliance with financial regulations.

As companies create more technology-based products and services, their initiatives are likely to come within the scope of a broader range of laws and regulations. Health-monitoring products might fall within the purview of the Food and Drug Administration; companies thinking about using drones for photography need to think about Federal Aviation Administration requirements.

Privacy

As privacy regulations continue to grow in complexity and reach, many organizations need to comply with multiple requirements at local, regional, and national levels. Legal specialists across the organization can help us understand what's required in each geography, align policies and controls for protecting personal information, and decide how to manage responses in the event of a breach.

Even local regulations can have implications across the enterprise. For example, citizens of European countries are subject to European and national privacy laws and regulations. The simple transfer of European employee personnel information to a US-based server will trigger a need to comply with the EU data privacy laws regarding such transfer of employee information.

Litigation

As one might expect, it's essential to partner with legal specialists in situations where litigation is possible or already in process. Examples are investigations of security breaches, particularly when law enforcement is involved. Another area of partnership is in responding to subpoenas and litigation discovery orders; a legal group may need to work with the information security team in order to collect the required information. To ensure that data is available for discovery when needed, we may also need to collaborate with the legal group to implement appropriate data retention policies.

Intellectual Property

Many organizations use a data classification structure to protect intellectual property, with the most highly classified information receiving the greatest protection. We work with legal groups to specify the classification structure and then implement controls on management and distribution of such information to provide the appropriate level of protection. We also partner to respond to suspected or known IP thefts. Suppose an employee loses a laptop storing the designs of future products; a dialogue with IP attorneys is essential to understand the implications and decide how to respond.

Contracts

Almost every contract with a supplier or customer contains a confidentiality provision, which sets expectations about how each party will maintain the confidentiality of the business transaction and any shared confidential information. We partner with the procurement organization as well as the legal group to define and implement these requirements into contracts.

If our company decides to outsource a business application to an external supplier, we'll typically work with the procurement organization and legal team to define these confidentiality and data security expectations, as well as the evidence we'll need to validate that those controls are operating properly. For example, when hiring a company to manage health benefits, we set expectations about how they must protect our employees' personal health information.

Our customers have expectations, too. Another company may need to share some IP with us to help us integrate our technology into their product. We need to understand their requirements and ensure that appropriate controls are implemented.

A security technology supplier has to meet customer expectations that go beyond the product's ability to provide protection. As I mentioned in Chapter 1, one of the irrefutable laws of security is that even a security feature can be used for harm. So suppliers must be able to discuss their security development lifecycle, privacy by design, and overall state of internal controls, all of which could ultimately affect the efficacy of the product.

Financial Compliance

In the United States and other countries, public companies are legally required to disclose “material events,” those likely to have significant financial impact that could affect investor decisions, including IT-related incidents. An important aspect of risk governance, therefore, is partnering with legal groups to understand the types of events and specific incidents that must be reported.

Guidance from the US Securities and Exchange Commission specifically discusses the obligation to disclose the impact of cyber attacks, including those that result in IP thefts. Companies are also required to disclose material increases in security spending in response to an attack, even if the attack didn’t result in a loss of IP (SEC 2011).

The legal team cannot do this alone because it lacks the security context of the event: the frequency of specific types of attack, the potential impact, and the cost of response. Therefore, the security team must be involved.

In 2010, Google disclosed that it had been breached in the widely publicized Operation Aurora attack. At around the same time, Intel also experienced an incident of similar sophistication. This was before the SEC issued its guidance in 2011, but as I pondered the potential ramifications of a cyber breach one sleepless night, I realized that I should call our SEC legal experts to discuss the incident. Subsequently, we disclosed the incident in our financial report for the first quarter of 2010 (Intel 2010).

Legal Specialists Within Business Groups

At large companies, each business group may have embedded legal experts. We need to work with them for issues directly related to their group. In addition, because of their connections within the group, these legal professionals can be extremely helpful in influencing the group’s controls and expectations.

Marketing groups, for example, usually include individuals who want to explore new ways to communicate with users via social media. This appetite for adventure is a good thing; it can benefit the business. But at the same time, we have to ensure that content is adequately protected and includes appropriate privacy protection and statements. If we bring up the issue directly with marketers, we may receive a lukewarm response, as they tend to view any controls as restrictions on their ability to move quickly. But the legal professionals within the marketing group understand the need for controls. So a good way to raise our concerns is to have a conversation with the business group’s attorney, who can help persuade others in the group that controls are needed.

While I was Chief Security and Privacy Officer at Intel, we implemented a program that reviewed all new externally facing online projects and monitors for potential problems (see sidebar). The projects ranged from web sites to more sophisticated tools, such as an application that users can download and use in conjunction with external social media sites.

As part of the review, we asked the project group who their legal contacts were so that we could verify that they’d received legal approval. We also asked whether trademark and branding teams had reviewed the initiative, which was essential in many cases—especially if the project was planning to register a new web site. Sometimes the answer was no, in which case we facilitated a dialogue with the trademarks and brands team. This enabled the trademark and brand people to manage the risk and helped forge yet another important relationship within the company.

SECURING INTEL'S EXTERNAL ONLINE PRESENCE

Intel's business groups use hundreds of web sites and third-party solutions, including social media platforms, to communicate and conduct business with customers and business partners. Collectively, these externally facing Intel-branded solutions were known as Intel's *external presence*.

Until 2006, these web sites proliferated rapidly in response to business needs, without centralized oversight. Given this growth and following a number of security incidents and the identification of several significant risks, we established the Intel Secure External Presence (ISEP) program to provide appropriate security for Intel's external presence (Leon 2011).

The goals of ISEP, which was a part of Intel's information security group, were to protect Intel's information assets and customers against threats such as loss of personal information and malware attacks, and to maintain compliance with laws, regulations, and standards. By achieving these goals, we also helped to protect Intel's corporate image.

We helped ensure this protection and compliance by reviewing all planned new external presence projects and by monitoring existing Intel-branded web sites. ISEP review and approval was mandatory for new externally facing online projects. We worked with Intel business groups to review planned projects before launch, whether they were to be hosted within Intel or by a third party.

Any ISEP-like process for reviewing a company's external presence should include several key aspects:

- **Ensure notification of new projects** by working closely with business groups and other stakeholders within the company. For example, the information risk and security team should be notified when business groups request new Internet domain names or seek approval to land a new application in the externally facing IT environment.
 - **Work with the business group on each project** to review details of the planned approach to maintaining security and privacy compliance. Verify that the project includes any required mitigating controls before giving approval.
 - **Establish an overarching governance board**, including senior managers from multiple stakeholder groups. This board should have enforcement powers including the ability to shut down web sites for noncompliance.
-

Human Resources

The human resources group is the organization's center of expertise on employee procedures, include legal specialists who are the organization's experts on employee-related laws. Because of its responsibilities, the HR group also tends to be heavily involved in insider risk considerations and applying action in any cases that are discovered. In some organizations, HR is also responsible for other functions, including internal and external communications. Because of this broad charter, the security team may form valuable partnerships with HR in several areas, including employee policies related to appropriate use and protection of information assets, internal communications, and investigations.

Setting Employee Expectations in Security Policies

Employees are part of the security perimeter, as I'll discuss in Chapter . Their behavior can have as much impact on security as the technical controls we use—particularly since a growing number of user interactions with the outside world take place on external web sites and networks, and on personal devices such as smartphones.

It is therefore critical to create employee policies that set expectations for secure behavior. If we can influence employees to behave in more secure ways, we can reduce risk for the business overall. However, the security team cannot write these policies without partnering with HR, including HR legal specialists, to ensure that they comply with employment laws and the organization's existing rules. Then, if an employee disregards the policies, we need to work with HR to take disciplinary action.

Careless behavior can have highly damaging consequences. Imagine an IT employee who decides to store some corporate data on a server at his home so that he can more easily work on projects when out of the office. But his home system is open to the Internet, and thus the data may be broadly exposed to anyone worldwide.

The employee's action has created a significant security risk. To explain the potential impact to HR, it may help to use analogies. We could say it's like an engineer taking critical product designs home and showing them to her neighbors. Or a factory employee taking dangerous chemicals home to experiment with them, and creating the danger of an explosion in his garage. If we have a good relationship with HR, we can have this kind of discussion and determine the appropriate consequences for the employee.

Employee Communications

The responsibilities of the employee communications group often include employee training, employee awareness, and internal distribution of other corporate information. This group's expertise can be very useful when we want to communicate security messages to the workforce. The group already has established communication channels and knows how to align messages with corporate style guidelines. A good employee communications group also knows how to present information in ways that engage employees rather than intimidate them.

In my prior roles running security and privacy, I always worked extensively with the employee communications group to create engaging security awareness messages, including interactive content that helps encourage secure practices when using social media and the Web.

Investigations

Partnership with HR is also essential in internal investigations, including investigations into insider threats responses. In other cases, we may already be pursuing an investigation and need help from HR legal specialists to access employee information.

Finance

The finance group typically takes the lead in managing enterprise-level risk and controls for the organization overall. Therefore, we need to partner with the finance group to assess the business impact of damage to information assets—a loss of confidentiality, integrity, or availability. This applies not only to internal systems that support business operations, but also to information technology-based products and services that generate revenue. We also work together to determine the required controls.

Sarbanes-Oxley Compliance

The corporate finance team usually has overall responsibility for Sarbanes-Oxley (SOX) infrastructure. We also work with the finance group, as well as legal groups, to determine whether we should categorize specific events as material and report them as required by SOX. This also includes product- or service-related vulnerabilities and controls that could have a material effect on revenue or corporate liability.

Working with Business Groups

Each sizeable business group is likely to have a group controller or other financial specialist responsible for financial controls. These finance experts can become important partners for the security team.

Because financial specialists focus on risk and controls, the culture among finance specialists has some similarities with the culture of the information risk and security teams. This shared focus can make it easier for us to communicate our concerns, particularly since the impact of information risk is often measured in financial terms. Therefore, the financial specialist can be a key contact point when we need to discuss information risk with business groups.

Sometimes these risk conversations can evolve into productive multi-way partnerships. A recent example: an IT team presented plans for new systems to support one of Intel's new businesses. As we assessed the information risks, we noticed that the plan didn't include fully redundant systems to ensure business continuity. When we asked why, it emerged that the business group hadn't requested redundancy because it would add cost. Revenue from this new business was initially expected to be modest, so the group's budget was limited.

However, when we discussed the revenue projections with the finance specialists who worked on the project, they expected the business to grow rapidly. This growth would also increase the information-related risk because a system failure would have a much bigger impact on revenue. As we discussed the implications, it became clear that it would make more sense to prepare for the anticipated growth by including redundancy

from the start. So we suggested that the business group negotiate a higher budget—and that’s what happened through a partnership between the business group managers, the information security team, and IT finance and business system specialists. The business group allocated increased funding that allowed IT to implement a redundancy safety net that would protect the growing business.

Internal Audit

Financial groups are often also responsible for an internal audit, which typically includes an IT auditing function—a job with considerable potential for overlap with the information security group’s role. If the security team and internal auditors duplicate each other’s efforts, we’ll waste resources and annoy business groups. Imagine that we contact a business manager to say that we need to conduct a risk evaluation of the group’s systems. The next day, internal auditors contact the same group and say they’re planning to do an audit, which some business managers might perceive to be essentially the same as a risk evaluation. What kind of reception do you think the auditors would receive?

We can minimize the overlap by partnering with internal auditors. This partnership becomes a mechanism for effectively allocating risk management resources. If the information security team has already assessed a system, auditors may be able to increase the efficiency of an audit by leveraging the work that the security team has already performed.

For effective partnership, our work must be thorough, transparent, and well documented so that auditors can see what we have done. We may also swap resources: sometimes security experts may act as guest auditors for specific projects because they have skills that the financial group lacks. The partnership can also be used for valuable dialogue and mutual support. If we’re concerned about a system that internal auditors have previously examined, we can ask for their opinion. We’ll sleep better knowing that another group of objective, risk-focused specialists has analyzed the system.

Corporate Risk Management

Most large organizations employ people whose job includes purchasing insurance for general business risks, including property and casualty insurance to protect the organization in the event of damage to a data center or another facility. When buying insurance, the corporate risk management team may need information from us about the organization’s IT business continuity and disaster recovery plans. Insurers ask for this information in order to set premiums.

Today, the corporate risk management team usually focuses on physical risks. But their scope is rapidly expanding to include IT-related risks as well as risks associated with products and services. Privacy breaches or other compromises can have a major impact on a company’s revenue, cost, and brand image. Because of this trend, insurance against cyber risks is a rapidly growing category, and we can expect a growing need to partner with the corporate risk management team to ensure adequate coverage of information risks.

Consider the case of Sony, which suffered a breach of its PlayStation Network—estimated by the company to cost at least USD 200 million (Perlroth 2011)—and then became embroiled in a legal dispute with its insurer, which claimed Sony’s insurance policy did not cover cyber risk. The breach at Target, in which hackers stole the payment

card accounts and personal information of millions of customers, is estimated to have cost the company roughly \$250 million. Reportedly, the insurance payout of \$90 million left the company \$158 million in the hole, plus what it paid for cyberattack insurance.

Privacy

Privacy and security are closely linked. However, increasing security doesn't always enhance privacy. In fact, it can have the opposite effect. Unfettered bulk collection and monitoring of the information and activities of users and their machines may be capable of increasing security, but it may also intrude on personal privacy. This data store may also be an attractive target for intruders.

This creates inherent tension between security and privacy interests. This tension is apparent at a national level in the way that privacy advocates respond to the use of surveillance and data mining. Government security organizations may feel that they protect data extremely well, but privacy advocates still object to the fact that information is collected and the way it is used.

Similar concerns apply at the enterprise level. We need to carefully manage the relationship between security and privacy, ensuring that we apply the appropriate level of controls to protect information without infringing on personal privacy.

The structure of this relationship varies between organizations. While at Intel, the information risk group that I managed for over a decade included the privacy team, which reported to me as the CISO. Then as we began to see growing confluence of the risks shown in Figure 3-1, I was promoted to a broader role as Chief Security and Privacy Officer, to give us an integrated governance and accountability structure. At other organizations, privacy is the responsibility of a separate group headed by a Chief Privacy Officer who is the CISO's peer. This arrangement necessitates careful management of the relationship between security and privacy teams to manage tension, align policies, and control breaches. In organizations with this structure, the security team sometimes complains that the privacy team is "getting in their way," which usually means that the security team wants to collect specific information and the privacy team objects.

Regardless of the organizational structure, it is the security team that is logically responsible for implementing IT controls. It is the product security team that is responsible for security development lifecycle (SDLC) and product security incident response processes (PSIRT). Laws define privacy rights; the organization's interpretation of those laws drives compliance requirements. It is the security team's responsibility to determine how to implement controls to support those requirements.

Corporate Security

The corporate security team focuses on physical security concerns ranging from door locks and guards to break-ins, fires, and natural disasters. By partnering with this team, we can make sure we're aligned on protection of key information assets. It wouldn't make sense to implement sophisticated data-protection tools on the servers in the data center and then leave the data center doors unlocked.

We also need to coordinate on other issues, including incidents that involve law enforcement. Not so long ago, assaults and harassment were almost always physical incidents handled by corporate security and the police. Today, there's a much bigger

overlap with information security. More crime is moving online, and we may encounter other problems, such as cyber bullying. Because of these trends, we may need to help assess the impact and drive the response.

Business Group Managers

Each business group has its own processes and applications, whether it's a product-focused unit responsible for generating revenue or an internal group managing finance or human resources. The information security team needs to partner with each group to implement security controls that protect the group's applications and information.

As the business acumen of our information security team increases, we can better fulfill our Protect to Enable mission by focusing on controls that improve security without impeding the business. This applies not only to the systems that support business operations, but also to the technology-based products and services the business unit creates. For example, we may discover product vulnerabilities through our security development lifecycle processes. We can partner with the business group to correct vulnerabilities before shipment, and we can work on training to prevent future mistakes due to poor coding, design, or architecture.

By working with business groups, we can also leverage their strengths. Business group managers can help drive decision-making and incident response. They can also help improve security by setting the "tone at the top," publicly setting expectations for their employees' security behavior. Suppose we notice that an increasing number of the employees at a specific facility are experiencing laptop thefts. We discuss the trend with the general manager and explain that we want to increase employees' awareness with messages about how to prevent theft. The business manager may offer to help by bringing up the topic at a site meeting or otherwise directly communicating with employees. This management request may exert a more powerful influence on employee behavior than messages sent by the security group.

HOW TO RESPOND TO EMERGENCIES

Defining a clear IT incident response process is an essential aspect of IT governance. Similarly, a clear PSIRT is an essential aspect of risk governance for technology-based products and services. Over time, while I was at Intel, we developed a clearly defined crisis management process for responding to emergencies and other significant incidents that affect IT infrastructure or services (Fleming and Tomizawa 2012). The goal of the process was to prevent material impact to the organization and its employees. Similarly, the goal of a PSIRT process is to prevent material impact to customers or even to society in general, depending on the nature of the risks.

Incidents that may trigger the process include cyber events and other information security incidents; physical incidents such as fires, leaks, and major outages that affect IT systems; and major disease outbreaks. A useful starting point for developing the process is the incident management principles based on the US Federal Emergency Management Agency's response to disasters.

Once initiated, an IT emergency response process (ITERP)• operates with a command-and-control structure, led by an incident commander who has overriding authority to make decisions across IT for the duration of the emergency. The structure consists of a virtual organization staffed on a volunteer basis by people from every discipline within IT. When an incident occurs, all team members perform their response roles instead of their normal duties until all issues are resolved.

Following an incident, the team should quickly identify the state of critical business processes that must continue during the crisis. It determines the current status of the key steps in the product cycle: design, build, order, ship, pay, and close. It assesses the physical state of the infrastructure, and analyzes the legal and other impacts if intellectual property or personal information is compromised. Decisions about response and remediation are driven by the incident commander and determined by business priorities.

PSIRT and privacy response processes should be structured along similar lines, focused on their respective mission-critical priorities.

While I was at Intel, the ITERP team, the PSIRT team, and the privacy incident response team proved to be essential components of the successful resolution of every crisis management, coordination, control, and communication activity across the company during my 13.5-year tenure.

Conclusion

Information risk has become a major concern for the entire organization. Managing information risk therefore requires a clear governance structure that enables the organization to make the right security decisions quickly and effectively.

Building the right governance structure can sometimes seem like a complex challenge. I've found that a good way to simplify and focus the thought process is to consider the following two cardinal rules. In my experience, these rules apply to all organizations, whether large or small, public, private, or non-profit.

- **Rule 1: *Structure drives behavior.*** Thinking about the behaviors that you want to see in the areas of security and privacy will help lead you to a structure that encourages those behaviors.
- **Rule 2: *You get what you measure.*** Thinking about the desired outcomes will help you determine how you should measure your organization's success in managing risk.

Think about how your own organization manages information risk. Do you develop strategies in close collaboration with business groups? Do you feel that you communicate well enough with every group to understand their priorities and implement controls that reflect them? Have you clearly defined all of the processes required to respond to a major breach or denial-of-service attack? If you answered "no" to any of these questions, you may need to improve your information risk governance.

Effective governance relies on partnerships between the information security team and other internal groups across practically every part of organization. In this chapter, I've described some of the most important partnerships and the value we can derive from them.

To develop these partnerships, CSPOs as well as Chief Security Officers and Chief Privacy Officers need more than just technical skills. We need to communicate in terms business people understand and build relationships that enable us to influence people at all levels across the organization. As the scope of information security expands, we also need extensive management and leadership skills, both to operate at an executive level and to coach and inspire our risk and security team.

CHAPTER 4



External Partnerships: The Power of Sharing Information

Chance favors the connected mind.

—Steven Johnson

After spending a day at a conference, I was having dinner with a dozen or so peers when a debate began about the dangers and benefits of sharing security information with other companies. One person turned to me and asked me whether, if I had information about a specific new threat, I would share it with him.

“You bet,” I said.

“But what if I was your competitor? Would you still share?” he asked.

“Our companies might compete for business,” I replied, “but in the security arena, my real competitors are the malicious actors who want to harm my company’s information systems. Those are my competitors, and they’re your competitors, too.”

As soon as I’d said this, several people at the table agreed. This agreement was gratifying—and not just because I felt that I had support for my views. The bigger implication was that my peers saw the value of sharing information outside their companies.

This hasn’t always been the case. Historically, many organizations frowned on the idea of sharing security information externally, and more than a few had policies forbidding it.

However, attitudes are changing. Although there is still resistance at some companies, many organizations now see the value of sharing information and have begun doing so. Evidence includes the growth of industry-specific information-sharing communities, such as the retail-industry group that formed after Target’s massive customer-information breach in 2013. There are also innovative partnerships that have a regional rather than industry-specific focus, such as the Arizona Cyber Threat Response Alliance.

Supportive actions by the US Government have also helped encourage information sharing. In 2014, the Federal Trade Commission and Department of Justice issued a policy statement indicating that sharing threat information was unlikely to raise antitrust concerns. This addressed a key reason that some big organizations had been reluctant to

share information. “Cyber threats are increasing in number and sophistication, and sharing information about these threats, such as incident reports, indicators, and threat signatures, is something companies can do to protect their information systems,” said Bill Baer, an Assistant Attorney General in charge of the DoJ antitrust division (U.S. Department of Justice 2014).

In 2015, the White House issued a statement encouraging information sharing as a way to help safeguard national and economic security, and directing the Department of Homeland Security to support the formation of information-sharing groups under the umbrella term Information Sharing and Analysis Organizations (ISAOs). And in 2015, legislation was proposed to promote sharing of threat information, although the effort stalled in Congress.

Despite the overall shift in attitude, some organizations still have reservations about sharing information. There are three major areas of concern. First, organizations worry about the legal and regulatory implications of revealing information about threats. A second, related concern is the public relations aspect. Both of these fears have a valid basis. Information security has become an enterprise risk management issue of board-level interest because of the potential effects. Information leaks revealing potential intrusions and data breaches can have legal consequences: the organization may be required to report the problems in order to comply with financial and privacy regulations, for example. If security issues become public, they may also damage the way the organization is perceived by customers and by the business community, potentially affecting a company’s profitability and its stock price. The third major area of concern is privacy. This also has a valid basis. For example, sharing information that identifies the victim of an attack, as some security specialists would like to do, clearly can expose machine data that can potentially compromise the victim’s privacy. Some people also see a risk, following the revelations of National Security Agency eavesdropping, that legislation could be used to enable government surveillance. For these reasons, I believe that any cybersecurity legislation must include appropriate privacy protection.

What’s the payoff from sharing information? My personal experience is that I have obtained real value: information shared by others has helped me understand threats and take action. I have also seen that it’s possible to share useful information while avoiding the issues mentioned above. Companies can share information about attacks without revealing personal information about the victim. They can share indicators of compromise without revealing confidential information. They can alert other trusted contacts during the early stages of investigating a threat, before it’s been determined whether a compromise has occurred that requires regulatory disclosure.

The growth of information-sharing groups shows that many other organizations now share my belief in the value of sharing information about threats and best practices. As I’ll explain in this chapter, sharing security information can provide considerable benefits in managing the risk of moving into new business relationships and adopting new technologies. We just need to find ways to reduce the risk of sharing. The solution lies in creating trusted information-sharing relationships with other organizations. The more we trust the relationship, the more sensitive the information that can be shared.

The need to share security information is being driven by rapidly changing business, technology, and threat landscapes. Increasingly, companies are collaborating with a broad variety of business partners. We share business information, and often we also use the same technology, or we sell or share technology with each other. As we do so, we also share risks. Understanding the risks faced by our partners, and the way they manage those risks, can help us protect our own organizations.

Looking more broadly across the technology landscape, all systems and devices are to some extent connected, whether they are owned by enterprises, individuals, or service providers. Almost every aspect of society depends on a worldwide, rapidly evolving, highly complex network of devices and services. This provides the central nervous system that supports innovation, economic development, and social interaction worldwide. But because we are all inherently interconnected, we share common risks. The threat landscape is dynamic, global, and increasingly complex. Threats may originate in any country and then spread rapidly across national and enterprise boundaries, causing extensive damage to organizations and individuals worldwide.

Because threats spread so quickly and the threat landscape is so complex, it is hard for any single organization to gain a clear view of all potential vulnerabilities, threats, and attacks. External partnerships can help. They provide additional intelligence that we can use to improve our own security posture. By exchanging information with other organizations, we gain what I call *outsight*, or a better understanding of what happens outside our own environment. We learn about new threats before they hit us directly. We see how other organizations are managing those threats. We learn about best practices for managing security operations. Using the information we gather from external relationships, we can increase the organization's ability to sense, interpret, and act on risk.

The Value of External Partnerships

Sharing security-related information can require initiative and courage. The idea of sharing information externally may run counter to the culture of the organization overall, including the culture within the security group. Organizations may view security information as proprietary and confidential, like intellectual property. Many still have policies against sharing information.

It's true that much security information is sensitive, and sharing it can introduce risks. Because of this, we need to be careful about what we share and with whom.

But think about the broader context of how organizations are increasingly sharing information. Most organizations have already recognized that they need to share sensitive business information with partners in order to develop, manufacture, and market new products. Collaboration with other companies is becoming an integral part of many other business processes. As organizations share information, they benefit from their partners' insights and expertise. As noted by Steven Johnson, author of *Where Good Ideas Come From: The Natural History of Innovation* (Riverhead Books 2010), many of the best ideas have emerged not through the inspiration of a single mind, but through the exchange of ideas. "You have half of an idea, somebody else has the other half, and if you're in the right environment, they turn into something larger than the sum of their parts," Johnson said in a speech at the 2010 TEDGlobal conference (Johnson 2010). "We often talk about the value of protecting intellectual property—building barricades, having secretive R&D labs, patenting everything that we have, so that those ideas will remain valuable ... but I think there's a case to be made that we should spend at least as much time, if not more, valuing the premise of connecting ideas and not just protecting them."

I believe that there's similar value in sharing security information. As we collaborate with business partners, we need to understand the threats to their environment, and how they manage risk, in order to determine what we need to do to protect our own organizations. Each partner in a value chain needs to protect information to a level

that is adequate to protect the other partners; the weakest link in the chain can impact everyone. Note that throughout this chapter, I use the terms “partner” and “partnership” in the colloquial sense, not to imply any specific type of formal legal relationship.

There are many other examples of how sharing information can benefit all organizations involved. If we are entering new markets through business partnerships, we need to understand the nature of the threats in those markets from the companies currently operating there. The same logic applies to using new technologies. Organizations are extending their environment to customers and becoming suppliers of mobile apps and web services in the process. As they do, they can learn from other companies’ experience how to manage the risks. Companies are increasingly sharing cloud capacity or other data-center infrastructure supplied by external providers, and can all benefit by sharing feedback with the provider about risks within the environment.

Despite these trends, some organizations still have policies stipulating that employees shouldn’t share internal information about risks and threats with anyone outside the company. This is sometimes the case even when the same organization willingly shares other IT-related information such as helpdesk or e-mail management best practices.

Without wishing to discount the real fears driving these policies, the value of sharing information often outweighs the risk of doing so. Let’s imagine that a CISO learns of a new threat affecting companies in his industry sector. He shares information about the threat with a peer at another company and, by doing so, gains insight that helps the organization mitigate an attack that has caused massive damage at other companies. By sharing information against company policy, the CISO took a personal risk. Yet by doing so, he averted the bigger risk of business disruption and damage to the organization’s reputation.

Failure to share information with others introduces its own risks. If we don’t share with peers, they won’t share with us, so we won’t benefit from their information and insights. I’ve seen cases in which information security professionals wanted to participate in communities, but weren’t allowed by their companies to share any internal security-related information. So they attended meetings but couldn’t contribute. Ultimately, their peers wouldn’t tolerate a situation in which these people were receiving information but giving nothing in return, and they were effectively voted off the island.

External Partnerships: Types and Tiers

Much of the publicity about information-sharing initiatives has focused on public-private partnerships related to critical infrastructure and national security. However, there are many other types of formal and informal external information-sharing relationships, including 1:1 partnerships and groups comprised solely of private-sector organizations.

External partnerships are most often used to share information about specific threats and best security practices. But some partnerships focus on other types of information. For example, security specialists within the high-tech sector share information in order to develop security standards, which are then implemented in various products.

Much of this security information is sensitive. Because of this, we need to be able to trust that the partners with whom we share information will treat it appropriately. The more sensitive the information, the greater the level of trust required. In general, the level of trust can be higher in relationships with fewer people, allowing more-sensitive information to be shared. As the number of people increases, there’s a greater chance that information will leak, so the level of trust tends to decrease and only less-sensitive information is shared.

Relationships therefore naturally tend to fall into a tiered pyramid model, as shown in Figure 4-1 (Willis 2012). At the top of the pyramid are the most-trusted relationships with the fewest partners; these are 1:1 partnerships between two individuals at different organizations, or between two security teams.

Information-sharing relationships between more than two partners are often referred to as communities. Because more people are involved, a legal or peer-enforced agreement is usually needed to define the level of trust and confidentiality expected among community members.

The two middle tiers of the pyramid include groups with intermediate levels of trust, sharing information with varying levels of sensitivity. The *targeted tier* typically consists of public-private partnerships aimed at protecting critical infrastructure. The *confidential tier* includes many private-sector communities, including regional communities and those focused on specific industry sectors.

At the bottom of the pyramid is the *public tier*, comprised of the largest communities with the lowest level of trust. At this level, information is often public and may be broadcast via the Internet. This tier might include groups that develop educational information about threats for public distribution, or CISOs who share their insights via public webcasts.

I should note that there is considerable overlap between these tiers. A group may have characteristics of both the targeted and confidential tiers, for example. Also, the number of members in groups within each tier (shown in Figure 4-1) is just a guideline: communities at all levels tend to grow over time as more organizations see the value and join.




Figure 4-1. Tiered pyramid model for trusted information-sharing partnerships and communities (adapted from Willis 2012). Source: Intel Corporation, 2012

How can you get involved in information-sharing partnerships? One good method is to start by participating in communities in the public tier, where the information shared has a relatively low level of sensitivity and therefore involves little risk. In these communities, you're likely to meet peers with whom you can begin to engage in 1:1 partnerships. As you become more knowledgeable about the communities that reflect your organization's key interests, you may then become involved in relationships in the middle tiers of the pyramid, where more confidential information is exchanged. I have always made sure that my teams and I actively participate in partnerships at all the tiers of the pyramid.

1:1 Partnerships Tier

- **Community structure:** Direct communication between CISOs at two organizations or between their teams
- **Typical number of partners:** 2
- **Example partnership/community:** Any two organizations who choose to share information
- **Example goal:** To mitigate shared threats by exchanging information with a business partner more quickly and in greater detail than would be possible within a larger group
- **Trust framework:** Personal trust and existing business relationships

Targeted Tier

- **Community structure:** A relatively small number of critical information infrastructure owners and operators sharing information to protect the infrastructure. Also includes key security ecosystem influencers, such as large security service providers or vendors.
- **Typical number of partners:** Up to about 50
- **Example partnership/community:** Information Sharing and Analysis Centers (ISACs)
- **Example goal:** To prevent advanced persistent threats (APTs) within the industrial base by sharing APT signature information
- **Trust framework:** Strong information-sharing frameworks, such as national security clearances and nondisclosure agreements, are required. Trusted sharing mechanisms, such as encrypted web portals with multifactor authentication, are also required.

Confidential Tier

- **Community structure:** Communities that represent industry sectors or other groupings, such as the banking sector and Internet service providers (ISPs), or regional forums
- **Typical number of partners:** Up to about 100

- **Example partnership/community:** BITS (financial services industry), Bay Area CSO Council (regional), Regional CSO Summits
- **Example goal:** To enable members to protect against common threats and vulnerabilities affecting their industries. For example, ISPs might share the command and control Internet addresses that botnets use.
- **Trust framework:** Communities typically use trust frameworks such as nondisclosure agreements or memoranda of understanding.

Public Tier

- **Community structure:** A broad range of communities that represent all user categories, including consumers, small- and medium-sized businesses, and industry in general
- **Typical number of partners:** 100s to 1,000s
- **Example partnership/community:** Forum for Incident Response and Security Teams (FIRST), National Cyber Security Alliance
- **Example goal:** To share best practices or informational bulletins about widely known threats and vulnerabilities that affect a large cross-section of users.
- **Trust framework:** Trust frameworks are not necessary; communities typically distribute information broadly through mechanisms such as e-mail distribution lists or public web sites.

Let's look at these tiers in more detail.

1:1 Partnerships

In my experience, 1:1 partnerships are some of the most valuable security relationships. They may be formal or informal, established at a corporate level or between individuals.

As I explained, a key advantage of a trusted 1:1 partnership is that we can more safely share highly confidential information. We can often create a stronger bond with a single individual than with a larger group. As a result, the shared information often has a depth and richness that's lacking in information shared within larger communities.

Another advantage is speed. Communication is often fastest in 1:1 partnerships, partly due to logistics. It's much easier to set up a meeting between two people than it is to organize a meeting with a dozen people. To exchange information about the latest developments, a CISO may be able to simply pick up the phone and have a conversation with his or her peer. Quickly sharing information enables a faster response to threats—and in the security arena, timeliness is often critical.

Here's an example showing how 1:1 partnerships can develop and benefit both partners. Through my participation in a larger security community, I got to know the CISO at a fast-growing e-commerce company whose customers were primarily consumers. We both would contact each other periodically for advice and information as we puzzled over the latest security challenges. Over time, these conversations evolved into open dialogues about best practices and benchmarking.

The relationship eventually evolved to a point where we both realized we could learn a great deal more by bringing our teams together in a face-to-face meeting. The resulting half-day meeting proved incredibly valuable to both teams. Our team was able to provide insights and experiences about managing security in a large, complex enterprise environment. This was helpful to the security team at the fast-growing e-commerce company, which was in the process of building an enterprise environment to support its fast-growing business. In return, the team at the e-commerce company was able to share the security challenges and experiences of operating a large consumer business with millions of online customers. This was extremely valuable to us at Intel because we were in the process of expanding our external online presence and were beginning to encounter some of the same challenges.

The partnership thus expanded from ad hoc conversations to a productive relationship between teams sharing experiences and best practices at multiple levels. It's hard to imagine that this extensive information exchange could have occurred within a larger community.

Another example: I met the CISO of a large manufacturing company at an industry event, and we stayed in touch through occasional e-mails. Then, during a period of especially large-scale industry attacks, our communications suddenly became much more frequent and detailed. It was extremely valuable to be able to pick up the phone and simply call a peer to share the latest knowledge about the attacks and responses.

I have frequent 1:1 meetings with peers at other companies, sometimes as often as several times a week. These meetings can serve several purposes. A few years back, I met with a team from a key supplier to discuss our strategy for securing employees' personal (bring-your-own) devices. I shared our best practices with this team, and during the question-and-answer discussion, team members also provided information about how they were addressing the same problem. The meeting served as a helpful benchmarking exercise for all of us.

At the same time, the discussion clearly demonstrated each company's commitment to protecting its partner's business information. It showed the depth of each company's strategy for protecting information—revealing a commitment that extended far beyond the desire to comply with contract confidentiality clauses. I felt more confident that if a security issue ever arose, I could talk directly to my counterparts at the supplier company because their commitment to protecting information would enable a productive approach to resolving problems.

Another recent discussion, this time with a potential customer, focused on the cloud. The organization was concerned about our use of the cloud as part of our infrastructure, and also as a part of the service connected to our product. Rather than respond to the lengthy survey they had put together, we met with them to discuss how Cylance uses the cloud and which data we store there. We discussed the risks that could exist in the cloud infrastructure, the potential implications of those risks, and how we manage those risks. We also discussed other precautionary steps the customer could take to further mitigate the potential risks. This discussion helped develop a relationship that built the most customer trust.

Communities

Participating in larger communities may not provide information that's quite as rich and deep as the information you'd obtain from a 1:1 partnership with a peer. But communities provide value in other ways.

Because they contain more people, communities provide breadth and diversity of perspective that help us make balanced risk decisions. With a larger number of participants, there's a better chance that one of them will have developed a solution to a problem, or can provide valuable new information about an industry attack.

Some communities focus on sharing threat-related information; others on benchmarking and best practices, influencing legislation, developing security standards, or public education.

Communities can also present great networking opportunities. Through participation in communities, I've met several people with whom I've subsequently developed closer 1:1 partnerships.

Community Characteristics

Like all groups, communities require a structure and a set of ground rules to be effective. Successful communities typically have the following characteristics:

- **Clear goals:** The community shares clearly defined common goals that benefit members, such as mitigating an industry-wide threat. A community may have several goals.
- **A strong framework of trust, such as a legal or peer-enforced agreement, that addresses risks related to information sharing among community members:** For example, the Industry Consortium for the Advancement of Security on the Internet (ICASI) has a strong multilateral nondisclosure agreement, while other communities, such as the Bay Area CSO Council, rely on a peer-enforced trust framework.
- **Trusted communications channels:** Members can safely contribute and access shared information using an effective trusted communications channel or mechanism, such as a secure web site. These channels are not always electronic; some regional groups conduct face-to-face meetings to further reduce the risk of compromise.

An organization is most likely to benefit from joining communities if those communities align with the organization's security goals. This means it's important to first clearly define those organizational security goals. To do this, some organizations have found it helpful to use a structured approach; they can more clearly categorize their goals by mapping them to a standard risk management model, such as the "defense in depth" model. Once an organization clearly understands its own security goals, it can identify communities whose objectives align with these goals.

Because there is such a diverse range of organizations, security threats, and goals, it is unlikely that any single information-sharing community structure meets all the needs of a large organization. For example, a company might participate in one community for benchmarking and another to tackle industry-specific threats.

Information-sharing communities thrive only when the participating organizations feel they're receiving valuable information, creating incentives to continue to share information with others.

What constitutes valuable information? A common definition is that information should be timely, specific, relevant to participants' concerns, and provides a suitable level of detail while protecting individual privacy (ENISA 2010). In practice, "valuable" usually means the information helps you achieve your security goals, whether those goals are long-term and strategic, or short-term and operational. Information useful for strategic goals might include an early warning that attackers are expected to target a specific industry. This helps members of the community plan their defenses. Information useful for operational goals typically includes more specific details, such as an attack signature. This helps organizations more quickly identify an attack and respond when it occurs.

As shown in Figure 4-1 (the targeted tier), some communities consist of government agencies working alongside an industry in what are usually known as public-private partnerships (PPPs). These PPPs can be particularly important for protecting critical information infrastructure. Internationally and within many nations, this infrastructure is largely owned and operated by the private sector, including carriers and network service providers. Sharing information about threats and attacks among public and private agencies therefore can help ensure security and resiliency of this infrastructure. Because the shared information is highly sensitive, these PPPs usually have strong trust frameworks including national security clearances.

An example of a much broader public-private community is InfraGard, a partnership between the FBI and private- and public-sector organizations that shares information and intelligence to prevent hostile acts against the U.S.

Other communities are primarily comprised of private-sector organizations. Some are industry-specific: members of an industry get together to share threat information and best practices, helping to reduce risk for each company while enhancing the industry's reputation overall. Others involve sharing across industries, such as Evanta's CISO Coalition, a cross-industry group of executives from large organizations. The Coalition is designed to facilitate secure, real-time interaction among members to vet critical information security issues, and then share best practices for resolving them. As a part of my efforts to expand my external partnerships, I was fortunate enough to become a founding member of this group's advisory board. Another cross-industry group is the Security Advisor Alliance, a cybersecurity nonprofit dedicated to aligning CISOs to help one another, supporting the information security community (including startups), and giving back to schools and nonprofits.

Some communities are regional, aimed at security professionals from private and public-sector organizations located within a specific area. These regional communities offer the advantage of convenience. It takes less time, effort, and expense to attend a regional event, which makes participation more attractive. Examples of regional groups and forums include ACTRA (see sidebar) and the San Francisco Bay Area CSO Council, described shortly.

New communities arise frequently. A community may form in response to a specific threat because companies are strongly motivated to share information about the threat in order to develop effective defenses. For example, the Conficker Work Group was formed specifically to address the risk posed by the Conficker worm.

ARIZONA CYBER THREAT RESPONSE ALLIANCE

Innovative new models for information-sharing communities are springing up as the value of sharing security-related information becomes more widely recognized. An example is the Arizona Cyber Threat Response Alliance, Inc., a regional public-private partnership. This cross-sector group shares information about threats and other issues among partners from industry, academia, law enforcement, and intelligence.

ACTRA grew out of relationships developed with FBI's InfraGard, the public-sector Arizona Counter Terrorism Intelligence Center (ACTIC), and the U.S. Department of Homeland Security. A key difference is that ACTRA is a nonprofit company with a full-time president in addition to voluntary participants including a board and technical subject matter experts. The goal is to improve security for members with a flat, responsive organizational structure and without adding a burdensome layer of process. The group disseminates information ranging from alerts in near real time to white papers that provide insights and highlight best practices. ACTRA has grown to include representatives from 14 critical infrastructure sectors. The group has found, based on discussions with its members, that multi-sector sharing improves threat visibility beyond the single-sector focus of industry-specific groups.

Community Goals

Communities may focus on narrowly defined goals, such as mitigating a specific threat, or they may have broader information-sharing goals, such as benchmarking security techniques. A single community may pursue several goals. The most well-known types of goals are sharing information about threats (to help member organizations mitigate those threats) and sharing best practices (to improve efficiency). I'll describe sharing goals next.

Sharing Information about Threats and Vulnerabilities

Perhaps the best-known function of communities is to provide a trusted mechanism for sharing information about threats and vulnerabilities. Members of the community can use this information to improve their tactical and strategic situational awareness.

I'm often asked by peers how I measure the value of the information obtained from external partnerships. A key metric is whether the early threat information has helped enable us to reduce risk. A single piece of information might make participation worthwhile if it helps us better mitigate risk and protect the company.

Information from the community can also be useful for corroborating evidence that we've already identified internally. If we observe a potential new threat within our environment, we may not feel that we have enough evidence to justify taking action. But we can often discuss the issue within a community. If others are experiencing the same problem, we can be more confident that it's a real issue. This gives us enough reason to act.

Some examples of communities that share threat information include

- **Information Sharing and Analysis Centers (ISACs):** ISACs are trusted industry-specific communities established by owners and operators of critical infrastructure resources. ISACs exist for a number of industry sectors, including communications, retail, electrical utility, health, and public transit. Services provided by ISACs include risk mitigation, incident response, and alert and information sharing.
- **Bay Area CSO Council:** This is a regional community that focuses on improving the sharing of intelligence and best practices among CISOs in the San Francisco Bay Area. The Council serves as a vehicle for CISOs to safely and securely share their attack experiences. Members may share artifacts, such as attack signatures, that they can then build into their organizations' detection and defense mechanisms (Jackson Higgins 2010). The forum uses a peer-enforced trust model rather than a formal legal framework. The group also creates subgroups to work on more highly classified information.

Sharing Best Practices and Benchmarking

Many communities also serve as a forum for exchanging best practices and for benchmarking operations. By sharing security best practices, we may be able to increase the efficiency and effectiveness of our own operations.

Tapping into the expertise of others can help us avoid reinventing the wheel. A typical example: A CISO is trying to create a bring-your-own device policy for her own organization. So she sends a message to community members and receives detailed advice from others who have already been through the process. This gives the CISO a head start in creating a policy that meets her organization's needs.

Besides enabling informal exchanges, communities may also operate formal benchmarking exercises. Some of the best-known examples are the security-related programs run by benchmarking firm CEB, Inc., which conducts studies and generates reports that compare companies in a variety of areas, from user security awareness to controls maturity (CEB 2015; also see the discussion of security awareness programs in Chapter 5). Benchmarking information generated by communities can also be useful for demonstrating the efficiency of security operations to other internal groups within your organization, such as an audit committee.

Some benchmarking information is sensitive and closely held because organizations feel that it could reveal too much information about their security operations. Other information is more general and is sometimes publicly available, such as the webinars and presentations published online by Intel and others. Even this general benchmarking information may yield risk insights. Observing what other companies are focusing on, and how they are allocating resources, can help security professionals think about how they need to manage risk within their own organizations.

One of the most established communities is the Forum for Incident Response and Security Teams (FIRST). This international group focuses on sharing best practices among computer security incident response teams. Trust relationships are peer-enforced. The group publishes a series of detailed best-practices guides and other documents for public use. Other activities involve the exchange of information for cooperative incident management.

Technology is helping to make information exchange more automated and therefore easier and faster, due in part to the adoption of standards for representing (STIX) and communicating (TAXII) information about threats. Platforms are emerging that use these standards for rapid, secure information sharing.

BENCHMARKING: WHO SHOULD YOU COMPARE YOURSELF WITH?

Many years ago, I was asked to manage Intel's first major IT benchmarking activity. It was a big task that entailed analyzing cost, quality, and other aspects of operations across our entire IT environment.

One of the first challenges was determining which organizations we should benchmark ourselves against. At the time, the conventional wisdom at most organizations was that you should compare yourself with similar businesses. The logic was that because these businesses were the most directly comparable, this approach would yield the most meaningful results. So the expectation was that I'd benchmark our operations against a collection of other big high-tech companies.

But I didn't want to benchmark our operations against only high-tech companies. Instead, I wanted to benchmark against a broad base of companies in industries such as retail, banking, manufacturing, consumer goods, and utilities.

The time came to present my selection of peer groups in a meeting with senior IT management. By this time, I'd already started the benchmarking process, and as I described the diversity of the companies included in the benchmark comparison, I could sense the atmosphere becoming increasingly hostile. Practically everyone felt that my approach was completely wrong. In fact, if there had been rotten tomatoes in the room, a few people would have been throwing them at me.

So I asked for a moment of quiet so that I could explain. If we were an airline that wanted to benchmark operations, who would we compare ourselves with?" I asked. Several people said they'd benchmark against other airlines.

"What do you think we would learn from that comparison?" I continued. "My guess is not much. We'd all have grown up in the same industry, and we'd probably have similar business processes. Many of our employees would have worked for the other companies and vice versa, so they'd probably implement similar practices. We might learn about minor efficiency improvements, but I wouldn't expect any breakthroughs."

“If I really wanted to dramatically improve the way I manage airline gate operations, I’d benchmark against a Formula 1 pit crew. Those crews can service a car and get it back on the road in 20 seconds or less. I’d think about what we could learn from studying their processes, their technologies, and their ability to communicate and organize, and I’d try to figure out which aspects could cross over into airline data operations. If we want to make dramatic improvements, we need to look at people who operate in an extreme operational environment—not at other airlines.”

I’m happy to say that the managers in the room recognized that there might be value in the approach I was suggesting, even if many of them still disagreed with it. Ultimately, benchmarking against companies in a broad range of industries did help us achieve some dramatic improvements, and I received an internal award for the initiative. The lesson is that sometimes we can learn more by looking outside a narrowly defined, traditional peer group. People in the same industry may be facing the same problems as we are and dealing with them the same way. For a fresh perspective, it can be worth looking farther afield.

Influencing Regulations and Standards

All of us operate within an increasingly complex regulatory environment, and we’re all affected by evolving technology standards.

It’s important to stay abreast of legislative developments. That can be a difficult and time-consuming job for any single organization, and so it may be helpful to become involved in a community whose goals include tracking regulatory activity.

In addition, communities can sometimes help influence public policy more effectively than a single organization can do alone. There’s strength in numbers, and communities often include some of the biggest companies in an industry.

An example of a community that focuses on policy is BITS (www.bits.org), the technology policy division of The Financial Services Roundtable, which represents 100 of the largest integrated providers of consumer financial services. Members of BITS cooperate on issues such as critical infrastructure protection, fraud prevention, and the safety of financial services. The organization works to influence public policy by communicating with public agencies. It also publishes reports for use across the industry, including a financial services security assessment. Thus, communities that focus on policy may help all participating companies and the reputation of the industry overall.

Businesses who offer services in multiple countries have a particular interest in the international regulatory environment. These include multinationals, of course, which are directly affected by the complex web of regulations at international, national, and local levels.

However, these regulations affect a surprisingly large number of other companies, including many that don’t have employees or facilities physically located in other countries. Today, almost any business with a web-based service consumed in multiple countries is effectively operating in a multinational environment. Regulations in those countries have impacts that stretch beyond geographical boundaries. For example, regional and local regulations such as the California data breach bill (SB1386) and European privacy guidelines require compliance by any company that stores information about residents of those areas, no matter where the company is located.

Corporate Citizenship

At many companies a large number of employees volunteer in ways that benefit their neighborhood or a wide variety of worthy causes. Businesses often provide support to help employees do this. There's a growing trend to leverage the organization's talent and expertise in volunteer corporate citizenship initiatives that are more closely related to the organization's goals and employees' technical expertise. Examples might include offering expert security advice to nonprofits or helping security initiatives in other countries.

Security-related corporate citizenship initiatives include the National Cyber Security Alliance, whose mission is to educate and empower society to use the Internet safely and securely (see staysafeonline.org). The sponsors of the alliance include large high-tech companies such as Intel. Senior managers at those companies also are among the directors of the organization.

Conclusion

The knowledge we acquire via external partnerships can help us protect our own organizations. I've experienced this first hand; indicators of compromise shared by others have helped me understand and respond to threats. The growth of information-sharing groups shows that many organizations are coming to the same conclusion. As Ken Athanasiou, Global Information Security Director at American Eagle Outfitters, said in a statement supporting the formation of the new retail ISAC: "Cyber-criminals work non-stop, and are becoming increasingly sophisticated in their methods of attack ... by sharing information and leading practices and working together, the industry will be better positioned to combat these criminals" (Retail Cyber Intelligence Sharing Center 2015).

Industry-specific groups such as the financial and industrial control ISACs have been widely acknowledged as helping companies quickly learn about threats and specific measures for combating them. Other groups provide different kinds of valuable information. The Evanta CISO Coalition has published metrics that its members can use for security benchmarking and dashboarding. Members of IASAP share information that helps them improve their awareness programs.

The security landscape has become increasingly complex and dynamic, and it's difficult to track and manage the risks without help from others. Sharing security information is also becoming more important as organizations increasingly collaborate with business partners and adopt new technologies. Understanding the risks faced by our partners, and the way they manage those risks, can help us protect our own organizations. As businesses move into new markets and use technology in new ways, we need to understand our biggest exposures and how to allocate resources most effectively to minimize business risk. Therefore, sharing information can help businesses remain competitive and successful.

Organizations have often been reluctant to share security information, but if we want help from other people, we have to be prepared to share information ourselves. By carefully using trusted partnerships that align with our security goals, we can increase our organization's ability to sense, interpret, and act on risk.

CHAPTER 5



People Are the Perimeter

There's a difference between interest and commitment. When you're interested in doing something, you do it only when circumstances permit. When you're committed to something, you accept no excuses, only results.

—Art Turock

A few years ago, a senior manager began bringing his corporate laptop into the cafeteria at lunchtime. Typically, he'd find an empty table, set down the laptop, and then walk out of sight to get his lunch. As he perused the salads and main courses, made selections, and paid for his food, his laptop sat unattended in plain view of hundreds of people using the large cafeteria.

My security team noticed the neglected laptop and pointed it out to me. I discussed the issue with the manager a few times, but he continued leaving the laptop unattended. So eventually, I began taking the laptop and leaving my business card in its place.

Not surprisingly, the manager became somewhat annoyed. "Nobody's going to steal the laptop because there are all these people around," he said.

"Okay," I responded. "I'll never take your laptop or complain again on one condition. If you really trust everybody here, you'll take off your wedding ring and leave it on top of the laptop. If you do that, you'll never hear from me again."

He thought about this for a while. Then he said, "You made your point." And he never again left the laptop unattended.

The Shifting Perimeter

This incident helped crystallize in my mind a new perspective about how we should approach information security. It demonstrated how each person's daily decisions can affect the risk dynamics of the company overall.

The traditional enterprise security paradigm, often expressed in castle-and-drawbridge terms, described a wall of technology that isolated and completely protected the workers behind it. To protect our people and information assets, we focused our efforts on fortifying the network perimeter and the physical perimeter of our buildings.

Today, however, a growing number of user interactions with the outside world bypass the physical and network perimeters and the security controls these perimeters offer. They take place on external web sites and social networks, on laptops in coffee shops and homes, and on personal devices such as smartphones. As the Internet of Things unfolds, those interactions will also take place on many more “things,” such as wearables, cars, and even household appliances.

This changing environment doesn’t mean the security perimeter has vanished. Instead, it has shifted to the user. The laptop left unattended in the cafeteria was clearly inside the physical perimeter, but the corporate information it contained was still potentially at risk due to the manager’s actions. People have become part of the perimeter. Users’ decisions can have as much impact on security as the technical controls we use.

Over the past few years, the idea of the people perimeter has won wider recognition and acceptance. Accordingly, organizations are placing more emphasis on employees’ security awareness and behavior.

One reason for this is the rash of high-profile insider exploits, such as the leaks by National Security Agency contractor Edward Snowden. Another is that technical controls have not kept pace with the attackers. Many exploits are reaching users because technical controls, particularly those on endpoint devices, have failed to prevent them. We are therefore more reliant on the user’s ability to detect suspicious activity. We also have been forced to deploy more back-end detection and response tools and staff to handle the flow of malware penetrating the corporate infrastructure. These ever-growing security operations teams, which become another layer of the people perimeter, typically are unable to keep up with the flood of malware and commit errors due to “alert fatigue.”

There’s a continuing emphasis on phishing attacks; the *2015 Data Breach Investigations Report* found that the percentage of users deceived by phishing actually increased from previous years, with 23% opening phishing messages and 11% clicking on attachments (Verizon 2015).

Older social-engineering techniques are also still effective, apparently. At hedge fund Fortelus Capital Management in London, the chief financial officer received an alarming phone call one Friday afternoon. The caller said he was from the company’s bank, and warned of possible fraudulent activity on the account. The CFO reluctantly agreed to generate codes enabling the caller to cancel 15 suspicious payments. When he logged into the firm’s bank account the following Monday, \$1.2 million was gone. The CFO lost his job and was sued by his firm for failing to protect its assets (Chelel 2015).

As almost every company becomes a technology developer as well as a technology consumer, employee security awareness behavior will become an even bigger issue. Security lapses by the employees working on technology-based products can have far-reaching impacts, creating vulnerabilities in the digital services and physical products delivered to millions of customers.

Compliance or Commitment?

Each day, employees make decisions that can affect the company’s information risk. Do I leave my computer unattended or not? Do I post this information on social media? Do I install this software on my device? Do I report this suspicious looking e-mail? When I’m in a coffee shop, do I connect to the corporate infrastructure via a secure virtual private network, or do I engage directly over the Internet?

We could view each of these decisions purely in terms of the potential for increased risk. However, there's also a positive side. If users become more aware of security and make better decisions, they can strengthen the organization's defenses by helping identify threats and prevent impact. Among CISOs surveyed recently by best-practices firm Corporate Executive Board, 50% said that insecure behaviors cause more than half of all breaches; but they also said employees are key to uncovering suspicious activity (CEB 2015).

Therefore, as information security professionals, we are in the behavior modification business. Our goals include creating a more security-conscious workforce so that users are more aware of threats and vulnerabilities, and make better security decisions. Furthermore, we need to influence employees' behavior both within the workplace and when they are home or traveling.

If the manager was comfortable leaving his laptop unattended in our cafeteria, would he also leave it unattended at the local coffee shop? At the airport? Or somewhere else where the risk of loss was even greater? My belief is he probably would. When trying to influence this person's behavior, I wanted to achieve more than a level of compliance. I wanted to initiate a feeling of commitment.

The term *compliant behavior* implies making the minimum effort necessary to achieve good performance to a predefined standard. It's like checking boxes on a list of security compliance items. Ultimately, employees feel they are being compelled to follow someone else's list of instructions. Because of this, compliance requires supervision and policing, and employees may sometimes engage in lengthy recreational complaining. If employees are simply following a checklist, what happens when they encounter a situation that's not on the list? They stop and await further instructions, or perhaps they are even unaware of the threat or ignore it.

In contrast, *committed behavior* is intrinsically motivated and self-directed. Being committed implies that people are emotionally impelled to invest in security; they take responsibility and ownership. When people feel committed, they tend to deliver above and beyond the bare minimum. Rather than simply following a predefined list of instructions, they are empowered to make decisions and judgment calls in real time, with a focus on how their actions affect others as well as themselves.

If we can create this sense of commitment in our users, we can implement security not as a wall but as a collective security force that permeates the entire organization. Individually and as a group, every person in the corporation uses their skills in security to protect the organization, handling known attacks today as well as quickly adapting to new threats tomorrow.

When I needed to influence the manager's behavior, I looked for a way to establish this level of commitment. I sought to change the way he felt about the laptop, and to do this I tapped into his emotional connection to his wedding ring.

Creating a culture of self-motivated commitment rather than compliance can make a big difference, as shown in studies by management guru Dov Seidman. His group looked at behavioral differences between businesses with a culture of self-governance, in which an organization's purpose and values inform employee decision-making and behavior, and those with a culture of blind obedience based on command-and-control and coercion. Organizations based on self-governance experienced three times more employee loyalty and half as many incidents of misconduct, compared with organizations based on blind obedience (Seidman 2011).

The implications for enterprise security are clear. As the boundaries between personal and corporate computing dissolve, employees may be accessing information from any location, on any device. If users behave in an insecure way while they are in the office, it's likely they will also exhibit insecure behavior when they're elsewhere. Conversely, if we can create a feeling of commitment that causes them to own responsibility for security, there's a better chance they will behave more securely both within the workplace and when they are outside our physical perimeter. This change in behavior improves the security of the device they are using, the information they are accessing, their personal lives, and the enterprise.

Examining the Risks

Before discussing ways that we can modify user behavior, I'd like to briefly mention some examples of what can happen if we don't influence the ways that users think and act.

As an experiment, the US Department of Homeland Security secretly dropped disks and thumb drives in the parking lots of government and private contractors' buildings. Their goal was to see whether people would pick them up and plug them into their computers. As reported by Bloomberg News (Edwards et al. 2011), up to 60 percent of the people who picked up the items inserted them into their office computers. That number rose to 90 percent if the item included an official-looking logo. Clearly, the security behavior of employees at these facilities left quite a bit to be desired.

Insider threats unfortunately continue to make the news. A former JPMorgan Chase & Co. employee was arrested by the FBI on charges of stealing customer data and trying to sell it to an undercover informant. As noted by *CSO*, similar incidents have occurred multiple times at the bank over the past few years, illustrating the company's inability to account for insider threats despite its substantial annual spending on security technology (Lambert 2015).

Think about what can happen with newer, more sophisticated exploits. A sophisticated attack targeted government departments using fake voice-mails to distract users while malware downloaded in the background. Using social engineering and targeted e-mails, the attackers tricked users into visiting web sites harboring self-extracting archives. The archives contained a recording media file purporting to be a voice-mail from a female journalist seeking information for a news story, alongside other files that downloaded malicious content (CNET 2015).

As in the example above, today's threats may arrive in the form of carefully personalized spearphishing communications designed to win the trust of targeted users. These users then unwittingly provide access to the information the attackers want. In essence, trust—in this case, the organization's trust in the user—has become the attack surface.

Let's say a company is looking to hire a credit analyst with a very specific set of skills. Attackers notice this and apply online, using a résumé that lists the exact skills required for the job and contains the terms the company's résumé-scanning software is likely to be looking for. Suitably impressed, the company's human-resources specialists forward the application to the company's credit-department manager, who has access to all the systems storing customer financial data. The manager trusts this communication because it has been sent from another department within the same company. So she clicks on the link to the résumé. Unfortunately, that action triggers the execution of malicious code. The human-resources team effectively acted as an infection agent, ensuring the attack reached its real target.

Social media accounts can become sources of risk even when they haven't been compromised. There have been several examples in which senior executives accidentally revealed information that was confidential or problematic in other ways. In November 2014, Twitter's CFO accidentally publicly tweeted a plan to buy another company, including the fact that he wanted help to make the deal happen at a meeting the following month. The CFO was apparently trying to send the message privately (Frier 2014).

At Houston-based fashion retailer Francesca's Holdings, a former CFO frequently shared his thoughts via a personal blog, Facebook page, and Twitter feed (Silverman 2012). Unfortunately, he also shared information that caused problems for his employer. The company fired him because he "improperly communicated company information through social media."

Users frequently post information on external social-media sites that attracts the attention of competitors or the media. To boost their job prospects, interns mention product features they helped develop during their summer job at a well-known company; sales representatives reveal the names of major clients; even senior executives have been known to unintentionally disclose key corporate strategies. In fact, services exist that specialize in aggregating apparently minor snippets of information from social-media and other web sites to build an accurate view of a company's size, geographical distribution, and business strategy, including hiring patterns that indicate whether the company is expanding and which new areas it is moving into.

Adjusting Behavior

To counter these risks, we need to make employees aware and empowered so they act as an effective part of the security perimeter. Increasing recognition of this need has led to the development of a small ecosystem focused on increasing security awareness and changing behavior, ranging from companies offering best practices to groups focused on internet safety for children. This includes companies that train users to avoid phishing exploits, using simulated phishing scenarios and other tools. Security awareness professionals have come together to share best practices (see sidebar).

While I was Intel's CISO and then Chief Security and Privacy Officer, we focused on building security and privacy protection into the corporate culture, getting employees to own responsibility for protecting enterprise and personal information. Achieving this required a lot of effort, and we realized that it took just as much work to maintain a culture of security and privacy as to build it.

Training is a key part of security efforts at most companies, and Intel is no exception. We supplemented general training, which fulfills most legal requirements, with specialized training for employees who have specific roles or access to sensitive information. Another effective technique was to embed security and privacy training into business processes. When an employee requested access to an application that handles sensitive information, they were automatically prompted to take training that focused on the related security and privacy concerns. We also used online training including video and other visually stimulating material as well as entertaining, interactive tools to help engage users (see Figure 5-1).

Find the Phish: See if you can tell why these 5 Web sites are scams. Not sure? Click on the “Phish Clue” button to reveal the answers.

- 1. Scary
- 2. Got your number
- 3. Word to the wise
- 4. Key to success
- 5. Character flaw



Figure 5-1. Intel’s internal “Find the Phish” interactive training tool helps employees spot web scams. Source: Intel Corporation, 2012

However, it is not enough to create good training. If nobody takes the training, the effort is wasted. We found that incentives such as public recognition helped ensure employees underwent training and absorbed the lessons. Ultimately, if people continued to avoid security training, we escalated compliance efforts by directly contacting them and their managers.

We also found we could help maintain and increase awareness by publishing security-related articles on Intel’s primary employee portal. Many of these articles included a personal aspect, such as preventing identity theft, keeping children safe online, and home wireless security tips. The focus on personal concerns recognized that the way employees behave outside the office is as important to enterprise security as their behavior in the office.

How did we know our security efforts paid off? We accumulated a variety of evidence, including independent benchmark results from Corporate Executive Board (2011), which indicated that Intel employees consistently ranked in the top 10 percent of companies for secure behavior. We also observed that employees acted as part of the security perimeter by alerting us to suspicious text messages or e-mails they’d received.

A Model for Improving Security Awareness

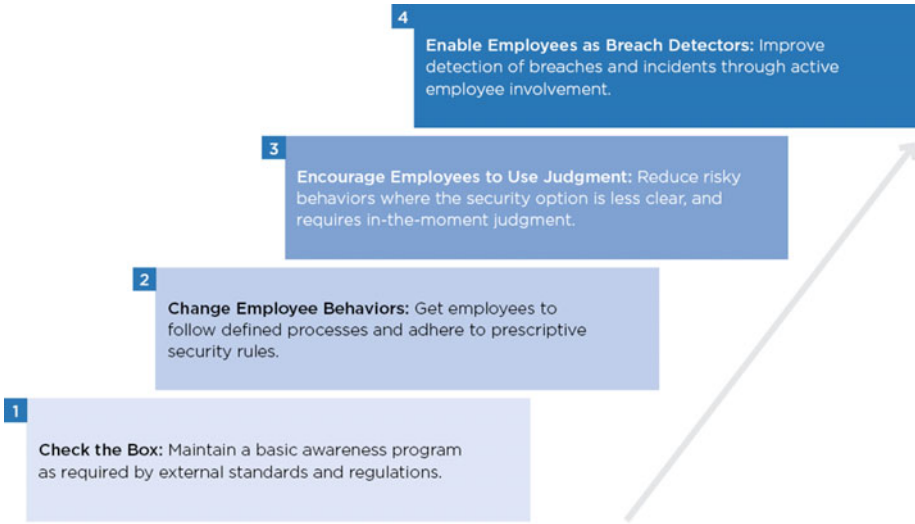
Some of the industry's most valuable work to advance organizational understanding of security awareness comes from best-practices firm CEB. The company offers a security awareness service that includes surveying key security-related employee behaviors at member organizations, and benchmarking the results against other organizations. The program attempts to understand the psychological reasons for insecure behavior; it then focuses on those psychological drivers when suggesting tactics to change employees' behavior. To date, CEB has collected some 300,000 employee responses from over 400 organizations.

The program has found that despite the importance of secure employee behavior, most organizations deliver only moderate amounts of training: just over an hour per year, on average, and only three to four employee communications per month. Survey results suggest that organizations can increase training time to as much as six hours a year before experiencing diminishing returns.

CEB emphasizes that organizations need to use an understanding of psychology to tailor their awareness efforts; awareness programs must target the specific root causes of employees' risky behavior in order to be effective. The company initially identified five psychological factors influencing security awareness and behavior: lack of knowledge of policy, lack of self-interest in security, inadequate perception of the risk to the organization, a low emotional commitment to security, and a perception that secure behavior imposes a high burden. It recently added a sixth factor: the ability to display good judgment.

CEB's findings suggest that the perceived burden of secure activities affects employee behavior more than any other psychological driver. That's the bad news. The good news is organizations can fix the perception of the burden both by reducing the burden itself and by addressing the other drivers. For example, employees' emotional commitment to security increases if their managers engage with them directly to emphasize the risks. Clear enforcement of policy compliance increases employees' self-interest in secure behavior and heightens their perception of risks.

Based on survey data collected over the years, CEB has developed a model that organizations can use to help plan and assess their security awareness programs. The model presents a four-stage progression toward higher security awareness and commitment, from basic check-the-box compliance to active involvement in security. It recognizes that in a complex threat environment, we can no longer rely only on policies that prescribe specific employee behaviors: we also need to enable employees to actively support security activities including breach detection.



Source: CEB analysis.

Figure 5-2. A four-stage model for programs seeking to improve security awareness and behavior. Source: CEB Inc., 2015

Employee awareness programs at Level 1 (Check the Box) simply respond to external regulatory requirements and don't explicitly aim to change specific employee behaviors. At Level 2, programs try to encourage users to adopt specific, simple behaviors, such as avoiding sharing passwords and sending sensitive information to their own personal e-mail addresses.

The third and fourth levels display greater levels of judgment and commitment. At Level 3, employees are able to make good judgment calls in the moment, especially in situations where the right answer is not immediately obvious. For example, they remember to pause before clicking to check whether an e-mail comes from a legitimate source or contains a phishing link. At Level 4, employees become an extension of the information security organization; they not only avoid security risks, but also notify information security when they see something suspicious. A key behavior here is an increase in reporting events such as spearphishing attempts.

The encouraging news is that that CEB surveys show a gradual improvement; the percentage of employees avoiding insecure behaviors such as password sharing has slowly increased over the past six years. Resistance to phishing has improved faster, though from a lower base. "I think at progressive companies the aspiration is changing," says CEB practice leader Jeremy Bergsman. "Most companies have been moving from Level 1 to Level 2 over several years, and are starting to think about Level 3. But progressive companies are moving beyond employee behavior as a risk to be reduced, and working on ways to make employees a control—an early warning system (Level 4)."

For example, a large telecommunications firm that participated in the CEB security awareness program wanted to empower employees to act as controls supporting information security. It provided each employee with a weekly report tracking his or her behavior, including the documents they accessed, and the devices and external locations used to connect to corporate systems (Figure 5-3). Employees were responsible for reading the reports, thus sharing responsibility with information security for detecting

breaches. The firm found that employees detected suspicious activities faster than would have otherwise been the case; users also proactively improved security by suggesting other activities that should be tracked and added to future reports.

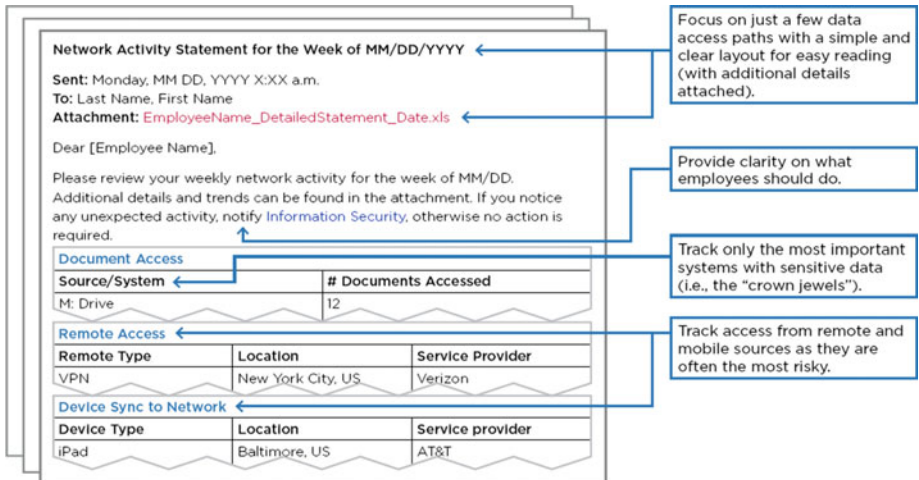


Figure 5-3. Example of weekly tracking report showing employees their activity. Source: CEB Inc., 2015

INTERNATIONAL ASSOCIATION OF SECURITY AWARENESS PROFESSIONALS

Exchanging ideas with other security professionals can help improve security awareness programs. The International Association of Security Awareness Professionals (www.iasapgroup.org) is an independent association of corporate security specialists who are seeking to do just that. IASAP is a non-profit, fee-based association dedicated solely to security awareness programs. Its goal is to serve as a trusted forum of security awareness professionals collaborating to improve employee security behavior. "Clear guidance has not been as available for employee security behaviors as it has been for technology solutions," says IASAP board member Michael Diamond. "Several awareness professionals noticed this gap, and that ultimately led to the formation of IASAP."

Some members have built programs from scratch; others inherited established programs in need of fresh ideas and new energy. Members meet in person two to three times per year to learn about other members' security programs and present their own. There's also a members-only sharing platform supporting Q&A, feedback, benchmarking surveys, member polls, guest speaker webinars, and teleconferences. Some members feel comfortable posting program resources that are available to other members for re-branding within their own programs.

Broadening the Awareness Model

I think that the CEB four-stage model shown in Figure 5-2 is a very useful tool. One limitation, in my opinion, is that the model is based on the traditional organizational view of information security. I believe that we need to expand the model to capture a more complete view of information risk. You might think of the following two additions as Levels 3a and 4a, respectively, of the model.

- Privacy awareness, which has become a critical concern. Just consider the number of breaches that have targeted personal information at retail, healthcare, and government organizations.
- A specific focus on engineers and other technology professionals, including those creating technology-based products and services. If engineers don't have a foundational understanding of privacy and security, they cannot design privacy and security into the technology they produce. As a result, a company's products may contain vulnerabilities that introduce significant risks for the business and its customers.

The Security Benefits of Personal Use

Employees use an ever-growing variety of personal devices every day, both inside and outside the physical workplace. This trend started with smartphones and laptops; it also includes wearable devices such as smartwatches and fitness monitors. Information security specialists naturally tend to focus on the security risks of using these devices for business purposes. As I discussed earlier in the book, I've found that the productivity benefits of personal devices often outweigh the risks. But even the security implications are not as one-sided as they might seem at first glance. I believe that, in some respects, allowing personal use may actually encourage better security.

In general, people are likely to take better care of their own possessions than someone else's. They feel a stronger connection to their own car than to one provided by their employer. If people are using their own computing device, they may take better precautions against theft or loss. And they may feel the same way if they are storing personal information on a corporate device. At Intel, we allowed reasonable personal use of corporate laptops, and therefore many employees stored personal as well as corporate information on their laptops. Because of this, they had a personal stake in ensuring the devices didn't get lost or stolen. I believe this sense of ownership contributed to our lower-than-average laptop loss rates.

Another company's experience provided some empirical evidence supporting this idea. The company conducted a tablet pilot deployment in which, for the first time, it allowed personal use of corporate devices. The company found that breakage and loss rates were dramatically reduced compared to its past experience with mobile devices. The CIO's conclusion was that employees simply take better care of devices when they use them for personal purposes.

Perhaps we should be similarly open-minded when considering the security implications of wearable devices. I met with managers at a large company who were pondering the security implications of smartwatches and fitness monitors, which

employees were already bringing into the workplace. Understandably, some people at the company wanted to make sure the devices could not interact with the corporate network. I observed that in the future, wearables could be harnessed to help identify users in ways that are less cumbersome for users than traditional controls such as passwords. Fitness devices, including some smartwatches, count the user's steps and monitor heart rate, and could therefore be used as biometric security devices in the future. As the devices evolve and accumulate more user data over time, they may become more adept at identifying each user's physiological and behavioral "signature." In addition, some smartphones include fingerprint recognition, which in itself can be a powerful authentication mechanism if the technology has been properly designed and implemented.

As security professionals, shouldn't we think about taking advantage of the benefits these technologies offer? We should seek to integrate into security strategies the broader variety of existing devices, which have useful features such as cameras and voice recognition and also contain data about our use patterns. Many of these devices already communicate with each other; why not take the next step and use the technology to eliminate the pain of using passwords? Why not find a way to reduce risk and cost, while providing a much better user experience, by using these devices to authenticate us automatically?

It may also be worthwhile to reexamine other assumptions about the security implications of personal devices. Some companies have policies forbidding the use of cameras in their offices. However, a smartphone includes a camera that employees can use to capture the off-the-cuff design sketches often scrawled on whiteboards during brainstorming sessions. This intellectual property can then be stored and encrypted on a hard drive within the enterprise. Is it safer to allow employees to photograph the image, or to copy it onto a piece of paper, or to leave it on the whiteboard where anyone might see it? Companies may come to different conclusions, depending on their culture and appetite for risk. But this is another illustration of the importance of considering all the possible business benefits as well as the risks when making technology decisions.

Roundabouts and Stop Signs

To try to reduce driving accidents at a dangerous curve in Chicago, the city painted a series of white lines across the road. As drivers approached the sharpest point of the curve, the spacing between the lines progressively decreased, giving the drivers the illusion they were speeding up, and nudging them to tap their brakes. The result was a 36 percent drop in crashes, as described by Richard Thaler and Cass Sunstein in their book *Nudge* (Yale University Press, 2008).

This traffic-control method succeeded in making drivers more aware and improving safety while keeping the traffic flowing with minimum disruption. I think this example provides a useful metaphor for information security. Some security controls are like stop signs or barriers: we simply block access to technology or data. But if we can shape the behavior of employees rather than blocking them altogether, we'll allow employees, and therefore the company, to move faster.

To use another traffic metaphor, a roundabout at an intersection typically results in more efficient traffic flow than an intersection with stop signs, because drivers don't have to come to a complete halt. The roundabout increases drivers' awareness, but they can proceed without stopping if the way is clear. Statistics have shown roundabouts are often safer than intersections.

Of course, we need to block access in some situations, such as with illegal web sites. But there are cases where it's more efficient and productive to make users aware of the risks, yet leave them empowered to make the decisions themselves.

Consider the case of a large multinational company whose business relied heavily on its significant intellectual property. To protect that proprietary IP, the company implemented data-loss protection software, including an application on employees' laptops. But instead of simply blocking transmission of information flagged as sensitive, the company configured the software to warn employees whenever it detected potentially insecure behavior. If an employee tried to transmit a confidential document, the software displayed a message that explained the potential risks and suggested ways to protect the information, such as encryption. After all, users may have good reasons for sending confidential documents, and preventing transmission could be detrimental rather than beneficial to the business. The company found that this warning caused 70% of users to change their behavior, representing a major reduction in risk. Yet because of the way the software was configured, users didn't complain about the security burden. The roundabout approach reduced risk without interfering with users' productivity.

Here's another hypothetical example. It may make sense to warn users visiting certain countries that they may be accessing material that is considered unacceptable. A US employee traveling on business might be working in a local office of a country with strict religious guidelines. The employee has a daughter who's in a beauty pageant, so it would be natural to check the pageant web site from time to time. But the images could be offensive in the country, so it makes sense to warn the employee to exercise caution. At Intel, we found that when we warn users in this way about potentially hazardous sites, the vast majority heed the warnings and don't access the web sites.

In the case of information security, there's an additional benefit of making controls as streamlined as possible. We all know if controls are too cumbersome or unreasonable, users may simply find ways around them. We kept this concern in mind when developing a social media strategy at Intel IT (Buczek and Harkins 2009). We were well aware of the risks associated with social media, but attempting to stop the use of external social media web sites would have been counterproductive and, in any case, impossible. We realized that if we did not embrace social media and define ways to use it, we would lose the opportunity to shape employee behavior.

As part of our initial investigation, we conducted a social media risk assessment. We found social media does not create new risks, but can increase existing ones. For example, there's always been a risk that information can be sent to inappropriate people outside the organization. However, posting the same information on a blog or forum increases the risk by immediately exposing the information to a much wider audience.

So we developed a social media strategy that included several key elements. We determined that we could reduce risk by implementing social media tools within the organization, so we deployed internal capabilities such as wikis, forums, and blogs. Initially, employees used these tools mainly to connect socially rather than for core business functions; we later integrated the tools into line-of-business applications to achieve project and business goals. We also worked with Intel's human-resources groups to develop guidelines for employee participation in external social media sites, and developed an instructional video that was posted on a public video-sharing site. The video candidly explained that Intel wanted to use social media to open communications channels with customers, partners, and influencers, to encourage people to adopt the technology, and to close the feedback loop. The information also included guidance

about how to create successful content and general usage guidelines. We also used technology to help ensure that employees followed the guidelines. We monitored the Internet for posts containing information that could expose us to risks, and we also monitored internal social media sites to detect exposure of sensitive information and violations of workplace ethics or privacy.

The Technology Professional

So far, I've focused mainly on the security roles of end users. But think about the broadening roles that technology professionals play at many companies. Historically, technology professionals have performed back-office IT roles at most companies, such as managing infrastructure and internal applications. Many also work on web sites and online services. We're now moving into a future in which companies in all industries will become creators of technology embedded in physical as well as digital products, and they'll hire developers to create that technology. These technical professionals are also part of the people perimeter, and their actions can have major positive or negative effects.

We've already seen several well-publicized problems caused by vulnerabilities in products. Fiat Chrysler recalled Jeeps in 2015 after researchers showed they could hack into a 2014 model and hijack its steering, brakes, and transmission. The researchers used an unsecured communications port to execute the attack (Dark Reading 2015). Similar concerns prompted the FDA to order organizations to stop using older drug infusion pumps made by Hospira when it was found that an unauthorized user could hack into them and change the dosage the pump delivers.

In traditional IT roles, technical professionals manage almost every element of the technology spanning our networks, data centers, and users' computing devices. They develop and install software. They configure, administer, and monitor systems. Their actions or inaction can make the difference between a system that is vulnerable and one that is reasonably secure.

Those systems include servers, which are still the IT assets most commonly attacked and robbed of data. An attacker may initially gain access to your company by compromising a user's laptop, but the biggest prize—databases of corporate intellectual property and personal information—still reside on the enterprise servers. To steal that information, attackers now typically often use a compromised end-user device to search the network for servers with inadequately configured access controls. Surveys show many attacks continue to exploit security holes that organizations could easily have fixed. Among organizations surveyed for the 2015 *Data Breach Investigations Report*, more than 30 of the exploited vulnerabilities had been identified as long ago as 1999, yet presumably not addressed at the victim organization. As the report notes, "Apparently, hackers really do still party like it's 1999."

Similar trends can be seen in the incidence of software errors. Many of the most serious, frequently exploited vulnerabilities in software are due to well-known errors that are "often easy to find, and easy to exploit," as noted in the 2011 CWE/SANS Top 25 Most Dangerous Software Errors (CWE/SANS 2011). Furthermore, the situation does not seem to be improving. As David Rice, author of *Geekonomics* (Addison-Wesley Professional 2007), puts it, most software is not sufficiently engineered to fulfill its designated role as the foundation for our products, services, and infrastructure (Rice 2007). This is partly due to the fact that incentives to improve quality are "missing, ineffectual, or even distorted," he concluded. To compete, suppliers focus on bringing products to market

faster and adding new features, rather than on improving quality. Rice estimated, based on government data, that “bad” error-ridden software cost the United States a staggering USD 180 billion even back in 2007.

Not surprisingly, the typical recommendations for improving IT security often sound remarkably familiar. That’s because they address problems already known to most organizations, but not fully addressed. As the Data Breach Investigations Report notes, the question is not which vulnerabilities should be patched (all of them should): “The real decision is whether a given vulnerability should be patched more quickly than your normal cycle or if it can just be pushed with the rest.” Previous editions of the report have recommended basic precautions such as ensuring passwords are unique; regularly reviewing user accounts to ensure they are valid and properly configured; securing remote access; increasing employee awareness using methods such as training; and application testing and code review to prevent exploits such as SQL injection attacks and cross-site scripting, which take advantage of common software errors.

The fact that these measures do not appear to be rigorously applied at many organizations takes us back to a key theme of this chapter: that the commitment of employees is as important as the policies and procedures you have in place. If administrators and developers are committed rather than just following directives, if they feel personally responsible for the security of the enterprise, and they will be more conscientious about ensuring the right technical controls are in place.

Insider Threats

High-profile national security breaches by insiders such as Edward Snowden and Chelsea Manning have made insider threats a considerably more prominent issue during the three years since the first edition of this book was published.

Among the 557 organizations participating in the 2014 Cybersecurity Watch Survey (CSO et al. 2014), 28 percent of cybercrime events were attributed to insiders. Furthermore, insiders accounted for the highest percentage of incidents in which sensitive or confidential information was stolen or unintentionally exposed.

Insider attacks also cause additional harm that can be hard to quantify and recoup, such as damage to an organization’s reputation. Insiders have a significant advantage because they can bypass physical and technical security measures such as firewalls and intrusion detection systems that were designed to prevent unauthorized access. The organization’s trust in the insider is used as the attack surface. In at least one case, the insider was the person one might least suspect: the head of information security at the Iowa state lottery, who hacked his employer’s computer system, and rigged the lottery so he could buy a winning ticket in a subsequent draw. By installing a rootkit on a lottery system, he could secretly alter the lottery’s random number generator, enabling him to calculate winning numbers in advance and buy a winning ticket in advance (Thomson 2015).

Unfortunately, even security firms are not immune to compromise; well-known cybersecurity company FireEye hired an intern who was later discovered to be a top Android malware developer. Unfortunately, his job at the security firm involved researching and analyzing Android malware, which raises the concern that he could have used his inside knowledge to develop malware capable of evading technical controls (Fox-Brewster 2015).

Yet surveys have also suggested that many insider attacks are opportunistic, rather than highly planned affairs. Many insiders take data after they've already accepted a job offer from a competitor or another company, and steal data to which they already have authorized access. In some cases, misguided employees may simply feel they're entitled to take information related to their job.

Clearly, all organizations need to be aware of the insider threat. It may not be possible to thwart all insider exploits, but we can take actions to reduce their likelihood and impact. Perhaps the biggest step we can take is to instill a culture of commitment. User behavior analytics technology can also help by detecting behaviors or access privileges that are outside the norm; perhaps technology could have prevented the case in which a former nursing assistant at an Orlando health network inappropriately accessed about 3,200 patient medical records, with no apparent motive. Besides disclosing the breach, the health network had to notify the affected patients and offer support, fire the employee, reeducate the workforce, and increase its efforts to audit and monitor access (Brinkmann 2015).

To help manage insider threats, consider a three-part approach: deter, detect, and discipline. Remember that successful implementation will require the involvement of the entire organization.

Deter

- Build security awareness and instill a culture of commitment, using the techniques discussed in this chapter.
- Make your company a great place to work. Employees are less likely to get disgruntled, and therefore less likely to seek ways to harm the company.
- Let people know you're watching. Technology can help monitor users' activity. Showing users their activity reports can help involve them in protecting the business. It also lets potentially malicious insiders know they're being watched.

Detect

- See something, say something. A committed workforce will tell you if they see something suspicious.
- User behavior analytics tools are becoming more and more effective at finding anomalies in access permissions and user activity, and identifying whether a user's actions are far enough outside the norm that they merit investigation.
- Form a team that focuses on insider threats and investigations. This should operate as a cross-functional team with involvement from human resources, legal, physical security, and information security groups.

Discipline

When an insider incident occurs:

- If it's an honest mistake without a big impact, immediate remedial training may be the best remedy.
- If the impact was low, and the incident seems more an error of judgment than a malicious act, a less heavy-handed approach may be appropriate—perhaps a written warning or a comment in the person's performance review.
- If the intent is clearly malicious, or the impact is significant, consider the options of termination and even engaging law enforcement.

Finding the Balance

One reason that organizations are focusing more attention on security awareness is that their technical controls have failed to prevent attacks from reaching employees and thus the core of the enterprise. Rapidly evolving new exploits, often involving social engineering as well as malware, have outstripped the capabilities of the security tools companies have relied on in the past.

Now, innovative security technology is becoming available that uses machine learning and artificial intelligence techniques to prevent malware much more effectively, on every type of device. This is great news for all consumers of technology. The adoption of this technology should result in a substantial reduction in risk, due to a precipitous drop in malware. The danger is that some will see this as an opportunity to dial back their security awareness efforts. I think this could be a mistake. We will always need to maintain a level of diligence and discipline in security and privacy awareness. However, we may be able to shift the emphasis of training toward prevention and future risks, and focus on how we should design, develop, and deploy technology that better protects privacy and resists attacks.

No matter how good our technical controls are, we will still need people to act as part of the perimeter. We need to create a sense of personal commitment and security as well as privacy ownership among our employees. If we succeed in this goal, we will empower employees to help protect the enterprise by making better security decisions both within and outside the workplace.

CHAPTER 6



Emerging Threats and Vulnerabilities: Reality and Rhetoric

Curiosity is lying in wait for every secret.

—Ralph Waldo Emerson

These days it's hard to read an online news source, pick up a newspaper, or watch TV without seeing reports of new threats: cybercrimes, data breaches, industrial espionage, and potential destruction of national infrastructure. These reports inevitably leave the impression that we are drowning in an inexorable tide of new and terrifying threats.

One has to question how much of this is rhetoric, and how much is reality. There are political and profit-driven motives for making threats seem bigger and more imminent than they really are. US government officials have warned that cyber attacks potentially can be “devastating, approaching weapons of mass destruction in their effects” (Levin 2010). Such warnings have been used to justify requests for increased national cybersecurity funding, as well as proposed restrictions on private networks. It's not surprising, therefore, that some experts have expressed skepticism about the real extent of the threat. In fact, academics at the George Mason University Mercatus Center have warned, “the United States may be witnessing a bout of threat inflation similar to that seen in the run-up to the Iraq War” (Brito and Watkins 2012).

On the other hand, common sense tells us new cyber threats really are emerging and growing. More data is online and vulnerable to attack, and millions of new Internet-connected devices are inevitably introducing new risks. Malware production has matured into a sizable industry. Government agencies and businesses have suffered real attacks attributed to nation-state actors: in 2014, for example, the US Government charged five members of the Chinese military with stealing information from SolarWorld and other companies, during a trade dispute over solar-energy products.

Given the flood of often-conflicting information, how can we get an accurate picture of the threat landscape so that we can develop an appropriate security strategy? How do we determine which threats directly affect our organizations, and distinguish them from those that are irrelevant? How do we decide which threats require immediate defensive measures, as opposed to those that attract attention but don't yet present significant risks?

In this chapter, I'll describe methods for identifying the real threat and vulnerability trends among the rhetoric. I'll also discuss some key areas of threat activity that have been analyzed using these methods. My goal is to help information security groups stay ahead of the attackers and focus their limited resources on mitigating the most important threats.

Structured Methods for Identifying Threat Trends

To identify the real trends in emerging threats among the mass of news and speculation, we need to carefully examine the available information using a structured, analytical approach. Unfortunately, many security groups absorb information about emerging threats using methods that are unstructured and sometimes almost haphazard.

A typical process looks something like this. The security team relies on external sources, such as news feeds and alerts, as well as informal anecdotes, to gather information about emerging threats. Based on this information, the team holds brainstorming sessions to review the threat landscape. The output from these sessions is a list of "top risks." Security resources are then focused on mitigating the items on the list.

There are several problems with this approach. Information comes from a narrow, limited range of sources, resulting in a blinkered security perspective that tends to stifle creative thinking. Also, the information is usually fragmented, making it difficult for the team to identify trends and gaps in the data. These deficiencies continue through security planning and implementation. Because the team lacks a full view of the threat landscape, it's hard to determine which threats require immediate attention and how much of the limited security budget they deserve. As a result, risks are incorporated into plans on an ad hoc basis, and not all risks are adequately mitigated. Finally, security teams often don't have a structured process for communicating threat information to other people within their organizations. Because of this, people outside the security group remain unaware of emerging risks and don't know how to respond when they experience an attack.

I realized the limitations of this approach several years ago, and began trying to inject more rigor into the risk-sensing strategy. Over time, those efforts progressively developed into a more structured risk-sensing process that helps identify threats, prioritize them, plan responses, and deliver actionable information to those who may need it. Through continued use, risk sensing can become a systemic process within any organization.

The process for analyzing emerging threats includes several valuable techniques that may be unfamiliar to some security groups. I have used a product life cycle analogy to track threats as they mature from theoretical risks into full-blown exploits. I have also used nontraditional analysis techniques, such as war games and threat agent profiles, to encourage creative thinking and identify threats that might otherwise be missed. I'll discuss these methods in more detail later in this chapter.

The process can be managed by a small core team, supplemented by a broad set of experts (including people outside the security group) across an organization. This arrangement ensures continuity while enabling the team to mine a diverse variety of sources to get a more complete picture of immediate and future threats.

Security team members should research a wide range of security topics in depth. This diversity of perspective and discussion essentially creates a crowd-sourcing of intelligence and reduces the influence of any single person's bias. Team members use typical sources, such as external feeds and analysis; they also mine academic research and hacker discussion forums, and connect with security professionals at other organizations. Other team members may scan the regulatory horizon to identify upcoming laws and regulations with potential impact, or analyze internal investigations and other near-miss incident data.

The team should hold regular meetings to analyze the threat landscape. At these meetings, each security domain expert explains his or her findings to other members of the security team. For each security topic, the discussion should include a review of recent events and a look ahead to the future. This helps identify the key trends and the factors driving those trends, provides context that can be used to analyze the current state, and predicts the likely evolution of each threat. The structured evaluation uncovers emerging risks that the team might otherwise miss. It's also useful to look back at previous predictions to see which ones were accurate, and to analyze the reasons why threats may not have materialized in the way that was expected.

It's important to communicate the findings to stakeholders across your organization in regular reports and briefings, including a wide-ranging annual assessment of the threat landscape. This communication provides further opportunities to get feedback from across the organization and its business units, which can then be used to refine your risk-sensing analysis.

The Product Life Cycle Model

I have found that a product life cycle model is a useful way to track and prioritize emerging threats as they evolve and begin to present real risks to the enterprise. Almost all security groups have a limited budget, so they need to focus their resources on effectively mitigating the highest-priority threats.

This model, shown in Figure 6-1, recognizes that many threats initially emerge as theoretical risks, but are on a path to exploitation, and we need to evaluate and monitor them.




Figure 6-1. The product life cycle model for tracking the evolution of threats.
 Source: Intel Corporation, 2012

Often, researchers or hackers first reveal a possible attack or vulnerability at a security conference or publish information about it online. Next, attackers begin testing the use of this technique and making their results publicly available. Once the method has been proven, the threat enters the production phase as attackers start exploiting it in earnest. Ultimately, the threat becomes a mature commodity—source code is often freely available, many variants exist, and organizations treat the threat as part of the everyday landscape and build defenses accordingly.

This life cycle model enables security teams to systematically track the evolution of threats. It helps us determine when to allocate resources to fighting each threat. As each threat approaches maturity, we can examine how it is likely to affect our organizations and plan appropriate mitigation.

In addition, this model provides a great way to communicate actionable information to business groups using terminology they already understand (the product life cycle). When we provide regular threat landscape assessments to stakeholders, each security topic should include a description of the activity at each life cycle phase, thus providing a context that helps the security team inform business groups about how they should act on each of these emerging risks.

Let’s examine some examples showing how this model can be used in real life. Figure 6-2 illustrates the evolution of threats targeting smartphones and other handheld devices. Researchers and hackers began to take notice of handheld devices almost a decade ago, demonstrating weaknesses and theoretical avenues of exploitation. Initially, they focused on what were then known as personal digital assistants. As smartphones

took off, attackers shifted their attention to this bigger market, which rapidly became a major area of threat activity. Monitoring trends at these earlier stages enables organizations to prepare. As threats mature and employees begin using smartphones more widely at work, well-prepared organizations are in a better position to develop risk mitigation measures including technical controls and incident response plans.




Figure 6-2. How an organization could use the product life cycle model to track and respond to smartphone security threats

By visually comparing activity across multiple threat areas, as shown in Figure 6-3, we can quickly identify major areas of activity and see the likely timing and extent of their impact. This chart also shows us areas in which there are numerous proof-of-concept tests and other activities that suggest major problems in the near future. And it indicates areas of focused research that may ripen into active exploitation over the long term.




Figure 6-3. A visual comparison of security-related activity across different technology areas. Data are for illustration purposes only. Source: Intel Corporation, 2012

Although the depth of detail in Figure 6-3 is valuable to the security team, I have found a simpler, consolidated view such as the chart in Figure 6-4 can help communicate the essential trends to a broader audience, supplementing other threat analysis materials. These simpler charts are based on the activity identified using the product life cycle model, but add further trend analysis and group the activity areas into four main clusters, depending on their level of activity and maturity potential and on their potential impact to the company. These clusters are

- *Sustained drivers:* These are areas that already have a high impact or otherwise cause considerable concern. Typically, they are characterized by commoditized distribution and active exploitation by multiple threat agents. Today, examples include malware and web attacks.
- *Critical trends:* These areas have begun undergoing active exploitation, with growing adoption beginning to shift toward commoditization. Current examples include social computing and smartphones.
- *Emerging trends:* These areas have a low current level of exploitation, but considerable research and proof-of-concept activity. Examples include embedded and cloud computing.

- *Disruptive trends:* These are areas with little or no active exploitation, but significant research activity and the disruptive potential to cause a major security problem. Frequently, they are discussed as theoretical risks, and because of this, many people in the industry would be caught off guard by a significant event. Examples include virtualization, an area in which potential threats and vulnerabilities have been exposed and a successful exploit could cause far-reaching damage.

I have found that clustering threat analysis information in this way enhances communication with stakeholders. Representing the information in easy-to-understand charts helps to convey the key trends and their potential impact to a broad cross-section of people, helping them quickly assess whether they need to make adjustments to security strategy.

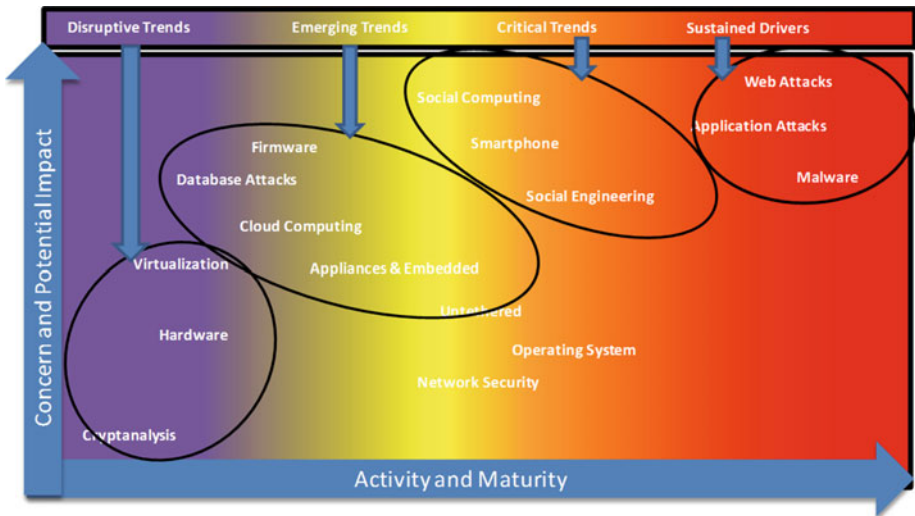


Figure 6-4. Clustering areas of threat activity to highlight trends.
Source: Intel Corporation, 2012

ASSESSING HOW TO RESPOND TO A NEW THREAT REPORT

A continuous stream of new threat reports emerges from agencies, intelligence services, and vendors. It can be hard to determine what to do with all the new information—especially since most security organizations have limited resources. Here are five questions you can ask yourself the next time you see a published threat report.

1. Are we immediately affected? Are the indicators of compromise shared in the report found in our environment? If so, we have an incident that we must deal with.
2. If we're not already affected, what is the likelihood that we will be a future target? We're more likely to be targeted if we work in the same industry as a previous victim, or if we are connected to them in another way (as a supplier, customer, or partner). If the attackers are hactivists or politically motivated threat actors, we are more likely to be targeted if we align with the victim's philosophy. Note that we may be a target even if there's no obvious linkage to the victim.
3. How were the victims attacked? What compensating controls do I have in my security stack to mitigate the risks across the kill chain of a similar attack?
4. Have we seen the same malware used, or families derived from it, against our assets?
5. Were any interesting tools, techniques, or procedures used that I should capture and share with my security team? This part of the report can be used to educate responders, architects, and risk managers so they can make better decisions.

Based on a blog post by Steve Mancini, Director of Information Security at Cylance (Mancini 2016).

Understanding Threat Agents

Besides the product life cycle analogy, there are other techniques that can help us think creatively about threats and identify risks we might otherwise miss.

Behind every threat is a human agent. To effectively plan defenses, it helps if we can understand why and how these agents operate: their motives, typical methods, and targets. However, I realized several years ago that we lacked agreed-upon definitions of threat agents, as well as a clear understanding of which agents actually pose the biggest risks to us.

Some agents and their activities attract considerable publicity, resulting in the "TV news effect" in which the most-publicized agents appear to be the biggest threat, so they often receive a disproportionately large percentage of limited mitigation resources. In reality, a wide spectrum of threat agents exists, some of which may be less well-known but pose bigger threats. For example, hactivists often want to publicize their activities as much as possible to draw attention to their cause. This publicity makes them appear to be a bigger threat than other groups, such as organized crime syndicates, which try to conceal their exploits.

In addition, terms often are used without clear agreement about what they mean. The phrase *advanced persistent threat* has become a buzzword whose exact meaning depends on who is using the term. It usually implies adaptive, long-term strategies employing a

variety of stealthy techniques and used by attackers with considerable resources. However, it's important to remember that a variety of agents may be capable of generating this type of threat. One thing that all these threat agents have in common is the use of malicious code to achieve their goals. But to understand and predict their likely motives and methods, it is useful to clearly define the agents, whether they represent nations or other powerful groups, such as organized crime. To solve this problem, Tim Casey, a member of my security team at Intel at that time, developed a standard threat agent library that provides a consistent, up-to-date reference describing the human agents that pose threats to our information assets (Casey 2007). The library helps risk management professionals quickly identify relevant threat agents and understand the importance of the threats.

The library acts as a collection point for information about each agent, making it easier to share information across your organization. It includes profiles of agents such as disgruntled employees, opportunistic employees, industrial spies, and politically motivated attackers. The library also catalogs agents' typical targets, objectives, skill levels, current activity, and exploit outcomes. When used as part of regular threat assessments, this model can help determine which agents pose the biggest risks to your organization. The security team can then use the information about their typical methods and exploits to help plan its strategy. The library helps the team understand why specific events and attack trends occur and what might happen next.

NSA'S CHIEF HACKER EXPLAINS HOW TO DEFEND AGAINST THREATS

It's hard to imagine someone who is better placed to provide advice about defending against advanced adversaries than Rob Joyce, who heads the National Security Agency's Tailored Access Operations (TAO) elite hacking unit. So the audience listened closely when he took the stage for an eye-opening talk at the 2016 Usenix Enigma conference. "My talk is to tell you, as a nation-state exploiter, what can you do to defend yourself to make my life hard," he said.

Joyce said that six intrusion phases comprise what is typically referred to as the "kill chain:" reconnaissance; initial exploitation; establish persistence; install tools; move laterally; and collect, exfiltrate, and exploit the data. Organizations can thwart attackers by disrupting the transition between any of these phases. For example, to help prevent reconnaissance turning into initial exploitation, you can reduce the attack surface by locking down or disabling devices that are unused or don't need to be open to access. "Don't assume a crack is too small to be exploited," he said. "We will look for that esoteric edge case."

Contrary to popular belief, advanced adversaries don't rely exclusively on zero-day exploits, Joyce added. Most intrusions occur via easier vectors: e-mail, web sites (using techniques such as waterholing—infesting web sites that are frequently accessed by users at the target organization), and removable media like USB drives. Joyce noted that you can't rely on users not to click, even with the best security policies and education (see my Irrefutable Laws of Information Security in Chapter 1), so you need technical controls that will prevent the execution of malicious code.

Once advanced attackers have established a beachhead, they try to steal credentials that enable them to maintain a presence, install tools, and move laterally to the prized assets they seek. Techniques such as segmenting the network, limiting administrator privileges, and forcing two-factor authentication can help make this more difficult. Joyce also said that he liked some of the new ideas emerging from the industry such dynamic privileges, which is analogous to the granular trust model described in Chapter 7: the level of access provided depends on factors such as the device you're using and your location.

Finally, he stressed the need to continually evaluate and improve your defenses. An organization with static defenses will drift to the back of the herd, where it is easily picked off by a predator (see Irrefutable Law #6). "Don't be that easy mark," he said.

Playing War Games

I like to conduct war games a few times a year. War games are intense role-playing exercises in which employees take on the role of attackers and attempt to compromise key assets using any feasible methods. I have found war games are particularly valuable for analyzing threats that may have major consequences but whose vulnerabilities are not well understood.

This technique provides the most comprehensive method of assessing threats to key assets because the people playing the role of our adversaries are essentially allowed to use any method to achieve their goals. However, because of this, it is also resource-intensive and should be used selectively.

Typical war games that I have overseen take one and a half days and may involve eight to ten staff from a variety of roles, such as factory workers, business process leads, salespeople, and technical experts. Some war games can take much longer; in *Wargaming for Leaders*, written by wargaming experts at management-consulting firm Booz Allen Hamilton, (Herman, Frost, and Kurz 2009), the authors discuss games that may last weeks and involve many more players across an organization.

A typical game focuses on a specific target or scenario, such as disabling a key facility or stealing trade secrets. You can use war games to examine potentially catastrophic events that have a low probability of occurrence, but a high probability of causing damage if they do occur. Team members are instructed about the threat agents involved and draw on archetypes from a threat agent library or descriptions provided by the game architect. Led by a facilitator, the team takes on the attacker's perspective and postulates ways to achieve the attack's objectives.

Because the team can propose any attack method, they often identify risks that might be overlooked using conventional methods. As the authors of *Wargaming for Leaders* put it, "We create the environment, the players engage, and what comes out of team play often surprises and even stuns everyone involved." For example, a malicious group might attempt a devastating attack by purchasing a small but essential technology provider and inserting malware into their products in order to infect their customers. After each game, security analysts examine the results to determine how to address newly identified vulnerabilities.

I also like to examine the cyber consequences of large physical events as part of disaster recovery planning. These could include earthquakes and tsunamis that damage data centers, or even solar flares that disrupt the communications that the business relies on. Exercises can include drills that last a day or more.

A large organization can justify the considerable effort involved in conducting these exercises because of the enormous potential benefit of mitigating the threats. In fact, some organizations hire professionals to create and facilitate these games. Booz Allen Hamilton, for example, has an extensive war gaming practice covering diverse subject areas including market dynamics, cybersecurity, geo-political events, and even real war scenarios.

But smaller organizations can also benefit by considering extreme events and formulating response plans. If you prepare for the extreme, you'll be more prepared to deal with everyday events. Planning doesn't need to be as resource-intensive as a full-blown war game. It can be as basic as bringing team members together to discuss likely scenarios and responses in a shorter tabletop exercise lasting just a few hours. This method enables members to get a feel for what it would be like to work together in the event of a real disaster. Considering these extremes can also provide motivation for introducing simple yet effective measures to reduce the risk that catastrophes will occur. You might realize it is worth increasing investment in user education to reduce the risk of social engineering compromises, or becoming more diligent about analyzing logs and network traffic to identify patterns that indicate suspicious activity.

Trends That Span the Threat Landscape

I've described some of the methods that can be used to analyze emerging threats. Now I'd like to turn to some key themes that have emerged from such threat analysis. These themes paint a broad-brush picture of threat and vulnerability trends spanning multiple technologies across the threat landscape.

Trust Is an Attack Surface

As the technology industry erects new technical defenses, attackers seek to bypass these controls by exploiting user trust, typically using social engineering techniques such as phishing.

If an attacker can win a user's trust with a sufficiently convincing e-mail or fake web site, the user will make it easy for the attacker by clicking a link or downloading a file. These actions usually undermine even the most rigorous system-level controls, initiating a chain of compromises that ultimately can result in major damage.

Whenever users place their trust in a new technology, attackers quickly follow. Studies have shown that users trust social media services more than other information sources. A user is more likely to click a link if it appears to have been sent by a social media "friend." Exploiting this trend, attackers have spread malware via social computing circles of trust such as friend networks.

Attackers have also been quick to take advantage of the trust users place in their smartphones and in other appliances such as game consoles. The exploitation of trust also extends to the relationships between systems. Once configured, communications

between systems often operate autonomously, without manual oversight. Smartphones are set to automatically update applications from trusted app stores; other systems blindly trust firmware updates and dutifully install them. This automation provides convenient opportunities to insert malicious code, abusing trust without the need to directly involve the user.

In the near future, I anticipate trust will become a commodity that is bought and sold. The digital reputation of systems and services will become critically important. In the past, tokens of trust, such as digital certificates and social computing credentials, were stolen for immediate use. In the future, they will be stolen so they can be sold in underground markets. The value of these tokens depends upon the access they grant and the other circles of trust they can be used to penetrate. Already, attackers are using stolen digital certificates to sign their malware in an attempt to avoid detection by operating system defenses.

I expect social engineering attacks will continue to present significant risks because they exploit human weaknesses and will adapt to take advantage of new technologies. So we, as security professionals, need to focus on the role of users as part of the security perimeter, as I discussed in Chapter 5. To reduce the risk to the enterprise, we need to make users more security-aware and influence them to act in more secure ways. But it's also important to note that a successful phishing exploit is also ultimately a technology failure that allowed malicious code to execute.

Barriers to Entry Are Crumbling

Our adversaries gravitate toward the path of least resistance. They tend to select targets that are easy to access and analyze, and they typically use the most readily available and cheapest tools.

They are much less likely to use methods with high barriers to entry such as the need for specialized expertise, expensive hardware or software, or access to extensive compute capacity. However, several of these barriers have begun to crumble as a result of trends such as cloud computing, lower-cost communications components, and commodity malware toolsets. This trend ultimately is likely to result in new types of attack.

A key factor is that security researchers are sharing not only their knowledge but also the tools they design as part of their research. Recently publicized tools, such as rogue base stations and Bluetooth sniffers, provide attackers with more accessible, low-cost ways to intercept network traffic. Researchers have uncovered vulnerabilities in femtocell devices (miniature, low-cost cell towers) that can be used to take control of the devices, lowering the barriers to attacks targeting cell phone data traffic.

Ultimately, lower barriers to entry mean increased risk to enterprises. However, because several of these areas are still at the research stage, it will take time for them to mature into active exploitation.

The Rise of Edge Case Insecurity

Each day, the environment becomes more complex with millions of new devices, each running its own operating system and collection of applications. This complexity generates new edge cases—problems or situations that occur only in unexpected or extreme situations.

Edge cases can include unlikely interactions between two familiar objects. A hacker team recently demonstrated that, with a popular smartphone, a paperclip (used to pop out the phone's SIM card at the critical moment), and a little patience, it's possible to gain access to contact information, phone call logs and voice mail, e-mails, and other information stored on the phone.

Overall, the growing number of third-party plug-ins and widgets introduce edge cases that are hard for developers to anticipate even if they use secure design techniques.

Interoperability between programs has resulted in a new category of hybrid attacks where malicious objects are concealed in innocent-looking ones to thwart detection. One proof of concept in 2011 demonstrated it was possible to conceal a fully functioning Trojan in an e-mail plug-in.

Some of these hybrid attacks have shown they can circumvent new security features. As web browsers and search engines try to protect users from malicious links, attackers are responding by hiding links in image search results, where they cannot be detected using standard tools. Research into network intrusion methods has discovered over a hundred methods of evading detection by manipulating traffic to remain functional but undetectable by typical tools.

There is no silver-bullet solution for eliminating edge-case insecurities. It's unlikely that even the most rigorous testing could ever uncover them all. The best approach may be to exercise caution when adopting new technologies with the potential to generate edge cases.

The Enemy Knows the System

The technology industry has often relied on security through obscurity: the idea that if attackers can't see the insecurities in code or other technology, they won't exploit them.

Over time, it has become clear that security through obscurity is poor security. To quote the maxim coined by Claude Shannon, one of the founders of modern computing, "The enemy knows the system."

It's now relatively easy for attackers to get access to the same tools enterprises use, such as web hosting services and smartphone application development tools. Hackers can now more easily engineer malware and attacks that take advantage of these elements. The fact that static platform controls tend to become less effective over time (one of the Irrefutable Laws of Information Security noted in Chapter 1) is partly due to the ability of malware authors to pretest their malicious code against technical controls. They can do this by obtaining code from malware repositories that have already been tested against existing controls, or by actually purchasing the technical controls.

Even the success of social engineering demonstrates that the attackers' knowledge of the target greatly increases the likelihood of successful deception. Today, competitors and other threat agents learn a great deal about a company and its employees by simply searching information publicly available on web sites or social media accounts.

Because we cannot assume insecure technology is safe just because it is hidden, we need to design with security in mind. The ineffectiveness of security through obscurity is also an argument in favor of standards and open-source solutions. This idea may initially seem counterintuitive, but the fact that open source is exposed to public scrutiny requires it to be secure. At a minimum, we should ensure devices are rigorously tested against industry standards because the attackers will do so.

Key Threat Activity Areas

Threats are evolving in many technology areas, from embedded systems to cloud computing. I'd like to discuss a few areas experiencing significant developments with implications for enterprise IT.

The Industry of Malware

Malware has become a profitable industry that increasingly resembles the legitimate software market, with market leaders, mergers, licensing agreements, real-time support, and open source. The organized business activity in this market reflects the extent to which well-crafted malware has become a viable career pursuit for members of the criminal underground.

Today, malware development and malware use may in some cases be distinct activities carried out by different groups or individuals. Malware authors are producing standardized toolkits, which have made life much easier for would-be attackers. These attackers can now simply buy or acquire a toolkit rather than expending the effort to identify vulnerable web sites and develop their own exploits.

The Zeus malware family provides a useful case study showing how complex this industry has become and how hard it is to accurately track developments. Sold mainly in underground forums, Zeus has been used extensively for theft by creating botnet nodes. During 2011, a code merger was reported between Zeus and another popular crimeware kit, complete with assurances of future support for the customers of both products. Around the same time, Zeus toolkit source code was made publicly available. Since then, multiple new variants have appeared and been used for a variety of attacks. At one point, security researchers attempting to monitor Zeus exploits discovered a server they believed was the hub of a Zeus botnet. However, the server was the equivalent of an espionage honey pot, allowing the botmasters to turn the tables by spying on the researchers who were attempting to analyze the hub.

Ransomware has also become a profitable activity for some organized crime elements. Ransomware was mostly at the validated proof of concept stage when I wrote the first edition of this book in 2012; it has since progressed to active exploitation with some commoditization. Today's ransomware exploits typically exploit system vulnerabilities using Trojans and other methods, then lock or encrypt information so users cannot access it and hold people and organizations hostage until they pay. In February 2016, a Los Angeles hospital paid a ransom in bitcoin after staff were locked out of the hospital's own network for more than a week; during the same month, one ransomware variant was reported to be infecting more than 90,000 PCs per day (Fox-Brewster 2016).

The Web Expands to the Internet of Things

The Web continues to present a huge attack surface. And this attack surface is growing rapidly as it expands to include the Internet of Things, encompassing nontraditional devices such as appliances and control systems, cars, wearable and medical devices, and the "smart" grid. Each of these is a potential source of risks.

Recent headlines have highlighted the growing threat activity focused on IoT. Researchers hacked into a Jeep via its Internet-connected entertainment system and remotely controlled the vehicle's functions, including turning off the transmission and brakes while someone was driving (Greenberg 2015). Other researchers showed that thousands of devices in hospitals are vulnerable to attack, including x-ray machines, MRI scanners, and drug infusion pumps, partly because medical equipment is increasingly connected to the Internet so that data can be fed into electronic patient records systems (Pauli 2015a). Yet another researcher demonstrated the ability to hack into FitBit fitness trackers via Bluetooth (Pauli 2015b). Many IoT devices, including cars, wearables, and home appliances, include wireless capabilities, so exploitation doesn't require a physical network connection.

Clearly, we should expect continued growth in IoT threat activity. However, should be noted that the activity to date has generally been at the research or early proof-of-concept phase (see Figure 6-1). As the IoT expands and matures, we will see a progression to advanced active exploits over the next few years; given the rapid pace at which IoT is evolving, if companies don't use good privacy and security design principles when building their products, the time from research to active exploitation could be much shorter than has typically been the norm.

Many embedded devices that are already installed in businesses are similarly vulnerable. Companies have a history of deploying specialized devices without adequate security controls, often because of the perception that specialized devices are "dumb" and do not have a full set of capabilities. In reality, the opposite is often true: devices marketed for a specific function are often capable of much more. Printers contain processors, use wireless connections, and may be capable of acting as file servers, for example. As a result, embedded devices can introduce as much risk, or more, to an organization as a traditional computing device since they lack security controls and administrators are generally unaware of the danger. New devices may be vulnerable to new attack methods: recent research showed that the sounds 3D-printer nozzles make as they cross the machine bed can be recorded by smartphones, analysed, and then used to duplicate prototypes (Nelson 2016).

The vulnerabilities in embedded industrial control systems were exposed by the widely publicized Stuxnet malware, which was used to sabotage the systems that supported Iran's uranium enrichment capabilities. The incorporation of computer-based control and automation technology into the existing electrical power infrastructure—resulting in the "smart grid"—is another source of potential vulnerabilities. The US government has warned of increasing threats to the grid, noting that many embedded systems lack adequate security controls and are susceptible to known techniques such as cross-site scripting attacks (US GAO 2012).

We might also see logical attacks as precursors to physical attacks. On a macro scale, a nation state might attack another nation's cyber infrastructure before staging a physical attack. This approach might also be applied at a more personal level. A burglar might remotely disable an Internet-connected alarm system before sneaking into a house, or perhaps even use the system's video cameras to watch the owners and note when they leave the house unattended.

Here are two more potential future IoT scenarios in which innovative technology designed to do good could be exploited for harm, unless designed with strong security and privacy protection. Last year, doctors for the first time inserted an artificial "eye" that enabled a blind person to see. The device is a retinal implant that receives signals from a video camera integrated into eyeglasses. Think ahead a few years, to a time when the implants are more sophisticated and can see in much higher resolution, and also include

software to automatically interpret visual information, such as QR codes. Then imagine that a malicious actor creates a QR code that triggers the vision system to download malware. Like the PC malware that paralyzed Sony's network in 2014, the malware then demands a ransom to re-enable the person's vision. Now consider the example of a cement company that's embedding sensors in the concrete mix used to build a new road, thus enabling local authorities to monitor traffic patterns and adjust signals to optimize the flow of vehicles. If the technology is not securely designed and implemented, all that a malicious person needs is the ability to execute malicious code, in order to falsify the traffic pattern in such a way that vehicles converge on the scene of a planned bomb attack.

Smartphones

Smartphones are attracting almost as much malicious interest as desktop and laptop platforms. However, even though smartphone sales have outstripped PC sales, smartphone malware isn't yet as prevalent as PC malware and doesn't cause the same kind of widespread damage. That's partly because most valuable corporate and personal data is still held on PCs and servers. Another factor is that smartphone vendors have somewhat greater control over applications, since users generally access them via vendor-controlled app stores.

Just as in legitimate software markets, malware authors are likely to maximize the value of their code by using tools that allow their software to run on multiple devices. They are increasingly targeting applications, a trend also seen on other platforms. Attackers have purchased copies of applications, incorporated their malicious content into the otherwise legitimate software, and then redistributed their code under a new name or as a "free" version of the original. On one smartphone platform, autodialing malware was found in more than 20 applications. Variations of a Trojan were found in dozens of applications and are believed to have been downloaded by at least 30,000 users.

A further development is the use of smartphones as bridges to traditional networks, resulting in the potential for enterprise network attacks that originate from within mobile networks.

In the future, we could see greater exploitation of location-based services to deceive users. Because smartphones contain location sensors such as Global Positioning System (GPS) chips, knowledge of the phone's location can be used to present targeted ads and useful information. For example, a user in a supermarket aisle might be presented with online coupons for products on nearby shelves. But this information could also be exploited to present fake coupons that are all the more convincing because they suggest that the sender knows the user's preferences.

Attackers could also exploit other smartphone capabilities to take advantage of the fact that the devices are carried into confidential meetings and other highly sensitive situations. Imagine being able to remotely control a device that has a microphone, a camera, or other recording capabilities. Or think about a vulnerability in any of the popular web-conferencing services that people use for confidential discussions and to exchange information.

Current trends in the mobile platform space indicate that attackers are most interested in stealing personal data. This trend is partly due to the increasing use of smartphones for financial and banking transactions, which provides new opportunities for identity thieves and other criminal groups. As a result, it is now important that smartphone hardware and software developers focus on protecting personal data.

Software developers should adopt the same discipline and commitment to following secure design principles as traditional platform developers. Today, more and more people are becoming app developers, creating software, and posting it online for others to use. One has to question how much security testing and validation has been applied to these applications. As users move more of their everyday activities onto smartphones and other small devices, the consequences of poor or insecure designs will have greater impact on individuals and their employers.

Web Applications

Web applications, primarily comprising client browsers and server-based applications, continue to be heavily attacked. Threat analysis indicates that this area is experiencing full exploitation activity and moving toward commoditization. There is also considerable research in this area, suggesting the number of attacks will continue to grow.

Attackers have adopted new techniques to hide their intentions and deceive users long enough to achieve their aims. As web browsers and search engines try to protect systems from malicious links, attackers are instead obfuscating their links in image search results, where they may not be detected.

Techniques for hiding messages within images have been used within the security realm since long before the invention of information technology. Now, this technique, known as *steganography*, is being used to hide malware and botnets on publicly used image hosting sites.

Search poisoning has also become a common method. Attackers using search poisoning tend to focus on events and topics of popular interest, optimizing their web pages to achieve high search engine rankings. After a search query, the victim clicks a link among the search results. They are redirected multiple times and eventually land on a page that is used as a vector to deliver malware.

Conclusion

In this chapter, I've outlined some of the real threat trends and described methods information security groups can use to analyze the threat landscape as it continues to evolve.

No doubt, new and more sophisticated types of exploitation will continue to emerge, and we need to stay aware of them. As Mustaque Ahamad, director of Georgia Tech Information Security Center, noted in 2011, "We continue to witness cyber attacks of unprecedented sophistication and reach, demonstrating that malicious actors have the ability to compromise and control millions of computers that belong to governments, private enterprises, and ordinary citizens."

Yet, as we try to make sense of the deluge of news about attacks and vulnerabilities, it's essential to retain a sense of perspective. Most threats do not take place using exotic, obscure methods. Instead, they take the path of least resistance, exploiting well-known vulnerabilities. Therefore, business can mitigate many of these threats by implementing basic, established security measures. To put it another way: when you hear hoof beats, think horses—not zebras.

Social engineering will continue to be a key attack method because it takes advantage of user trust and is hard to prevent using technical controls. Therefore, as I discussed in Chapter 5, we need to continue to focus on educating users to become more security-aware. By doing so, we can reduce the risk to the enterprise.

Ultimately, while doing our best to prevent compromises and breaches, we must remember we cannot control the threat actors and their exploit attempts. Because all threat categories use malicious code in some way, advanced preventive tools that effectively stop the execution of malicious code can greatly reduce the potential of compromise. But all organizations face the possibility of some level of compromise, making defense in depth as essential as ever. Losers ignore the trends. Winners survive by being able to predict, prevent, detect, and respond.

CHAPTER 7



A New Security Architecture to Improve Business Agility

An organization's ability to learn, and translate that learning into action rapidly, is the ultimate competitive advantage.

—Jack Welch

Some *Star Trek* episodes feature suspense-filled battles in which adversaries use sophisticated phase-shifting weapons that can be rapidly adjusted until they find a way to penetrate static force-field defenses. For a beleaguered starship, the only effective response is to use similarly adaptable and fast-changing shields.

As information security professionals, we also need extremely agile defenses that quickly adapt to meet new demands. Attackers are continually adapting, and defenders also need to continually adapt. But rapidly evolving threats are only part of the challenge. We also need to continually adapt our defenses to a rapidly changing technology landscape.

As information risk and security groups consider the future, it's clear that we need to radically change our approach in order to face the challenges ahead and support the Protect to Enable mission.

One problem in recent years has been that most of the protection offered by the industry has not kept up with the attackers. Because these tools have failed to prevent harm, many companies have defaulted to a detect-and-respond approach. This means they continue to expose themselves to high risks and higher long term costs since they are reactively responding to attacks that have already breached the organization's defenses.

We also need to consider whether our existing control architecture improves or impedes business agility and velocity. It's important to recognize that controls can place a "drag coefficient" on the business. By hindering users, they can stifle business velocity and innovation. Users react to this control friction by circumventing the controls whenever possible; as a result, the controls can actually introduce new risks, as discussed in Chapter 2.

As we move forward, we will need an agile security architecture that quickly and automatically learns and adapts to new challenges as they emerge. A learning system is harder to defeat because it can more quickly predict and thus prevent new attacks. The pace of change is so rapid that we cannot predict all the challenges we will face, and manual or semi-manual processes will not be enough to keep up. We will need solutions that can learn to manage what we don't know.

The right control architecture will enable flexibility that helps the business move more quickly, allowing us to rapidly adopt new technologies and emerging usage models while continuing to provide security in the ever-evolving threat landscape.

A few years ago, after intense brainstorming sessions, the information risk and security team I led at Intel devised a new security architecture for the company. This architecture represented our implementation of the Protect to Enable strategy, using the technologies that were current at that time. With the benefit of hindsight, I believe that we got many things right—but there were also some omissions because we didn't have a full understanding of the controls that would be needed.

In this chapter, I'll provide a high-level overview of a new security architecture and describe how it meets some key challenges. Some of this overview is based on the work at Intel a few years ago, but I have added a new perspective on controls that I have realized is lacking in the industry. An important aspect of this new perspective is the concept of control friction. As I'll explain later in the chapter, I've developed a simple framework called the 9 Box of Controls, which takes control friction into account when assessing the value of security controls.

I believe that the architecture includes some novel approaches that may be valuable to many organizations facing these universal challenges. My conversations with peers at other companies have validated this view. Many of them are considering similar strategies and in some cases have begun implementing them.

Any future security architecture must provide better prevention, and it must also be more flexible, dynamic, and granular than traditional enterprise security models. This will help us all accommodate future evolving usage models. We can provide users with different levels of access depending on factors such as the devices they are using and their location. To achieve this, the architecture dynamically adjusts a user's access privileges as the level of risk changes. For example, an employee should have more limited access to our systems when using a less-secure device than when using a more hardened or perhaps fully managed enterprise-class system.

The new architecture greatly improves threat management. As new attacks appear, we need to be able to recognize good from bad in milliseconds, so that we can stop the bad and allow the good. For any attack that gets past the preventive controls, we need to learn as much as we can without compromising the user's computing performance or privacy. This information enables us to investigate what occurred, so we can quickly take action to mitigate the risk and also learn how to prevent similar attacks in the future. A control architecture should assume that attempts at compromise are inevitable, but we should also understand that it's possible to achieve real prevention for 99% or more of malicious code. We can apply artificial intelligence and machine learning to analyze the features of files, executables, and binaries to stop malicious code prior to execution. For the remaining attacks, representing less than 1% of malware, we need to focus heavily on survivability.

The 9 Box of Controls, Business Trends, and Architecture Requirements

Before diving into the specifics of the architecture, I'll explain the 9 Box of Controls. Then I'll recap some of the key business and technology trends, focusing on how they drive the need for specific capabilities in security technology.

9 Box of Controls

There are three primary types of security controls: prevention, detection, and response. Prevention occurs when an action or control prevents a risk before it affects users or the environment. Detection is identifying the presence of something malicious that has already entered the environment. Response is a reaction. From a risk perspective, prevention focuses on minimizing vulnerability and the potential for harm, while detection and response focus on minimizing damage.

There are also three primary levels of control automation: automated, semi-automated, and manual. Automated control occurs entirely through machines. Semi-automated involves some level of human intervention. Manual controls are managed entirely by hand.

The combinations of these control types and automation levels comprise the cells of the 9 Box, as shown in Figure 7-1. Risk increases as we move from prevention to detection to response. Cost increases as we move from automated to semi-automated to manual controls.




Figure 7-1. The 9 Box of Controls

However, there is a third dimension to the 9 Box: control friction. As we know, friction is the force that causes a moving object to slow down when it is in contact with another object. Similarly, controls can impose a “drag coefficient” on business velocity—they can slow the user or a business process. However, friction is not a fundamental, immutable force like gravity or electromagnetism. Instead, we have the ability to determine how much control friction we apply. Apply too much control friction, and business users will go around IT and its security controls. This adds cost: IT is no longer managing the technology; data and business silos are created, and the organization loses its volume purchasing power. It also adds risk: because the security team lacks visibility into the technology, it cannot prevent compromises, detection is difficult, and in many cases response after the fact becomes the only option. If a business adheres to high-friction controls, the effect can be to generate systemic business risk. High-friction controls can hinder business velocity; the organization can lose time to market and the ability to innovate, and over the long term it may even lose market leadership.

IT Consumerization

As I discussed in Chapter 5, consumerization is a major IT theme with ever-broadening impact. It includes several trends, including the adoption of new applications and support for consumer devices.

Many highly mobile employees want to use their own consumer devices, such as smartphones, wearables, and tablets, for work. This increases productivity by enabling employees to collaborate and access information from anywhere, at any time. To support this, organizations provide access to corporate e-mail and other applications from employee-owned smartphones and tablets.

Some people believe that in the future, all devices will be consumer-owned, and that enterprises will no longer purchase devices for their users. I believe this might be the case in some work environments, but I doubt that it will suit all organizations. For a company providing call center services, with most employees working from home, it might make sense that employees exclusively use their own personal systems for work. But this strategy could be more risky for a financial services company whose employees handle highly sensitive information that is subject to extensive regulatory requirements.

Nevertheless, the consumerization trend continues to grow at almost all organizations. Accordingly, we’ll need to provide employees with a level of access to resources from an expanding continuum of client devices, some of which may have much weaker security controls than today’s enterprise clients (see sidebar).

CONSUMERIZING ENTERPRISE IT AND “ENTERPRISING” THE CONSUMER

Discussions of IT consumerization tend to draw a clear line between business devices that can be managed and trusted, and personal consumer devices that are essentially unmanaged and untrusted.

However, not all consumer devices are created equal. From a security standpoint, it may be more valuable to think about a device’s capabilities than to categorize it based solely on whether it’s marketed as an enterprise device or a personal device. The security of a device depends on the inherent features of the hardware, operating system, and applications, and on whether it enables us to add further security and manageability capabilities that mitigate the risks of enterprise use.

As the variety of consumer devices, such as smartphones and wearables, continues to expand, users may choose from dozens of models with different levels of security capabilities. Greater security and manageability means that IT can place greater trust in the device and provide a correspondingly greater level of access to enterprise resources.

Extending this idea further, the information security group could evaluate the security of available consumer devices and provide guidance about the level of enterprise access that users will be allowed with each device. Users may prefer to buy a more secure device because it will provide them more access. With greater access, they can use the device for more of their daily work activities. This ability in turn enables them to be more productive.

At the same time, employees increasingly expect to have available to them at work the types of consumer services and cloud applications that they use in their personal lives. These include social computing applications such as blogs and wikis, video-sharing sites, and file-sharing services.

We need a security architecture that enables us to more quickly support new devices and provide access to a greater range of applications and data, without increasing risk. We need to be able to dynamically adjust the levels of access we provide and the monitoring we perform, depending on the security controls of the client device.

New Business Needs

Nearly all companies now rely on a growing network of business partners, and conduct many of their interactions with those partners online. Many organizations are also expanding into new markets through both organic growth and acquisitions. Because of these business trends, most organizations will need to provide access to a broader range of users, many of whom are not employees. Many organizations also need to be able to smoothly integrate acquired companies and provide them with access to resources. In general, we need to quickly provide new users access while minimizing risk and providing selective, controlled access only to the resources they need.

Cloud Computing

Most organizations are already using cloud services in some form to achieve benefits such as greater agility and lower cost. Some are also implementing a private cloud based on virtualized infrastructure while using external cloud services for noncritical applications. In the future, I expect greater use of hybrid clouds that use both internal and external resources, especially for organizations that are anchored to legacy environments. Organizations able to let go of their legacy environments will predominantly use the cloud, with limited internal infrastructure.

This trend means that IT services at many organizations will be provided by a mixture of traditional and cloud-based internal and external services. During a typical day, employees may access a variety of different services, some of which are internal and some external. Ultimately, they should be able to easily move between these services without needing to log in multiple times or even know where the services are located.

Securing access to cloud-based services presents challenges that aren't easily addressed using conventional security controls. In cloud environments, systems and their data are virtualized and may migrate dynamically to different network locations. This makes it difficult to effectively restrict access using traditional security controls such as firewalls, which rely on fixed locations of systems and a more static nature of the data. We need much more granular and dynamic controls that are linked to the resources themselves rather than just their network location.

Changing Threat Landscape

The threat landscape is evolving rapidly. Increasingly, attackers have taken a stealthy approach, creating malware that quietly gains access and attempts to remain undetected in order to maintain access over time. This has been possible because the security solutions deployed on endpoints in most organizations today do not adequately prevent malicious code from executing. As the number of threats increases and new types of malware emerge, we need to focus on the 9 Box of Controls and seek new prevention methods reduce risk, reduce cost, and reduce control friction.

Traditional enterprise security architectures have relied largely on protective controls such as firewalls located at the network perimeter and signature-based antivirus at the end points. At the same time, our focus has shifted to providing controlled access to a broader range of users and devices, rather than simply preventing access. Combine this with a continually changing threat landscape, and we can assume that attempts to compromise the environment are inevitable. Although existing perimeter controls such as firewalls will continue to have some value, we need tools that can dramatically increase the ability to prevent attackers from gaining access to the environment, but in way that does not introduce a cost burden or a high degree of control friction.

Privacy and Regulatory Requirements

The growing emphasis on privacy requirements and the increasingly complex regulatory environment have many implications for the way we manage information. Some regulations create the need for more control over where information is stored and require specific levels of protection and tracking. Our architecture must provide this assurance, allowing us to build a high-security environment and access controls appropriate for the protection of highly regulated information. In addition, the security controls themselves must not introduce privacy risks.

New Architecture

To meet these rapidly changing requirements, we need a highly flexible and dynamic architecture. The architecture should enable us to more quickly adopt new devices, use models, and capabilities; provide security across an increasingly complex environment; and adapt to a changing threat landscape.

Key goals include helping increase employee productivity while supporting new business requirements and technology trends, including IT consumerization, cloud computing, and access by a broader range of users. At the same time, the architecture should be designed to prevent risk, reduce our attack surface, and improve survivability—even as the threat landscape grows in complexity and maliciousness.

The architecture moves away from the traditional enterprise trust model, which is binary and static. With this traditional model, a user is in general either granted or denied access to resources; once granted, the level of access remains constant. The new architecture replaces this with a dynamic, multi-tiered trust model that exercises more fine-grained control over identity and access control, including access to specific resources. This means that for an individual user, the level of access provided may vary dynamically over time, depending on a variety of factors—such as whether the user is accessing the network from a highly secure managed device with advanced anti-malware capabilities or an untrusted and perhaps unmanaged device.

The architecture's flexibility allows us to take advantage of trust based on real proof that malware execution is being prevented. Increasingly, devices may include some level of hardware-enforced security designed to ensure the integrity of the applications and data on the device. The architecture takes this into account when determining whether to allow access to specific resources—a more-trusted platform can be allowed greater access than a less-trusted one. The architecture is based on four cornerstones:

- **Trust calculation:** This unique element of the architecture handles user identity and access management, dynamically determining whether a user should be granted access to specific resources and, if so, what type of access should be granted. The calculation is based on factors such as the user's client device and location, the type of resources requested, and the security controls that are available.

- **Security zones:** The infrastructure is divided into multiple security zones that provide different levels of protection. These range from trusted network zones containing critical data, with tightly controlled access, to untrusted zones containing less-valuable data and allowing broader access. Communication between zones is controlled and monitored; this helps ensure users can only access the resources for which they have been authorized and prevents compromises from spreading across multiple zones.
- **Balanced controls:** To increase flexibility and the ability to recover from a successful attack, the model emphasizes the need for preventative controls but also to balance them with detection and response.
- **User and data perimeters:** Recognizing that protecting the enterprise network boundary is no longer adequate, we need to treat users and data as additional security perimeters and protect them accordingly. This means an increased focus on the endpoint device and prevention of malicious code, in addition to increasing user awareness and building data protection into the information assets.

I'll describe each of the four cornerstones in more detail.

Trust Calculation

The trust calculation plays an essential role in providing the flexibility required to support a rapidly expanding number of devices and usage models. The calculation enables us to dynamically adjust users' levels of access, depending on factors such as the devices and networks they are currently using.

It calculates trust in the interaction between the person or device requesting access (source) and the information requested (destination). The calculation consists of a source score and a destination score, taking into account the controls available to mitigate risk. As shown in Figure 7-2, the result of this calculation determines whether the user is allowed access and the type of access provided.

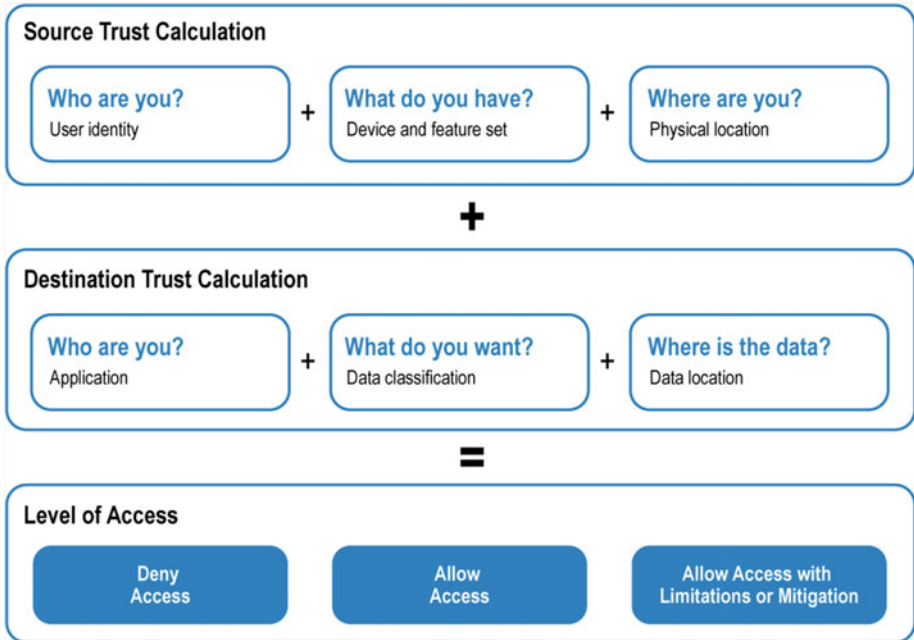


Figure 7-2. Trust calculation. Source: Intel Corporation, 2012

Source Score

Trust in the source, or requestor, is calculated based on the following factors:

- **Who:** The identity of the user or service requesting access and our confidence level in the authentication mechanism used; how confident are we that users are who they say they are?
- **What:** The device type, its control capabilities, our ability to validate those controls, and the extent to which IT manages the device.
- **Where:** The user's or service's location. For example, a user who is inside the enterprise network is more trusted than the same user connecting through a public network. There may also be other considerations, such as the geographical region where the user is located.

Destination Score

This is calculated based on the same three factors, but these are considered from the perspective of the destination (the information the source is trying to access):

- *Who*: The application that stores the requested data. Some applications can enforce greater controls, such as enterprise rights management (ERM), and therefore provide a higher level of trust.
- *What*: The sensitivity of the information being requested and other considerations, such as our ability to recover it if compromise occurs.
- *Where*: The security zone in which the data resides.

Available Controls

The trust calculation also takes into account the security controls available for the zone. If the only controls available are controls that simply block or allow access, we might deny access due to lack of other options. However, if we have extensive preventative controls with highly granular levels of access, detailed logs, and highly tuned security monitoring—as well as the ability to recover from or correct problems—then we can allow access without creating additional risk.

Calculating Trust

The trust calculation adds the source score and the destination score to arrive at an initial trust level. The available controls are then considered to make a final decision about whether access is allowed and, if so, how. This calculation is performed by a logical entity called a *policy decision point* (PDP), which is part of the authentication infrastructure and makes access control decisions based on a set of policies.

Based on the results of this calculation, the PDP makes a decision, allocating a trust level that determines whether the user can access the requested resource and the type of access that is allowed. Broadly, the decision will fall into one of the following categories:

- Allow access
- Deny access
- Allow access with limitations or mitigation

This trust calculation therefore allows us to dynamically apply granular control over access to specific resources. For example, employees using IT-managed devices with additional hardware features such as a trusted platform module (TPM), global positioning system (GPS), and full disk encryption would be allowed access to more resources than when using devices that lack those features.

Employees directly connected to the network would typically get greater access than when using a public network. If we are unable to verify the location of a high-security device such as a managed PC, we would allow less access.

The trust calculation also can be used for more fine-grained distinctions between different device models. For example, we could provide different levels of access based on manageability, hardware-enabled authentication and encryption, and installed applications.

We anticipate situations in which the trust level is not adequate to allow any access, but there is still a business requirement to allow a connection or transaction to occur. In these conditions, the result of the trust calculation could be a decision to allow access with limitations or with compensating controls that mitigate the risk. For example, a user might be allowed read-only access or might be permitted access only if additional monitoring controls are in place.

Today, the trust calculation makes decisions based on information gathered from components at multiple levels of the infrastructure, such as network gateways, access points, and user devices. Once the trust calculation mechanism is in place, we can extend it to include information from a broader range of sources.

The trust calculation can be used to determine access to internal systems by business partners as well as employees. Let's say we're collaborating with another company on the design of a new product. An engineer at that company wants access to a specific document. We can add a variety of criteria to the trust calculation for deciding whether to grant access. Did the engineer's request originate within the business partner's enterprise network? Is it consistent with the type of request that we'd expect from an engineer? If so, we have a higher level of trust in the requestor.

If we cannot establish an adequate level of trust in the user's device, but other factors provide enough confidence to grant access, we might provide one-time access for a specific job. We could do this by allowing a document to be downloaded, but only within a container that ensures the document is completely removed from the user's device once the job is completed.

Longer term, the trust calculation could become a mechanism that is used to determine access to both internal and external resources, including cloud-based applications.

Security Zones

The architecture divides the IT environment into multiple security zones. These range from untrusted zones that provide access to less valuable data and less important systems to trusted zones containing critical data and resources.

Because the higher-trust zones contain more valuable assets, they are protected with a greater depth and range of controls, and we restrict access to fewer types of devices and applications, as shown in Figure 7-3. However, devices allowed access to higher-trust zones also have more power; they may be able to perform actions that are not allowed within lower-trust zones, such as creating or modifying enterprise data.



Figure 7-3. As the value of an asset increases, the depth and span of controls increase, while the number of allowed devices, applications, and locations decrease. Source: Intel Corporation, 2012

Aligning the infrastructure in this fashion provides an excellent way to right-size security controls so that security resources are utilized effectively. It also helps improve the user experience by enabling employees to choose from a wider range of devices, such as smartphones, for lower-risk activities. However, all devices should have, at a minimum, advanced endpoint capabilities that prevent more than 99% of malware from executing.

Access to zones is determined by the results of the trust calculation and is controlled by *policy enforcement points* (PEPs). PEPs may include a range of controls, including firewalls, application proxies, intrusion detection and prevention systems, authentication systems, and logging systems.

Communication between zones is tightly restricted, monitored, and controlled. We separate zones by locating them on different physical or virtual LANs; PEPs control communication between zones. This means that if one zone is compromised, we can prevent the problem from spreading to other zones or increase our chances of detection if it does spread. In addition, we can use PEP controls, such as application proxies, to provide devices and applications in lower-trust zones with limited, controlled access to specific resources in higher-trust zones when required.

The architecture includes three primary categories of security zones: untrusted, selective, and trusted. Within the zones, there are multiple subzones.

Untrusted Zones

These zones host data and services (or the interfaces to them) that can be exposed to untrusted entities. This allows us to provide widespread access to a limited set of resources from non-managed consumer devices, without increasing the risk to higher-value resources located in other zones. Untrusted zones might provide access to enterprise resources, such as corporate e-mail and calendars, or they might simply provide Internet access.

These zones are regarded as “shark tanks,” with a high risk of attack and compromise. Therefore, detective and corrective controls are needed to mitigate this risk. These controls might include a high level of monitoring to detect suspect activity and correction capabilities such as dynamic removal of user privilege.

We anticipate a need to provide controlled access from these zones to resources in higher-trust zones. For example, an employee using an untrusted device might be allowed limited, read-only access to customer data located in a trusted zone; or their device might need access to a directory server in a trusted zone to send e-mail. We expect to provide this controlled access using application proxies. These proxies act as secure intermediaries, evaluating the request from the device, gathering the information from the resource in a trusted zone, and passing it to the device.

Selective Zones

Selective zones provide more protection than untrusted zones. Examples of services in these zones include applications and data accessed by contractors, business partners, and employees, using client devices that are managed or otherwise provide a level of trust. Selective zones do not contain critical data or high-value intellectual property. Several selective subzones provide access to different services or users.

Trusted Zones

Trusted zones host critical services, data, and infrastructure. They are highly secured and locked down. Examples of services within these zones are administrative access to data center servers and network infrastructure, factory networks and devices, enterprise resource planning (ERP) applications, and design engineering systems containing intellectual property. Accordingly, we might only allow direct access to these resources from trusted systems located within the enterprise network, and all access would be monitored closely to detect anomalous behavior.

Many organizations have implemented secure high-trust zones as part of their transition to an enterprise private cloud. Implementing these zones is a key step in allowing several types of applications to be moved onto virtualized cloud infrastructure, including applications requiring high security. The security features in these trusted zones include application hardening and increased monitoring.

NEW SECURITY ARCHITECTURE IN ACTION: A DAY IN THE LIFE OF AN EMPLOYEE

This example (illustrated in Figure 7-4) describes how the new security architecture could enable an organization's sales force to access the information they need in the course of a day. At the same time, the architecture protects security by dynamically adjusting the level of access provided, based on the user's device, its location, and its capabilities for preventing malicious code, and by monitoring for anomalous behavior.

The employee travels to a customer site. The employee is using a personal smartphone with limited security features and so is allowed access only to services in untrusted zones. From here, the employee can view limited customer information, including recent orders, extracted from an enterprise resource planning (ERP)

system in a trusted zone—but only through an application proxy server, which protects the trusted zone by acting as an intermediary, evaluating information requests, accessing the ERP system, and relaying the information to the user.

If a smartphone requests an abnormally large number of customer records—an indication that it may have been stolen—further access from the smartphone is blocked. To help understand the reason for the anomalous access, there is increased monitoring of the employee’s attempts to access the system from any device.

The employee reaches the customer site and logs into the enterprise network from a company-owned mobile business PC. Because this device is more trusted, the employee now has access to additional capabilities available in selective zones, such as the ability to view pricing and create orders that are relayed by an application proxy to the ERP system in a trusted zone.

The employee returns to the company’s office and connects to the corporate network. Now the employee is using a trusted device from a trusted location and has direct access to the ERP system in a trusted zone.

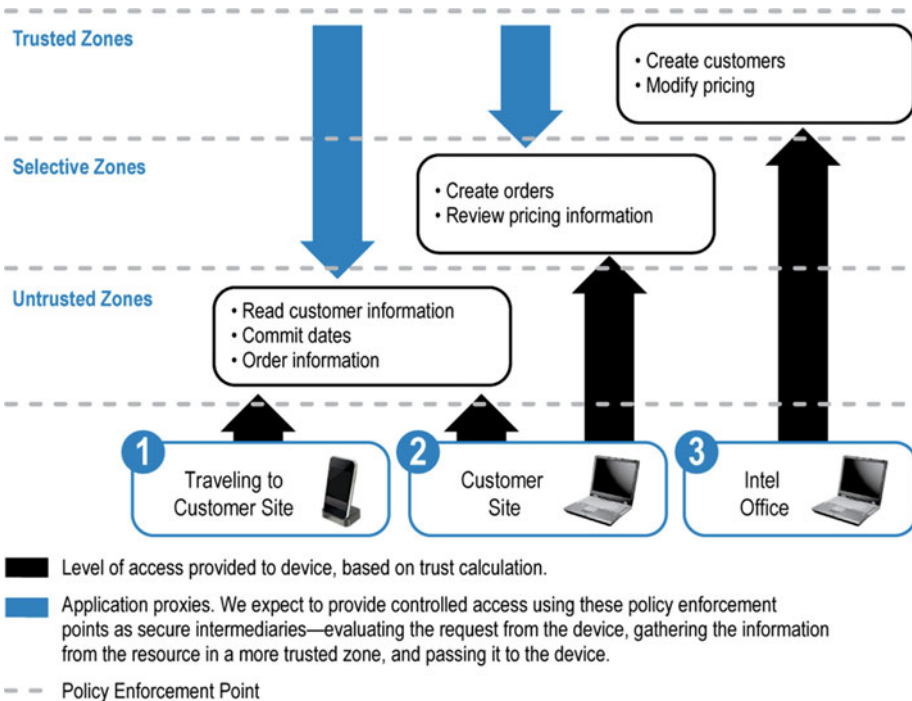


Figure 7-4. The new security architecture dynamically adjusts the user’s access to information, based on factors such as the user’s device and location. Source: Intel Corporation, 2012

Balanced Controls

Over the past decade, enterprise security has focused heavily on controls such as firewalls, signature based antivirus, and intrusion detection systems such as behavior-based anomaly detection tools. As we have seen so often in the past few years, this approach is not working. At many companies, the default belief is now that prevention is not possible and we can only correct problems after they have occurred.

However, the new security model requires that we understand the implications of the 9 Box. Preventative controls should not only stop malicious code from executing but do so in a way that lowers our overall cost of controls and with low friction. More effective prevention will reduce the alert fatigue within due to the “whack-a-mole” effect associated with over-reliance on detective (monitoring) and response controls. Detection capabilities will also be more effective because effective prevention reduces the “noise” in the environment. Over the long term, this approach will free up resources that can then be applied to other corrective controls.

By using the 9 Box to guide the control philosophy, and demanding solutions that continually shift down and to the left (reducing risk, cost, and control friction), we will be able to change the risk dynamics in the industry.

USING SECURITY ANALYTICS TO DETECT SUSPICIOUS BEHAVIOR

Almost all organizations have experienced security issues involving both external attackers and insiders, including attempts to steal intellectual property. Investigations have identified markers and indicators that are frequently associated with these events. If we can spot these indicators sooner, we can respond and mitigate the threats more quickly.

Security analytics technology can be used to detect suspicious behavior as the environment becomes more complex and attackers become more adept at concealing compromises. The technology automates the process of analyzing large volumes of data to detect and monitor anomalous activity, allowing companies to detect problems that they might otherwise miss. These capabilities are similar to those already implemented by financial institutions to prevent fraudulent credit-card transactions, and by online consumer services to prevent theft of user data.

On a large scale, logging data generated by servers and sensors across the network can be collected into a database for analysis. Security business intelligence can also be applied at the level of individual users and devices, as long as we are careful to protect users’ privacy.

The balance between preventative, detective, and corrective controls will vary, depending on the security zone. In high-trust zones, we implement extensive monitoring to detect possible attempts to steal data or compromise critical systems. Redundancy within each type of control can be used to provide additional protection.

The following includes possible examples of using detective and preventative controls:

- An employee attempts to send a confidential document to an external e-mail address. Monitoring software detects the attempt, prevents the document from being sent outside the firewall, and asks the employee if he or she really intended to do this. If the employee confirms that this was intended, the document may be transmitted, or if the document is highly sensitive, a redacted version may be sent.
- Inappropriate use of a document protected with enterprise rights management technology results in revocation of access to the document.
- The system allows access to specific documents but tracks the activity. A user can download a few documents without causing concerns. However, if the user attempts to download hundreds of documents, the system slows down the speed of delivery (for instance, only allowing ten to be checked out at a time) and alerts the user's manager. If the manager approves, the user is given faster access.
- The detection of an infected system places the system on a remediation network, isolating the system and restricting access to enterprise information and applications. The system may retain some ability to access corporate assets, but all activity is closely logged to enable incident response if necessary.
- When a system is found to be compromised, we examine all its recent activities and interactions with other systems. Additional monitoring of those systems is automatically enabled.

USING MACHINE LEARNING TO IMPROVE ANTI-MALWARE TECHNOLOGY

Traditional antivirus software relies on recognizing characteristic signatures to identify specific malware families. But today, adversaries have access to off-the-shelf malware toolkits that make it easy to create custom malware variants that signature-based antivirus products cannot recognize. This custom malware sails past traditional antivirus products as if they didn't exist.

Machine learning technology provides a solution to the problem. Rather than relying on humans to identify malware signatures, machine learning technology can automatically analyze and classify hundreds of thousands of characteristics per file, breaking them down to an atomic level to discern whether an object is "good" or "bad" in real time.

The process works like this. A machine learning platform continuously collects vast amounts of data from many sources. It analyzes the data and extracts DNA-level features that the machine learning platform itself determines are unique characteristics of good and bad files. Most of these characteristics are so microscopic that human malware researchers and reverse engineers don't understand their importance. The software constantly adjusts to the real-time threatscape, thus learning to make higher-fidelity decisions. For each file, the platform assigns a threat score that is used to automate policy-based protection decisions—ignore, alert, block, or terminate file/process execution. A mathematical model encapsulating the platform's intelligence is then periodically extracted and incorporated into an anti-malware solution that is installed on endpoints. Using this solution, it's possible to stop more than 99% of malware before execution.

Users, Data, and the Internet of Things: The New Perimeters

The concept of balanced controls also extends to the protection of users and data. Traditional network security boundaries are dissolving with the proliferation of new devices and users' expectations that they should be able to access information from anywhere at any time. Users are under direct assault from a barrage of attacks designed to trick them into taking actions that can compromise the information on their devices or on enterprise systems. These trends mean that we need to think more broadly about how we protect information, as well as the users of this information.

While we continue to implement enterprise network controls, such as perimeter defenses and the detective controls described earlier, we need to supplement these controls with a focus on the users and on the primary assets we are trying to protect such as intellectual property. The new architecture therefore expands our defenses to two additional perimeters: the data itself and the users who have access to the data.

Data Perimeter

Important data should be protected at all times: when it is created, stored, and transmitted. This becomes increasingly challenging as we move data to more and more devices and let more people access it. How do we protect information when it's located outside the physical perimeter on a personal device?

One approach is to use technologies that closely integrate protection with high-value data so that the data remains protected as it moves to different devices and locations. Technologies such as enterprise rights management and data leak prevention can be used to watermark and tag information so that we can track and manage its use. With enterprise rights management, the creator of a document can define exactly who has access rights throughout the life of the document and can revoke access at any point. Data loss prevention is used to tag documents, track their movements, and prevent transfer outside the organization if necessary.

User Perimeter

As I described in Chapter 5, people are part of the security perimeter, and we need to treat them as such. Users can become security risks for a variety of reasons. They are targeted more frequently in social engineering attacks, and they are more vulnerable to these attacks because their personal information is often readily available on social networking sites. They may also click malicious links in e-mail, download malware, or store data on portable devices that then are lost. A combination of training, incentives, and other activities can help instill information security and privacy protection into the corporate culture and successfully encourages employees to own responsibility for protecting enterprise and personal information.

Internet of Things

The Internet of Things can be viewed as an extension of the user and data perimeters into new connected devices and systems such as cars, wearables, and smart buildings. IoT devices should be included in the security architecture; for example, the trust calculation could be applied to access from IoT devices, so that the security of the device is a factor in determining the level of access provided. For machine-to-machine communications, each communicating machine can be considered conceptually as analogous to a user; the security architecture focuses on preventing, detecting, and responding to behavior that it identifies as anomalous.

Conclusion

This chapter describes a new control architecture designed to support the Protect to Enable mission. With this approach we can lower risks, lower costs, and lower control friction. It will also allow for faster adoption of new services and capabilities because it helps prevent risk and improve survivability. I believe that this architecture can be used to meet a broad range of evolving requirements, including new usage models and threats. The architecture's flexibility and granular trust model should also make it easier for the security team to identify and contain anomalous activity that signals potential insider threats. By publishing information about the architecture, I hope to encourage others to take advantage of these ideas. I also hope that making this information available will stimulate more discussion and ideas, and that others will build on these concepts to create further innovations that benefit all of us.

CHAPTER 8



Looking to the Future: Emerging Security Capabilities

Learn from yesterday, live for today, hope for tomorrow. The important thing is not to stop questioning.

—Albert Einstein

The Web has existed for two decades, yet it's only in the last few years that we've gained a clearer picture of what the Internet may become, and how the emerging capabilities may shape the future.

As early as 1993, companies like AOL started offering access to online newsgroups, soon followed by dial-up Internet access using early web browsers. As laptops became more affordable, many people started accessing the Internet while on the move. The rise of smartphones introduced built-in sensors, such as cameras, global positioning system receivers, and touch-sensitive screens, into consumers' everyday computing experiences. Businesses began using the information gathered from users' devices to offer personalized experiences, ranging from location-based driving directions to selected advertisements. The variety of Internet-connected devices rapidly expanded to include tablets, home DVRs, appliances, and cars. Devices also became smarter, with improved voice and gesture recognition.

We're now entering a world in which these elements will be combined to create much richer context-aware experiences for users and new opportunities for businesses. Our devices will know us, and they will know other devices. In fact, devices may almost become part of us: many companies are already shipping wearable computers, including smart athletic garments that work with smartphone apps to monitor your biometrics and suggest ways to improve your performance.

Each day, billions of computing devices will perform functions on our behalf, often communicating among themselves to get the job done. Much more information will be collected from sensors such as cameras, microphones, and GPS receivers embedded into the user devices. This data will be combined with other information to create context-aware experiences that are far more personalized and compelling. Already, cameras and image recognition technology, combined with behind-the-scenes analytical software, can be used to identify a user's age bracket and gender, and tailor their experience accordingly. Early applications based on this technology are being piloted and in some cases deployed by large companies, including retailers (see sidebar).

Estimates of the projected size of the context-aware computing market continue to grow. When the first edition of this book was printed, Gartner, Inc. (2011) expected context-aware technologies to create huge business opportunities affecting an estimated \$96 billion in annual consumer spending worldwide by 2015. In 2013, forecasts suggested the market would reach \$120 billion by 2018 (MarketsandMarkets 2013). And a report in late 2015 forecast the market will swell to \$185 billion by 2020 (Global Industry Analysts 2015). During this period, it's expected that a significant percentage of all payment transactions will be validated using contextual information.

RICHER EXPERIENCES IN THE RETAIL ENVIRONMENT

As people buy more goods online, retailers are seeking to entice shoppers into brick-and-mortar stores by using technology to create richer, context-aware experiences.

Macy's and some other big-name stores are already using beacons, which detect the smartphones of nearby shoppers and, if they have opted in, send them targeted offers or mobile games with gift-card prizes (Tode 2015). Brands including Kate Spade and Levi's use smart display tables and shelves that sense when customers pick up a product and engage them with relevant videos and product information. The technology tracks every interaction, so stores can analyze shopper behavior and measure the impact on sales (Perch Interactive 2016). LEGO stores use augmented reality video screens to show kids what they can build with each LEGO box. The screens recognize each box and display a 3D image of a toy that can be created from it, blended into a real-time video of the child in the store. Canadian sports retailer Sport Chek's flagship stores integrate hundreds of screens in displays up to 16 feet tall, using gesture, touch, and RFID to sense customer input and display customized interactive content.

As an Advertising Age column noted, technology may ultimately help transform the physical store into a venue for interactive experiences that increase brand affinity—acting as an event space, gallery, help desk, or even a test kitchen. If that happens, online sales may work in tandem with, rather than as a substitute for, a physical store (Fulford 2015).

These new technologies also introduce new risks, as I described in the discussion of emerging threats and vulnerabilities in Chapter 6. The sensors and other new capabilities embedded into millions of intelligent new devices can be exploited for dangerous purposes. Malicious individuals might be able to remotely access home security surveillance systems to determine when you're not at home. Researchers have already demonstrated the ability to remotely control the brakes and other functions of an Internet-enabled car. After remotely hijacking a Jeep Cherokee driven by a reporter, researcher Charlie Miller commented, "Right now I could do that to every [Chrysler] car in the United States on the Sprint network (Pagliery 2015)." The hack prompted Fiat Chrysler to recall 1.4 million vehicles to fix the issue (Greenberg 2015b).

As security professionals, we may tend to focus obsessively on this darker side of the picture. Looking for threats and vulnerabilities is part of our role. We've seen that attackers find ways to exploit new technologies almost as soon as they appear. Analysis of emerging threats by many firms indicates that this trend will continue. As attackers adapt, we must adapt, too. Our role will be more important than ever. As more aspects of people's daily lives are based on technology, it will become increasingly important to secure the technology. The Protect to Enable mission will expand accordingly; in fact, it is becoming a corporate social responsibility, as I will explain further in Chapter 9.

The positive news is that new technologies can also be used to enhance security. As information risk becomes an even more high-profile concern, suppliers are building more security into their products and services. Devices will include a greater level of baseline security hardening to reduce the likelihood of compromise and minimize the impact.

Context-aware computing also introduces new privacy concerns. By definition, context-awareness involves taking advantage of information about the user to create personalized experiences. This makes it even more important to appropriately protect users' information and privacy. A clear organizational commitment to privacy will be important to ensure this protection. A growing number of other organizations have formally committed to complying with a single set of privacy principles worldwide—although this is becoming difficult due to the proliferation of localized privacy laws and the elimination of the EU safe harbor agreement (see Chapter 1).

An organization's privacy commitment must also extend to applications and systems. Suppliers are becoming increasingly aware of this, and some are already taking additional steps to ensure user data is collected anonymously. The new baseline security capabilities built into products, such as hardware-enforced protection and accelerated encryption, may also help enhance privacy by protecting user data. In addition, the information provided by sensors can be used to create context-aware security. Today, some cars can automatically adjust seat, mirror, and pedal positions to suit different drivers. They adjust these settings when they detect the presence of the driver's personal car key. In the future, as cars become more intelligent and include more sensors, they might identify the driver using a camera and microphone. If they don't recognize the driver, they might disable the car and alert the owner via their built-in wireless Internet connection. Cars might include a maintenance mode that lets mechanics drive it while when it's being serviced, but only within a radius of a few miles. Similarly, as I'll discuss later in this chapter, the sensors in an enterprise-class device, such as a business laptop PC, could be used to prevent theft and help protect the information it contains.

From the perspective of the enterprise information security team, these emerging capabilities will allow increased trust in users and their devices. When we have a higher level of trust, we can provide the user with greater access to sensitive enterprise information and other resources.

I believe that this dynamic evaluation of trust is a key capability that new security architectures should include, as I discussed in Chapter 7. Employees may want to access our systems from a variety of devices and locations, including personal smartphones and tablets as well as business PCs. When a user requests access to enterprise systems, the architecture should dynamically calculate trust based on contextual information such as the user's identity, the security features of the device they're using, their physical location, and the resources they're trying to access. The architecture then will decide whether to grant access and the level of access that should be allowed. As manufacturers increase the security capabilities in their devices, the model will be able to take this into account. We'll have increased trust in a device, and we'll be able to provide a correspondingly greater level of access.

In this chapter, I'll take a closer look at some of the emerging security capabilities that we can expect in products and services. First, though, I'd like to set the stage by examining some of the key underlying trends that make these security capabilities both necessary and possible.

Internet of Things

Many everyday objects are becoming more intelligent. They're acquiring processors, sensors, software, and the ability to communicate. This trend is made possible by Moore's Law: processors and other hardware components continually become faster and less expensive, and, therefore, ubiquitous as a result. This accelerating trend is creating the Internet of Things, a massive expansion of the Internet as it swells to include billions of devices and household objects. Intelligent devices in cars, home electronics, and other "things" will far outnumber those in more conventional computing platforms and even those in mobile devices such as smartphones. Gartner, Inc. estimates that during 2016, 5.5 million new "things" will be connected every day. Juniper Research expects 38.5 billion connected devices by 2020 (Loechner 2015); Cisco expects an even higher number of 50 billion (Cisco Systems 2015b).

Gartner, Inc. (2011b) identifies several key technologies and capabilities contributing to this trend, including sensors, image recognition, and wireless payments using *near field communications* (NFC) technology. Sensors that detect and communicate changes in their environment are being embedded not just in mobile devices, but in an increasing number of places and objects. Emerging applications will take advantage of this information. For example, camera-based image recognition technologies are expanding from mainly industrial applications to broad consumer and enterprise uses. These systems gather information about users and then analyze this information to personalize the user experience. Wireless NFC, based on a communications standard analogous to the Radio Frequency Identification (RFID) technology used for product-tracking, lets users make payments by waving a mobile phone or smartwatch in front of a compatible reader.

With technologies such as NFC, the concept of the Internet may broaden to include an even wider variety of “dumb” objects, like drink cans or fertilizer bags (Gartner 2011b). This trend will provide opportunities for innovations that were not previously possible. Today, items in stores may include 2D bar codes that can be read by smartphones. In the future, store items may include NFC on the packaging or shelf label allowing them to wirelessly identify themselves to nearby devices, such as a shopper’s smartphone. The shopper will then be able to learn not only about the product, but also alternatives, and could even view cross-selling and up-selling suggestions.

Devices such as the Nest Learning Thermostat have provided a glimpse of the future. This home heating controller is designed to be intuitive and simple to operate, replacing complex menus and instructions with a single big button and a dial. Users can remotely monitor and set the temperature from their smartphones, so they know the house will be warm by the time they get home. But perhaps the most interesting capability is that, as its name suggests, it can learn. The Nest monitors use of the heating system and attempts to learn the user’s preferences—when the heating is switched on and off, and the desired temperature. After studying the use patterns for a while, the Nest begins to predict and autonomously set the temperature and timing itself. Since Nest launched many other companies have followed suit with similar devices not only for home heating but also for other sensors and alarms, including water sensors, motion sensors, and do-it-yourself internet-based home security systems.

I believe that devices like this are early examples of a much larger trend. As the Internet of Things grows, more interactions will occur directly between devices, rather than between people and device. Devices and objects will interpret and act on information provided by other objects. This will enable much more intuitive and streamlined experiences in many different fields. Consider the following scenario, described by Plantronics CTO Joe Burton (2012). A doctor visits a patient in a hospital room. A smart device the doctor is wearing turns on the doctor’s workstation in the room, then authenticates the doctor to the patient management system, detects which patient is near the doctor, and pulls up the patient’s record. When the doctor leaves the room, the information accumulated during the visit is saved and the workstation powers down.

Consistent User Experience Across Devices

Users now demand the same quality of experience in the workplace that they’ve become accustomed to in their personal lives. This includes the ability to access information across a continuum of devices, including PCs, smartphones, and tablets. They expect to be able to move from one device to another. They also expect intuitive applications on all of these devices, with the application’s features tailored to the device’s size and capabilities.

IT therefore needs to provide users with a consistent experience across devices and the ability to seamlessly transition between them. As enterprise information security professionals, we need to focus on the user experience and on enabling this broader range of devices while managing the risks.

Cloud Computing

The cloud is as much a new business model as it is a technology shift. The ability to obtain flexible IT services on demand lets businesses operate more dynamically—quickly taking advantage of business opportunities and growing or shrinking infrastructure capacity to meet demand. Cloud services can also potentially reduce cost.

However, cloud computing can also add new security complexities and data-protection concerns. Organizations may use multiple cloud providers, while also operating a private cloud for the most sensitive applications. Users need to be able to easily access services delivered from any of these multiple environments. From the enterprise perspective, we need to enable a seamless user experience while minimizing risk. This implies a federated model in which the user needs to log in only once; the user's credentials can then be used to access multiple applications. However, this also means that an attacker may only need to gain access once in order to compromise several environments.

Big Data Analytics

Businesses have quickly realized the value of analytical tools for real-time analysis of massive amounts of unstructured data. In the future, these analytic capabilities will increasingly be used to interpret data from sensors as well as from databases, social media, and other sources. The analysis of this information will then be used to create new personalized experiences, like the retail examples discussed in the “Richer Experiences in the Retail Environment” sidebar.

This analysis can also be integrated with existing enterprise systems to create sophisticated customer-focused services. Here's a scenario described by Accenture (2012): a rental car company automatically detects when an accident with one of its cars has happened, initiates emergency services if needed, and issues a replacement rental car to meet the renter at the scene, greatly improving the chances of creating a loyal customer for life.

Artificial Intelligence

Artificial intelligence is rapidly maturing, and it's now clear that AI will help all of us in a variety of ways, both in business and our personal lives. AI is already used to identify meaningful patterns in data for many purposes, including information security, and to understand and translate speech. AI will certainly play a role in self-driving cars. Over time, AI will become capable of taking on broader and greater responsibilities. As Alphabet Inc. executives Eric Schmidt and Jared Cohen put it: “Eventually it will be possible to give a computer unstructured data—say, spreadsheets used to manage business records—and receive quality advice on improving operations.” (Schmidt and Cohen 2015) In our personal lives, perhaps we'll have a helper like Jibo, a “social robot” that recognizes your face, converses with you, helps manage your calendar and basic tasks, and learns your preferences so it can adapt and help you better.

Business Benefits and Risks

By now, it should be apparent that the richer experiences enabled by these capabilities are as important to businesses as they are to users. New, context-aware experiences may attract customers and create new revenue. Furthermore, focusing on the user experience may be essential for business survival. If we don't provide rich and appealing user experiences, customers may gravitate toward competitors that do.

Our challenge is to manage the risks associated with these new experiences. The good news is that new security capabilities are emerging to help us do so.

New Security Capabilities

The IT ecosystem is increasingly focusing on building security into hardware, software, and services. We'll all be able to take advantage of this security to protect users and the enterprise. I think of these capabilities as the equivalent of termite-resistant building materials used in construction. They may not prevent termite attacks altogether, but they can stop some of them and minimize the impact of others. For example, Dell is using technology from Cylance to protect the BIOS firmware in its business PCs. The technology is designed to check if systems are secure when users boot them up; after the PC boots, the software checks a hash of the BIOS against a known good version stored in a secure cloud.

Suppliers will need to frequently enhance these defenses to ensure they remain effective. As I noted in Irrefutable Law #6 in Chapter 1, security controls operate in a dynamic environment in which attackers are constantly learning and adapting their approach. Unless the defenses also adapt, they will lose their effectiveness over time.

I expect the ecosystem will increasingly view these security features as a way to differentiate products to meet the needs of distinct categories of customers. As a parallel, think about how the auto and other consumer industries developed. Initially, manufacturers focused on getting the public to buy cars en masse. Accordingly, the focus was on mass-producing just a few models at the lowest cost. As Henry Ford famously said, "Any customer can have a car painted any color that he wants so long as it is black" (Ford and Crowther 1922). Ford's mass-production strategy was enormously successful in popularizing cars among the American public. By 1918, half of all cars in the United States were Model Ts (The Henry Ford Museum 2003). But once consumers became more familiar with cars, they started demanding models that met specific needs. As manufacturers responded, the industry began to develop the huge variety of models that we see today.

In the same way, suppliers will offer a range of products or services with differing levels of security, including higher-security versions for the most sensitive enterprise uses and less-secure versions for consumers. This trend has already been evident for some time in products such as servers and PCs, and we're beginning to see it in cloud services.

In a closely connected trend, we'll see increasing use of contextual information to improve security. Some of this context will be provided by the sensors built into devices, such as cameras and GPS receivers. In addition, analytical and monitoring tools will be able to gather valuable contextual information from the environment. For example, they may examine databases containing information about users' access history and other relevant data.

Baseline Security

A greater level of baseline, hardware-enforced security features will be important in all categories of devices, from smartphones to full-featured PCs. These capabilities will protect the information on the device itself, and the information that is accessed from the device. They'll enable greater trust in the device, and because of this trust we'll be able to provide users of the device with access to more resources, as I described in Chapter 7. The potential business benefits include increased user satisfaction and productivity.

I believe that these features will become particularly valuable as the Internet of Things takes shape. Many new, connected devices and objects won't be powerful enough to run traditional software security controls. Do I expect the computers that control my car or my home to run full intrusion prevention systems or traditional antivirus suites? No, but it is possible to run lightweight AI-based agents that can determine good from bad in milliseconds. This capability has already been demonstrated: in the summer of 2015, Cylance showed its AI-based anti-malware agent running on a Raspberry Pi platform, which is based on the ARM processors that are in many appliances and other IoT devices (Bradley 2015). I also believe that many of these new devices should include protection that limits their functions to the desired purpose, reducing the risk that they could be successfully attacked and manipulated via the Internet or a wireless network.

For enterprise security, these baseline hardware security capabilities will provide help in key focus areas, including threat management, ID and access management, data protection, and remote monitoring. Some expected baseline capabilities include protected environments, encryption, hardware acceleration, enhanced recovery, and integration with security software, as described next.

Protected Environments

Increasingly, hardware will provide protection for essential functions and data in the form of trusted layers and execution environments. I think of this approach as analogous, at the hardware level, to the way organizations are implementing network security zones within an enterprise environment (as described in Chapter 7). The most valuable and critical functions receive the greatest protection, as well as increased monitoring and recovery capabilities.

Attackers have become increasingly adept at compromises using tools, such as rootkits, that operate at or below the operating system level, making them harder to detect and prevent by most traditional security applications. Implementing protection at the hardware level can help prevent compromise of firmware, operating systems, hypervisors, and other fundamental system components. Hardware-level protection can also help alert security professionals to attempted attacks and aid in system recovery. However, hardware-level protection must be designed, developed, and implemented correctly or it could actually do more harm than good, because compromise at this level can give attackers wide-ranging access to the software and data on the system. Concerns have already begun to surface and are growing. Researchers demonstrated the ability to hack the microcontroller inside flash cards, enabling the execution of code that can be used to perform a man-in-the-middle attack (Paganini 2014). Networking equipment supplier Juniper Networks found that its firewall operating system contained "unauthorized code" that surreptitiously decrypted virtual private network traffic (Goodin 2015). MIT researchers suggested there are weaknesses in the implementation of key provisioning for Intel Software Guard Extensions (SGX), a set of hardware instructions designed to improve security by sealing software into hardware-protected enclaves (Chirgwin 2016).

Encryption

Many organizations already use disk encryption to protect data against loss or theft. But in a world where devices are always on and always connected, traditional software-based hard disk encryption is not sufficient. New capabilities will make encryption an even more pervasive technology used to protect information throughout its life, both when it is stored and when it is transmitted. Devices will include self-encrypting drives that maximize protection while minimizing the performance impact; encrypted input-output will help protect data during communications. Capabilities that currently exist in larger systems, such as total memory encryption, will become common in PCs and other end-user devices.

Hardware Acceleration

There's often a trade-off between security and performance. Controls, such as software-based encryption and malware scans, certainly help increase protection, but the performance impact can also increase frustration for users, to such an extent that some may avoid using the security features altogether (see the discussion of control friction and the 9 Box of Controls in Chapter 7). Accelerating functions in hardware can shift the balance in favor of security by decreasing the impact, both on users and on enterprise systems. For example, complex calculations required by standard encryption algorithms can be accelerated using hardware instructions rather than executed entirely in software.

Enhanced Recovery

As I've discussed in previous chapters, we must assume that attempts to compromise are inevitable, despite our best efforts. As attacks become increasingly sophisticated, the ability to recover from compromises will become even more important. Future capabilities will help organizations recover from low-level attacks that target fundamental system components such as firmware or the BIOS. The system will be able to detect changes in these components, whether due to malicious attacks or accidental corruption. It will then be able to take steps to restore the components to a known good state, alerting users and the security team when necessary. Other anticipated recovery features include enhanced capabilities to revoke cryptographic keys to reduce the spread and impact of compromise.

AI-Based Security and Automation

AI-based security applications will play valuable roles in preventing attacks. Today, for example, Cylance uses AI-based agents to distinguish good from bad in milliseconds. These applications will be able to provide an even greater level of protection when they are integrated with hardware-based security, as exemplified by the Dell-Cylance BIOS protection agreement described earlier in this chapter. This kind of integration will enable software to more closely monitor the underlying hardware and firmware for attacks that might otherwise go undetected. For example, security software could use hardware features to detect symptoms, such as memory state changes, caused by specific types of attack. Companies are also researching better ways to authenticate users by employing behavioral biometrics: identifying users based on a combination of hard-to-duplicate characteristics such as they way they swipe characters on a smartphone or even how they walk when carrying the device.

AI will be used more broadly over time to enable a greater level of automation in threat detection, prevention, and response. In the future, AI might be used to dynamically evaluate trust and the corresponding level of access that's provided to a user (see the granular trust model in Chapter 7).

Context-Aware Security

The theme of context awareness underlies many of the rich user experiences described in this chapter. Context awareness can also enhance security: the same sensors and analytical tools that help organizations create personalized experiences can also be used to mitigate risk.

In the home, TVs might be able to recognize when a child is watching, and show only appropriate channels. In supermarkets, cameras that are already used for physical security could help increase the efficiency of automated checkout stations. As I described, image recognition technology can determine a shopper's approximate age. By using this information, perhaps in conjunction with data from a scanned driver's license, the system could help avoid the need for cashiers to manually approve alcohol sales, leading to faster checkouts for consumers and reduced costs for stores.

The sensors in portable devices, such as mobile PCs and smartphones, may also be used to help protect against theft and unauthorized use. A simple case might utilize the device's camera, microphone, and GPS receiver to help authenticate you as the device's owner. If the user looks and sounds like you, and the PC is at your house, we have more confidence that the person using it is really the owner.

Additional technologies in portable devices, such as NFC, will allow more sophisticated examples of context-aware security. Devices will know when they're no longer in proximity of their owner, and may enter a protected state to prevent data loss. If your phone is near your laptop, we have greater confidence that you are the user trying to access the information on the laptop. When your phone moves away, the laptop deduces that you have moved away, too, and begins to armor itself by locking the screen. As you move progressively farther away, the laptop first goes into standby to save power, and then begins encrypting its contents for protection.

The GPS receiver in a portable device can also be used to geofence the device and the data it contains. If the receiver detects that a PC has moved outside a specific area, the device could alert the owner and the enterprise support team. The same capabilities could help protect data whose movement is restricted by specific geography-related requirements such as export controls. The device could detect when it's in a country subject to these controls, and encrypt the data it contains to protect it.

Cloud Security and Context Awareness

Cloud service providers recognize that some organizations are still reluctant to move critical data to external clouds due to security, regulatory, and privacy concerns. Suppliers have been working to add security capabilities designed to address these concerns. As they do so, we can expect more cloud services that are differentiated based on the level of trust they offer.

Suppliers might offer a “plain vanilla” cloud service for noncritical applications, along with a more expensive high-trust cloud service. Besides offering additional technical controls, secure clouds might include guarantees that the supplier will meet specific privacy and other data-protection regulatory requirements. This tiered strategy resembles the zoned approach to network security that organizations are implementing as part of their evolving security architecture. Zones that host critical applications are protected by a variety of controls, ranging from network segmentation and hardened virtualization host servers to additional monitoring.

In the future, client-aware cloud services will be able to tailor the access they provide based on the security capabilities of the client in order to mitigate risk. A fully managed device that includes hardware-based enterprise security features and a full software security suite may get more access than an unsecured personal device. At the same time, a cloud-aware client will be able to validate that the cloud service it is accessing is genuine, and that it offers the required level of security.

As businesses use a growing number of cloud services, security requirements become more complex. A single enterprise may use multiple external cloud services while also operating a private cloud and a traditional computing environment. It will be important to streamline access for users. We can expect more emphasis on technology that eliminates the need for users to authenticate to each individual service.

Security Analytics and Data Protection

Security context can be provided not only by sensors, but also by analyzing information about the enterprise environment and the threat landscape. As attackers become stealthier, this analysis will become an increasingly important part of an organization’s defenses. Within the industry, many are moving toward the use of security analytics tools to analyze patterns of network traffic and system use. I expect to see increasingly sophisticated external services that analyze a broad range of information in order to thwart attacks.

As information is used on more devices outside the enterprise network perimeter, it will also be increasingly important to focus on controls that are integrated with the data itself. Many organizations are already protecting information with technologies such as enterprise rights management. In the future, these capabilities are likely to become more sophisticated and automated, allowing businesses to define policies that automatically store sensitive data in highly secured locations.

Conclusion

New technologies bring challenges, but they also bring opportunities for the CISO and for the organization overall.

The rich context-aware experiences that I’ve described in this chapter are entirely dependent on IT. To deliver these experiences, organizations will need to understand and manage the risks. As the experts in information risk, CISOs and other security professionals should have opportunities to become closely involved in the development and implementation of key business initiatives. This will result in a higher profile for the information risk and security team across the entire organization.

To fully take advantage of these opportunities, CISOs will need broad business and people skills as well as a thorough knowledge of security controls. With the addition of these skills, I believe the role will evolve into the chief security and trust officer (CSTO), with broad responsibilities to enable the business through trusted infrastructure, applications, and business processes. As this transition occurs, the CSTO becomes the essential enterprise architect, with the IT organization becoming a peer or perhaps a subordinate. I'll discuss these skills further in the next chapter.

CHAPTER 9



Corporate Social Responsibility: The Ethics of Managing Information Risk

Be the change you wish to see in the world.

—Gandhi

In the past year or so, we have passed a major inflection point; it has become clear that almost every powered device will compute, communicate, and have an IP address. As technology becomes embedded into the fabric of our lives, exploits that take advantage of technology vulnerabilities may increasingly impact the well-being of almost everyone in society. This makes it particularly important that we apply the right ethical values to shape the way we design, develop, and implement these technologies.

The past few years have seen an escalating cycle of risk, with correspondingly greater impacts for businesses and individuals. If that trajectory continues as technology becomes more pervasive, the implications for society could be catastrophic. This means we should all, as security professionals, contemplate our ethical responsibilities not only to the organizations we work for, the customers we serve, and the company's shareholders, but also to society. To put it another way, I believe that information security and privacy are issues of corporate social responsibility.

Yet even as it becomes even more important to consistently apply an ethical approach to managing information risk, business demands and other challenges can make it increasingly difficult to do so. Companies' continuous efforts to drive growth and accelerate time to market translate into demand for faster implementation of internal systems and new technology-based products. At the same time, implementing effective security and privacy is becoming more difficult due to a more complex threat landscape and the expanding, fragmented regulatory environment.

These factors result in increasing pressure on technology and business professionals to take risky short cuts. In some cases, there may be clear conflicts between business priorities, such as the deadline for launching a new product, and "doing the right thing" in security and privacy terms. There are also many gray areas in which the right course

of action is not immediately clear; whether to expend resources on protection against a threat that's still on the distant horizon, for example. I'll explore these ethical dilemmas, and offer suggestions about how to find solutions to them, later in this chapter.

WHAT IS CORPORATE SOCIAL RESPONSIBILITY?

Definitions of corporate social responsibility typically focus on the idea that companies look beyond their profits and legal obligations to their broader role in society. A common theme is that a company should take into account the social, ethical, and environmental effects of its activities on its employees and the community around it. Here are three definitions that summarize some of the key concepts:

“The notion of companies looking beyond profits to their role in society is generally termed corporate social responsibility (CSR)... It refers to a company linking itself with ethical values, transparency, employee relations, compliance with legal requirements, and overall respect for the communities in which they operate. It goes beyond the occasional community service action, however, as CSR is a corporate philosophy that drives strategic decision-making, partner selection, hiring practices, and, ultimately, brand development.” (McComb 2002)

“CSR is about businesses and other organizations going beyond the legal obligations to manage the impact they have on the environment and society. In particular, this could include how organizations interact with their employees, suppliers, customers, and the communities in which they operate, as well as the extent they attempt to protect the environment.” (Lea 2002)

“The continuing commitment by business to behave ethically and contribute to economic development while improving the quality of life of the workforce and their families as well as of the local community and society at large.” (World Business Council for Sustainable Development 2007)

The Expanding Scope of Corporate Social Responsibility

Despite the obvious societal implications of security and privacy risks, most companies don't consider them to be CSR issues today. That may change over time, as public and corporate awareness of the risks continues to expand. Already, some major technology companies include descriptions of how they manage security, privacy, and business continuity in their CSR reports (see sidebar). That trend may spread as companies in other industries add more technology-based products and services.

Consumer data protection is one area of information risk that is already widely treated as a CSR issue; it is even included in the International Standards Organization corporate social responsibility standard (ISO 26000). As Forrester Research analyst Heidi Shey put it, “It’s time to start thinking of protecting customer data as a corporate social responsibility, and not to check off boxes for compliance or a thing that must be done so you can avoid some nasty breach costs.” (Shey 2014).

In terms of the potential impact on society, security and privacy could be considered a digital extension of consumer safety, which companies have viewed as a CSR issue for many years. Furthermore, a quick review of the history of CSR shows that its scope has continually evolved and broadened to include new issues, typically as public awareness of those issues has increased. For example, it’s not so long ago that rivers and oceans were used not only as human sewage dumps but also as a convenient method for disposing of industrial waste; as late as 1969, one large river in Ohio was so polluted that it regularly caught fire. Yet today, discussions of environmental impacts are typical in CSR reports, and in the last few years have further evolved into a focus on climate change: in 2015, 82% of the world’s largest companies included data about carbon emissions in their reports (KPMG International 2015).

While early social-responsibility efforts were often philanthropic in nature (such as the funding for public libraries and education provided by Andrew Carnegie, founder of US Steel), corporate social responsibility reporting is now a mainstream business practice worldwide, undertaken by more than 90% of the world’s largest companies.

TECHNOLOGY COMPANIES THAT TREAT INFORMATION RISK AS CSR

Some large technology companies—including Cisco, Microsoft, and Intel—already position information risk areas such as security, privacy, and business continuity as corporate social responsibility items, and discuss them in their CSR reports. While the reports devote space to the companies’ achievements, they also describe corporate positions and principles on key issues such as data protection and transparency. Cisco’s 2015 CSR report, for example, notes the company’s commitment to produce a twice-yearly transparency report that includes data requests or demands for customer data received from law enforcement and national security agencies around the world (Cisco 2015).

Apple CEO Tim Cook has also spoken out about his company’s commitment to privacy and security, particularly when protecting user data. In a letter published on the company’s web site, he said: “We don’t “monetize” the information you store on your iPhone or in iCloud. And we don’t read your e-mail or your messages to get information to market to you.” Cook has argued vociferously that government should not have “back door” access to systems in order to thwart terrorism. “The reality is if you put a back door in, that back door’s for everybody, for good guys and bad guys,” he said on CBS’ 60 Minutes (Rose 2015). “I don’t believe that the tradeoff here is privacy versus national security. I think that’s an overly simplistic view....we should have both.”

The Evolution of Technology and Its Impact

To continue the exploration of why I believe security and privacy is a matter of corporate social responsibility, here’s another quick historical perspective, this time examining the emergence of information risk in the context of technology’s evolution.

The march of technology can be viewed as a succession of major waves, each lasting roughly 100 years (Rifkin 2013). Each wave has brought transformative benefits to society, but also significant challenges. The first wave, starting in the 1760s, included steam power, railways, and early factories as well as mass education and printing. The second wave, starting roughly in the 1860s and continuing well past the mid-1900s, included automobiles, electricity, mass production, and had an even bigger effect on society. Many of today’s corporate social responsibility issues today are the negative impacts of those first two waves of technology: examples are environmental impacts due to industrial production, mining, and oil drilling; factory working conditions; and the safety of mass-produced items.

Table 9-1. *The March of Technology*

Version 1.0: 1760s	Version 2.0: 1860s	Version 3.0: 1990s
Steam and coal	Electric lights	The Internet
Railways	Communications	Molecular biology
Factories	Oil and gas	Renewable energy
Printing press	Mass production	“Smart” everything
Mass education	Automobiles	

The third wave began in the 1960s, with early computers, but only really gained momentum in the 1990s. It includes the Internet and smart “things,” molecular biology and genetic engineering, and renewable energy. Arguably, this technology wave may have the broadest impact on society of any to date. Each previous wave lasted about 100 years, so history suggests that we are far from reaching the crest. If this wave was a movie, we’d still be watching the opening credits.

If the opportunities presented by this third wave of technology are unparalleled, so are the risks to society. As I’ve argued in earlier chapters, as technology has spread exponentially, so have the threats and their impacts, while security controls have progressed at a more linear, incremental rate. As a result, there’s a continually growing gap between the capabilities of the controls and the impact of exploits. If the impact of security breaches seems big now, consider what the impact will be in 10, 20, or 50 years, when technology is even more pervasive throughout society.

Let’s consider some of the potential impacts by reiterating two examples from Chapter 6. Last year, doctors for the first time inserted an artificial “eye” that enabled a blind person to see. The device is a retinal implant that receives signals from a video camera integrated into eyeglasses. Think ahead a few years, to a time when the implants are more sophisticated and can see in much higher resolution, and also include software to automatically interpret visual information, such as QR codes. Then imagine that a malicious actor creates a QR code that triggers the vision system to download malware.

Like the PC malware that paralyzed Sony’s network in 2014, the malware then demands a ransom to re-enable the person’s vision. Now consider the example of a cement company that’s embedding sensors in the concrete mix used to build a new road, thus enabling local authorities to monitor traffic patterns and adjust signals to optimize the flow of vehicles. If the technology is not securely designed and implemented, all that a malicious person needs is the ability to execute malicious code, in order to falsify the traffic pattern in such a way that vehicles converge on the scene of a planned bomb attack.

Here’s example of a real-life attack that unfortunately has already occurred. Over a four-day period during November 2008, members of an Islamic militant organization carried out a series of 12 coordinated shooting and bombing attacks across Mumbai. The attacks killed 164 people and wounded at least 308. Of the funding that enabled the attack, \$2 million was raised by cyber crime (Goodman 2015). Think about how cyber crime works. Typically, the cybercrime cycle starts with stealing someone’s identity by installing malicious code on a device or by taking advantage of insecure behavior. So ask yourself: If I don’t keep my systems up to date, if I don’t design and implement them well, and educate employees to ensure they are security-aware, am I indirectly contributing to terrorism? The answer is that you might be—although in most cases, you won’t even know it.

As I discussed in Chapter 6, four motivations account for the majority of serious exploits. Terrorism is one. The others are financial gain, warfare, and hacktivism. Each of these motivations can result in consequences with broad impacts across society: economic damage, loss of services, damage to morale, degradation of government services, and even human casualties.

As all companies become technology companies, the technology they create and deploy may be exposed to exploits with potential impact on society. The same applies, of course, to public-sector organizations. Even though this idea is becoming more widely accepted, I occasionally encounter people who don’t believe it applies to their organization. Recently, as I fielded questions after giving a talk, an audience member commented that she was on the board of a local school and definitely didn’t see the school as a technology organization. “Does your school have a web site that parents and kids can use to view and update information?” I asked. She said yes. Then I asked “Does your school have an app that lets parents check whether their kids attend class?” No, she said, but the school was considering it. “Let’s imagine you have a web site that’s not well designed, and a malicious person decides to take advantage of that with a zero-day exploit,” I said. “He can compromise the site and the personal information of the parents and children that use it.” I added that if a school takes its technology to the next level by making an app available to parents or kids, it becomes even more clearly a technology supplier—and its security concerns now include product vulnerabilities. By the time I’d finished explaining, the audience member asked me if I could come and explain the issues to her board, which of course I agreed to do.

Here’s another school example, one that highlights the risks of failing to consider all the ethical implications: A Pennsylvania school district issued laptops to some 2,300 students, then remotely activated the laptops’ webcams—without informing the students—and used the webcams to secretly snap students at home, including in their bedrooms. Surveillance software on the laptops also tracked students’ chat logs and the web sites they visited, and then transmitted the data to servers, where school authorities reviewed and shared the information and in at least one case used it to discipline a student. Ultimately, the school district was forced to settle a class-action lawsuit that charged it had infringed on the students’ privacy rights (Bonus 2010).

Maintaining Society's Trust

The third wave of technology offers opportunities for all organizations. But as the opportunities increase, so does the obligation to use technology responsibly. If we don't implement appropriate security and privacy protection, consumers won't trust the technology. If they don't trust the technology, they will be reluctant to use it. This could potentially affect any company that supplies technology, and impact the growth of the digital economy overall.

Unfortunately, the privacy and security breaches that have hit the headlines in recent years have weakened that trust. As a result, consumers' trust in technology sank last year in 70 percent of countries surveyed worldwide, according to the Edelman Trust Barometer, a widely used indicator of trust in business and government. Worse, the rapid implementation of new technologies that are changing everyday life, "from food to fuel to finance," emerged as a new factor depressing trust overall. "By a two-to-one margin, respondents in all nations feel the new developments in business are going too fast and there is not adequate testing," the study concluded (Edelman 2015).

Top US regulators have urged companies to expand and clarify their privacy efforts. Federal Communications Commission chairman Tom Wheeler said Internet service providers have a responsibility to make sure personal data is held securely and that companies are transparent about the data that's being captured. "There's no question that with connected devices, data is becoming today's currency, and we need to be aware of the impact of that on consumers," added Federal Trade Commission Chairwoman Edith Ramirez, noting a recent Pew Research Center survey found that 47% of Americans lacked confidence that they understand what companies will do with their personal information, and had mixed feelings about whether or not to share it (Hamblen 2016). The weakening of trust is a dangerous trend. Breaking someone's trust is like crumpling up a perfect piece of paper: you can work to smooth it over, but it will never be the same again.

All organizations inevitably experience security and privacy issues. The question is how we respond to them. We can manage them in way that focuses on limiting our liability, or we can focus on doing the right thing for those who may be impacted. I recently participated in a peer group discussion that evolved into an intense debate on this very issue. The discussion was prompted by the major breaches that occurred in 2014 and 2015; as a group, we discussed how we might jointly develop the concept of a "minimum standard of care" for security and privacy. Some people wanted to focus on limiting corporate liability for a breach. I believed that was the wrong goal, and argued that the primary focus should be on protecting our customers. My reasoning was that if we protected our customers, we would limit our liability as a natural consequence. But if we focused only on limiting liability, we would likely fail to take the necessary steps to protect our customers. Furthermore, I believed that the lens we chose to view the problem with would bias strategy and outcomes over the long term. A liability-focused standard would inevitably cause us to direct our efforts into seeking ways to limit our responsibility for the technology we create and manage. But if the standard focused on protecting the people who might be impacted, we would direct our efforts to thinking about how best to prevent, detect, and respond to risks.

The Ethics of Managing Information Risk

Some professions, such as certified public accountants and doctors, have ethical standards that may require them in some cases to break ranks with their organizations, such as if they see signs of illegal activities or financial manipulation. We expect doctors to be personally accountable for decisions that affect the lives of their patients, rather than simply deflecting responsibility for health decisions onto someone else within the organization. If CPAs or doctors fail to meet these professional and ethical standards, they may lose their ability to practice.

Although there are many professional certifications for security and privacy professionals, there's currently no equivalent to these medical or legal qualifications. Security and privacy managers are not automatically barred from practicing their trade if they fail to meet professional standards. However, we should all assume a similar level of personal accountability for our decisions—especially since our actions may have broader implications for society. Regrettably, not all of us do. Some security and privacy managers see their role as simply managing a risk register: they identify the risks, and perform the analysis and associated cost estimates, but then they take the register to other executives who then make the decisions. By doing so, they are abdicating responsibility and deflecting accountability onto someone else.

As the senior security and privacy professional within the organization, CSPOs should share responsibility for information risk decisions equally with the other corporate executives and the board. People are often told that they need to “think like an owner;” we need to act like an owner too. And ultimately, we need to think about our responsibility to all the people we work for—including customers and anyone else in society impacted by our actions—as well as our responsibility to the executives we report to. If you don't think your manager is right, think hard about the possible consequences of not speaking out and where your responsibility ultimately lies.

The recent events at automakers have shown all too clearly what can happen when corporate culture produces a system in which professionals are driven to behave unethically in order to meet business goals, or fail to take responsibility for their actions, while senior executives apparently remain ignorant. In the Volkswagen emissions-testing scandal, engineers included software specifically to deceive test equipment so that cars could meet the emissions targets required for sale in the US. An investigation into General Motors ignition-switch problems that caused at least 13 deaths described the “GM Salute,” in which employees sat in meetings, with their arms folded and pointing outward at others, as if to say that the responsibility lay with those other people, not with the employees (Maynard 2014). At both automakers, top executives said they were unaware of the actions of the lower-ranking employees who were directly involved in the issues.

In our daily lives, we encounter many situations in which we need not only to decide on the right course of action, but also to take responsibility for voicing our opinions so that they are considered by the company as a whole. Suppose that a business manager is proposing an action that's legal but conflicts with our security values and approach to protecting customers' information. Or imagine that implementing the right level of protection risks the target dates for a critical product launch. Or that failing to tell customers or suppliers about a potential vulnerability becomes the equivalent of a lie.

In the book *Giving Voice to Values*, author and educator Mary Gentile discusses the ethical dilemmas that many people face in businesses today. Her assumption, as she observes in the book, is that “in many if not most of the managerial and financial misbehaviors we have seen in the past, there were enough people who recognized the lapses in ethics and judgment to have stopped them. The problem was that they did not believe it was possible to do so.” Gentile then focuses on providing techniques to help people voice their concerns and take action at “those times and situations when we believe we know what is right and want to do it, but we experience external pressures—from our boss, our colleagues, our customers—to do otherwise. As a result, we are not sure how to raise our concerns.”

DISCLOSING SECURITY ISSUES: A TALE OF TWO COMPANIES

Questions about how to deal with the discovery and disclosure of security issues are likely to generate difficult ethical discussions for many companies. The following examples show how two companies dealt with security issues in very different ways.

In December 2015, networking vendor Juniper Networks disclosed that an internal code review had discovered “unauthorized code” in its firewall operating system that could allow hackers to gain administrative access and decrypt encrypted VPN traffic. The company said it had not received any reports of exploits using the vulnerability; it said it had released patches to fix the problem and urged customers to update their systems (Worrall 2015). This is a case in which a company appears to have managed a difficult issue well, in my opinion. It highlights the tough questions and discussions that companies face when managing potential security issues. How deeply do you test and review your code, knowing that the deeper you dig the more likely you are to find vulnerabilities? If you do find a problem, how do you handle it? Do you disclose it, quietly fix it, or even ignore it? Does your company have the right value structure to ensure that decisions reflect its responsibilities to customers and to society?

Now consider a contrasting example. In 2015, a vendor of dental practice-management software agreed to pay \$250,000 to settle US Federal Trade Commission (FTC) charges that it falsely advertised the level of encryption it provided to protect patient data (Federal Trade Commission 2016). According to the FTC, the company spent two years touting its “encryption capabilities” for protecting patient information and meeting “data protection regulations”—yet at the time, it was well aware that its software didn’t provide the encryption required by HIPAA. It seems clear that a company that makes deceptive claims of this kind lacks a value structure capable of ensuring ethical security and privacy decisions.

The challenges described in *Giving Voice to Values* probably seem familiar to many of us who are responsible for managing information risk (see sidebar article). First, how do we decide what is the ethical course of action? Then, how do we take action by voicing our opinions when it really matters?

One starting point is to define the organization's critical security and privacy principles, which then can serve to guide our decisions. These principles should be derived from the organization's corporate values. For example, a company that prioritizes customer service should also be committed to protecting customer information, and therefore its critical principles should include privacy by design.

We then need to think about how to focus the company on those principles: how we create the right language to express the principles to others, and how we enroll our organizations in principle-based decision making. We need to make security and privacy clearly visible in the decision-making process, not just within the information security organization but across the entire organization. That sends a message to everyone, including customers as well as people within the organization, that security and privacy are corporate priorities. By demonstrating our commitment to these principles, we can create trust in our organization and its technology.

We can use our security and privacy principles as a compass to guide us through the dilemmas we encounter. We can approach these dilemmas using the same framework that we apply to any source of information risk: sense, interpret, and act (see Chapter 3).

- **Sense:** Are changes on the way that that could conflict with our security and privacy principles? What is the dilemma that we will face?
- **Interpret:** Analyze the issue to determine the following: Can I make this decision? Which of our principles can guide my decision? Who do I need to talk to? What actions can I take, and what are the direct and indirect consequences of each?
- **Act:** Will my action align with the organization's best interests? What about the interests of our customers, and of society in general? Will my action or lack of action create embarrassment for the company? Is my action practical? Who should I tell?

Conclusion

As we progress through the third wave of technology, and our reliance on technology expands even further, so does the potential societal impact of security and privacy issues. Our professional and ethical responsibilities require that we hold ourselves accountable for doing what we know is right. This is true today, and will be even more so in the future. This means that we will have to take career risks to make sure that security and privacy are appropriately handled within the organization, including ensuring that issues are discussed at board level. I'll discuss how to do this in more detail in the next chapter on the 21st Century CISO.

CHAPTER 10



The 21st Century CISO

Leadership is the art of mobilizing others to want to struggle for shared aspirations.

—Jim Kouzes and Barry Posner,
The Leadership Challenge

The finance director sounded frustrated and exhausted. Our IT auditors had been trying to tell her about an obscure yet important data backup problem that affected SOX compliance. But her background was in accounting, not technology, and as the IT experts presented page after page of technical information elaborating the intricacies of backup processes, her eyes glazed over. The more they tried to explain by adding yet another layer of detail, the more confused and frustrated she became.

That's when I thought of a solution. "Imagine," I said, "we've got a passenger train running from station A to station B. That's what our backups are like; they're carrying data from our servers to tape."

"We know the train arrived at station B, so we know the backup occurred," I said. "But we don't know how many passengers got on at station A, and we don't know how many got off at station B. So we can't definitively say we actually backed up all the information, and to comply with SOX, we need to be certain."

The finance director sat up. For the first time since the start of the presentation, she seemed alert and engaged. And from that point on, we made progress. She asked how we planned to solve the problem, we briefly mentioned a couple of the possible solutions, and the meeting ended on an upbeat note.

My storytelling, using an off-the-cuff metaphor, succeeded where the more traditional approach had failed. It communicated a technical security issue in terms that a senior businessperson could understand and remember. And it illustrates one of the key skills of the 21st century CISO. We need to extend our reach outside the security organization to communicate with and influence people at all levels, from all backgrounds.

Chief Trust Officer

In this chapter, I'll explain some of the skills and traits I believe CISOs need in order to fulfill their changing role. To set the stage, I'd like to step back for a moment and briefly recap the changing focus of information security overall.

As I've discussed earlier in the book, every company is becoming a technology company. And as the potential impact of information risk expands, it is becoming essential to manage security and privacy as a corporate social responsibility. The CISO's role should therefore expand to span the full breadth of information-related risks, as described in Chapter 1. At many organizations, this is already happening. CISOs are taking on responsibility for privacy, regulatory compliance, and product and service security, in addition to more traditional IT security functions.

This is a huge opportunity for CISOs to step into a more valuable, high-profile role within the organization. The core skills of information security professionals—evaluating and mitigating risk—are as essential for mitigating new risks associated with product security, privacy, and regulatory compliance as they are for more traditional IT-related threats. But perhaps this broader role requires a different title that more accurately reflects the convergence of risk responsibilities, such as Chief Trust Officer or Chief Information Risk Officer.

Taking on a larger role requires a broader view and a corresponding set of skills. We need to communicate in terms that business people understand, and build relationships that enable us to influence people at all levels across the organization. We also need extensive management and leadership skills, both to operate at an executive level and to inspire our expanded risk and security team.

The ability to manage the full range of information-related risks is a necessity, not just for the CISO, but for the organization. If we do not step into a broader role, the organization must acquire these abilities elsewhere. Because of this, CISOs who do not adapt to this role run the risk of becoming irrelevant to the organization. Alternatively, these risk areas will be managed in a stove-piped, fragmented way, in which case the organization may never discuss the aggregation of risks and the controls necessary to manage them. If this occurs, organizations will certainly generate unmanaged risks to themselves, their customers, and to society.

Until recently, one of the CISO's biggest challenges was obtaining funding for security initiatives. Today, due to the prevalence of large breaches, it's often easier to find funding. But more funding doesn't always lead to greater security or a better outcome for the organization. Sometimes the fear of breaches drives organizations to invest heavily in controls that generate a high degree of control friction, restricting users' ability to do their jobs. For example, some organizations have installed controls that prevent users from downloading apps or files, or even accessing some web sites. These controls threaten to stifle users' ability to innovate and hinder overall business velocity. Furthermore, determined users will find ways around the controls, such as using less-secure personal systems to access "forbidden" resources.

CISOs need business acumen to understand the impact of security controls on others in the organization. As I discussed earlier in the book, our approach to security architecture should start with an understanding of the 9 Box of Controls, including the friction that controls can generate. Business acumen is also necessary to communicate technical risks in language that nontechnical people in the business can grasp, and to understand that some risks are worth taking. Risk-taking is fundamental to business. Without it, no business value would be created.

The Z-Shaped Individual

If we don't already have the skills required of the 21st century CISO, we need to acquire them.

To some extent, this trend parallels what is happening in most technology-related professions: IT professionals need to acquire business acumen as well as depth of IT knowledge. The concept of “T-shaped” individuals has been widely used to describe the idea that IT professionals need to be able to provide value horizontally, across business groups in the organization, as well as vertically at all levels within IT.

This concept is useful, but it doesn't fully encompass the skills of the 21st century CISO. The unique role of CISOs and other security professionals might be better represented as a “Z-shaped” individual, as shown in Figure 10-1. Adding the third dimension of core security skills, such as risk assessment and understanding of controls, allows us to deliver value across the business and all areas of IT.

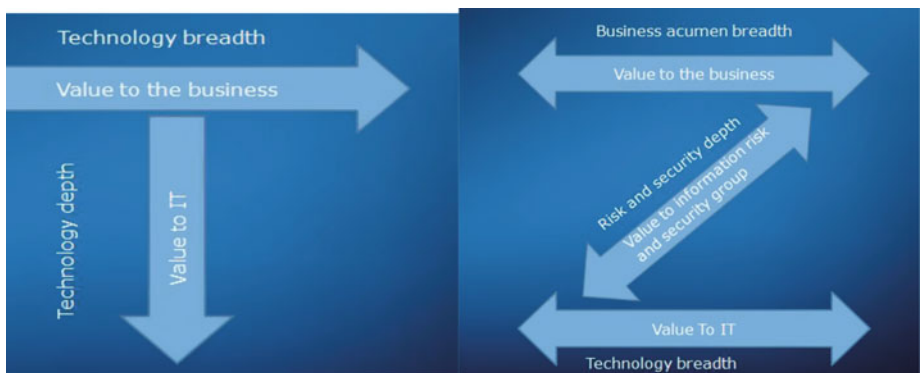


Figure 10-1. *The T-Shaped IT professional (left) and the Z-shaped CISO (right)*

The 21st century CISO needs to understand business priorities and processes well enough to identify how security controls help or constrain the business. To gain this level of understanding, he or she has probably gained experience in areas that are central to the company's business, which, of course, vary depending on the company's core focus. For example, the CISO might previously have worked in manufacturing operations, services, or mergers and acquisitions.

The CISO needs technical knowledge too, although the depth of technical knowledge required remains a subject of intense debate among my peers. I've observed CISOs at smaller and less-complex organizations who feel they need deeper technical skills to do their jobs. This is not surprising. With much smaller security teams, CISOs at smaller companies may need to be more involved in day-to-day technical details as well as managing people. At larger and more complex organizations, CISOs are less likely to spend time delving into technical detail.

However, all CISOs need to be able to understand enough about the technology to absorb the important issues and communicate these issues to other managers outside the security group. This means that our technical knowledge must be broad, ranging from devices to data centers. We need to know enough about devices, such as smartphones, PCs, tablets, and new evolving device types such as wearables, to understand the security implications as well as the benefits. At the other end of the scale, we need to know enough about data centers and physical access controls to understand and communicate the important security requirements and challenges.

Our core risk management and security skills provide the link that completes the “Z” by connecting technology and business. We understand how to assess and manage risk by applying procedural, technical, and physical controls to meet the organization’s legal, privacy, and security requirements.

Foundational Skills

Becoming a Z-shaped individual is the foundation for one of the 21st century CISO’s essential traits: establishing credibility across the organization. We must be credible in order to build trusted relationships with executives and specialists across the organization and to discuss the vast range of issues that affect the business. This credibility is built on the competence that comes from understanding the business and technology as well as possessing core security skills. By becoming Z-shaped, we will also be better positioned to influence risk management for the company’s product and service strategy, as opposed to having those risks managed independently by another group.

Our ability to influence the organization also springs from a clear mission. I use the term *centered* to describe this. We can effectively present our case because we have a strong sense of purpose and a clear understanding of why the security group exists and what we are trying to achieve.

This idea returns us to the theme of this book: Protect to Enable. In our global economy, most companies operate in highly competitive markets. As the security organization, our mission is to enable the free flow of information and rapid implementation of new capabilities to ensure success and long-term competitive survival. Other CISOs may work at more risk-averse organizations, and therefore some aspects of their mission may differ. However, the mission always needs to be aligned with the organization’s business priorities. It is essential that this mission becomes a part of who we are and why we exist. It provides a sense of purpose that lends authenticity and consistency to our actions and helps us build credibility across the organization.

As we all know, security can be a particularly distracting profession, with a constant barrage of day-to-day emergencies and diversions. So we need a clear mission in order to retain a strong sense of direction. Like expert sailors, we can progress toward our goal amid the day-to-day distractions and diversions, making continual adjustments and corrections to stay on course as the winds shift.

We also need to retain a sense of curiosity. To engage with others, we need to be genuinely interested in what they do. This curiosity enables us to continue to learn, building on and broadening the competencies that then enhance our credibility.

Another major reason we need to be learners is to stay ahead of the enemy. Threat agents are always learning because they must. As new threats emerge, we put in place new controls. But once implemented, these controls tend to be static, while threat agents

are dynamic, coming up with new techniques to bypass the controls. Therefore, our thinking must also be dynamic, and we must continually learn in order to protect against ever-evolving threats.

Becoming a Storyteller

We cannot influence people unless we communicate with them. And as the scope of information risk expands, we need to communicate with a wider range of people across the organization.

Communicating with people isn't always easy, as most of us have discovered. If we start relaying technology details to those who aren't technologists, we won't capture their interest. In fact, we run the risk of doing the opposite, as I described in the example at the start of this chapter.

To communicate, CISOs must become chameleon-like, with the ability to blend into a variety of environments. We need enough knowledge of each business domain to be able to communicate with different groups using language they understand. And we need to discuss these subjects at different levels. A CFO may only want to hear a high-level summary expressed in terms of financial impact and return, which is often not easy when discussing security investments targeting hard-to-quantify threats. Product group managers want to hear security issues expressed in terms that relate to sales, marketing, and operational efficiency.

I've found storytelling to be a powerful tool for communicating with diverse people across the organization. When I frame security issues as stories and images that people can understand, they relate better to the issues even if they lack a background in technology.

I like to tell stories using metaphors and analogies. They are easily remembered, and they translate complex subjects into simple terms everyone can understand. In fact, the metaphors I've used throughout this book, such as the perfect storm in Chapter 1, the train backup in this chapter, and the roundabouts and traffic lights in Chapter 5, have helped me communicate security issues to many people. To use yet another analogy, orchestra conductor Benjamin Zander said, "The conductor of the orchestra doesn't make a sound. His power comes from awakening possibility in others." (Zander and Zander 2000). In the same way, I believe the power of the CISO comes from awakening the awareness of risk among people across the organization. I use stories based on metaphors to create that awareness.

For example, employees often find it hard to understand the dangers of stealthy threats. This is because the threats are unobtrusive, concealing themselves so they can steal information over the long term. Users are usually not even aware that a problem exists on their system. They still associate malware with obvious, annoying symptoms such as screen messages and system crashes. So when we tell them we've detected dangerous software on their machine, they have a hard time believing that it matters. That is why we have to focus on prevention using low-friction controls. If we do not achieve this as a profession, we will perpetuate the worsening cycle of risk we are experiencing today.

To communicate the danger, and the need for effective preventative controls, I sometimes use the analogy of ants and termites. "Malware used to be like food-eating ants in the kitchen," I explain. "You'd know when you had an infestation because you'd see ants crawling over the countertops and walls. Once you knew about them, you'd spray or set traps to eliminate them."

“But today, threats are more like the termites that can live in your walls. You can’t see them, and you may not even know they are there. But they’re doing much more damage than ants ever did. In fact, they may be destroying the structural integrity of your house.”

I’ve found using analogies helps quickly drive home messages. People immediately understand that these invisible threats can undermine the structure of the computing environment, just as termites undermine houses. This makes them more likely to accept the next step, which is that we have to perform the digital equivalent of tenting their computer to eradicate the vermin, but without toxicity to users or the computing environment.

THE NIST FRAMEWORK: A COMMON LANGUAGE FOR RISK MANAGEMENT

To discuss information risk management across the organization, it’s helpful to use a common language that everyone, including non-technical people, can understand. I’ve found the National Institute of Standards and Technology (NIST) Cybersecurity Framework to be a helpful tool for communicating the issues. Development of the framework was triggered by a 2013 presidential executive order on improving the security of critical infrastructure. This led to a year-long private-sector-led effort to develop a voluntary how-to guide for organizations. Many companies contributed input about standards, best practices, and guidelines to that effort. I was one of the first security leaders among the Fortune 500 companies to engage the framework.

The framework creates a common taxonomy and terminology for managing risk, making it easier for security teams and others to communicate. It fosters collaboration. In addition, each organization can measure its risk management maturity level against the framework. As the framework is used by more people, including business executives, it may help to increase the overall understanding of information risk and how to manage it, which would be a good thing for all organizations.

Fear Is Junk Food

Just as building trusted relationships is essential to influencing the organization, I also think we need to transcend the doom-and-gloom that can pervade discussions of security topics.

The security industry has a tendency to use fear to sell products. Unfortunately, this tendency reflects the fact that many people in the security industry profit from insecurity: their revenue grows when more breaches and other incidents occur. Internally, as security professionals, we sometimes share this tendency to use fear as a tool to obtain additional budget or other resources. Of course, security really is about scary things: threats, vulnerabilities, and risk. But focusing on fear as the primary motivator is like living on a diet of junk food. It may provide immediate gratification, and it’s somewhat addictive, but ultimately it’s not healthy for either the CISO or the rest of the organization.

In the short term, fear can scare people into action and help drive funding for security projects. However, relying on fear alone can only work for so long. Eventually, it has the opposite effect. It causes the CISO to lose credibility. In fact, I think relying on fear may even contribute to the high rate of job turnover among CISOs. Those who rely too much on selling fear are snacking on an unhealthy diet, and eventually the organization realizes this and rejects them.

Ultimately, fear doesn't work for other reasons too. Most people don't want to listen to a continuous stream of negativity. If we are always seen as the source of negativity, we will lose our audience. If we are continually viewed as the group that says no, we will be ignored. People will bypass security restrictions in order to meet their business needs.

Even within the security organization, fear can become a gravitational force, a black hole drawing ever-increasing attention to the negative side of security issues and draining energy that should be directed to enabling the business. This is why we need to focus on solutions that deliver the three key benefits I discussed earlier in the book: a demonstrable and sustainable bend in the curve of risk; the ability to lower the total cost of controls; and low control friction to improve business velocity and the user experience.

Accentuating the Positive

So how do we take a more positive approach? By focusing on our mission, which is to Protect to Enable. This mission shifts the emphasis from the negative to the positive: how we can help the business achieve its goals by solving these information risk and security problems. It puts hope and optimism before the challenge.

This mission is aligned with the business. Rather than being antagonistic, it is based on common values. It sets an optimistic tone, and, in the long term, optimism is a far better motivator than pessimism. Threats may be frightening, but our goal is to see past the threats and identify the opportunities. To paraphrase the noted Stanford University behavioral scientist Chip Heath, there's no problem that cannot be solved without a new framework. Therefore, if we can't see a solution, we have the wrong framework. Protect to Enable provides a new framework. So does the 9 Box of Controls, with its focus on cost efficiency and control friction as well as effectiveness. These tools help us focus on finding solutions.

Imagine you're invited to attend a meeting to discuss whether the company should start using a specific cloud-based business application from a new supplier. Clearly, this product introduces risks: it comes from an unfamiliar supplier, it's accessed over the Internet, and it means sensitive data will be stored outside the enterprise.

A narrow security view might focus solely on minimizing the risk. However, this narrow view can lead to a Catch-22 situation, as discussed in Clayton Christensen's book *The Innovator's Dilemma* (Harvard Business School Press 1997). Typically, it goes something like this. To minimize the risk, the organization initially restricts the use of a new technology. For example, the technology can only be used for low-risk data, or by a narrow segment of employees. The problem with this approach is that it also reduces the business benefit to the point that the benefit of the technology cannot justify the expense and effort of adopting it. So we reach an impasse. To make the technology a viable proposition, we need to be able to show a business benefit, but we can't show a business benefit because we won't allow viable use of the technology.

Protect to Enable provides the new framework that frees us from the innovator's dilemma. It allows us to focus on the opportunity and identify benefits that outweigh the risks. For example, introducing a new supplier increases competition for our existing suppliers, leading to future savings for our organization. This benefit aligns with the business and is one that everyone in the organization understands. Perhaps less intuitive, but equally important, the savings can be used to fund security controls to mitigate the risk of using the technology more widely. Now our benefit/risk equation has a positive result rather than a negative one. By enabling the technology to be used more widely, we realize bigger business benefits that outweigh the additional cost of controls. This example also underlines the need for CISOs to build business acumen that enables us to see the opportunity and how it can be used to overcome the challenge of funding security initiatives.

Let's look at another example, this time from my experience at Intel in the days before I had defined our Protect to Enable mission. Several years ago, a highly damaging worm was discovered in our environment, requiring a significant emergency response from our team. Upon investigating, we traced the origin of the worm to an employee's personal system.

Our immediate response was that of a stereotypical security group. We shut down this usage to eliminate the risk of future infections. We immediately tightened security policy to ensure only corporate-owned PCs could access the network, and we ruthlessly went through the environment and cut off access by any devices not managed by IT.

Our response was successful in the sense that it reduced the risk of infection. But it led to other risks we hadn't foreseen. Eliminating personally owned PCs from the network meant we now needed to issue corporate PCs to contract employees. This meant that we had to provide more people with devices that allowed full access to the corporate environment. It also, of course, increased capital costs. The broader impact was that it eliminated the potential business benefits of letting people use their own personal devices for work.

Subsequently—driven largely by employee demand, as well as the massive proliferation of new consumer devices—we revisited this issue. This time, we examined it from the perspective of Protect to Enable. We looked at the business opportunities if we allowed personally owned systems on the network, and then how we could mitigate the risks. As I mentioned in Chapter 1, we rapidly discovered that the business value is enormous. Helping employees communicate and collaborate at any time can drive significant productivity gains. It also helps make employees happy. They love using their personal smartphones, PCs, and tablets, and they appreciate that we enable them to do so.

These benefits easily outweigh the cost of the technology required to reduce the risk of allowing access by personal devices. True, some of this technology wasn't available at the time we experienced the original security problem. But if we had focused on the opportunity first, perhaps we could have found ways to provide some level of access while mitigating the risk, and experienced at least some of the benefits we enjoy today.

Demonstrating the Reality of Risk

Of course, the security organization's role still centers on managing risk, which includes discussing the negative consequences of people's actions. If we frame this discussion carefully, I believe we can inform without fearmongering. By describing possible

outcomes and solutions without using emotional language, in terms listeners can understand, we create a context in which the organization can make the decisions that are best for the business.

Even when we have to highlight unpleasant outcomes, we're not fearmongering if our information is based clearly on reality. Here's another example from my experiences at Intel. As our customers' use of the Internet expanded, Intel's marketing groups naturally wanted to expand their external online presence by creating new web sites. So we, as Intel's information security group, began assessing the risks and the security controls required. Some of our marketing teams didn't find this an appealing prospect. They needed to move quickly, with the freedom to communicate however they thought best, and they viewed security procedures as bureaucracy that slowed them down and hindered their ability to communicate with customers and partners.

What happened next was far more persuasive than any of our initial efforts to forestall potential problems. A few web sites were launched without rigorous quality control. Hackers found the weaknesses in these sites, but they didn't crash the sites or steal information. Instead, they inserted links to porn sites.

When this unfortunate fact was discovered, it provided the leverage we needed to improve security procedures. I realized this was a case where a picture spoke a thousand words. So, to illustrate the impact, I simply showed the links to people within the company. This wasn't fearmongering. It was simply demonstrating the real consequences of their actions on the brand. Everyone could understand the implied question: Do we want our brand to look like this? This ended, once and for all, any discussion about whether we needed to apply rigorous quality control to external web sites.

The CISO's Sixth Sense

In the book *Blink: The Power of Thinking Without Thinking*, author Malcolm Gladwell (Little, Brown & Co. 2005) describes an interesting experiment. Researchers asked subjects to play a game in which they could maximize their winnings by turning over cards from either of two decks. What the subjects didn't know was that the decks were subtly stacked. They could win by selecting from one of the decks, but selecting from the other deck would ultimately lead to disaster. After about 80 cards, the subjects could explain the difference between the decks. But they had a hunch something was wrong much sooner, after only 50 cards. And they began showing signs of stress and changing their behavior even sooner, after only about 10 cards, long before they cognitively understood a difference existed.

As CISOs, we develop a sixth sense about security issues. Often, my instincts suggest a need to act or begin investigating a specific direction long before our group is able to fully understand or explain what is happening. This sixth sense is particularly relevant in the security realm, where our information is almost always imperfect or incomplete. When a threat strikes, we do not have time to conduct extensive research or wait for evidence to accumulate. Therefore, we need to act decisively based on imperfect information.

I think we develop this sixth sense from the diverse experiences and skills we've acquired during our careers. We can also foster this sixth sense by being aware. Some security professionals tend to be inwardly focused, looking only at the data and systems they need to protect. As described in Chapter 4, I have directed my teams to try to be

more open and outward-looking, sharing information and seeking input from a variety of sources, including peers across our company and at other organizations. This can help CISOs spot early warning signals and correlate information to quickly identify threats. Like secret service agents scanning a crowd, our experience helps us spot anomalies, to see the signals and ignore the noise.

By identifying future risks early, we may be able to prevent them entirely, or at least minimize their impact. We may also reduce the overall effort needed to deal with the risk. Early action may avoid the need for emergency response and a potentially major cleanup effort.

Taking Action at the Speed of Trust

A sixth sense is only of value if the organization can act on it quickly. This requires two things. First, we need the courage to take a leap of faith based on what we believe. This courage is rooted in the attributes I discussed earlier in this chapter, such as being centered and credible, with a clear sense of our mission.

The second requirement is that the organization responds quickly when we inform them about a security issue. This rapid response is only possible if we have established trusted relationships with people across the organization. Because of these relationships, the organization can act at the *Speed of Trust*, as Stephen M. R. Covey describes it in the book of the same name (Free Press 2008). Faster, frictionless decisions are possible because people know, from experience, that our information is reliable and that our focus is on enabling rather than spreading fear.

The CISO as a Leader

Above all, 21st century CISOs must become effective leaders who can inspire their teams to enable and protect the organization

Over the years, I've identified three essential themes I try to instill in my team and constantly reinforce in our day-to-day interactions. Our security team members must believe in our mission; they must feel they belong within the security group and the company as a whole; and they must feel they matter.

If I can make people feel that they believe, they belong, and they matter, they will tackle any challenge. As Kouzes and Posner put it in *The Leadership Challenge* (Kouzes and Posner 2012), "leadership is the art of mobilizing others to want to struggle for shared aspirations." If people understand the greater goal, it helps establish an emotional connection that guides their everyday actions. This is a key reason that I have thought so much about defining the mission, and that I have spent so much time helping the teams I have led to see how their jobs are connected to the business's objectives and concerns.

For example, a typical operational goal might be to patch all systems within a week of a new software release. This goal is more meaningful if we establish the links to the business using the *I believe, I belong, and I matter* mantra: "I believe in the mission of Protect to Enable. If I'm not protecting to enable, the other employees at the organization I belong to cannot do their jobs effectively. The company doesn't achieve its results, and the company doesn't execute its vision. Patching systems quickly matters because it helps our users do their jobs, which in turn helps the business achieve its goals."

Learning from Other Business Leaders

As leaders, we can learn a lot from how other business leaders work. Today, managers are moving away from command-and-control to a more collaborative approach that takes advantage of the diversity of employee ideas and strengths. I'm not talking about a consensus process, which can lead to endless debate and indecision. Rather, a leader's goal is to ensure alignment to a common mission and accelerate decisions. Within this framework, differing viewpoints and debate spark creativity, generating new ideas and a productive tension that can drive results.

Because security can be frustrating, even daunting, it's vital to find ways to help employees stay motivated. It's important to help employees feel they are making progress, not just when they achieve major milestones, but in solving the smaller problems they face every day. A key study found that even small wins boost motivation, productivity, and creativity. In the *Harvard Business Review* article describing the study, authors Teresa Amabile and Steven Kramer (2011) determined that the feeling of making progress is the most important contributor to an employee's emotions, motivations, and perceptions.

Opportunities to lead occur continually, in every interaction with our teams, with other people in IT, and with business partners. The question we need to ask ourselves is whether we are seizing these opportunities to reinforce our mission and ultimately to help the organization achieve success.

In highly technical jobs and organizations, we have a tendency to focus on technical challenges while overlooking the "people factor." I think it's important to remember the need for personal connections, which foster the sense of belonging. When we know a little more about each other, we care more as a result. I think about this in my day-to-day interactions. If a team member is making a presentation, are we paying attention and asking thought-provoking questions, or are we distracted? And if so, do we think they will feel they belong?

When we meet with a team member to discuss their struggles with a project, are we helping them think through the issues and come up with solutions? Are we helping them believe they can overcome the challenges and that the results will matter to the company and to us? Or are we just taking them to task? Each interaction is an opportunity for coaching and helping employees improve their performance.

A final requirement of effective leadership is the ability to develop other leaders within the security group. Otherwise, the group's strengths in managing risk for the business will last only as long as the current CISO's tenure. By building competence in depth, the CISO can ensure that the organization delivers sustained performance over time. We will discuss this in more depth in the next chapter.

Table 10-1 shows research by executive-search firm Korn Ferry suggesting that cybersecurity leaders need a unique set of attributes, including the ability to think outside the box, dig deeply into issues, exercise judgment at board level, and be a credible business partner (Alexander and Cummings 2016).

Table 10-1. *Attributes of Cybersecurity Leaders*(Alexander and Cummings 2016)

Key Attributes for Cybersecurity Executives			
Competence	Experience	Traits	Drivers
Strategic, global thinker (sees big picture)	Depth of technical experience	Learning agile (can adapt to the new and different)	Seeks high visibility and accountability roles
Thinks outside the box	Understands the evolving legal and regulatory environment	Flexible	Strives to be agent of change (not agent of “no”)
Analytical (digs deeply into issues)	Has successfully handled security incidents in the past	Tolerance for ambiguity	Must “thread the needle to balance driving change with managing enterprise risk”
Possesses business savvy (understands how information is used in daily operations)		Intellectually curious	Pursues close engagement with business leaders (works to add business value)
Balances competing priorities		Bias for action	
Communicates and influences broadly (board, senior management)			
Attracts, builds, and leverages talent			

Voicing Our Values

Obviously leadership means taking responsibility. Yet some CISOs seem to forget this, at least occasionally. A typical situation goes something like this. The CISO warned of a security issue but couldn’t obtain the budget or resources to address it. So the CISO abdicated responsibility because someone else had made the decision not to fund a solution. I take a different view. I believe even if we disagree with the decision, we should do our best to voice our values. We need to articulate the potential impact to the organization, to our customers, and to society, as I discussed in Chapter 9.

As partners in the organization’s strategy, we should commit to the decision and share full accountability and responsibility with our peers. Having said that, we also need to clearly express our personal values and stay true to our principles. Adhering to our values may mean taking career risks, as discussed in Chapter 9. Therefore it is critical that we take the time to reflect on what our principles and values really are. This personal

journey, which we all need to take, adds another dimension to the Z-shaped individual, a dimension of values (Figure 10-2). As Mary Gentile, the author of *Giving Voice to Values* (2010), puts it, “We are more likely to voice our values if we have decided that the costs of not doing so, and the benefits of trying, are important enough to us that we would pursue them even though we cannot be certain of success in advance. In order to get to this place of clarity, we need to spend serious time thinking about our own identity, our personal and professional purpose, and our own definition of success and failure.”

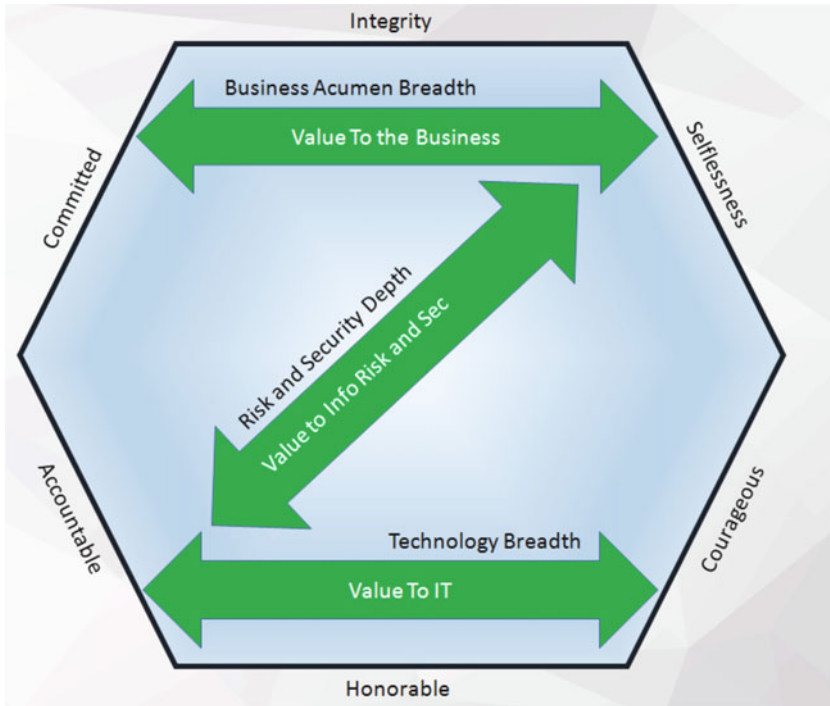


Figure 10-2. Another dimension of the Z-shaped individual: the personal values that guide our actions

Discussing Information Risk at Board Level

Clearly, corporate discussions of any topics that have such far-reaching potential impact on society should include participation by the executive board. Board awareness of security has increased somewhat due to the spate of well-publicized breaches. Yet surveys show that the majority of boards are still not aware of major security and privacy issues. A recent study found that only 32% of boards review security and privacy risks, and only 45% have any involvement in security strategy (PWC 2015).

In contrast, a significant number of security professionals believe that the CEO and executive boards are responsible to society for the sometimes disastrous impact of security and privacy issues. In another recent survey, one sixth of security professionals

said they advocate arrest and a prison sentence for the CEO or board members after a breach (Websense 2015). That seems to indicate that they feel their management is not taking the problem seriously enough, or perhaps even chooses to look the other way, and that they are concerned about the broader consequences to society.

Given the broad and ever-growing importance of security and privacy, boards need be much more involved in than they have been in the past. It is the CISO's responsibility to bring important security and privacy issues to the board, and initiate a debate about the potential impacts of those issues and the right response. Even with the current heightened awareness of security issues, it may not be easy to get the board's attention, because board members have so many other business issues to worry about. It can help to hone in on the handful of risks with the largest potential financial impact or other major implications such as damage to the company's brand. Key areas for boards to consider include

- **Security and privacy strategy:** Is it cohesive and complete?
- **The security and privacy leadership:** Do they act with a level of independence? Do they take ownership of issues, or do they simply manage a risk register?
- **Incident response planning and drills:** Do they occur? Are they integrated across the organization?
- **“Tone from the top:”** Is the executive team engaged? Do their actions match their words?
- **Security and privacy governance:** Does it have the appropriate decision-making structure, including the right level of “tension” between different stakeholders? Is it set up to ask the “high contrast” questions (as discussed in Chapter 2)?

The CISO must take responsibility for determining which issues merit the board's attention. That determination will depend on the potential impact of an exploit conducted against the company's internal systems or technology-based products and services.

C-I-S-O ATTRIBUTES

In this chapter, I have covered a range of abilities and characteristics that the 21st century CISO requires. Many of these probably sound familiar, but it's all too easy to forget them amid the demands of hectic daily schedules. I've found a good way to remind myself of some of the key attributes is simply to look at my job title. The letters in CISO help me remember that we all need Character, Intuition, Skills, and Objectivity. So if you're struggling to remember all the details in this chapter, just remember you're a CISO. You need Character to ensure your actions demonstrate integrity; Intuition to anticipate what's needed and act accordingly, taking risks when necessary; Skills that span business, technology, and a wide variety of risk areas; and Objectivity in order to avoid falling prey to fear-mongering.

Conclusion

As the technology environment continues to evolve, many people believe we're moving toward a future in which organizations outsource much of the delivery of IT services. If this trend continues, what does it mean for the CISO?

In this view of the future, the organization shifts away from IT implementation to procurement and management of suppliers and services, while setting direction and establishing an overall IT architecture.

In addition to this, the organization will need to retain the core competency of the security group: the management of information risk. Essentially, organizations cannot outsource risk. We can hire companies to deliver our business systems, but we're still responsible for compliance with regulations that affect our companies, such as SOX and HIPAA. And if a breach results in theft or leakage of personal information or critical intellectual property, we're still responsible for reporting it. Furthermore, we still suffer the damage to our brand, even if the breach was due a failure of the supplier's systems. As regulations proliferate and more and more personal information is stored in business systems, the risks can only increase.

Therefore the CISO's abilities will remain essential, even if the job title changes. The organization must retain the management of information risk as a core competency. As CISOs, we are poised to continue providing that core competency as long as we can effectively work within this new environment by developing the abilities I've described in this chapter and throughout this book. These abilities enable us to work with others to support the Protect to Enable mission.

I'll close this chapter with an excerpt from a speech by Teddy Roosevelt; the sentiments seem as relevant today as when he made the speech back in 1910. "It is not the critic who counts; not the man who points out how the strong man stumbles, or where the doer of deeds could have done them better. The credit belongs to the man who is actually in the arena, whose face is marred by dust and sweat and blood; who strives valiantly; who errs, who comes short again and again, because there is no effort without error and shortcoming; but who does actually strive to do the deeds; who knows great enthusiasms, the great devotions; who spends himself in a worthy cause; who at the best knows in the end the triumph of high achievement, and who at the worst, if he fails, at least fails while daring greatly, so that his place shall never be with those cold and timid souls who neither know victory nor defeat." (Roosevelt 1910)

We need to be in the arena, and so do our teams. Our mission, as information security and privacy professionals, is a worthy cause. With our efforts to prevent harm to our organizations, our customers, and to society, we can ensure that tomorrow is better than today.



Performance Coaching

If your actions inspire others to dream more, learn more, do more, and become more, you are a leader.

—John Quincy Adams

Over the years I have attended and taught many management and leadership classes. I have also received and written countless performance reviews. I have overseen the ratings and reviews for literally thousands of employees, starting when I ran a call center for a large retailer back in the late 80s, before I attended graduate school. One thing that is clear to me, after so many years participating in these annual and semi-annual corporate rituals, is that there is the potential for considerable ambiguity, particularly when assessing soft skills, those that cannot be measured using hard metrics such as the ability to meet deadlines or deliver revenue commitments.

This ambiguity makes it hard for employees to understand how to meet their manager's expectations. It makes it hard for them to understand the factors that may be limiting their progress from a junior player in the organization to a more senior role. I believe this ambiguity can be clarified, although there will probably always be some qualitative differences in perspective between employee and manager, and even among different managers.

For these soft skills, I believe performance *coaching* needs to be emphasized over performance *management*. This is because at many organizations, performance management focuses primarily on promoting the fittest and eliminating the weak. The process looks at who is getting the best ratings and who is getting the worst. Managers then work to remove the lowest performers from the organization. This selection process is a natural cycle, and one that should continue to play a role. However, I believe that coaching can yield better long-term results, both for individuals and for the organization. Coaching focuses on helping everyone in the organization, including ourselves, reach their full potential. The ultimate goal is to create a high-performance organization in which everyone performs to the utmost of his or her ability.

To effectively coach people, we need to be able to define the soft skills that are required at each level of the progression from entry-level employee to executive. Then we can coach them about how to acquire these skills and move up the organization. The tables in this chapter are intended to provide those definitions, to provide some clarity in these areas of potential ambiguity. They are based on tables that I have used, adapted, tested, and refined over many years in a wide variety of roles. Although I created the

tables for my own employees, the skills listed in the tables are not specific to information risk professionals; they may be equally applicable to employees in other disciplines.

The soft skills in the tables generally describe *how* people work, which can be almost as important to the organization as *what* they do. How people behave and communicate affects not only their own ability to achieve goals but also the performance of those around them. An individual contributor who interacts poorly with others may impair the performance of his or her team, and cause interpersonal problems that the team's manager has to spend time fixing. A senior manager who lacks these soft skills can have an even broader impact, hindering the performance of the organization.

I have published older versions of these tables to my employees, in the belief that feedback should be multi-directional and that leaders as well as employees should be measured using the same publicly available criteria. I have also shared these tables with industry peers. I am providing them in this book in the hope that they will be beneficial to others, and that they will generate comments and feedback that I can use to improve future iterations of this living document.

How to Use the Tables

Each of the 11 tables in this chapter focuses on a specific area of soft skills, such as initiative, commitment, professionalism, or communication. Each table follows the same format, with five columns representing the skills required at progressively higher levels of the organization, from junior employees to emerging executives. The leftmost two columns represent individual contributors: entry-level employees and more seasoned intermediate professionals. The rightmost three columns represent increasingly senior management positions: a line manager responsible for a team; a senior manager who may be responsible for multiple teams, each headed by a line manager; and a leader who is responsible for an entire information risk organization and should be able to work directly with the company's board and top executives.

As one might expect when discussing soft skills, this is not an exact science. The columns show a progression, but they do not represent a precise scale, and there is overlap in some areas. An implicit assumption throughout the tables is that someone in a more senior role has already acquired the skills needed in less-senior positions (i.e. in the columns farther to the left). The skills required at more junior levels tend to be more narrowly defined and constrained; those required at more senior levels tend to be broader in scope, with more far-reaching impact. For these reasons, the tables may be easiest to absorb by reading down the columns (to see all the skills for each role) rather than across the rows.

Over the years, I have used these tables in various ways. I have used them to help employees understand where they need to enhance their skills and abilities if they want to move up to more senior positions. I've also used them to help employees self-assess. Here are some examples of ways to use the tables in everyday work situations:

- An employee believes he or she should be promoted to a more senior position. You ask them to assess their own skills in each area. You also do your own assessment of their skills. Then the two of you discuss any differences between those assessments, and pinpoint areas that the employee should work on in order to acquire the skills needed for a higher-level position.

- You provide an entry-level employee, enthusiastic but fresh out of college, with a roadmap of the skills they'll need to acquire if they want to progress to VP level in the future. This gives them a practical tool that they can use to guide their personal and career development.
- You use the tables to identify your own Achilles' heel, the weak spot that hinders your progression to an executive level. You notice that even though your skills mostly match those in the Emerging Executive column, the skills in a few areas correspond to those that you'd expect in a more junior manager. Those are skills that you need to improve.
- During a coaching session with an employee, you count roughly how many of their skills are already at the next most senior level, the next column in the table. If 80% of their skills match, they may be ready to move up. If there's only a 20% match, they need to work on bringing the rest of their skills up to scratch.

The tables cover the following areas: independence and initiative, efficiency and effectiveness, commitment, professionalism, discipline, teamwork, problem-solving, communication skills, and goal-setting.

Independence and Initiative

This category, as its name suggests, is all about someone's ability to act independently and take the initiative. As you'd expect, the expectations increase dramatically as one progresses up the organization. An entry-level employee may require very specific direction for each new task. A more experienced employee (Intermediate) should be able to define action plans and complete small projects with minimal supervision. A line manager should take responsibility for leading his or her team. An emerging executive can deal with tough issues at executive level, and take responsibility for risky independent decisions that he or she believes are in the best interest of the organization. See Table 11-1.

Table 11-1. *Independence and Initiative*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Takes direction and turns it into results; assumes ownership of deliverables	Acts independently with a specific charter	Embraces role as manager to lead his/her team; sets direction in support of higher level goals	Seeks, identifies, and solves problems while taking responsibility for the outcome	Makes risky independent decisions and takes responsibility for the outcome

(continued)

Table 11-1. (continued)

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Handles multiple simultaneous tasks with some supervision	Responds creatively to customer needs	Effectively summarizes and reports team’s activity	Takes unpopular positions and makes them happen	Deals with tough issues, with no “air cover,” at an executive staff level
Identifies roadblocks and resolves or escalates	Shapes problem statements and defines action plan to complete assignments	Holds self accountable for work he or she doesn’t directly control		Can foresee and take action on problems that do not yet exist
Works with manager to establish workload priorities, clarify expectations, and get feedback	Requires only minimal direction for small projects	Assumes responsibility for work that requires attention, even if it is outside direct scope of his/her role		
Identifies value-added activities and sometimes initiates actions	Seeks buy-in from manager on workload timing and prioritization	Drives risk and security charter among other managers across the organization		

Efficiency and Effectiveness

Efficiency and effectiveness are both important, related skills. An efficient employee works quickly and uses fewer resources. An effective employee is highly productive. A company that combines effectiveness and efficiency achieves better results faster, using fewer resources. Table 11-2 shows the progression from an entry-level employee’s ability to follow efficient processes to a manager’s ability to manage the resources of a group or an entire organization.

Table 11-2. *Efficiency and Effectiveness*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Works at consistent and predictable pace	Schedules work and communicates timelines for output	Delegates appropriately; gets results by working through others and with others	Manages strategic planning and organizational scheduling, and makes good tradeoffs for the organization	Recognizes that what you say “no” to is as vital for driving organizational efficiency as what you say “yes” to
Demonstrates effective work habits enabling timely completion of tasks	Demonstrates ability to manage to multiple work items with inter-dependencies	Plans, schedules, and balances resources among projects to avoid crises and minimize fire fighting	Manages administrative resources to increase personal efficiency	
Learns from mistakes and applies learnings to subsequent tasks	Remains calm and in control of work demands while maintaining work/life balance	Devotes time to improving group’s efficiency	Dispositions items and issues quickly	
Works with manager to prioritize workload	Understands priorities, plans accordingly, and makes real-time adjustments	Uses project management tools and stakeholder input to maximize output and leverage resources	Communicates, and demonstrates through his/her own actions, that people are rewarded for results, not hours worked	
Begins to question time spent on routine tasks with low added value	Networks with others to identify shortcuts and efficiencies	Actively prioritizes by weeding and feeding the project list to maximize organizational effectiveness		

Commitment

Commitment reflects someone’s loyalty to the organization and their willingness to devote time and energy to the cause. In an entry-level employee, commitment is demonstrated by personal work ethic and willingness to take on more work. As people move up the organization, they demonstrate commitment by taking ownership of bigger issues and focusing on driving the best outcome for the organization. See Table 11-3.

Table 11-3. *Commitment*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Demonstrates strong personal work ethic	Aligns individual goals with organizational goals	Drives issues for the benefit of multiple groups across the organization	Holds self accountable for company’s performance	Becomes a role model, demonstrating strong sense of “company first” with the right corporate social responsibility
Readily takes on more workload within job scope	Takes ownership of problems	Recognizes what is best for the organization versus what might be best for the department	Demonstrates a high level of dedication and personal commitment to the success of all employees	Tolerates the indirect control and influence that result from matrix management
Answers the specific questions asked (doesn’t drift)	Provides complete answers to questions; anticipates doubtful areas and works to eliminate concerns	Demonstrates commitment to work/life balance: creating a good home life as well as a good work life	Demonstrates that growth never stops and that we all need to continually learn in order to improve	
	Makes specific requests for necessary information; asks only for what is needed	Knows when to quit on a losing decision but willing to risk self to do the right thing	Subordinates ego to the needs of others and of the company	

Professionalism

Professionalism is the extent to which someone demonstrates the attitudes, skills, and methods required to execute their professional role. For an entry-level employee, this includes adhering to established company policies. For senior managers, it involves demonstrating broader and deeper adoption of the company's values and principles. See Table 11-4.

Table 11-4. *Professionalism*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Demonstrates pride in his/her craft	Sets high personal standards	Modifies behavior to embrace corporate values	Demonstrates unquestioned confidentiality and adherence to the organization's code of conduct, values, and principles	Becomes a role model exemplifying corporate values, growth, consistency, integrity, composure, respect for others, and accountability
Has a courteous and businesslike manner, demonstrating understanding of basic values, role, and appropriate behavior	Holds self accountable for his or her actions	Matches actions with words	Demonstrates strong integrity and motivation with the most honorable intentions	
Respects confidentiality of information, with strict adherence to confidentiality policies	Maintains composure and is not defensive		Aggressively seeks feedback and coaching to grow into a role model	

Discipline

Discipline is the ability to remain focused and execute consistently despite the many distractions of everyday working life. As employees rise to higher-level positions, the distractions and demands increase, requiring greater focus and discipline. See Table 11-5.

Table 11-5. *Discipline*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Approaches work in an orderly fashion	Consistently maintains high standards of accuracy and thoroughness	Stays on point, even with heavy distraction	Demonstrates the stamina and fortitude to remain focused and not succumb to premature conclusions	Prevents the organization from getting distracted
Consistently meets routine deadlines and executes well	Consistently documents intentions and results	Can discern urgency from importance, and prioritize accordingly		Understands the value of “silver bullets” and uses them wisely
Overcomes basic snags and remains focused to stay on course and deliver expected outputs	Does not initiate or perpetuate wasteful communication	Doesn’t waste energy on rhetoric or reactions that lead to no meaningful conclusions		
Demonstrates progression to greater discipline over time				

Teamwork

Individuals must be able to recognize the need to work with others as a team, share expertise, and take on suitable team roles. Managers need to create, inspire, and lead teams, utilizing each member’s talents in the best way. See Table 11-6.

Table 11-6. *Teamwork*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Comprehends the importance of teamwork	Willingly shares knowledge and leverages expertise with others in team	Recognizes and assembles appropriate team players; encourages diversity and utilizes each member's unique talents	Sponsors and leads teams across broad entities	Commissions teams to solve broad, long-term problems
Requires some coaching on appropriate level of team involvement	Independently determines and executes appropriate team role and level of involvement	Provides training and coaching to his/her team Actively engages team members and others to generate win-win solution Recognizes when a team needs course correction Willing to make personal sacrifices for the sake of the team	Nurtures multiple teams within an organization Inspires teams to achieve an extraordinary level of performance	Becomes a key player within the executive team

Problem-Solving

Problem-solving is an important skill for any information risk management professional. Individual contributors need to be able to analyze and solve problems. Managers need to help their teams solve problems and focus on broader issues including those that involve other organizations. See Table 11-7.

Table 11-7. *Problem-solving*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Solves problems with coaching	Takes ownership of problem resolution	Coaches teams to solve problems	Resolves complex problems across the organization	Resolves strategic problems, particularly those involving external parties
Understands cause and effect	Drives analysis of cost, benefit, risk, and probability of success	Identifies and resolves problems not obvious to others, including those beyond his/her previous experience	Champions enduring improvements through structured approaches such as task forces	Identifies proactive and predictive processes to identify the consequences and solve the problems of broad business initiatives
	Uses available resources and solid methodology to solve problems within charter	Uses consultative and consensus processes with ease		Acts as role model for commitment to previously agreed process improvements designed to systematically solve problems

Communication

Good communication helps organizations thrive. It is essential in almost any role, from entry-level team members who must communicate with their colleagues and managers to executives who must communicate messages to the entire organization. Because communications skills are so important, I’ve divided them into three areas, each with its own table: listening, style, and clarity.

Listening

Communication starts with listening. For junior employees, the ability to listen helps create a clear understanding of what’s required. More senior employees actively solicit multiple viewpoints, listen for the meaning behind the words, and intercept emotional outpourings that can overwhelm a situation. See Table 11-8.

Table 11-8. *Communication Skills: Listening*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Confirms understanding	Listens to the broader meaning of what is being said, and seeks opportunities to add value	Hears frustrations and seeks advice about how to respond	Can hear beyond emotion and respond with meaningful commitments and actions	Finds the practical solution amid the noise from team members and executives
Makes listening an overt activity	Listens to others' ideas, and incorporates them into the work; demonstrates respect for others by ensuring their entire message is heard	Seeks others' perspectives and listens to all viewpoints and ideas; encourages mutual understanding	Steps back during debates and identifies the key issues	Can listen to strong-willed or irrational requests and provide appropriate direction
Listens and responds to customers and stakeholders	Adds information or perceptions to expand the concept or the opportunity	Reinforces understanding through active listening; builds confidence in others that their message is being heard	Intercepts escalating emotion before it overwhelms a situation involving other employees or customers	
Before ending conversation, summarizes conversation and achieves closure and agreement	Comes to meetings prepared to review the data and communicate information in a logical fashion Knows when it is better to listen than to talk		Smoothly cross-references prior conversations to ensure truth and consistency	

Style

How you communicate can be as important as *what* you communicate. Each person’s communication style should develop to match their role as they progress through the organization. See Table 11-9.

Table 11-9. *Communication Skills: Style*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Communicates well with others without creating confusion or unnecessary conflict	Delivery of analysis is comprehensive, instructive, and easily understood	Recognizes the requirements of each situation and adapts style accordingly	Demonstrates patience, persistence, and polish in communications	Develops own motivating style
Responds willingly and capably to direct verbal or written questions	Detects when someone is trying to direct them in a conversation and can follow as opposed to veering off track	Remains composed under pointed fire	Maintains a professional demeanor under pressure; can deflect “fire”	Has perfect timing; times communication for maximum impact
Interactions with others are viewed positively; other people do not avoid working with this person	Recognizes and is not deterred by different communication styles	Seeks and responds effectively to feedback on own management behavior	Credibly responds to questions when he or she doesn’t know the answer; can bluff but remain directionally correct or say “I don’t know, but I’ll find out”	Can make and communicate decisions on the fly with high precision and without disrupting other activities
		Uses post-mortems effectively	Says the right thing at the right time	Creates and delivers “state of the union addresses” and “one voice responses” for medium-sized and large organizations

(continued)

Table 11-9. (continued)

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
				Knows when to tell (give direction) versus lead, and does both things well

Clarity

Clear communication helps ensure that information and ideas are accurately shared throughout the organization. Experienced staff should be able to summarize data and create clarity from a confusing mass of information. Senior managers create consistent and clear messages for diverse audiences. See Table 11-10.

Table 11-10. *Communication Skills: Clarity*

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Keeps messages clear and concise	Focuses on and highlights key points	Tells the story, not the facts; delivers the core meaning and the answer (what actions to take) when appropriate	Takes multiple messages from various sources and reconstitutes or links them into a larger, more meaningful message	Sends clear and consistent messages to a broad audience, including external parties
Presents facts accurately, using relevant data	Remains clear about the goal and does not meander—stays on point	Draws summary conclusions from large amounts of information	Creates consistent and clear messages despite complex scope of material	Helps people from different backgrounds quickly grasp complex subjects at a high level
Independently determines areas that need clarity, and seeks and adds appropriate details	Demonstrates awareness of target audience, and tailors message accordingly	Brings clarity to complex situations; asks the right questions to lead the conversation to results, and avoids stating opinion up front	Lean communication: uses the minimum number of words to express a point	Brings clarity to issues across multiple organizations who may have opposing interests

(continued)

Table 11-10. (continued)

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Keeps work neat and well organized		Is aware when he or she has confused senior management	Does not confuse executive management	

Goal-Setting

All experienced staff should be able to identify and set goals, from line managers setting goals for their team to leaders defining the organization’s mission. See Table 11-11.

Table 11-11. Goal-Setting

Entry-level	Intermediate	Line Manager	Senior Manager	Leader/Emerging Executive
Drafts individual goals and reviews with manager for approval	Identifies and declares opportunities	Set goals for team; ensures goals are clear and stated in terms of measurable results	Sets strategic as well as tactical goals	Creates missions
	Presents compelling data to support recommended goals	Aligns goals and expectations with upper management	Demonstrates ability to set goals when starting with a blank sheet	Challenges self, staff, and peers to take on increasingly higher leverage objectives
	Anticipates needs and requirements	Provides a degree of focus on strategic issues; demonstrates vision in areas of expertise	Can drive an organization to articulate commitments, maintain focus, adjust priorities, and raise the bar	Can drive consensus on vision
		Fosters innovation and creative thinking; encourages discussion and feedback in setting goals	Challenges existing paradigms and explores new possibilities	Helps others make the connection between the vision and the deliverables necessary to achieve the higher goals

Conclusion

I believe that performance coaching focused on soft skills can help everyone in the organization achieve their full potential, and thus contribute to the creation of a high-performance organization. I'd like to conclude by examining what makes a manager an effective performance coach. A good performance coach

- Develops and mentors managers and other employees, managing people to higher expectations and greater results.
- Stretches others and themselves to achieve beyond the norm, and rejects mediocrity.
- Creates more key players than he or she consumes, becoming a net developer of people for the organization.
- Holds people accountable for results and coaches them to achieve those results.
- Distinguishes motion from progress, and separates the means from the end.
- Responds positively to feedback about his or her own behavior as a manager or individual.
- Is sought out to provide performance coaching to senior players who report to other managers.
- Handles tough conversations with employees about their behavior or performance crisply, without creating a litigation risk.
- Saves senior players from self-destructing or falling short of their potential.
- Demonstrates empathy and can save employees who are struggling due to work-related or personal reasons and might otherwise leave the organization.

APPENDIX A



References

- Accenture. 2012. *Accenture Technology Vision 2012*. <http://www.accenture.com/us-en/technology/technology-labs/Pages/insight-accenture-technology-vision-2012.aspx>.
- Ahamad, Mustaque. 2011. *Georgia Tech Releases Cyber Threats Forecast for 2012*. Comment in Georgia Tech press release. <http://www.scs.gatech.edu/content/georgia-tech-releases-cyber-threats-forecast-2012>.
- Alexander, Aileen and Jamey Cummings. 2016. "The Rise of the Chief Information Security Officer." *People + Strategy* 39:1.
- Alperovitch, Dmitri. 2012. Comment in *Georgia Tech Emerging Cyber Threats Report 2012*. http://www.iisp.gatech.edu/sites/default/files/documents/2016_georgiatech_cyberthreatsreport_onlinescroll.pdf.
- Amabile, Teresa M., and Steven J. Kramer. 2011. "The Power of Small Wins." *Harvard Business Review* 89:5.
- Ashford, Warwick. 2015. "Security researchers disclose flaws in Kaspersky and FireEye products." *ComputerWeekly.com*. September 7, 2015. <http://www.computerweekly.com/news/4500253029/Security-researchers-disclose-flaws-in-Kaspersky-and-FireEye-products>.
- Bazerman, Max H. and Ann E. Tenbrunsel. 2011. *Blind Spots: Why We Fail to Do What's Right and What to Do about It*. Princeton: Princeton University Press.
- Ben-Shalom, Omer, Manish Dave, Toby Kohlenberg, Dennis Morgan, Stacy Purcell, Alan Ross, Timothy Verrall, and Tarun Viswanathan. 2011. "Rethinking Information Security to Improve Business Agility." Intel Corporation. http://www.intel.com/Assets/PDF/whitepaper/Rethinking_Information_Security_Improve_Business_Agility.pdf.
- Bonus, Angelia. 2010. "Pennsylvania school district settles laptop privacy lawsuit." CNN report, October 12, 2010. <http://www.cnn.com/2010/CRIME/10/12/pennsylvania.school.webcams.settlement/>.
- Bradley, Tony. "Run Cylance Infinity OEM engine on Raspberry Pi." *TechSpective* article posted online October 15, 2015. <https://techspective.net/2015/10/15/run-cylance-infinity-oem-engine-on-raspberry-pi/>.
- Breakwell, Glynis. 2007. *The Psychology of Risk*. Cambridge, UK: Cambridge University Press.
- Brinkmann, Paul. 2015. "Orlando Health reports data breach for 3,200 patients." *Orlando Sentinel*. Published online July 2, 2015. <http://www.orlandosentinel.com/business/brinkmann-on-business/os-orlando-health-data-breach-20150702-post.html>.

Brito, Jerry and Tate Watkins. 2012. "Loving the Cyber Bomb? The Dangers of Threat Inflation in Cybersecurity Policy." *Harvard Law School National Security Journal*, 3:39–83.

Buckley, Chris. 2015. "China Passes Antiterrorism Law That Critics Fear May Overreach." *New York Times*. December 27, 2015. <http://www.nytimes.com/2015/12/28/world/asia/china-passes-antiterrorism-law-that-critics-fear-may-overreach.html>.

Buczek, Laurie and Malcolm Harkins. 2009. "Developing an Enterprise Social Computing Strategy." Intel Corporation. <http://www.intel.com/content/dam/doc/white-paper/intel-it-developing-enterprise-social-computing-strategy-paper.pdf>.

Business Continuity Institute. 2016. "Cyber attack top business threat for second year running." Press release published February 8, 2016.

Carty, Matt, Vincent Pimont, and David W. Schmid. 2012. "Measuring the Value of Information Security Investments." Intel Corporation. <http://www.intel.com/content/dam/www/public/us/en/documents/best-practices/information-security-investments-paper.pdf>.

Casey, Timothy. 2007. "Threat Agent Library Helps Identify Information Security Risks." Intel Corporation. <http://www.intel.com/it/pdf/threat-agent-library.pdf>.

Casey, Tim and Brian Willis. 2008. "Wargames: Serious Play that Tests Enterprise Security Assumptions." Intel Corporation. http://www.intel.com/it/pdf/Wargames-Serious_Play_that_Tests_Enterprise_Security_Assumptions.pdf.

Chelel, Kit. 2015. "A London Hedge Fund Lost \$1.2 Million in a Friday Afternoon Phone Scam." *Bloomberg Business*. Published online July 7, 2015. <http://www.bloomberg.com/news/articles/2015-07-07/friday-afternoon-scam-cost-hedge-fund-1-2-million-and-cfo-s-job>.

Chirgwin, Richard. "Intel's SGX security extensions: Secure until you look at the detail." *The Register*. February 1, 2016. http://www.theregister.co.uk/2016/02/01/sgx_secure_until_you_look_at_the_detail/.

Christensen, Clayton M. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, Mass.: Harvard Business School Press.

Cisco Systems, Inc. 2011a. *Cisco Connected World Technology Report 2011*. <http://www.cisco.com/en/US/netso1/ns1120/index.html>.

Cisco Systems, Inc. 2011b. *Email Attacks: This Time It's Personal*. http://www.cisco.com/en/US/prod/collateral/vpndev/ps10128/ps10339/ps10354/targeted_attacks.pdf.

Cisco Systems, Inc. 2015. *2015 Corporate Social Responsibility Report*. http://www.cisco.com/assets/csr/pdf/CSR_Report_2015.pdf.

Cisco Systems 2015b. *Internet Of Things Will Deliver \$1.9 Trillion Boost To Supply Chain And Logistics Operations*. Cisco Systems, Inc. press release.

Clark, Sandy, Stefan Frei, Matt Blaze, Jonathan Smith. 2010. "Familiarity Breeds Contempt: The Honeymoon Effect and the Role of Legacy Code in Zero-Day Vulnerabilities." In *Proceedings of the 26th Annual Computer Security Applications Conference*. New York: Association for Computing Machinery. doi: 10.1145/1920261.1920299.

Colgan, William B. 2010. *Allied Strafing in World War II: A Cockpit View of Air to Ground Battle*. Jefferson, NC: McFarland.

Compeau, Joseph, Nicole Haggerty, Ramasastry Chandrasekhar. 2013. *Intel Corp. - Bring Your Own Device*. Ivey Business School case study.

- Corporate Executive Board Company, The (CEB). 2012. Information Risk Executive Council. Arlington, VA. <http://www.executiveboard.com/exbd/information-technology/it-risk/index.page>.
- Corporate Executive Board Company, The (CEB). 2015. CEB March 2015 Information Risk Peer Perspective Survey.
- Covey, Stephen M. R. with Rebecca R. Merrill. 2008. *The Speed of Trust: The One Thing That Changes Everything*. New York: Free Press.
- CSO Magazine, US Secret Service, Software Engineering Institute CERT Program at Carnegie Mellon University, Deloitte. 2011. *2011 CyberSecurity Watch Survey: Organizations Need More Skilled Cyber Professionals To Stay Secure*. Press release. http://www.sei.cmu.edu/newsitems/cybersecurity_watch_survey_2011.cfm.
- Culp, Scott. 2010. *10 Immutable Laws of Security*. Microsoft Corporation. <http://technet.microsoft.com/library/cc722487.aspx>.
- CWE/SANS. 2011. *CWE/SANS TOP 25 Most Dangerous Software Errors*. <http://cwe.mitre.org/top25/>.
- Department of Justice. 2014. "U.S. Charges Five Chinese Military Hackers for Cyber Espionage Against U.S. Corporations and a Labor Organization for Commercial Advantage." Press release. May 9, 2014.
- Department of Telecommunications, Government of India. 2009. Instructions to Internet service providers. Letter dated February 23, 2009, No. 820-1/2008-DS Pt. II.
- Edelman. 2015. 2015 Edelman Trust Barometer. <http://www.edelman.com/insights/intellectual-property/2015-edelman-trust-barometer/>.
- Edwards, Cliff, Olga Kharif, and Michael Riley. 2011. "Human Errors Fuel Hacking as Test Shows Nothing Stops Idiocy." Bloomberg News. Posted June 27, 2011. <http://www.bloomberg.com/news/2011-06-27/human-errors-fuel-hacking-as-test-shows-nothing-prevents-idiocy.html>.
- European Commission. 2011. *ePrivacy Directive: circumstances, procedures and formats for personal data breach notifications*. http://ec.europa.eu/information_society/policy/ecomms/doc/library/public_consult/data_breach/ePrivacy_databreach_consultation.pdf.
- European Commission. 2012. *Proposal for a regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)*. http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf.
- European Network and Information Security Agency (ENISA). 2010. *Incentives and Challenges for Information Sharing in the Context of Network and Information Security*. <http://www.enisa.europa.eu/activities/Resilience-and-CIIP/public-private-partnership/information-sharing-exchange/incentives-and-barriers-to-information-sharing>.
- Evered, Rob and Jerzy Rub. 2010. "Maintaining Information Security while Allowing Personal Hand-held Devices in the Enterprise." Intel Corporation. http://www.intel.com/Assets/PDF/whitepaper/Maintaining_Info_Security_Allowing_Personal_Hand_Held_Devices_Enterprise.pdf.
- Federal Trade Commission. 2016. "Dental Practice Software Provider Settles FTC Charges It Misled Customers About Encryption of Patient Data." Press release issued January 5, 2016. <https://www.ftc.gov/news-events/press-releases/2016/01/dental-practice-software-provider-settles-ftc-charges-it-misled>.

- Fleming, Virgil and Naoyuki Tomizawa. 2012. "Intel IT: Keeping the Business Running in a Crisis." Intel Corporation. http://media12.connectedsocialmedia.com/intel/03/7906/Intel_IT_Keeping_Business_Running_in_Crisis.pdf.
- Fong, David, Toby Kohlenberg, and Justin Philips. 2010. "Enterprise Security Benefits of Microsoft Windows 7." Intel Corporation. <http://www.intel.in/content/dam/www/public/us/en/documents/case-studies/intel-it-windows-7-upgrade-security-brief.pdf>.
- Food and Drug Administration. 2016. "FDA outlines cybersecurity recommendations for medical device manufacturers." FDA press release. <http://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm481968.htm>.
- Fox-Brewster, Thomas. 2015. "Darkode Shutdown: FireEye Intern Accused Of Creating \$65,000 Android Malware." *Forbes*. Published online July 15, 2015. <http://www.forbes.com/sites/thomasbrewster/2015/07/15/fireeye-intern-dendroid-charges/>.
- Fox-Brewster, Thomas. 2016. "As Ransomware Crisis Explodes, Hollywood Hospital Coughs Up \$17,000 In Bitcoin." *Forbes*. Published online February 18, 2016. <http://www.forbes.com/sites/thomasbrewster/2016/02/18/ransomware-hollywood-payment-locky-menace/>.
- Frier, Sarah. 2014. "Twitter CFO Noto Has an Oops Moment With Mistaken Tweet." *Bloomberg Business*. November 24, 2014. <http://www.bloomberg.com/news/articles/2014-11-25/twitter-cfo-noto-has-an-oops-moment-with-mistaken-tweet>.
- Fulford, Charles. 2015. "Retail Is About to Be Reinvented, Driven by Digital Technologies." *Advertising Age*. August 28, 2015. <http://adage.com/article/digitalnext/retail-reinvented/300129/>.
- Gartner, Inc. 2005. "Gartner Survey Shows Spending for Compliance and Corporate Governance to Account for 10-15 Percent of an Enterprise's 2006 IT Budget." Gartner Inc. Press release. http://www.gartner.com/press_releases/asset_141532_11.html.
- Gartner, Inc. 2011a. "Gartner Says Context-Aware Technologies Will Affect \$96 Billion of Annual Consumer Spending Worldwide by 2015." Gartner Inc. Press release. <http://www.gartner.com/it/page.jsp?id=1827614>.
- Gartner, Inc. 2011b. "Gartner Identifies the Top 10 Strategic Technologies for 2012." Gartner Inc. Press release. <http://www.gartner.com/it/page.jsp?id=1826214>.
- Gartner, Inc. 2015. "Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015." Gartner Inc. Press release. <http://www.gartner.com/newsroom/id/3165317>.
- Gartner, Inc. 2015b. "Gartner Says It's Not Just About Big Data; It's What You Do With It: Welcome to the Algorithmic Economy." Gartner Inc. Press release. <http://www.gartner.com/newsroom/id/3142917>.
- Gentile, Mary. 2010. *Giving Voice to Values*. New Haven: Yale University Press
- Gladwell, Malcolm. 2005. *Blink: The Power of Thinking Without Thinking*. New York: Little, Brown & Co.
- Global Industry Analysts. 2015. *Context Aware Computing – A Global Strategic Business Report*. Global Industry Analysts, Inc. research report. http://www.strategyr.com/MarketResearch/Context_Aware_Computing_CAC_Market_Trends.asp.
- Goodin, Dan. 2015. "'Unauthorized code' in Juniper firewalls decrypts encrypted VPN traffic." *Ars Technica*. December 17, 2015. <http://arstechnica.com/security/2015/12/unauthorized-code-in-juniper-firewalls-decrypts-encrypted-vpn-traffic/>.

- Goodman, Mark. *Future Crimes: Everything Is Connected, Everyone Is Vulnerable and What We Can Do About It*. New York: Doubleday.
- Greenberg, Andy. 2015. "Hackers Remotely Kill a Jeep on the Highway—With Me in It." *Wired*. Posted July 21, 2015. <http://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>.
- Greenberg, Andy. 2015b. "After Jeep Hack, Chrysler Recalls 1.4M Vehicles for Bug Fix." *Wired*. Posted July 24, 2015. <http://www.wired.com/2015/07/jeep-hack-chrysler-recalls-1-4m-vehicles-bug-fix/>.
- Gutierrez, Esteban, Toby Kohlenberg, Sridhar Mahankali, and Bill Sunderland. 2012. "Virtualizing High-security Servers in a Private Cloud." Intel Corporation. <http://www.intel.de/content/dam/www/public/us/en/documents/best-practices/virtualizing-high-security-servers.pdf>.
- Hamblen, Matt. 2016. "At CES, Feds prod companies to expand privacy efforts." *Computerworld* January 6, 2016. <http://www.computerworld.com/article/3019832/data-privacy/at-ces-feds-prod-companies-to-expand-privacy-efforts.html>.
- Information Risk Executive Council. 2011. *Security Controls Maturity Benchmark Summary*. Information published in *2011-2012 Intel IT Performance Report*. Intel Corporation. <http://www.intel.com/content/dam/www/public/us/en/documents/best-practices/intel-it-annual-performance-report-2011-12.pdf>.
- Intel Corporation. 2010. Form 10-Q for the quarterly period ended March 27, 2010; Filed May 3, 2010. <http://www.intc.com/secfiling.cfm?filingID=950123-10-42822>.
- Intel Corporation. 2011. *Worldwide Device Estimates Year 2020—Intel One Smart Network Work*.
- Intel Corporation. 2012a. "Thinking Differently About IT Value: 2011-2012 Intel IT Performance Report." <http://www.intel.com/content/dam/www/public/us/en/documents/best-practices/intel-it-annual-performance-report-2011-12.pdf>.
- Jackson Higgins, Kelly. 2010. "'Operation Aurora' Changing the Role of the CISO." *Dark Reading*. March 16, 2010. <http://www.darkreading.com/attacks-breaches/operation-aurora-changing-the-role-of-the-ciso/d/d-id/1133225>.
- Joffe-Walt, Chana and Alix Spiegel. 2012. "Psychology Of Fraud: Why Good People Do Bad Things." National Public Radio broadcast. Transcript accessed online May 28, 2012. <http://www.npr.org/2012/05/01/151764534/psychology-of-fraud-why-good-people-do-bad-things>.
- Johnson, Steven. 2010. *Where Good Ideas Come From: The Natural History of Innovation*. New York: Riverhead Books, a subsidiary of Penguin Books (USA).
- Johnson, Steven. 2010. Talk at TEDGlobal 2010. http://www.ted.com/talks/steven_johnson_where_good_ideas_come_from.html.
- Keteyian, Armen. 2010. "Digital Photocopiers Loaded With Secrets." CBS News article posted online April 20, 2010. http://www.cbsnews.com/2100-18563_162-6412439.html.
- Kouzes, James and Barry Z. Posner. 2012. *The Leadership Challenge*. San Francisco: Jossey-Bass, an imprint of John Wiley & Sons.
- KPMG International. 2015. *The KPMG Survey of Corporate Responsibility Reporting*. November 2015
- Kupperwasser, Yosef. 2007. *Lessons from Israel's Intelligence reforms*. The Brookings Institution. http://www.brookings.edu/~media/research/files/papers/2007/10/intelligence%20kupperwasser/10_intelligence_kupperwasser.pdf.

- Lambert, Leslie. 2015. "User behaviors can expose bad actors before it's too late." CSO article posted online July 27, 2015. <http://www.csoonline.com/article/2951814/cyber-attacks-espionage/what-can-we-learn-from-jpmorgans-insider-breaches.html>.
- Lea, Ruth. "Corporate Social Responsibility: IoD Member Opinion Survey." The Institute of Directors, November 2002.
- Leon, Fred. 2011. "Securing Intel's External Online Presence." Intel Corporation. <http://www.intel.com/content/dam/doc/white-paper/intel-it-securing-intels-external-online-presence-paper.pdf>.
- Levin, Carl. 2010. Opening Statement of Senator Carl Levin, Senate Armed Services Committee Hearing on Nominations of Vice Admiral James A. Winnefeld and Lieutenant General Keith B. Alexander.
- Lindstrom, Pete. 2008. "Five Immutable Laws of Virtualization Security." Burton Group blog entry posted online January 08, 2008. <http://srmsblog.burtongroup.com/2008/01/five-immutable.html>.
- Loechner, Jack. 2015. "IoT Connected Devices Triples To 38 Billion By 2020." Center for Media Research brief. August 27, 2015. <http://www.mediapost.com/publications/article/256678/iot-connected-devices-triples-to-38-billion-by-202.html>.
- LosHuertos, Gary. 2010. "Herding Firesheep in New York City" Blog entry posted online October 27, 2010. http://money.cnn.com/2010/12/14/technology/firesheep_starbucks/.
- MarketsandMarkets. 2013. *Context Aware Computing Market - Global Advancements, Emerging Applications, Worldwide Forecasts and Analysis (2013 - 2018)*. MarketsandMarkets research report. <http://www.marketsandmarkets.com/PressReleases/context-aware-computing.asp>.
- Massachusetts Institute of Technology Sloan School Center for Information Systems Research. 2012. IT Governance. <http://cisr.mit.edu/research/research-overview/classic-topics/it-governance/>.
- Maynard, Micheline. 2014. "'The GM Nod' And Other Cultural Flaws Exposed By The Ignition Defect Report" *Forbes*, June 5, 2014. <http://www.forbes.com/sites/michelinemaynard/2014/06/05/ignition-switch-report-spares-ceo-barra-but-exposes-gms-culture/#2715e4857a0b3ac65828408a>.
- McAfee, Inc. 2011. Press release. *McAfee Q2 2011 Threats Report Shows Significant Growth for Malware on Mobile Platforms*. <http://www.mcafee.com/us/about/news/2011/q3/20110823-01.aspx>.
- McComb, Michael. 2002. "Profit to Be Found in Companies That Care," *South China Morning Post*. April 14, 2002.
- Miller, Ron and Joe Varga. 2011. "Benefits of Enabling Personal Handheld Devices in the Enterprise." <http://www.intel.com/content/dam/doc/best-practices/intel-it-it-leadership-benefits-of-enabling-personal-handheld-devices-in-the-enterprise-practices.pdf>.
- Millman, Rene. "Updated: 97% of malicious mobile malware targets Android." *SC Magazine UK*. June 26, 2015. <http://www.scmagazineuk.com/updated-97-of-malicious-mobile-malware-targets-android/article/422783/>.
- Morgan, Steve. 2015. "Cybersecurity job market to suffer severe workforce shortage." CSO. July 28, 2015. <http://www.csoonline.com/article/2953258/it-careers/cybersecurity-job-market-figures-2015-to-2019-indicate-severe-workforce-shortage.html>.

- Nakashima, Ellen and Matt Zapposky. "U.S. charges Iran-linked hackers with targeting banks, N.Y. dam." *Washington Post*. March 24, 2016.
- Nelson, Patrick. 2016. "3D printers wide-open to hacking." *Network World*. March 8, 2016. <http://www.networkworld.com/article/3041436/security/3d-printers-wide-open-to-hacking.html>.
- Nest Labs. 2012. Nest Learning Thermostat web site. <http://www.nest.com/>.
- Paganini, Pierluigi. 2014. "Firmware vulnerability allows man-in-the-middle attack using SD Memory cards." *The Hacker News*. January 01, 2014. <http://thehackernews.com/2014/01/firmware-vulnerability-allows-man-in.html>.
- Pagliery, Jose. "Chryslers can be hacked over the Internet." *CNNMoney*. July 22, 2015. <http://money.cnn.com/2015/07/21/technology/chrysler-hack>.
- Pauli, Darren. 2015a. "Thousands of 'directly hackable' hospital devices exposed online." *The Register*. September 29 2015. http://www.theregister.co.uk/2015/09/29/thousands_of_directly_hackable_hospital_devices_found_exposed/.
- Pauli, Darren. 2015b. "'10-second' theoretical hack could jog Fitbits into malware-spreading mode." *The Register*. October 21 2015. http://www.theregister.co.uk/2015/10/21/fitbit_hack/.
- Perch Interactive. 2015. Perch Interactive web site. www.perchinteractive.com/.
- Perloth, Nicole. 2011. "Insurance Against Cyber Attacks Expected to Boom." *New York Times* blog post December 29, 2011. <http://bits.blogs.nytimes.com/2011/12/23/insurance-against-cyber-attacks-expected-to-boom/>.
- PhishLabs. 2016. *2016 Phishing Trends & Intelligence Report: Hacking the Human*. https://pages.phishlabs.com/2016-Phishing-Trends-and-Intelligence-Report-Hacking-the-Human_PTI.html.
- Rajab, Moheeb Abu, Lucas Ballard, Panayiotis Marvrommatis, Niels Provos, and Xin Zhao. 2010. "The Nocebo Effect on the Web: An Analysis of Fake Anti-Virus Distribution." In *Large-Scale Exploits and Emergent Threats*. Usenix. http://static.googleusercontent.com/external_content/untrusted_dlcp/research.google.com/en/us/pubs/archive/36346.pdf.
- Retail Cyber Intelligence Sharing Center. 2015. "Retailers Launch Comprehensive Cyber Intelligence Sharing Center." Press release published May 14, 2015. <https://r-cisc.org/2015/05/14/retailers-launch-comprehensive-cyber-intelligence-sharing-center/>.
- Rice, David. 2007. *Geekonomics: The Real Cost of Insecure Software*. Boston: Addison-Wesley Professional.
- Rifkin, Jeremy. *The Third Industrial Revolution: How Lateral Power Is Transforming Energy, the Economy, and the World*. New York: St Martin's Griffin.
- Roosevelt, Theodore. 1910. "The Man in the Arena." Speech at the Sorbonne, Paris, France. April 23, 1910. <http://www.theodore-roosevelt.com/images/research/speeches/maninthearena.pdf>.
- Rose, Charlie. 2015. "Inside Apple, Part 2." *60 Minutes* interview, December 20, 2015. CBS Interactive. <http://www.cbsnews.com/news/60-minutes-apple-tim-cook-charlie-rose/>.
- Schmidt, Eric and Jared Cohen. 2015. "Inventive artificial intelligence will make all of us better." *Time*. December 21, 2015. <http://time.com/4154126/technology-essay-eric-schmidt-jared-cohen/>.

- Schmidt, Michael, Nicole Perlroth and Matthew Goldstein. "F.B.I. Says Little Doubt North Korea Hit Sony." *New York Times*. January 7, 2015. <http://www.nytimes.com/2015/01/08/business/chief-says-fbi-has-no-doubt-that-north-korea-attacked-sony.html>.
- Scott, Mark. 2016. "U.S. and Europe in 'Safe Harbor' Data Deal, but Legal Fight May Await." *New York Times*. February 2, 2016. http://www.nytimes.com/2016/02/03/technology/us-europe-safe-harbor-data-deal.html?_r=1.
- Seidman, Dov. 2011. "Measuring HOW We Do Business." *Forbes* article posted online November 27, 2011. <http://www.forbes.com/sites/dovseidman/2011/11/27/measuring-how-we-do-business/>.
- Shey, Heidi. 2014. "Pet The Unicorns And Think Of Protecting Customer Data As A Corporate Social Responsibility." Forrester Research blog post April 23, 2014. http://blogs.forrester.com/heidi_shey/14-04-23-pet_the_unicorns_and_think_of_protecting_customer_data_as_a_corporate_social_responsibility.
- Silverman, Rachel Emma. 2012. "Facebook and Twitter Postings Cost CFO His Job." *Wall Street Journal* article posted online May 14, 2012. <http://www.wsj.com/articles/SB10001424052702303505504577404542168061590>.
- Sinek, Simon. 2009. *Start with Why: How Great Leaders Inspire Everyone to Take Action*. New York: Portfolio.
- Slovic, Paul. 2010. *The Feeling of Risk: New Perspectives on Risk Perception*. New York: Routledge.
- Solove, Daniel J. 2011. *Nothing to Hide: The False Tradeoff between Privacy and Security*. Yale University Press.
- Sunderland, Bill and Ajay Chandramouly. 2011. "Overcoming Security Challenges to Virtualize Internet-facing Applications." Intel Corporation. <http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/cloud-security-and-secure-virtualization-paper.pdf>.
- Taleb, Nassim Nicholas. 2007. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.
- Thaler, Richard H. Thaler and Cass R. Sunstein. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
- Thomson, Iain. 2015. "Lottery IT security boss guilty of hacking lotto computer to win \$14.3m." *The Register*. Published online July 22, 2015. http://www.theregister.co.uk/2015/07/22/lotto_infosec_director_guilty/.
- Tode, Chantal. "Macy's peps up Black Friday shopping via beacon-triggered mobile game." *Mobile Commerce Daily*. Published online October 30, 2015
- Verizon. 2015. *2015 Data Breach Investigations Report*. <http://www.verizonenterprise.com/DBIR/2015/>.
- US Department of Justice (DoJ). 2014. "Justice Department, Federal Trade Commission Issue Antitrust Policy Statement on Sharing Cybersecurity Information." Press release issued April 10, 2014. <http://www.justice.gov/opa/pr/justice-department-federal-trade-commission-issue-antitrust-policy-statement-sharing>.
- US Environmental Protection Agency (EPA). 2011. "Oil Pollution Act Overview." <https://www.epa.gov/laws-regulations/summary-oil-pollution-act>.
- US Government Accountability Office (GAO). 2012. "Challenges in Securing the Modernized Electricity Grid." <http://www.gao.gov/products/GAO-12-507T>.

US Securities and Exchange Commission. 2011. CF Disclosure Guidance: Topic No. 2. Issued October 13, 2011. <http://www.sec.gov/divisions/corpfin/guidance/cfguidance-topic2.htm>.

Van Derbeken, Jaxon. "S.F. officials locked out of computer network." *San Francisco Chronicle*. Published online Tuesday, July 15, 2008. <http://www.sfgate.com/bayarea/article/S-F-officials-locked-out-of-computer-network-3205200.php>.

Venables, Philip. 2008. Speech at RSA Conference 2008.

Websense. Inc. 2015. "Research: Penalties, Punishment & Prison for Serious Data Breaches say e-Crime Congress Respondents." Press release, March 23 2015. <https://community.websense.com/blogs/websense-news-releases/archive/2015/03/23/research-penalties-punishment-amp-prison-for-serious-data-breaches-say-e-crime-congress-respondents.aspx>.

Weil, Peter and Jeanne W. Ross. 2004. *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Boston, Mass.: Harvard Business School Press.

Willis, Brian. 2012. "Sharing Cyber-Threat Information: An Outcomes-based Approach." Intel Corporation. <https://www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/Intel%20-%20Sharing%20Cyber-Threat%20Information.pdf>.

World Business Council for Sustainable Development. 2007. *Corporate Social Responsibility: Meeting changing expectations*.

Worral, Bob. 2015. "Important Announcement about ScreenOS." Juniper Networks security announcement. Posted online December 17, 2015. <http://forums.juniper.net/t5/Security-Incident-Response/Important-Announcement-about-ScreenOS/ba-p/285554>.

Zander, Rosamund Stone and Benjamin Zander. 2000. *The Art of Possibility: Transforming Professional and Personal Life*. Boston, Mass.: Harvard Business School Press.

Index

■ A

Advanced persistent threats (APTs), 13

Architecture. *See also* Security architecture

balanced controls

anti-malware technology, 114–115

definition, 106

detective and preventative controls, 114

intrusion prevention systems, 113

security analytics, 113

employee productivity, 105

hardware-enforced security, 105

security zones, 109

architecture, 111–112

critical data and resources, 109

definition, 106

devices and application types, 109

PEPs, 110

selective zones, 111

trusted zones, 111

untrusted zones, 110–111

user's device and location, 112

traditional enterprise trust model, 105

trust calculation

access type, 106

allow access, 108

available controls, 108

business partners, 109

definition, 105

destination score, 108

devices and usage models, 106

internal and external resources, 109

policy decision point (PDP), 108

source score, 107

trust calculation, 108

user and data perimeters

defenses and detective control, 115

Internet of Things (IoT), 116

protect information, 115

security, 116

traditional network security, 115

user and data perimeters definition, 106

Arizona Counter Terrorism Intelligence Center (ACTIC), 59

Arizona Cyber Threat Response Alliance (ACTRA), 59

Assessment models, 25

■ B

Business benefits and risks

baseline security

AI-based security and

automation, 125–126

encryption, 125

enhanced recovery, 125

hardware acceleration, 125

hardware-enforced, 124

protected environments, 124

building security, 123

context-aware security

business intelligence and data protection, 127

cloud security and context

awareness, 126

experiences, 123

image recognition technology, 126

portable devices, 126

sensors and analytical tools, 126

contextual information, 123

mass-production strategy, 123

■ **C**

- Chief information security officer (CISO)
 - chief trust officer, 139–140
 - foundational skills, 142
 - junk food fear, 144
 - leader, 148
 - business leaders, 149–150
 - information risk-board level, 151–152
 - operational goal, 148
 - values, 150–151
 - Z-shaped individual, 151
 - managing risk, 146
 - organizations outsource, 153
 - positive approach, 145–146
 - sixth sense, 147
 - speed of trust, 148
 - storyteller, 143–144
 - T-shaped individuals, 141
 - Z-shaped individuals, 141
- Chief Information Security Officer (CISO), 35
- Chief Security and Privacy Officer (CSPO), 35–36
- Coaching performances
 - commitment, 160
 - communication, 164
 - clear communication, 167
 - goal-setting, 168
 - listening, 164–165
 - style, 166–167
 - definition, 155
 - discipline, 161–162
 - efficiency and effectiveness, 158–159
 - independence and initiative, 157–158
 - management, 155
 - problem-solving, 163–164
 - professionalism, 161
 - soft skills, 156, 169
 - tables
 - soft skills, 156
 - work situations, 156–157
 - teamwork, 162–163
- Context-aware computing, 119
- Corporate governance model, 36
- Corporate social responsibility (CSR), 130
 - definitions, 130
 - maintaining society trust, 134
 - managing information

- act, 137
- events, 135
- interpret, 137
- risk professional and ethical standards, 135
- security issues, 136
- sense, 137
- managing information risk, 135
- scope of, 130
- technology
 - control and impacts, 132
 - ethical implications, 133
 - evolution, 132
 - march of, 132
 - motivations, 133
 - potential impacts, 132
 - public-sector organizations, 133
 - real-life attack, 133
 - technology and business professionals, 129
 - treat information risk, 131
- Cybersecurity legislation, 8

■ **D**

- Distributed denial-of-service (DDoS)
 - threats, 21

■ **E**

- Emerging security capabilities
 - accelerated encryption, 119
 - artificial intelligence, 122
 - business benefits and risks (*see* Business benefits and risks)
 - cloud computing, 122
 - consistent experience across devices, 121
 - context-aware computing and security, 119
 - data analytics, 122
 - enterprise information security, 120
 - hardware-enforced protection, 119
 - information security, 120
 - Internet of Things (IoT), 120
 - enabled car, 119
 - Moore's law, 120
 - nest learning thermostat, 121
 - NFC, 120
 - RFID, 120
 - wireless NFC, 120

- malicious purposes, 119
- organization's privacy commitment, 119
- security professionals, 119
- shopper's smartphone, 121
- Enterprise resource planning (ERP)
 - system, 111
- External partnerships
 - advantage of, 55
 - ACTRA, 59
 - attacks and threads, 58
 - benchmarks
 - information, 61
 - operations, 60
 - CISO, 60
 - communities, 57
 - characteristics, 57
 - goals, 59
 - constitutes valuable information, 58
 - corporate citizenship, 63
 - enabling informal exchanges, 60
 - FIRST, 61
 - information sharing, 49
 - legal implications and revealing
 - security, 50
 - partnerships, 55
 - private-sector organizations, 58
 - public-relations aspect, 50
 - regional communities, 58
 - regulations and standards, 62
 - security landscape, 63
 - share security information, 50
 - share threat information, 58
 - technology landscape, 51
 - threat landscape, 51
 - tiered pyramid model
 - confidential tier, 53–55
 - information-sharing
 - relationships, 53–54
 - partnerships tier, 54
 - public tier, 53, 55
 - targeted tier, 53–54
 - types of, 52
 - value of, 51
 - vulnerabilities information and
 - threats, 59

■ **F**

- Federal Trade Commission (FTC), 136
- Forum for Incident Response and Security Teams (FIRST), 61

■ **G**

- Global Positioning System (GPS), 96, 108
- Governance
 - dictatorial approach, 33
 - information risk, 32
 - IT policies, 33
 - life-threatening consequences, 33
 - MIT CISR, 32–33
 - structure
 - archetypes, 34, 90
 - corporate information, 36
 - CSPO, 36
 - engagements, 36
 - hybrid governance model, 35
 - monitor (sense) risk, 35
 - operations, 36
 - oversight, 35
 - security and privacy, 35

■ **H**

- Health Insurance Portability and Accountability Act (HIPAA), 9
- Human resources (HR)-related processes
 - capabilities, 31
 - information risks, 31
 - internal partnership, 31
 - key areas, 32
 - security team, 32
 - technology transitions, 31

■ **I, J, K, L**

- Information security
 - business enable[®], 5
 - balancing act, 6
 - core competencies, 5
 - legal and human resources (HR)
 - groups, 6
 - personal smartphones, 6
 - primary variables, 7
 - responsibilities grew, 6
 - transformation, 6
 - trust, 8
 - tuned to target, 6–7
 - wireless network implementation, 6
 - businesses and organization, 3
 - central nervous system, 3
 - fast-moving environment, 4
 - interdependent risks, 4

Information security (*cont.*)
 management risk, 16
 dynamic and flexible, 16
 incorporate privacy and regulatory compliance, 16
 network boundary, 16
 marketers, 4
 perfect storm, 1–2
 rapid proliferation (information and devices), 12
 regulatory environment, 3
 regulatory flood (*see* Regulatory flood)
 security and privacy, 5
 smartphones, 3
 technology, 3
 threat landscape (*see* Threat landscape)
 traditional mission and vision, 2
 Information Sharing and Analysis Centers (ISACs), 60
 Information Sharing and Analysis Organizations (ISAOs), 50
 Insider threats, 78
 detect, 79
 deter, 79
 discipline, 80
 organization’s reputation, 78
 security firms, 78
 three-part approach, 79
 Intel IT Emergency Response Process (ITERP), 47
 Intel Secure External Presence (ISEP), 41
 Internal partnerships
 business group managers, 46
 corporate risk management, 44
 corporate security, 45
 far-reaching web, 37
 fellow travelers, 38
 finance group
 business groups, 43
 internal audit, 44
 SOX, 43
 human resources
 communications, 42
 procedures, 42
 internal investigations, 43
 security policies, 42
 information security group
 and teams, 37
 ITERP, 47

legal groups, 38
 business groups, 40
 contracts, 39
 data classification, 39
 financial compliance, 40
 intellectual property, 39
 ISEP-like process, 41
 litigation, 39
 privacy, 38
 privacy, 45
 response process, 46–47
 risk review boards, 37
 standing committees, 37
 Internet of Things (IoT), 94, 120
 enabled car, 119
 Moore’s law, 120
 nest learning thermostat, 121
 NFC, 120
 RFID, 120
 technologies, 3
 wireless NFC, 120
 Irrefutable Laws. *See* Information Security, 14

■ **M**

Massachusetts Institute of Technology Center for Information Systems Research (MIT CISR), 32
 Misperception. *See* Risk misperception
 Moore’s law, 120

■ **N**

National Institute of Standards and Technology (NIST), 144
 Near field communications (NFC), 120
 Non-Intel managed systems (NIMS), 20

■ **O**

Organization’s privacy commitment, 119

■ **P, Q**

Perimeters
 balance finding, 80
 changing behavior, 69
 compliance/commitment, 66, 68
 hypothetical web sites, 76
 insider threats, 78

- detect, 79
- deter, 79
- discipline, 80
- organization's reputation, 78
- security firms, 78
- three-part approach, 79
- interactive training tool, 70
- personal devices, 74
- phishing attacks, 66
- risk examination, 68–69
- roundabouts and stop signs, 75–77
- security awareness
 - awareness model, 74
 - CEB, 71
 - employee awareness programs, 72
 - encourage users and news, 72
 - four-stage model, 72
 - international association, 73
 - judgment and commitment, 72
 - telecommunications firm, 72
 - tracking report, 73
 - training, 71
- security benefits, 74–75
- shifting perimeter, 65
- social-engineering techniques, 66
- technology professionals, 77
- traffic-control method, 75
- training, 69
- user interactions, 66
- Playing war games, 90
- Policy decision point (PDP), 108
- Policy enforcement points (PEPs), 110
- Product life cycle model
 - clustering areas, 87
 - commodity-source code, 84
 - critical trends, 86
 - disruptive trends, 87
 - emerging trends, 86
 - evolution of threats, 84
 - highest-priority threats, 83
 - security-related activity, 86
 - smartphone security threats, 85
 - sustained drivers, 86
 - threat analysis materials, 86
 - threat report, 87–88
- Product security incident response
 - processes (PSIRT), 45
- Public-private partnerships (PPPs), 58

■ R

- Radio Frequency Identification (RFID)
 - technology, 120
- References, 171
- Regulatory flood
 - cyber-security legislation, 8
 - e-discovery, 11
 - financial regulations, 10
 - high-tech exports, 8
 - IT capabilities, 8
 - personalization *vs.* privacy, 9
 - protecting personal information, 9
 - scope, 11
 - storage and protection, 8
- Retail environment, 118
- Risk misperception, 17
 - assessment models, 25
 - communication
 - asymmetry information, 27
 - building credibility, 28
 - laptops, 28
 - pirating software, 27
 - risk perceptions changing, 26
 - decision makers, 23
 - employees, 18
 - inevitable bias, 25
 - lure of the shiny bauble, 20
 - mitigate, 24
 - perception, 18
 - economic and psychological factors, 18
 - organization security posture, 18
 - security professionals, 18
 - social-media site, 18
 - security professionals, 20
 - alert fatigue, 21
 - DDoS threats, 21
 - history of, 21
 - NIMS, 20
 - scrubbing data, 21
 - security and privacy, 22
 - set and forget error and security controls, 21
 - target fixation, 20
 - threat controls, 23
 - untrusted devices, 20

■ **S**

- Sarbanes-Oxley Act (SOX), 10, 43
- Security architecture
 - 9 box of controls, 101
 - business needs, 103
 - cloud computing, 104
 - control architecture, 100
 - IT consumerization, 102–103
 - novel approaches, 100
 - privacy and regulatory requirements, 105
 - threat landscape, 104
 - threat management, 100
- Sharing security information. *See* External partnerships
- Shifting perimeter, 65
- Smartphones, 3, 96
 - security threats, 85
 - web applications, 97
- Systems development life cycle (SDLC), 45

■ **T, U, V, W, X, Y, Z**

- Tailored Access Operations (TAO), 89
- Technology professionals, 77
- Threat landscape
 - APTs, 13
 - cybercrime online, 13
 - irrefutable laws, 14
 - spearphishing attacks, 13
 - stealthy malware, 13

- Threats and vulnerabilities
 - Malware industry, 94
 - often-conflicting information, 82
 - rhetoric, 81
 - smartphones, 96
 - structured methods
 - analyzing emerging threats, 82
 - blinker security perspective, 82
 - playing war games, 90
 - product life cycle (*see* Product life cycle model)
 - risk-sensing analysis and strategy, 82–83
 - security team, 82
 - threat landscape, 83
 - threat landscape
 - barriers, 92
 - broad-brush picture, 91
 - edge case insecurity, 92
 - obscurity, 93
 - phishing, 91
 - smartphones, 91
 - social engineering attacks, 92
 - threats and identify risks, 88–89
 - web, attack surface
 - embedded devices, 95
 - industrial control systems, 95
 - IoT, 95
 - nontraditional devices, 94
- Trusted platform module (TPM), 108