SECURITY

Security Operations Center

Building, Operating, and Maintaining Your SOC

Joseph Muniz
Gary McIntyre
Nadhem AlFardan, CCIE No. 20519

ciscopress.com

Security Operations Center

Joseph Muniz
Gary McIntyre
Nadhem AlFardan

Cisco Press
800 East 96th Street
Indianapolis, Indiana 46240 USA

Security Operations Center

Joseph Muniz, Gary McIntyre, Nadhem AlFardan

Copyright© 2016 Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

First Printing November 2015

Library of Congress Control Number: 2015950793

ISBN-13: 978-0-13-405201-4

ISBN-10: 0-13-405201-3

Warning and Disclaimer

This book is designed to provide information about building and running a security operations center (SOC). Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact international@pearsoned.com.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Business Operation Manager, Cisco Press: Jan Cornelssen

Acquisitions Editor: Denise Lincoln

Managing Editor: Sandra Schroeder

Senior Development Editor: Christopher Cleveland

Senior Project Editor: Tonya Simpson

Copy Editor: Keith Cline

Technical Editors: Dr. Fred Mpala, Matthew Waters

Editorial Assistant: Vanessa Evans

Cover Designer: Mark Shirar

Composition: codeMantra

Indexer: WordWise Publishing Services

Proofreader: Sarah Kearns

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Joseph Muniz is a consultant at Cisco Systems and security researcher. Joseph started his career in software development and later managed networks as a contracted technical resource. Joseph moved into consulting and found a passion for security while meeting with a variety of customers. He has been involved with the design and implementation of multiple projects, ranging from Fortune 500 corporations to large federal networks. Joseph is the author of and contributor to several books and is a speaker for popular security conferences. Check out his blog, http://www.thesecurityblogger.com, which showcases the latest security events, research, and technologies.

Gary McIntyre is a seasoned information security professional focusing on the development and operation of large-scale information security programs. As an architect, manager, and consultant, he has worked with a wide range of public and private sector organizations around the world to design, build, and maintain small to large security operations teams. He currently holds a Masters degree from the University of Toronto and has also been a long-time $(ISC)^2$ instructor.

Dr. Nadhem AlFardan has more than 15 years of experience in the area of information security and holds a Ph.D. in Information Security from Royal Holloway, University of London. Nadhem is a senior security solution architect working for Cisco Systems. Before joining Cisco, he worked for Schlumbeger and HSBC. Nadhem is CISSP certified and is an ISO 27001 lead auditor. He is also CCIE Security certified. In his Ph.D. research, Nadhem published a number of papers in prestige conferences, such as IEEE S&P and USENIX Security, mainly around cryptoanalysis topics. His work involved him working with organizations such as Google, Microsoft, Cisco, Mozilla, OpenSSL, and many others, mainly to help them assess and fix major findings in the Transport Layer Security/Secure Sockets Layer (TLS/SSL) protocol. His work is referenced in a number of IETF standards.

About the Technical Reviewers

Dr. Fred Mpala is a security professional with broad experience in security and risk management.

Matthew Waters is a seasoned security professional and chief information security officer within the financial sector, specializing in large-scale transformation programs.

Dedications

Joseph Muniz: I would like to give a huge "thank you" to my friends and family for supporting me for this and my other crazy projects. This book goes out to Irene Muniz, Ray Muniz, Alex Muniz, Raylin Muniz, Ning Xu, my friends at Cisco, and the many other great people in my life.

Gary McIntyre: For Candice and Winston, who paid the highest price to see this through.

Acknowledgments

Joseph Muniz: I will start by thanking Gary McIntyre and Nadhem AlFardan for including me on this project. I really enjoyed collaborating on the material and hope that they do not mind my input. If they do, it is probably too late by now anyway.

I had help with validating my content and would like to recognize Jeff Williams and Aamir Lakhani. Jeff is the NetFlow ninja and assisted with reviewing my Lancope contributions. Aamir is my good friend and co-authored two books with me prior to this project. I let Aamir beat up all my drafts and appreciate both his and Jeff's time on this project.

I also want to thank Jamey Heary for helping me line up this opportunity. Jamey is a brilliant engineer and the author of *Cisco ISE for BYOD and Secure Unified Access* (Cisco Press, 2013). I would not have been involved with this project without him.

Finally, I would like to thank Denise Lincoln and the rest of the Cisco Press team for their support while we determined and then polished the final content for this book. They are a very professional group and a pleasure to work with.

Gary McIntyre: As with any book like this, it's hard to acknowledge everyone who might deserve it. Thankfully, a book about security encourages a bit of obfuscation, silence, and mystery.

To Dean T., who first taught me about correlations using Cheap Trick and Queen. For Bev T., for never being boring. For Leslie C., for teaching me not to bore others. For Beth, who could not stay to see this. For my parents, who likely still don't understand what I do for a living and love me anyway. Finally, for all the customers, colleagues, and students I've had over my career: your fingerprints are all over this thing. Thank you.

Contents at a Glance

Introduction xx

Part I: SOC Basics

Chapter 1 Introduction to Security Operations and the SOC 1

Chapter 2 Overview of SOC Technologies 35

Part II: The Plan Phase

Chapter 3 Assessing Security Operations Capabilities 69

Chapter 4 SOC Strategy 91

Part III: The Design Phase

Chapter 5 The SOC Infrastructure 103

Chapter 6 Security Event Generation and Collection 123

Chapter 7 Vulnerability Management 189

Chapter 8 People and Processes 215

Part IV: The Build Phase

Chapter 9 The Technology 245

Chapter 10 Preparing to Operate 319

Part V: The Operate Phase

Chapter 11 Reacting to Events and Incidents 347

Chapter 12 Maintain, Review, and Improve 365

Index 397

Contents

Introduction xx

Part I SOC Basics

Chapter 1 Introduction to Security Operations and the SOC 1

Cybersecurity Challenges 1
- Threat Landscape 4
- Business Challenges 7
- *The Cloud 8*
- *Compliance 9*
- *Privacy and Data Protection 9*

Introduction to Information Assurance 10

Introduction to Risk Management 11

Information Security Incident Response 14
- Incident Detection 15
- Incident Triage 16
- *Incident Categories 17*
- *Incident Severity 17*
- Incident Resolution 18
- Incident Closure 19
- Post-Incident 20

SOC Generations 21
- First-Generation SOC 22
- Second-Generation SOC 22
- Third-Generation SOC 23
- Fourth-Generation SOC 24

Characteristics of an Effective SOC 24

Introduction to Maturity Models 27

Applying Maturity Models to SOC 29

Phases of Building a SOC 31

Challenges and Obstacles 32

Summary 32

References 33

Chapter 2 Overview of SOC Technologies 35

Data Collection and Analysis 35
- Data Sources 37

Data Collection 38
The Syslog Protocol 39
Telemetry Data: Network Flows 45
Telemetry Data: Packet Capture 48
Parsing and Normalization 49
Security Analysis 52
Alternatives to Rule-Based Correlation 55
Data Enrichment 56
Big Data Platforms for Security 57
Vulnerability Management 58
Vulnerability Announcements 60
Threat Intelligence 62
Compliance 64
Ticketing and Case Management 64
Collaboration 65
SOC Conceptual Architecture 66
Summary 67
References 67

Part II: The Plan Phase

Chapter 3 Assessing Security Operations Capabilities 69
Assessment Methodology 69
Step 1: Identify Business and IT Goals 71
Step 2: Assessing Capabilities 73
Assessing IT Processes 75
Step 3: Collect Information 82
Step 4: Analyze Maturity Levels 84
Step 5: Formalize Findings 87
The Organization's Vision and Strategy 87
The Department's Vision and Strategy 87
External and Internal Compliance Requirements 87
Organization's Threat Landscape 88
History of Previous Information Security Incidents 88
SOC Sponsorship 89
Allocated Budget 89
Presenting Data 89

Closing 90
Summary 90
References 90

Chapter 4 SOC Strategy 91
Strategy Elements 91
Who Is Involved? 92
SOC Mission 92
SOC Scope 93
Example 1: A Military Organization 94
Mission Statement 94
SOC Scope Statement 95
Example 2: A Financial Organization 95
Mission Statement 95
SOC Scope Statement 95
SOC Model of Operation 95
In-House and Virtual SOC 96
SOC Services 98
SOC Capabilities Roadmap 99
Summary 101

Part III: The Design Phase

Chapter 5 The SOC Infrastructure 103
Design Considerations 103
Model of Operation 104
Facilities 105
SOC Internal Layout 106
Lighting 107
Acoustics 107
Physical Security 108
Video Wall 108
SOC Analyst Services 109
Active Infrastructure 110
Network 111
Access to Systems 112
Security 112
Compute 115

Dedicated Versus Virtualized Environment 116
Choice of Operating Systems 118
Storage 118
Capacity Planning 119
Collaboration 119
Ticketing 120
Summary 120
References 120

Chapter 6 Security Event Generation and Collection 123
Data Collection 123
Calculating EPS 124
Ubuntu Syslog Server 124
Network Time Protocol 129
Deploying NTP 130
Data-Collection Tools 134
Company 135
Product Options and Architecture 136
Installation and Maintenance 136
User Interface and Experience 136
Compliance Requirements 137
Firewalls 137
Stateless/Stateful Firewalls 137
Cisco Adaptive Security Appliance ASA 138
Application Firewalls 142
Cisco FirePOWER Services 142
Cloud Security 152
Cisco Meraki 153
Exporting Logs from Meraki 154
Virtual Firewalls 155
Cisco Virtual Firewalls 156
Host Firewalls 157
Intrusion Detection and Prevention Systems 157
Cisco FirePOWER IPS 160
Meraki IPS 161
Snort 162
Host-Based Intrusion Prevention 162

Routers and Switches 163
Host Systems 166
Mobile Devices 167
Breach Detection 168
Cisco Advanced Malware Prevention 168
Web Proxies 169
Cisco Web Security Appliance 170
Cloud Proxies 172
Cisco Cloud Web Security 172
DNS Servers 173
Exporting DNS 174
Network Telemetry with Network Flow Monitoring 174
NetFlow Tools 175
StealthWatch 177
Exporting Data from StealthWatch 179
NetFlow from Routers and Switches 182
NetFlow from Security Products 184
NetFlow in the Data Center 186
Summary 187
References 188

Chapter 7 Vulnerability Management 189
Identifying Vulnerabilities 190
Security Services 191
Vulnerability Tools 193
Handling Vulnerabilities 195
OWASP Risk Rating Methodology 197
Threat Agent Factors 198
Vulnerability Factors 198
Technical Impact Factors 200
Business Impact Factors 200
The Vulnerability Management Lifecycle 202
Automating Vulnerability Management 205
Inventory Assessment Tools 205
Information Management Tools 206
Risk-Assessment Tools 206
Vulnerability-Assessment Tools 206

Report and Remediate Tools 206
Responding Tools 207
Threat Intelligence 208
Attack Signatures 209
Threat Feeds 210
Other Threat Intelligence Sources 211
Summary 213
References 214

Chapter 8 People and Processes 215

Key Challenges 215
Wanted: Rock Stars, Leaders, and Grunts 216
The Weight of Process 216
The Upper and Lower Bounds of Technology 217
Designing and Building the SOC Team 218
Starting with the Mission 218
Focusing on Services 219
Security Monitoring Service Example 220
Determining the Required SOC Roles 223
Leadership Roles 224
Analyst Roles 224
Engineering Roles 224
Operations Roles 224
Other Support Roles 224
Working with HR 225
Job Role Analysis 225
Market Analysis 225
Organizational Structure 226
Calculating Team Numbers 227
Deciding on Your Resourcing Strategy 228
Building Your Own: The Art of Recruiting SOC Personnel 229
Working with Contractors and Service Bureaus 229
Working with Outsourcing and Managed Service Providers 230
Working with Processes and Procedures 231
Processes Versus Procedures 231
Working with Enterprise Service Management Processes 232
Event Management 232

Incident Management 233
Problem Management 233
Vulnerability Management 233
Other IT Management Processes 233
The Positives and Perils of Process 234
Examples of SOC Processes and Procedures 236
Security Service Management 236
Security Service Engineering 237
Security Service Operations 238
Security Monitoring 239
Security Incident Investigation and Response 239
Security Log Management 240
Security Vulnerability Management 241
Security Intelligence 241
Security Analytics and Reporting 242
Breach Discovery and Remediation 242
Summary 243

Part IV: The Build Phase

Chapter 9 The Technology 245
In-House Versus Virtual SOC 245
Network 246
Segmentation 247
VPN 251
High Availability 253
Support Contracts 254
Security 255
Network Access Control 255
Authentication 257
On-Network Security 258
Encryption 259
Systems 260
Operating Systems 261
Hardening Endpoints 262
Endpoint Breach Detection 263
Mobile Devices 264
Servers 264

Storage 265
Data-Loss Protection 266
Cloud Storage 270
Collaboration 271
Collaboration for Pandemic Events 272
Technologies to Consider During SOC Design 273
Firewalls 273
Firewall Modes 273
Firewall Clustering 276
Firewall High Availability 276
Firewall Architecture 277
Routers and Switches 279
Securing Network Devices 280
Hardening Network Devices 280
Network Access Control 281
Deploying NAC 282
NAC Posture 284
Architecting NAC 285
Web Proxies 290
Reputation Security 290
Proxy Architecture 292
Intrusion Detection/Prevention 295
IDS IPS Architecture 295
Evaluating IDS IPS Technology 296
Tuning IDS/IPS 298
Breach Detection 300
Honeypots 301
Sandboxes 302
Endpoint Breach Detection 303
Network Telemetry 306
Enabling NetFlow 308
Architecting Network Telemetry Solutions 310
Network Forensics 312
Digital Forensics Tools 313
Final SOC Architecture 314
Summary 317
References 318

Chapter 10 Preparing to Operate 319

Key Challenges 319

People Challenges 319

Process Challenges 320

Technology Challenges 321

Managing Challenges Through a Well-Managed Transition 321

Elements of an Effective Service Transition Plan 322

Determining Success Criteria and Managing to Success 322

Deploying Against Attainable Service Levels 323

Focusing on Defined Use Cases 325

Managing Project Resources Effectively 328

Marching to Clear and Attainable Requirements 329

Staffing Requirements for Go-Live 329

Process Requirements for Go-Live 330

Technology Requirements for Go-Live 331

Using Simple Checks to Verify That the SOC Is Ready 332

People Checks 332

Process Checks 336

Technology Checks 340

Summary 346

Part V: The Operate Phase

Chapter 11 Reacting to Events and Incidents 347

A Word About Events 348

Event Intake, Enrichment, Monitoring, and Handling 348

Events in the SIEM 349

Events in the Security Log Management Solution 350

Events in Their Original Habitats 350

Events Through Communications and Collaboration Platforms 350

Working with Events: The Malware Scenario 351

Handling and Investigating the Incident Report 353

Creating and Managing Cases 354

Working as a Team 355

Working with Other Parts of the Organization 357

Working with Third Parties 359

Closing and Reporting on the Case 362
Summary 363

Chapter 12 Maintain, Review, and Improve 365
Reviewing and Assessing the SOC 366
Determining Scope 366
Examining the Services 367
Personnel/Staffing 369
Processes, Procedures, and Other Operational Documentation 371
Technology 372
Scheduled and Ad Hoc Reviews 373
Internal Versus External Assessments 374
Internal Assessments 374
External Assessments 374
Assessment Methodologies 375
Maturity Model Approaches 375
Services-Oriented Approaches 376
Post-Incident Reviews 378
Maintaining and Improving the SOC 381
Maintaining and Improving Services 381
Maintain and Improving Your Team 383
Improving Staff Recruitment 383
Improving Team Training and Development 384
Improving Team Retention 386
Maintaining and Improving the SOC Technology Stack 387
Improving Threat, Anomaly, and Breach-Detection Systems 388
Improving Case and Investigation Management Systems 391
Improving Analytics and Reporting 392
Improving Technology Integration 392
Improving Security Testing and Simulation Systems 393
Improving Automated Remediation 394
Conclusions 395

Index 397

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([ ]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

Many security books are available, but they focus on either products or on very high-level security best practices. We could not find a book about developing a security operations center. This lack of coverage meant that people interested in this topic would have to take the literature available from security books and interview existing SOC leaders to learn how it should be done. We identified this gap and decided to write this book.

In this book, we use a blend of industry experience and best practices of many of our customers to create a guide for those interested in how a SOC should be developed and managed. This book presents the collective view of its three authors (each a consultant at Cisco Systems). We have worked with hundreds of customers, ranging from Fortune 500 companies to large military organizations, giving us a broad and deep understanding about how SOCs are developed and run in the real world. We hope that our research and experience assists you with your existing SOC or with future security operations projects.

Who Should Read This Book?

This book is written for anybody interested in learning how to develop, manage, or improve a SOC. This book, which is based on a blend of our industry experience and best practices, can serve as a guide for those interested in creating and managing a SOC. A background in network security, network management, and network operations would be helpful, but it is not required to benefit from this book.

How This Book Is Organized

Chapter 1, "Introduction to Security Operations and the SOC": This first chapter provides an introduction to security operations. First, we examine challenges that organizations face that justify the investment in a SOC. This chapter also covers high-level SOC topics such as vulnerability management, threat intelligence, digital investigation, and data collection and analysis, thus setting the stage for the rest of the book.

Chapter 2, "Overview of SOC Technologies": This chapter explains how to deploy and customize SOC technologies. The goal is to gain a basic understanding of technology and services found in most modern SOC environments. Topics include data collection, data processing, vulnerability management, and case management.

Chapter 3, "Assessing Security Operations Capabilities": This chapter describes a methodology for developing SOC requirements and assessing security operations capabilities against those requirements. The purpose is to help define a SOC strategy and create a supporting roadmap.

Chapter 4, "SOC Strategy": This chapter explains how to develop a strategy and roadmap for SOC capabilities using results from tactics learned in Chapter 3. These can be used as a benchmark to validate progress and justify future investment in the SOC.

Chapter 5, "The SOC Infrastructure": This chapter covers design considerations for a SOC infrastructure. Topics include the SOC facility, analyst environment, and various technology considerations.

Chapter 6, "Security Event Generation and Collection": Most SOC environments focus heavily on collecting and analyzing data. This chapter describes common tools found in a SOC environment that are used to produce and collect such data.

Chapter 7, "Vulnerability Management": This chapter explains how a SOC can create and deliver a vulnerability management practice. Topics covered include identifying and managing vulnerabilities, calculating risk, and ranking threats to properly prioritize which vulnerabilities to address first.

Chapter 8, "People and Processes": This chapter focuses on developing an appropriate governance and staffing model based on the range of services the SOC will provide. The goal is to not only recruit the right team but also to retain those valuable resources. The chapter also explains how to leverage outsourcing or managed service providers when the situation calls for such resources.

Chapter 9, "The Technology": SOC technologies were briefly discussed in earlier chapters, but this chapter covers how common technologies found in the SOC fit into a SOC architecture. The focus is on design considerations for SOC network, security, collaboration, and storage technologies.

Chapter 10, "Preparing to Operate": This chapter walks you through common steps required to transition a SOC to a fully operational state. Topics include challenges faced before going live, transition plan development, and the impact of proper project management.

Chapter 11, "Reacting to Events and Incidents": In this chapter, you learn how the people, processes, and technology parts of an effective SOC come together to react to events and incidents. This chapter includes a storyline that walks you through how an effective SOC would react to an identified security incident.

Chapter 12, "Maintain, Review, and Improve": The final chapter of this book wraps things up by explaining how to review and maintain the SOC and its services. This chapter also covers how to continuously improve the SOC services to guide its evolution. Improvement topics include how to identify issues with the SOC so that they can be remediated to improve SOC service.

This page intentionally left blank

Chapter 1

Introduction to Security Operations and the SOC

"It's better to light a candle than curse the darkness."

This book focuses on the best practices to develop and operate a security operations center (SOC). The journey begins with a review of important concepts relevant to information security and security operations. This chapter opens with a discussion about the continuously evolving security landscape and how new cybersecurity challenges impact how we perceive security operations. The discussion then turns to information assurance and its link to both risk management and security operations. This chapter does not cover information assurance and risk management in depth, but instead provides you with sufficient background information on these topics while relating them to security operations. The chapter then covers incident response and how it is used by security operations. The chapter concludes by introducing a number of concepts associated with the core topics of the book: planning, designing, building, and operating a SOC.

Cybersecurity Challenges

Security attacks are becoming increasingly complex and exhibiting increasingly sophisticated capabilities. So, addressing the complexity and sophistication of such attacks must include not only investing in preventive measures but also the development of intelligent and integrated monitoring capabilities incorporated into an incident response program.

Arguably, getting compromised at some point is inevitable. As the previous CEO of Cisco Systems, John Chambers, said, "There are two types of companies: those who have been hacked and those who don't yet know they have been hacked." So, be warned: A security breach is not an *if* but a *when*. The good news is that a breach does not necessarily mean that the business will immediately experience negative impact, because attackers usually need time to accomplish their objectives beyond gaining unauthorized access to the network. Discovering and preventing this type of behavior is just one of the many reasons organizations develop a SOC.

One exercise to help people understand the inevitability of cyberthreats is to ask this question: If you knew you were going to be breached, would you perform security differently within your organization? This leads to a series of questions that you should ask leadership within your organization about how your business deals with a security compromise. Somebody needs to be responsible for answering these questions; for most organizations, those people run the SOC or have security operations responsibilities.

If you need to justify a SOC budget, here are some questions focused on dealing with a security compromise:

- **How** can you detect a compromise?
- **How** severe is the compromise?
- **What** is the impact of the compromise to your business?
- **Who** is responsible for detecting and reacting to a compromise?
- **Who** should be informed or involved and **when** do you deal with the compromise once detected?
- **How** and **when** should you communicate a compromise internally or externally, and is that needed in the first place?

The OODA Loop[1] is one methodology that addresses some of these questions. The concept was first developed by military strategist and USAF Colonel John Boyd.[2] Colonel Boyd created a four-step approach designed to determine the appropriate response to a problem. An OODA Loop, as shown in Figure 1-1, consists of the following steps: observe, orient, decide, and act. In the context of cybersecurity, the four steps are as follows:

- **Observe:** Monitor, collect, and store data from various points in your network as the first step in the OODA Loop.
- **Orient:** Analyze collected data in search of suspicious activities. This usually involves the use of tools to process and analyze incoming and stored data.
- **Decide:** Determine an action course based on the results of the analysis phase and the experience you have gained from previous loop iterations.
- **Act:** Execute the action course you determined in the preceding step.

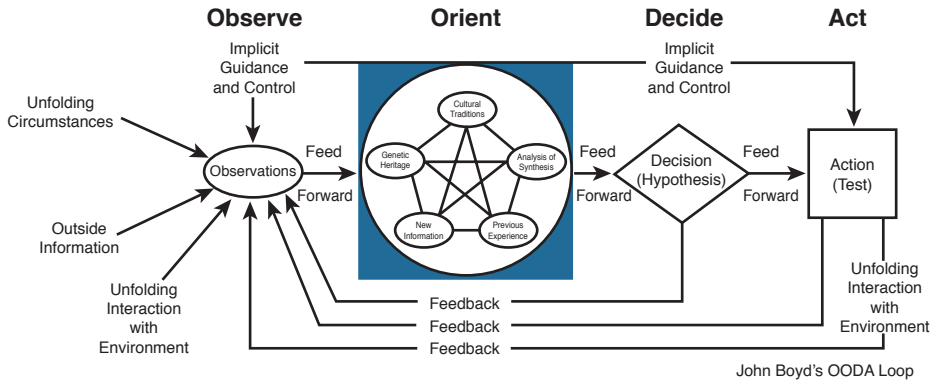

Figure 1-1 *The OODA Loop*

The OODA Loop assumes that continuous improvement is an integrated part of the process, allowing you to learn from your previous experiences, feeding lessons learned into the loop activities to achieve better performance every time you complete the four steps. Continuous monitoring and enforcement concepts are an ongoing theme throughout this book.

The OODA Loop may have been designed to deal with military attacks, but the concepts apply to defending any form of attack, including cyberthreats. For our next model, let's switch to the attacker's viewpoint by reviewing the *cyber kill chain*. The cyber kill chain,[3] developed by Lockheed Martin's Computer Incident Response Team and shown in Figure 1-2, describes the progression an attacker follows when planning and executing an attack against a target. This model helps security professionals identify security controls and actions that can be implemented or improved to detect, deny, and contain an attack scenario.

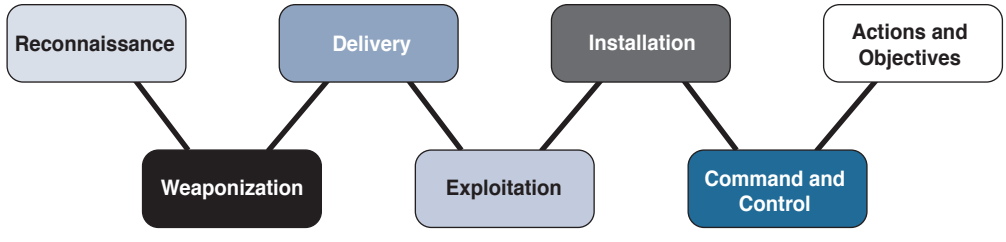

Figure 1-2 *The Cyber Kill Chain*

The various phases of the cyber kill chain are as follows:

- **Phase 1, Reconnaissance:** Research, identification, and selection of targets, often represented as crawling Internet websites such as conference proceedings and mailing lists for e-mail addresses, social relationships, or information on specific technologies.
- **Phase 2, Weaponization:** Coupling a remote-access Trojan with an exploit into a deliverable payload, typically by means of an automated tool also known as a weaponizer.

- **Phase 3, Delivery:** Transmission of the weapon to the targeted environment.
- **Phase 4, Exploitation:** Triggers the intruder's code. Most often, exploitation targets an application or operating system vulnerability, but it could also more simply exploit the users themselves or leverage an operating system feature that auto-executes code.
- **Phase 5, Installation:** Installation of a remote-access Trojan or back door on the victim system enables the adversary to maintain persistence inside the environment.
- **Phase 6, Command and Control:** Beacon outbound to an Internet controller server to establish a command and control channel.
- **Phase 7, Actions and Objectives:** Intruders take actions to achieve their original objectives. Typically, this objective is data exfiltration that involves collecting, encrypting, and extracting information from the victim environment; violations of data integrity or availability are potential objectives as well. Alternatively, the intruders may only desire access to the initial victim box for use as a hop point to compromise additional systems and move laterally inside the network.

These models make up a high-level way to look at cybersecurity attack and defense concepts. According to the cyber kill chain, attackers perform reconnaissance to identify the easiest and most effective way to breach a network. Defense teams using the OODA Loop may catch this behavior during the observing or orientation stage and decide to react by locking down security or raising awareness of a possible threat from an identified source. An attacker may move to Phase 4 of the cyber kill chain by weaponizing and delivering an exploit against a target. The OODA Loop may once again identify the attack during the observe or orientation stage and react through containing the threat and patching the weakness to avoid future attacks. So essentially, the OODA Loop is a defense strategy against every phase of the cyber kill chain.

Developing security operations and security incident response capabilities is critical to breaking or reducing the impact of an attacker executing the cyber kill chain against your organization. The OODA Loop is just one conceptual model that you can use to break the chain or contain the chain of events during an incident. The important question to ask is this: How does your organization apply the OODA Loop or similar concepts against different stages of the cyber kill chain when an attack against your organization occurs? Most SOCs execute a variation of the OODA Loop by using controls associated with these three elements: people, processes, and technology.

Threat Landscape

The Verizon 2015 Data Breach Investigation Report (DBIR)[4] showed that 60 percent of businesses being breached happened within minutes or less. The report also showed that half of these incidents took anywhere from months to even years before being

uncovered. So in summary, breaches tend to happen very quickly and on average take a long time to be detected by the targeted organization. These numbers demonstrate the importance of having an effective security operations program in which a mature SOC plays a significant role. Figure 1-3, taken from the DBIR, compares the time spans of when organizations are compromised versus when they are able to discover the breach.

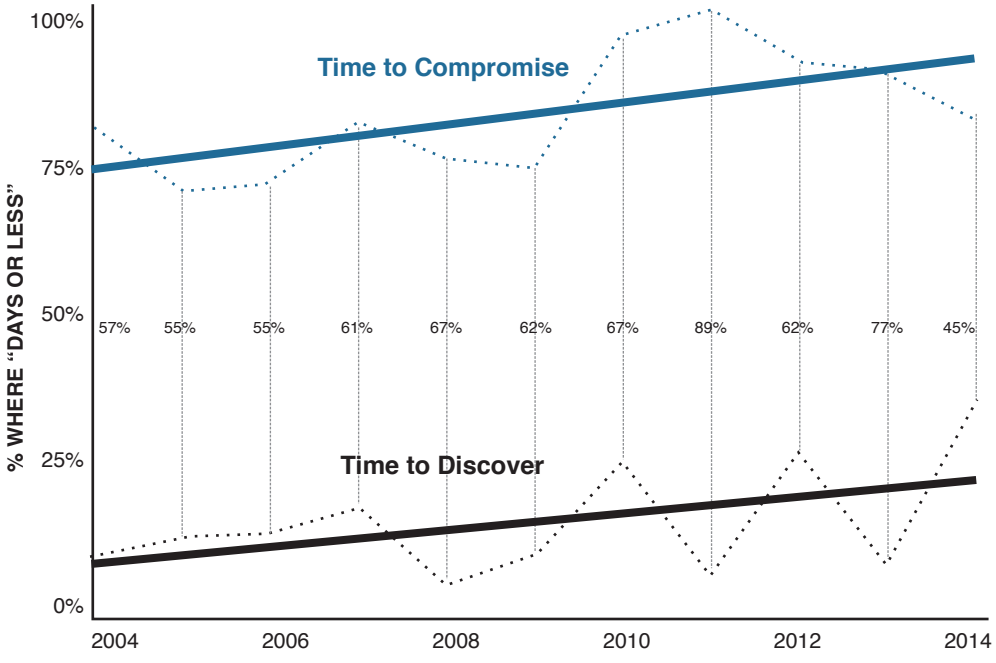

Figure 1-3 *2015 DBIR Comparing Breach and Discovery Time*

As the security landscape continues to evolve, new and automated offensive tools become readily available for a larger audience. Well-organized libraries of offensive tools are now packaged on free-to-download-and-use Linux distributions such as Backtrack[5] and Kali,[6] making it possible for almost anyone to test and develop tools and exploits. Attackers known as Script Kiddies might not understand the details of how the attack works, but packaged offensive tools make launching sophisticated attacks as simple as point-and-click execution. An analogy to this scenario is that the average person cannot build a telephone but can make a phone call as long as the system is prebuilt and ready to use. Many attack tools are developed in a point-and-click fashion, such as shown in Figure 1-4, which showcases a weaponized denial-of-service (DoS) attack application known as Low Orbit Ion Cannon (LOIC).

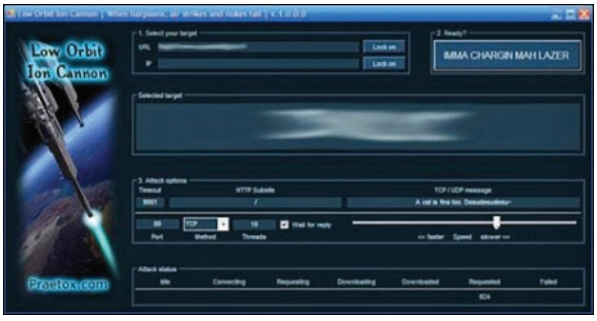

Figure 1-4 *LOIC DoS Attack Application*

Why are security products unable to stop these threats? Aren't governments and big businesses investing significant amounts of money into developing countermeasures to cyberattacks? The answer is yes; however, there is a reason why defense technologies need to continue to evolve. Just like solution providers, attackers have labs for researching products. If a "silver bullet" or single solution hits the market and is effective at preventing cyberattacks, that product will become popular and purchased by many organizations, along with their adversaries. Hackers will test various exploits against the new product until one succeeds, and will then either sell the exploit on underground markets, weaponize it for less-skilled attackers, or use it for some other malicious purpose.

Security researchers saw this pattern when the sandbox concept was introduced to detect malicious software. Administrators put a sandbox security product hosting various vulnerable applications at the edge of their inside network with the intention of making it the first target malware would strike during a breach. Malware would launch inside that monitored sandbox, giving away its identity, and thus the other security products could be made aware of the attack. Attackers learned about this technology and developed techniques such as targeting specific systems versus attacking the most vulnerable system, delaying the attack long enough to bypass the sandbox and so on. Figure 1-5 shows an article[7] from Networkworld.com explaining how hackers are now bypassing sandbox technology. This is why the attack and defense innovations continue to be a cat-and-mouse game, meaning both sides will develop a technique that eventually the other side will adapt to.

Figure 1-5 *Networkworld.com Article on Bypassing Sandbox Technology*

Where are these exploits being sold once developed? Many new exploits can be found on underground and notably growing commercial zero-day vulnerability markets operated by private brokers. Examples of boutique exploit vendors that sell *known-unknown* zero-day flaws include the French firm VUPEN Security,[8] Netragard,[9] ReVuln,[10] and Endgame Systems.[11] NSS Labs[12] estimates that on a yearly basis more than 100 *known-unknown* zero-day vulnerabilities are available for sale. To find examples of underground markets selling zero-day exploits, download the Tor browser from https://www.torproject.org/download/download, access the Darknet, and go to the Hidden Wiki at http://wikitjerrta4qgz4.onion/. Figure 1-6 shows a list of websites marketing anything from contract killers to the infamous Silk Road marketplace.

Figure 1-6 *Darknet Marketing Sources*

Warning Access the Darknet at your own risk! Because of the anonymous nature of hosting and accessing websites via Tor, many onion network resources contain malicious techniques not found on the Internet.

Business Challenges

In addition to the technical security threat landscape, legal and business-imposed decisions impact the way organizations operate information security. Examples of such decisions include moving services and information to the cloud; meeting compliancy requirements; the proliferation of bring your own device (BYOD); and the rise of the Internet of Everything (IoE), which brings people, processes, things, and data together by combining machine-to-machine (M2M), person-to-machine (P2M), and person-to-person (P2P) connections. For example, the IoE networked connection of people, processes, data, and things, shown in Figure 1-7, introduces challenges related to collecting, processing, storing, and analyzing large volumes of data generated at high velocity and with a variety of formats (that is, big data analytics, as discussed later in this chapter and in Chapter 2, "Overview of SOC Technologies."

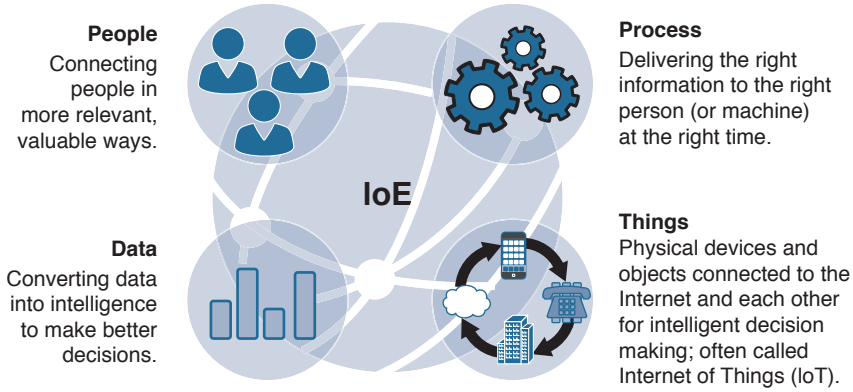

Figure 1-7 *Networked Connection of People, Process, Data, Things (Source: Cisco Internet of Everything Report[13])*

The Cloud

Cloud services are here to stay and evolve. Over the past decade, there has been a steady increase in the adoption of cloud services by small, medium, and large organizations from various industries. This trend maps to the economic gains expected from consuming cloud-based services.

According to various studies related to cloud adoption, security is one of the top concerns for chief information officers (CIO) for migrating their services to the cloud. For example, a study conducted by ChangeWave Research in April 2014[14] identified security concerns as being the leading reason why both large companies (more than 1000 employees) and small- to medium-sized companies (fewer than 1000 employees) are not moving to the cloud. Another study conducted by the firm McKinsey in conjunction with the World Economic Forum in May 2014[15] revealed that "some 70 percent of the respondents said that security concerns had delayed the adoption of public cloud computing by a year or more, and 40 percent said such concerns delayed enterprise-mobility capabilities by a year or more."

The exact security concerns vary from the ability of cloud service providers guaranteeing basic access control and confidentiality all the way to achieving audit and compliance requirements. Adoption rates also vary between the use of public and private cloud because private cloud tends to have more security based on being dedicated to a single client or company. What this book addresses is how to operate your SOC in an environment where cloud services are being consumed for business purposes regardless of the cloud model. You learn more about cloud services later in this book.

Compliance

Being compliant with mandatory or discretionary information security or privacy standards requires not only an investment in technology but also, in almost all cases, a fair amount of culture change. Required changes can also include a reform to the model by which information is managed and how security operations are conducted.

Typically, standards mandate reactive and preventive security measures to protect information. This means that all regulatory and discretionary security standards explicitly highlight the need for information security operations. In some cases, this is referred to as the need for continuous monitoring of the information security posture.

Examples of security standards many organizations must comply with include the following:

- The Payment Card Industry Data Security Standard (PCI DSS),[16] a standard set by the PCI Security Standards Council (PCI SSC) that applies to all entities that store, process, and/or transmit cardholder data, mandates a number of technical and operational security requirements. The standard mandates, under the "regularly monitor and test networks" control objective mapped to PCI DSS requirements 10 and 11, that organizations must regularly monitor and test networks to find and fix vulnerabilities.
- ISO/IEC 27001:2013, a standard that organizations can adopt to manage information security as a program, requires organizations to operate information security technical and nontechnical controls under the umbrella of an information security management system (ISMS) in which a plan, do, check, and act (PDCA) lifecycle is operated and maintained. Annex A.10.10 of the standard, titled "Monitoring," lists a number of important operation controls in the areas of audit logging, monitoring system use, and so on. The check phase of the PDCA lifecycle requires you to assess and, where applicable, measure process performance against ISMS policies, objectives and practical experience, and report the results to management for review.

PCI DSS and ISO/IEC 27001:2013, along with other information security and privacy standards, demand that security controls be properly monitored and that information security events and incidents be appropriately handled. Failure to comply with standards may lead to large fines or jail time for all persons held accountable for information assurance.

Privacy and Data Protection

In addition to business-centric standards, organizations must adhere to country-specific laws and regulations related to privacy and data protection. Examples of country-specific laws and regulations include the following:

- **United States:** US-EU Safe Harbor on Data Protection directive went into effect in October 1998 and prohibits the transfer of personal data to non-European Union countries that do not meet the EU-defined *adequacy* standard for privacy protection.

- **Germany:** The Federal Data Protection Act[17] (*Bundesdatenschutzgesetz* in German, BDSG) covers a range of data protection issues. For example, according to the act, organizations must have policies, procedures, and controls in place to protect all data types and categories that are under the BDSG umbrella. In addition, each German state has a data protection law of its own.
- **Japan:** The Act on the Protection of Personal Information[18] (APPI) prescribes the duties to be observed by entities handling personal information with the objective of protecting the rights and interests of individuals.
- **United Kingdom:** The Data Protection Act 1998[19] governs the protection of personal data in the UK.

Developing security operations capabilities is critical to supporting an organization's compliance state with such regulations; these same capabilities also allow organizations to react appropriately to security incidents that might result in an infringement of the law. Being compliant with these and other mandatory laws and regulations comes in direct support to ensuring that information is protected and monitored against potential attacks that would impact its availability, integrity, authentication, and confidentiality. In many cases, *reasonable effort* (as legally understood) must be used to protect sensitive data; otherwise, persons responsible for information assurance face potential fines or even jail time.

Let's now review the concepts underlying information assurance.

Introduction to Information Assurance

According to the U.S. Department of Defense (DoD) Directive 8500.01E,[20] information assurance (IA) refers to "measures that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for restoration of information systems by incorporating protection, detection, and reaction capabilities." The National Institute of Standards and Technology (NIST) SP 800-59[21] and other NIST documents use the same definition for IA.

Taking a United Kingdom (UK) perspective to IA, the UK Home Office defines IA[22] as "the practice of managing information-related risks around the confidentiality, integrity and availability of information, particularly sensitive information. It covers information whilst in storage, processing, use or transit; and the risks created by both malicious and non-malicious actions." The UK Cabinet Office has a slightly different definition for IA, stating that "the confidence that information systems will protect the information they carry and will function as they need to, when they need to, under the control of legitimate users" and that "information systems include any means of storing, processing or disseminating information including IT systems, media and paper-based systems." Most requirements for IA will contain similar characteristics in the language used for protecting

various aspects of data handled by your organization. It is extremely important to know exactly how IA is defined in your country and all associated requirements your organization must meet to avoid legal ramifications.

A question that is often asked is this: What is the difference between IA and information security? Taking NIST as our reference, NIST 800-53[23] defines information security as "the protection of information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide confidentiality, integrity, and availability." You should notice that a close relationship exists between IA and information security. Considering the NIST definitions of IA and information security, you can think of IA as a superset that covers information security. In general, IA tries to address the different risks associated with information systems or media, encompassing technical, procedural, and human factor aspects involved in delivering physical and digital protection for information systems or media.

Before trying to address IA or information security requirements, it is imperative to understand the security risks associated with your environment. This is where the risk management discipline becomes relevant. It is worth noting that establishing a SOC supports the basic requirement of managing the security risk of an organization. So, the next section introduces risk management concepts.

Introduction to Risk Management

Let's start by defining what *risk* is. A common definition of risk is the *probability* of a *threat* executing on *vulnerability* and the *impact* resulting from successful exploitation. This definition of risk is similar to the one listed in the NIST guide for conducting risk assessments.[24] Dealing with risk cannot be achieved without understanding the organization-specific threat landscape and vulnerability status. This can be achieved using various types of security services such as a risk assessment.

A risk assessment is the process used to assign some value to risk associated with assets. The output of a risk assessment can then be used to make better-informed decisions on how to deal with risk. Decisions include one or more of the following: *mitigate*, *transfer*, *accept*, or *avoid*. Risk management involves combining the output of risk assessment with the decision on how to address risk.

Examples of frameworks that try to formalize the topic of risk management include ISO/IEC 27005:2010, ISO/IEC 31000:2009, NIST SP 800-39, OWASP Risk Rating Methodology, and the DoD Risk Management Framework (RMF). These frameworks use similar processes for managing risks, as shown in Figure 1-8, taken from the ISO/IEC 31000:2009 standard. Taking the DoD RMF as an example, the framework assumes a six-step process. The process starts with risk categorization and ends with effectiveness measurement through security controls monitoring.

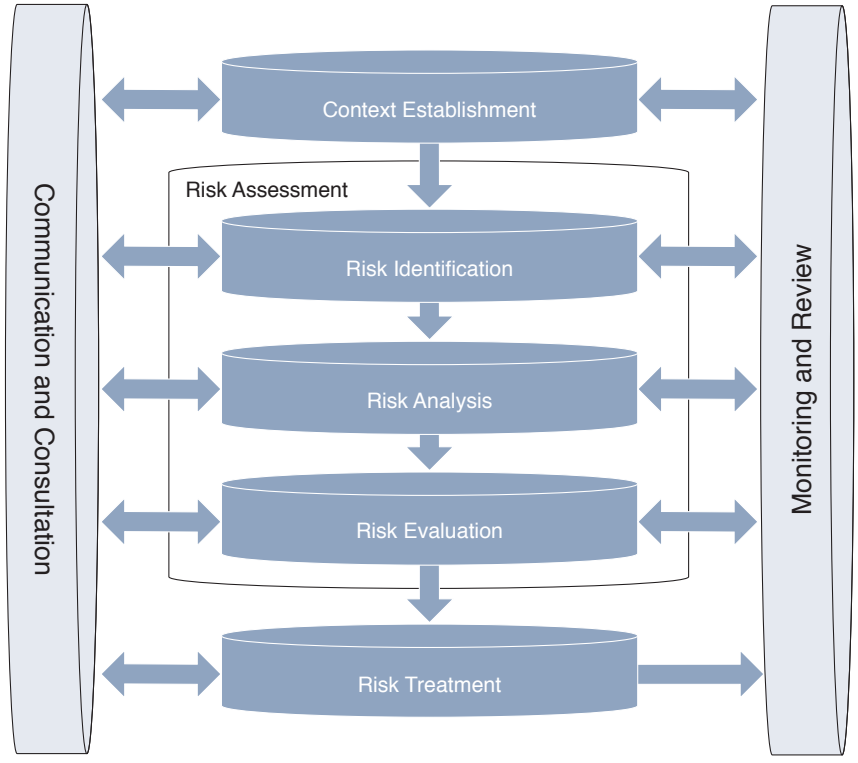

Figure 1-8 *Risk Management Framework (As per ISO/IEC 31000:2009)*

The task of assessing and managing risk should not be regarded as only a point-in-time exercise, but rather a continuous process that accepts input from various real-time and offline sources. Information generated by security operations is a vital source to the risk-assessment process. Consider the example in which security operations identified that a malformed IP packet destined for a critical system has caused this system to crash. The sequence of events should be investigated, and the risk associated with this system, and possibly other similar systems, should be reviewed. Changes made to the environment will have an effect on risk associated with individual assets and the overall security posture of the organization. For example, how should the organization react to the resignation of a system administrator who has responsibilities over privileged access systems? The proactive actions taken by the organization should be based on predefined processes that have been created to address new risk values introduced by this event. Another example is the public announcement of a new security vulnerability affecting a number of internal assets that are *not* publicly accessible. What should the new risk values associated with these internal assets be as a result of this vulnerability discovery? Let's consider the latter scenario to assess the risk associated with the internal assets based on the new security vulnerability announcement, using the Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) risk-assessment methodology that was first developed by the Software Engineering Institute (SEI)[25] at Carnegie Mellon University. Table 1-1 demonstrates going through an OCTAVE risk-assessment exercise.

Table 1-1 *Example of a Risk-Assessment Exercise*

Risk Component	Description
Vulnerability	A new vulnerability affecting internal assets has been announced. The analysis shows that a number of critical internal assets are indeed vulnerable.
Threat description	Vulnerable assets are classified as critical. The attack can be easily executed on an asset if the attacker can access the service over the network.
Existing controls	The internal assets are not connected to the Internet. The assets are protected by a firewall that allows internal users only. The assets are protected by an intrusion prevention system (IPS); however, the IPS vendor has not released signatures that can protect the assets from being exploited through the newly announced vulnerability.
Probability	Unlikely. The assets can be only exploited by internal users who have access to the assets over the internal network.
Impact	Critical. Exploiting the vulnerability results in the attacker gaining full administrative access to the system.

Next, we will map the information in Table 1-1 to a risk heat map that we can use to identify the priority of this event based on the *probability* of it occurring and *impact* to the business if the event occurs. The risk heat map is a simple color-based representation using a two-dimensional matrix calculating probability and impact levels to identify the organization's level of concern for the event. Figure 1-9 depicts a four-level severity risk heat map using low, medium, high, and extreme as the depth levels of concern.

In the next risk heat map example, we assessed the probability to be unlikely and the impact to be critical, resulting in a moderate priority course of action. It is important to note that the impact, probability, and the priority values presented in Figure 1-9 should be clearly defined and understood when designing your security assessment methodology. For example, a moderate action plan could be the implementation of the required system patches following the standard process of evaluating the patches in the test environment, following the standard change management process, and then applying the patches in production during a standard change window. You learn how to create different forms of heat maps and calculating different risk values in Chapter 7, "Vulnerability Management."

Probability		Impact			
		Negligible	Marginal	Critical	Catastrophic
	Certain	HIGH	HIGH	EXTREME	EXTREME
	Likely	MEDIUM	HIGH	HIGH	EXTREME
	Possible	LOW	MEDIUM	HIGH	EXTREME
	Unlikely	LOW	LOW	MEDIUM	EXTREME
	Rare	LOW	LOW	MEDIUM	HIGH

Figure 1-9 *Simple Risk Heat Map*

Information Security Incident Response

Detecting and responding to information security incidents is at the core of security operations. The team assigned to security operations is expected to monitor the organization's assets within scope and react to security events and incidents, including the detection and investigation of what would be considered indicators of compromise (IOC). IOCs are technical and nontechnical security compromise signals that could be detected with technology, processes, and people. For example, detecting a user accessing files from a USB memory device on an enterprise desktop machine can indicate that a policy related to restricting the use of USB memory devices has been violated and that a security control has been circumvented. Another example is detecting the IP address of an Internet botnet command-and-control server inside your network probably indicates that one or more of your systems have been compromised.

Responding to incidents starts by first detecting that an incident has actually occurred and is within the scope assigned to the security operations team. An example is capturing some IOCs with a monitoring system and either investigating the events or handing off the forensic tasks to another party. This process typically involves different internal and potentially external entities depending on factors such as the tools used, type of event, skills of parties involved, and location triggering the alarms.

Preparing a SOC to manage incidents extends to cover people, processes, and of course, technology. For example, a SOC is usually expected to educate users about how they must report security incidents and keep informed of the channels available for users to report what is perceived as a security incident. The exact sequence of steps to follow and the parties to involve depend on the nature of the incident. A typical incident-handling process follows the list of steps presented in the incident response (IR) timeline in Figure 1-10. Let's look at *detection*, which is the second stop on the incident response timeline.

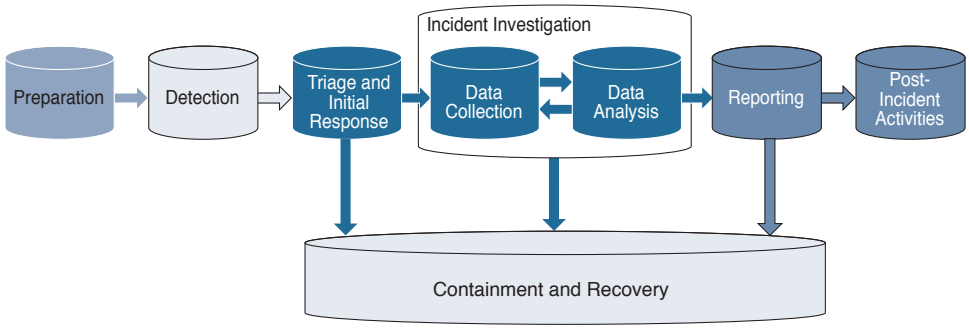

Figure 1-10 *Incident Response Timeline*

Incident Detection

Detection refers to the phase in which an incident is observed and reported by people or technology, and the process that handles the reporting aspects. A sample process for handling incoming incident reports is shown in Figure 1-11. For the process to be effective, the following must be documented and formalized:

- Identify the sources, such as people and technology, that are responsible for detecting and reporting computer security incidents.
- Identify the channels through which computer security incidents should be reported.
- Identify the steps that should be taken to accept and process computer security incident reports.
- Identify the requirements on people and technology for the process to work.

The steps described in the sections that follow outline the actions you expect to take after an incident has been reported, regardless of the channel through which the incident has been reported. Let's look at what follows detection, starting with triage.

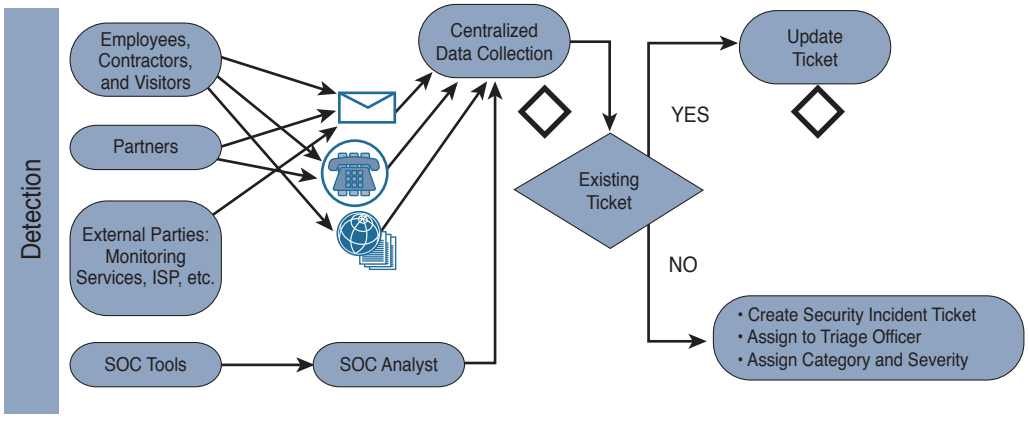

Figure 1-11 *Incident Detection Phase*

Incident Triage

Incident triage represents the initial actions taken on a detected event that is used to determine the remaining steps according to the incident response plan. The triage phase consists of three subphases: *verification*, *initial classification*, and *assignment*. An example of applying these phases is receiving an alert about an unauthorized system accessing a sensitive part of the network. The administrator would verify whether the alarm is something that should be investigated, classify the risk level for the event, and assign responsibility for handling the incident based on the level of associated risk.

The triage phase is concerned with answering questions such as the following:

- Is the incident within the scope of the program?
- Is it a new incident, or is it related to a previous reported one?
- What category should the incident be assigned to?
- What severity level should be assigned to the incident?
- Who should be assigned to analyze and investigate the incident (that is, the incident handler)?
- Is there a timeframe associated with the incident?

The triage process needs to be developed to prioritize incidents and move them along the incident response timeline to be analyzed and eventually conclude with some form of resolution. This usually involves placing the incident in a *category* for organization and assignment purposes along with applying a *severity level* so that incidents can be prioritized. Let's look at categorizing incidents and then at applying severity levels.

Incident Categories

All computer security incidents should be assigned a category. The category value identifies the type of the incident and its potential type of impact. Assigning a category value helps the SOC allocate the appropriate resources to analyze and investigate a computer security incident. Table 1-2 shows a sample list of categories that you could use for categorizing incidents.

Note that some incidents might have more than one category and that the category can change as the incident progresses or as the investigation of the incident unfolds new findings.

Table 1-2 *Computer Security Incident Categories*

Category Number	Name	Description
0	Exercise	This is used when conducting an approved exercise such as an authorized penetration test.
1	Unauthorized Access	This represents when an individual gains logical or physical access without permission to a client network, system, application, data, or other resource.
2	Denial of Service (DoS)	This is used when an attack successfully prevents or impairs the normal authorized functionality of networks, systems, or applications by exhausting resources.
3	Malicious Code	This identifies when there is a successful installation of malicious software, such as a virus, worm, Trojan horse, or other code-based malicious entity, that infects an OS or application.
4	Scans/Probes/ Attempted Access	This includes any activity that seeks to access or identify a client computer, open ports, protocols, service, or any combination for a future attack.
5	Investigation	This includes unconfirmed incidents that are potentially malicious or anomalous activity deemed by the reporting entity to warrant further review.

Incident Severity

Severity levels are based on the expected or observed impact of an incident. This is used for the prioritization of the incident, taking into account the amount of resources that should be assigned, and determines the escalation process to follow. Note that the severity level of an incident may change as the investigation unfolds. Chapter 7 covers calculating severity levels.

Table 1-3 shows a list of sample severity values. Levels can be more or less granular depending on your operational requirements.

Table 1-3 *Computer Security Incident Severity Levels*

Level	Description
High	Incidents that have severe impact on operations
Medium	Incidents that have a significant impact, or the potential to have a severe impact, on operations
Low	Incidents that have a minimal impact with the potential for significant or severe impact on operations

Once you categorize an incident and apply a severity level, you can decide how to resolve the incident. Let's review this stage of the incident response timeline.

Incident Resolution

The lifecycle of an incident should eventually lead to some form of resolution. This may include data analysis, resolution research, a proposed or performed action, and recovery. The objective of this phase is to discover the root cause of the incident, while working on containing the incident at the earliest stage possible.

During the analysis phase, the SOC team and the other teams should collaborate to achieve the quickest and best form of resolution. Access to some systems might be required to conduct necessary investigation activities. Figure 1-12 shows a sample process for handling the analysis tasks.

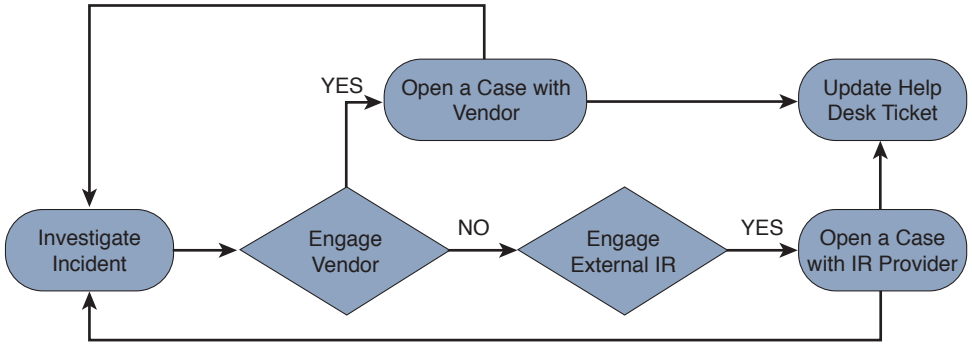

Figure 1-12 *Incident Analysis Phase Process*

The analysis and investigation phase involves the activities undertaken by SOC and by other teams for the purpose of

- Identifying compromised systems and accounts
- Understanding the impact of the computer security incident
- Identifying unauthorized access attempts to confidential data
- Understanding the chain of events that have led to the computer security incident

The containment phase involves the actions performed to quickly stop a computer security incident from escalating or spreading to other systems (that is, minimizing the potential damage caused by a computer security incident). Note that depending on the computer security incident and its impact, the containment phase can happen before, during, or after the analysis phase.

The process shown in Figure 1-13 is a sample containment process. You may, however, proceed with containment before or during the incident analysis.

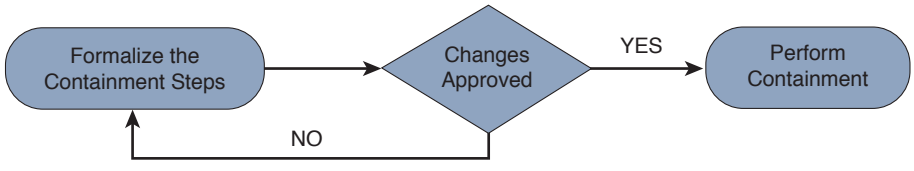

Figure 1-13 *Incident Containment Phase Process*

The exact steps to follow to contain a computer security incident vary depending on the nature of the incident and business criticality of the asset. Examples of containment actions include the following:

- Disconnecting a system from the network
- Moving an infected system to the quarantine network
- Stopping a service or a process
- Disabling an account
- Adding a firewall rule
- Adding an intrusion prevention system (IPS) signature/rule that would detect and block the attack's specific vector

Incident Closure

Closing a computer security incident refers to the eradication phase in which vulnerabilities that lead to the incident have been closed and all the incident traces have been

cleansed. The closure process also includes testing systems to ensure that the eradication steps and controls were effective and that the vectors used by the attack do not exist anymore or are ineffective. Predefined actions to consider include applying any final information about the event, its final classification, any external notifications, and archiving information about the incident.

If the incident involved violating regulatory requirements or resulted in the infringement of law, the SOC might be obliged to notify external entities. The SOC should have processes in place so that this communication is timely, accurate, and contains all information required for such notification.

The exact steps to follow when eradicating a computer security incident will vary depending on the nature of the incident. Examples of eradication actions include the following:

- System reconfiguration
- System re-imaging or rebuild
- Patching of systems
- Software update
- Deletion of accounts
- Deletion of files

In many cases, the closure and containment phases are closely coupled.

Post-Incident

This is the "lessons-learned" phase in which you seek to improve the IR processes and reflect on other people, processes, and technology controls. Post-incident activities will vary depending on the severity of the computer security incident. Valuable knowledge gained from computer security incidents can be useful to prevent/mitigate future incidents in the form of proactive services such as enhancing security features of functions within defenses.

In general, the process of post-analysis starts after incident closure and mainly includes documenting proposals of improvements in linkage to the different areas of the SOC: people, processes, and technology. In addition, recommendations for improving the SOC capabilities might result from a scheduled or an ad hoc risk-assessment exercise.

In summary, following an incident response timeline, the SOC team would handle a number of critical tasks, from incident detection to closure. Each stage of an incident can have its own processes and involve other groups within the organization. It is important that the plan, design, and build phases for incident response be defined, documented, and sponsored by the right authorities within the organization. Incident response is all about timing, and the worst time to figure out responsibilities and incident-handling process is during an active attack.

Now that we have covered a general overview of the responsibilities of the SOC, let's take a look at the history of SOC generations to get a better idea of how SOCs have matured over time.

SOC Generations

Our understanding of SOC components and expected services has changed over time. This is a reflection to the adjustment in our perception of the criticality of information assurance and security operations. This transformation comes in response to the ever-changing security threat landscape, in addition to our increasingly adopting formal information security standards, requiring the establishment and management of a formal security operations model and review processes.

The SOC's journey for the past 15 years can be broken to four *incremental* generations, shown in Figure 1-14. Ideally, an organization that uses technologies from the fourth generation, such as big data security analytics, should have adopted most of the SOC services from the previous generations. This might not always be the case in practice, though.

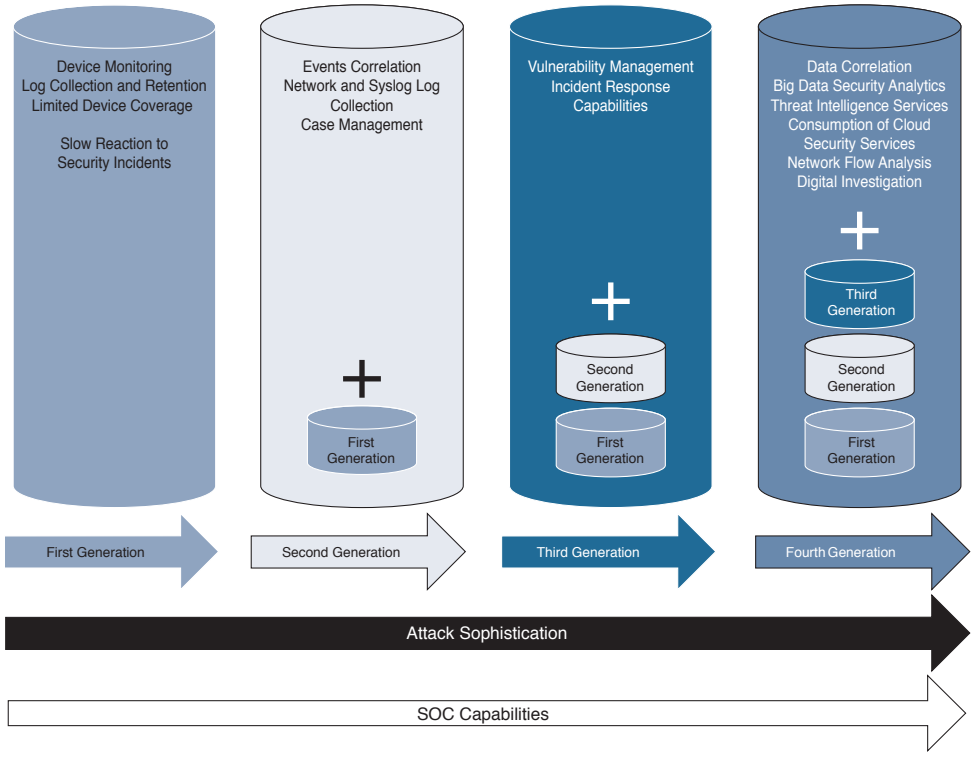

Figure 1-14 *The SOC: Four Generations*

The four generations reflect SOC capabilities in response to increasingly sophisticated attacks. Refer to the description of the generations to identify which SOC generation service your organization might be offering. Let's look at the details of the services offered in each SOC generation.

First-Generation SOC

In this generation, the wider IT operations team delivered what would be considered as SOC functions and services. This team was not necessarily skilled or trained to handle information security events and incidents. Security operations were not delivered by the establishment of a formal SOC, but in many cases by an IT operations individual or a team who focused on a blend of tasks. This could be responsibility for device and network health monitoring, managing antivirus security across the organization, and log collection. Log collection for the first-generation SOC was limited in the number of sources and types of devices capable of producing logs, such as firewalls. In many cases, storing logging messages was done locally. In other cases, a central logging facility was provisioned to receive log information, mainly in the form of unencrypted syslog or Simple Management Network Protocol (SNMP) messages.

In this generation, logging messages were rarely proactively analyzed, and were instead referred to if an incident was reported or some sort of troubleshooting was required. In addition, the concept of information security incident response was not formally established or appreciated. The process of identifying, communicating, and reacting to potential information security incidents was generally slow (and in many cases, ad hoc).

Let's look at an example of how a first-generation SOC would operate. Consider, for instance, that a number of systems have reported a substantial and relatively abnormal number of failed login attempts for the Microsoft Windows Active Directory domain administrator account within what is considered to be a short period of time. No evidence indicates suspicious activity or that a system compromise has been reported.

In the first generation of SOC, logging messages would most likely be locally saved on each system rather than stored on a centralized collection system such as a security information and event management (SIEM). The unsuccessful login attempt events would be saved to the local Windows security log store and buried under a large number of events generated by other various activities. Unless the Microsoft Active Directory system administrator manually accessed and analyzed the logs, the events in this example would have likely gone unnoticed, overlooking what could potentially be an account compromise and leading to what could be considered a major security incident.

Second-Generation SOC

This is the generation in which SIEM tools started to emerge. Early generations of SIEM providers such as netForensics, Network Intelligence (later acquired by EMC), and Cisco Security Monitoring, Analysis, and Response System (MARS) promised to detect network threats, releasing administrators from the complex and in many cases impossible task of manually analyzing huge amounts of log information. The early providers