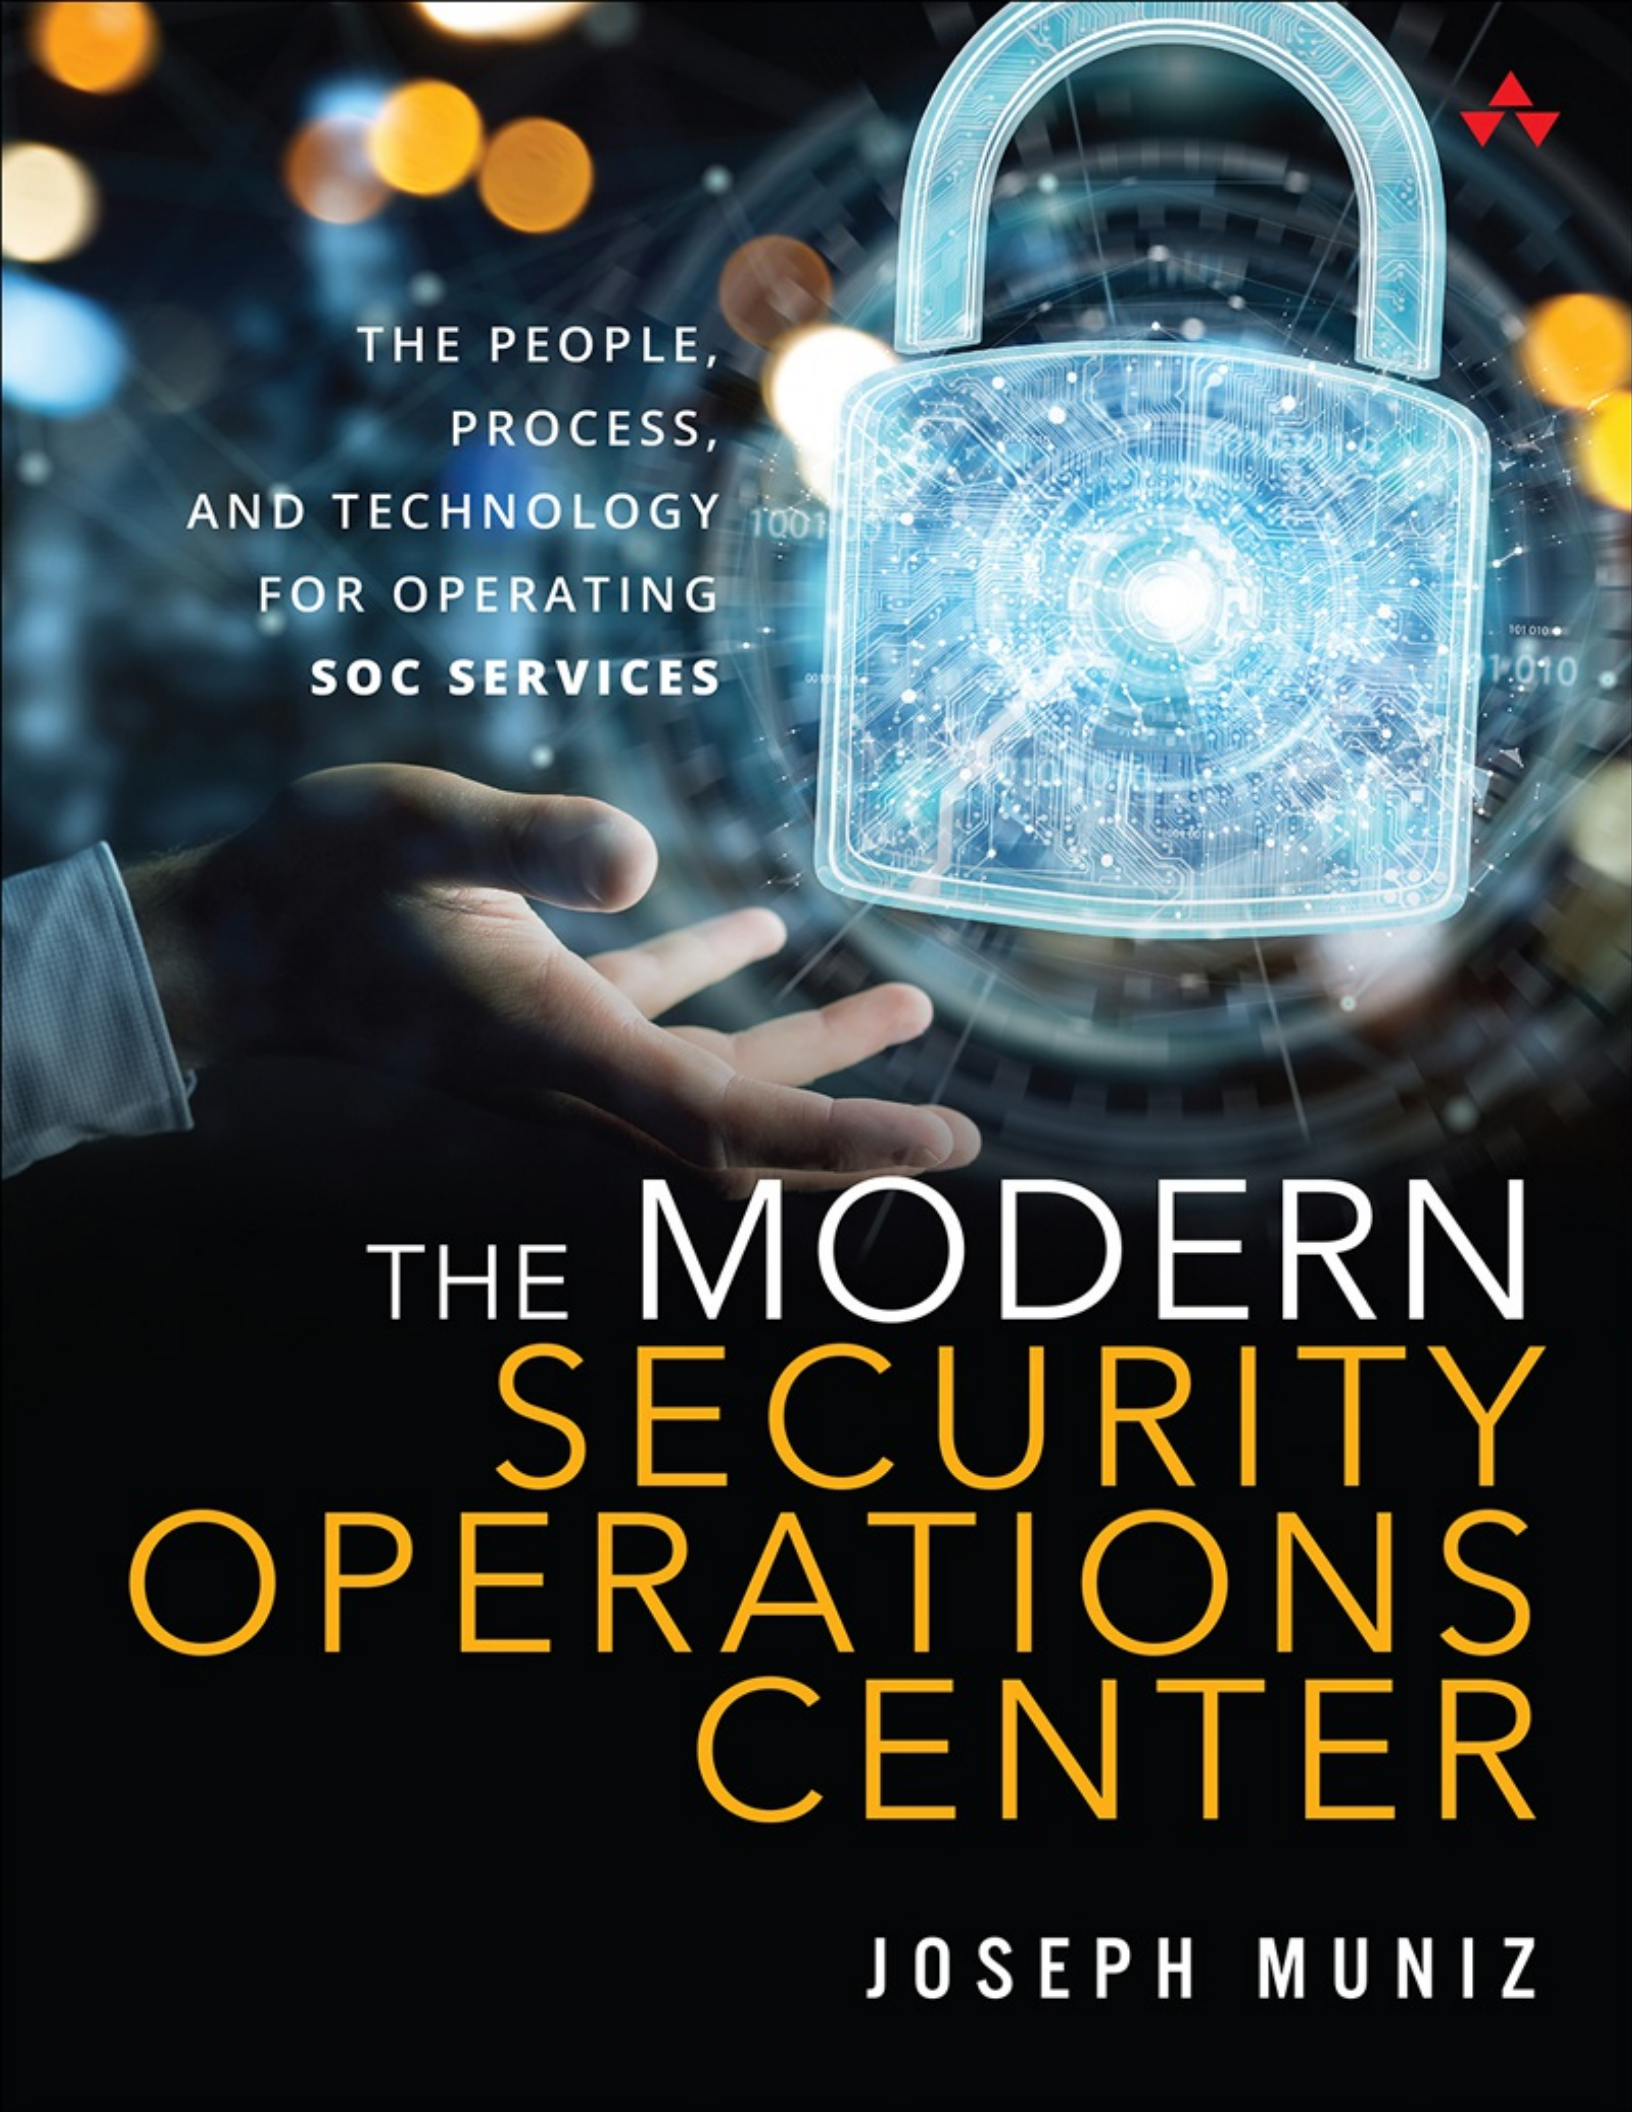

THE MODERN SECURITY OPERATIONS CENTER

JOSEPH MUNIZ

About This eBook

ePUB is an open, industry-standard format for eBooks. However, support of ePUB and its many features varies across reading devices and applications. Use your device or app settings to customize the presentation to your liking. Settings that you can customize often include font, font size, single or double column, landscape or portrait mode, and figures that you can click or tap to enlarge. For additional information about the settings and features on your reading device or app, visit the device manufacturer's Web site.

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a "Click here to view code image" link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

The Modern Security Operations Center

The Modern Security Operations Center

The People, Process, and Technology for Operating SOC Services

Joseph Muniz

Addison-Wesley

Boston • Columbus • New York • San Francisco • Amsterdam • Cape Town
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sidney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

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

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

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 intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2021930517

Copyright © 2021 Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions/.

ISBN-13: 978-0-13-561985-8
ISBN-10: 0-13-561985-8

ScoutAutomatedPrintCode

Editor-in-Chief
Mark Taub

Acquisitions Editor
James Manly

Development Editor
Christopher A. Cleveland

Managing Editor
Sandra Schroeder

Senior Project Editor
Tonya Simpson

Copy Editor
Bill McManus

Indexer
Ken Johnson

Proofreader
Betty Pessagno

Technical Reviewers
Anthony Giandomecino

Kevin Tigges
Willow Young

Editorial Assistant
Cindy Teeters

Cover Designer
Chuti Prasertsith

Compositor
codeMantra

Dedication

I would like to dedicate this book to two people. First, I want to dedicate it to Atticus Muniz, who can't walk, can't read, can't even understand how to use the toilet, but one day all of this will come. He is one year old and growing. Hopefully he will accomplish something great and while doing so make time to read this book. Second, I want to dedicate this book to Raylin Muniz, who is 11 and one of the most aggressive bookworms I've ever met. Hopefully she also will add this book to her reading list.

Table of Contents

Preface

Chapter 1: Introducing Security Operations and the SOC

Introducing the SOC

Factors Leading to a Dysfunctional SOC

Cyberthreats

Investing in Security

The Impact of a Breach

Establishing a Baseline

The Impact of Change

Fundamental Security Capabilities

Signature Detection

Behavior Detection

Anomaly Detection

Best of Breed vs. Defense in Depth

Standards, Guidelines, and Frameworks

NIST Cybersecurity Framework

ISO 3100:2018

FIRST Service Frameworks

Applying Frameworks

Industry Threat Models

The Cyber Kill Chain Model

The Diamond Model

MITRE ATT&CK Model

Choosing a Threat Model

Vulnerabilities and Risk

Endless Vulnerabilities
Business Challenges
In-House vs. Outsourcing
Services Advantages
Services Disadvantages
Hybrid Services
SOC Services
SOC Maturity Models
SOC Maturity Assessment
SOC Program Maturity
SOC Goals Assessment
Defining Goals
SOC Goals Ranking
Threats Ranking
SOC Goals Assessment Summarized
SOC Capabilities Assessment
Capability Maps
SOC Capabilities Gaps Analysis
Capability Map Next Steps
SOC Development Milestones
Summary
References

Chapter 2: Developing a Security Operations Center

Mission Statement and Scope Statement
Developing Mission and Scope Statements
SOC Scope Statement
Developing a SOC
SOC Procedures
Designing Procedures

Security Tools
- Evaluating Vulnerabilities
- Preventive Technologies
- Detection Technologies
- Mobile Device Security Concerns

Planning a SOC
- Capacity Planning
- Developing a Capacity Plan

Designing a SOC Facility
- Physical SOC vs. Virtual SOC
- SOC Location
- SOC Interior
- SOC Rooms
- SOC Computer Rooms
- SOC Layouts

Network Considerations
- Segmentation
- Logical Segmentation
- Choosing Segmentation
- Client/Server Segmentation
- Active Directory Segmentation
- Throughput
- Connectivity and Redundancy

Disaster Recovery

Security Considerations
- Policy and Compliance
- Network Access Control
- Encryption

Internal Security Tools
- Intrusion Detection and Prevention

Network Flow and Capturing Packets
Change Management
Host Systems
Guidelines and Recommendations for Securing Your SOC Network
Tool Collaboration
SOC Tools
Reporting and Dashboards
Throughput and Storage
Centralized Data Management
Summary
References

Chapter 3: SOC Services

Fundamental SOC Services
SOC Challenges
The Three Pillars of Foundational SOC Support Services
Pillar 1: Work Environment
Pillar 2: People
Pillar 3: Technology
Evaluating the Three Pillars of Foundational SOC Support Services
SOC Service Areas
FIRST's CSIRT
Developing SOC Service Areas
In-House Services vs. External Services
Contracted vs. Employee Job Roles
SOC Service Job Goals
Resource Planning
Service Maturity: If You Build It, They Will Come
SOC Service 1: Risk Management

Four Responses to Risk
Reducing Risk
Addressing Risk
SOC Service 2: Vulnerability Management
Vulnerability Management Best Practice
Vulnerability Scanning Tools
Penetration Testing
SOC Service 3: Compliance
Meeting Compliance with Audits
SOC Service 4: Incident Management
NIST Special Publication 800-61 Revision 2
Incident Response Planning
Incident Impact
Playbooks
SOC Service 5: Analysis
Static Analysis
Dynamic Analysis
SOC Service 6: Digital Forensics
SOC Service 7: Situational and Security Awareness
User Training
SOC Service 8: Research and Development
Summary
References

Chapter 4: People and Process

Career vs. Job
Developing Job Roles
General Schedule Pay Scale
IT Industry Job Roles
Common IT Job Roles

SOC Job Roles
- Security Analyst
- Penetration Tester
- Assessment Officer
- Incident Responder
- Systems Analyst
- Security Administrator
- Security Engineer
- Security Trainer
- Security Architect
- Cryptographer/Cryptologist
- Forensic Engineer
- Chief Information Security Officer

NICE Cybersecurity Workforce Framework
- Nice Framework Components

Role Tiers

SOC Services and Associated Job Roles
- Risk Management Service
- Vulnerability Management Service
- Incident Management Service
- Analysis Service
- Compliance Service
- Digital Forensics Service
- Situational and Security Awareness Service
- Research and Development Service

Soft Skills
- Evaluating Soft Skills
- SOC Soft Skills

Security Clearance Requirements

Pre-Interviewing

Interviewing
Interview Prompter
Post Interview
Onboarding Employees
Onboarding Requirements
Managing People
Job Retention
Training
Training Methods
Certifications
Company Culture
Summary
References

Chapter 5: Centralizing Data

Data in the SOC
Strategic and Tactical Data
Data Structure
Data Types
Data Context
Data-Focused Assessment
Data Assessment Example: Antivirus
Threat Mapping Data
Applying Data Assessments to SOC Services
Logs
Log Types
Log Formats
Security Information and Event Management
SIEM Data Processing
Data Correlation

Data Enrichment
SIEM Solution Planning
SIEM Tuning
Troubleshooting SIEM Logging
SIEM Troubleshooting Part 1: Data Input
SIEM Troubleshooting Part 2: Data Processing and Validation
SIEM Troubleshooting Examples
Additional SIEM Features
APIs
Leveraging APIs
API Architectures
API Examples
Big Data
Hadoop
Big Data Threat Feeds
Machine Learning
Machine Learning in Cybersecurity
Artificial Intelligence
Machine Learning Models
Summary
References

Chapter 6: Reducing Risk and Exceeding Compliance

Why Exceeding Compliance
Policies
Policy Overview
Policy Purpose
Policy Scope
Policy Statement
Policy Compliance

Related Standards, Policies, Guidelines, and Processes
Definitions and Terms
History
Launching a New Policy
Steps for Launching a New Policy
Policy Enforcement
Certification and Accreditation
Procedures
Procedure Document
Tabletop Exercise
Tabletop Exercise Options
Tabletop Exercise Execution
Tabletop Exercise Format
Tabletop Exercise Template Example
Standards, Guidelines, and Frameworks
NIST Cybersecurity Framework
ISO/IEC 27005
CIS Controls
ISACA COBIT 2019
FIRST CSIRT Services Framework
Exceeding Compliance
Audits
Audit Example
Internal Audits
External Auditors
Audit Tools
Assessments
Assessment Types
Assessment Results
Assessment Template

Vulnerability Scanners
Assessment Program Weaknesses
Penetration Test
NIST Special Publication 800-115
Additional NIST SP 800-115 Guidance
Penetration Testing Types
Penetration Testing Planning
Industry Compliance
Compliance Requirements
Summary
References

Chapter 7: Threat Intelligence

Threat Intelligence Overview
Threat Data
Threat Intelligence Categories
Strategic Threat Intelligence
Tactical Threat Intelligence
Operational Threat Intelligence
Technical Threat Intelligence
Threat Intelligence Context
Threat Context
Evaluating Threat Intelligence
Threat Intelligence Checklist
Content Quality
Testing Threat Intelligence
Planning a Threat Intelligence Project
Data Expectations for Strategic Threat Intelligence
Data Expectations for Tactical Threat Intelligence
Data Expectations for Operational Threat Intelligence

Data Expectations for Technical Threat Intelligence
Collecting and Processing Intelligence
Processing Nontechnical Data
Operational Data and Web Processing
Technical Processing
Technical Threat Intelligence Resources
Actionable Intelligence
Security Tools and Threat Intelligence
Feedback
Summary
References

Chapter 8: Threat Hunting and Incident Response

Security Incidents
Incident Response Lifecycle
Phase 1: Preparation
Assigning Tasks with Playbooks
Communication
Third-Party Interaction
Law Enforcement
Law Enforcement Risk
Ticketing Systems
Other Incident Response Planning Templates
Phase 1: Preparation Summary
Phase 2: Detection and Analysis
Incident Detection
Core Security Capabilities
Threat Analysis
Detecting Malware Behavior
Infected Systems

Analyzing Artifacts
Identifying Artifact Types
Packing Files
Basic Static Analysis
Advanced Static Analysis
Dynamic Analysis
Phase 2: Detection and Analysis Summary
Phase 3: Containment, Eradication, and Recovery
Containment
Responding to Malware
Threat Hunting Techniques
Eradicate
Recovery
Digital Forensics
Digital Forensic Process
First Responder
Chain of Custody
Working with Evidence
Duplicating Evidence
Hashes
Forensic Static Analysis
Recovering Data
Forensic Dynamic Analysis
Digital Forensics Summary
Phase 3: Containment, Eradication, and Recovery Summary
Phase 4: Post-Incident Activity
Post-Incident Response Process
Phase 4: Post-Incident Response Summary
Incident Response Guidelines
FIRST Services Frameworks

Summary
References

Chapter 9: Vulnerability Management

Vulnerability Management
- Phase 1: Asset Inventory
- Phase 2: Information Management
- Phase 3: Risk Assessment
- Phase 4: Vulnerability Assessment
- Phase 5: Report and Remediate
- Phase 6: Respond and Repeat

Measuring Vulnerabilities
- Common Vulnerabilities and Exposures
- Common Vulnerability Scoring System
- CVSS Standards

Vulnerability Technology
- Vulnerability Scanners
- Currency and Coverage
- Tuning Vulnerability Scanners
- Exploitation Tools
- Asset Management and Compliance Tools
- Network Scanners and Network Access Control
- Threat Detection Tools

Vulnerability Management Service
- Scanning Services
- Vulnerability Management Service Roles
- Vulnerability Evaluation Procedures

Vulnerability Response
- Vulnerability Accuracy
- Responding to Vulnerabilities

Cyber Insurance
Patching Systems
Residual Risk
Remediation Approval
Reporting
Exceptions
Vulnerability Management Process Summarized
Summary
References

Chapter 10: Data Orchestration

Introduction to Data Orchestration
Comparing SIEM and SOAR
The Rise of XDR
Security Orchestration, Automation, and Response
SOAR Example: Phantom
Endpoint Detection and Response
EDR Example: CrowdStrike
Playbooks
Playbook Components
Constructing Playbooks
Incident Response Consortium
Playbook Examples: Malware Outbreak
Automation
Automating Playbooks
Common Targets for Automation
Automation Pitfalls
Playbook Workflow
DevOps Programming
Data Management

Text-File Formats
Common Data Formats
Data Modeling
DevOps Tools
DevOps Targets
Manual DevOps
Automated DevOps
DevOps Lab Using Ansible
Ansible Playbooks
Blueprinting with Osquery
Running Osquery
Network Programmability
Learning NetDevOps
APIs
NetDevOps Example
Cloud Programmability
Orchestration in the Cloud
Amazon DevOps
SaaS DevOps
Summary
References

Chapter 11: Future of the SOC

All Eyes on SD-WAN and SASE
VoIP Adoption As Prologue to SD-WAN Adoption
Introduction of SD-WAN
Challenges with the Traditional WAN
SD-WAN to the Rescue
SASE Solves SD-WAN Problems
SASE Defined

Future of SASE
IT Services Provided by the SOC
IT Operations Defined
Hacking IT Services
IT Services Evolving
Future of IT Services
Future of Training
Training Challenges
Training Today
Case Study: Training I Use Today
Free Training
Gamifying Learning
On-Demand and Personalized Learning
Future of Training
Full Automation with Machine Learning
Machine Learning
Machine Learning Hurdles
Machine Learning Applied
Training Machine Learning
Future of Machine Learning
Future of *Your* SOC: Bringing It All Together
Your Future Facilities and Capabilities
Group Tags
Your Future SOC Staff
Audits, Assessments, and Penetration Testing
Future Impact to Your Services
Hunting for Tomorrow's Threats
Summary
References

Index

Preface

Defending your organization from cyberthreats is a cat and mouse game. Both sides are constantly changing their tactics. When the defense tools work, the adversaries acquire the defense technology, reverse engineer it, and develop strategies to bypass it. When the adversaries start to succeed at bypassing security tools, defense companies take note, research the attack being used, and adjust defense capabilities in their tools to prevent future successful exploitation. Somewhere in between all of this back and forth is your organization.

Security is about the combination of people, process, and technology working together to accomplish a goal. You don't just buy a few products, plug them in, and magically eliminate the risk of being exploited. Security is a journey, which you must continue to invest in. It is not a destination. You don't one day become secure and be done with it. You can't buy your way to being secure. It requires an investment in a team responsible for security, commonly referred to as the security operations center (SOC).

Vision

My purpose for writing this book is to help every organization regardless of size, budget, or mission understand how to turn those responsible for the security of their organization into a security operations center. I do believe security is the responsibility of everybody in the organization, but one or more people need to have security as their primary job, and they need to be recognized for that role.

In this book, I describe how to build security services to support your organization. Some organizations run their business from the cloud. Other organizations do not. Some organizations have a budget to build a new SOC, while others need to convert what they have into a SOC that can support the organization now. Wherever you are at in your security journey, I have

designed this book to incorporate industry guidelines, popular frameworks, and my own personal experience to give you an overview of how mature SOCs around the world run their security practice. I believe any organization can run a mature SOC as long as the organization recognizes its security team and what they do as a formal SOC.

My vision for this book is to take a vendor-agnostic approach to security with a focus on capabilities and best practices that will prepare you for the threats of tomorrow. I include tons of open-source and commercial product examples, but I always focus on the outcome of the recommendation so the vendor of choice won't matter. I reference specific guidelines to validate my recommendations and explain the risk of not performing what is covered in this book. I believe security professionals of all levels of experience can benefit from this book and I hope this book becomes a valuable asset in your journey against cyberthreats.

Who Should Read This Book?

I believe anybody with an interest in cybersecurity will benefit from this book. I explain concepts using different viewpoints ranging from what leadership expects to those behind the keyboard care about. Topics include building a SOC, risk management, vulnerability management, incident management, analysis of malware, compliance, digital forensics, situational and security awareness, and research and development. All of these topics correspond to services that are provided by mature SOCs around the world. Anybody who is interested in learning how to build these services into their security practice will benefit from this book.

How This Book Is Organized

This book can be used the day you start planning to build a SOC and can act as a resource to help mature an existing SOC. Chapter 1 starts with a general overview of all high-level SOC concepts. Chapter 2 focuses on how to build a SOC, including aligning the SOC to the business, mission statements,

scope, and everything that should be considered as you plan a SOC.

Chapter 3 introduces the fundamental SOC services I find in mature SOCs around the world. I work through each service in the remaining chapters, including how to deal with risk, vulnerabilities, compliance, and other challenges organizations rely on the SOC to handle. Chapter 8 provides different approaches to some topics, such as when to launch an incident response versus a forensic investigation. Chapters 2 and 11 cover technologies that are common today as well as technologies that are futuristic but will eventually become part of the average SOC. An example is how cloud technologies such as Secure Access Service Edge (SASE) will eventually become as common as Voice over IP (VoIP) has become in most businesses.

Throughout this book I include examples of tools and techniques used by both red teams and blue teams, meaning I show how to execute real-world exploitation as well as how SOCs around the world defend against modern attacks. Many of the tools are open source, including using Kali Linux for exploitation, Ansible for automation, and NIST publications (among others) as guidelines. When the topic requires referencing enterprise tools, I try to bounce between vendors to give you a general feel of the capabilities rather than the specifics of how a particular vendor's tool functions. My goal is to keep a vendor-agnostic approach to security, which is why I include examples from many different vendors and open-source options.

Chapter 11 concludes this book by making predictions about the future of the SOC. My predictions are based on industry trends, conversations with customers, and personal experience in the industry for 20+ years. I believe many of the topics in this book are fundamental security concepts and will be relevant for many years after this book's publication. I wrote this to prepare you for the threats of tomorrow regardless of how they look. I believe this book has something for every organization to benefit from, and hope it helps you in your journey to building and running a successful security operations center.

Book Structure

I have organized this book from general SOC concepts to detailed SOC services. The following is a short summary of each chapter and how it relates to building and maintaining a mature SOC.

- **Chapter 1, "Introducing Security Operations and the SOC":** This chapter introduces high-level SOC concepts. I provide ways for you to validate which capabilities you currently have as well as how to assess your existing processes so you can plan where you can improve your SOC, if one already exists. The purpose of this chapter is to serve as a primer for the remaining chapters and help you establish your current state of security so you can use this book to develop a mature SOC.
- **Chapter 2, "Developing a Security Operations Center":** Chapter 2 focuses on the fundamental business and operational requirements that need to be in place before your SOC can provide service. Topics include who should sponsor, manage, and support the SOC, what type of policies and procedures need to be developed, and other business objectives that are essential prerequisites for your SOC goes live. The second half of this chapter focuses on operational requirements, such as how to plan the SOC workspace, how to accommodate SOC team members with different responsibilities, and what type of technology needs to be considered depending on what services your SOC plans to offer. Addressing the topics in this chapter is essential before your SOC can provide any value to the organization.
- **Chapter 3, "SOC Services":** This chapter introduces many of the topics that are explored in depth in the subsequent chapters of the book. I introduce the fundamental SOC services that are common in organizations around the world. I cover how these services can be delivered by the SOC and everything you need to consider as you look to stand up a new SOC service. This includes when you should outsource a service versus when it makes sense to develop the service using in-house capabilities. Chapter 3 represents the point at which your SOC is moving to a go-live state and starting to provide value to the organization.
- **Chapter 4, "People and Process":** Every SOC service requires the right people and processes to be successful. This chapter introduces all of the different job roles that are common in mature SOCs around the world. It describes skill requirements for each of the roles as well as expectations

for daily duties. I cover how to find the right people for your SOC and groom them using different programs that tie directly back to the SOC's service success. Topics include job roles, recruiting, interviewing, onboarding, and outsourcing people and process.

- **Chapter 5, "Centralizing Data":** One fundamental SOC capability is being able to work with both the organization's data and external security data such as threat intelligence. Many new SOCs start off by offering log management and analysis services. This means the SOC is responsible for collecting logs from various tools, analyzing the logs, and providing a response when certain events are identified. Centralizing data is not an easy task as data comes in many formats and more advanced uses of data require different types of programming and automation skillsets. I cover everything related to how to collect and use data, which is a critical steppingstone to many of the other services offered in mature SOCs around the world.
- **Chapter 6, "Reducing Risk and Exceeding Compliance":** Dealing with risk and ensuring compliance are common requirements for many SOCs. The types of risk will vary between organizations, but in general, a SOC is responsible for reducing risk. Compliance can be mandated by local or federal government, mandated by service providers such as credit card companies, or mandated by an organization's leadership. In addition to these topics, this chapter covers some peripheral topics because everything security related is a form of risk management and has some form of compliance element. For example, managing vulnerabilities can be part of a risk management program as well as a requirement to be compliant with some regulation or policy. I believe risk management and compliance are the most critical of the services your SOC can provide.
- **Chapter 7, "Threat Intelligence":** Building on Chapter 5, which stresses the importance of collecting and using data for SOC services, this chapter focuses on the critical data source of threat intelligence. I believe threat intelligence represents the future of all security technologies and thus have dedicated this chapter to the topic. Data is becoming too massive to manually review, and concepts such as baselining normal behavior or comparing things against what others are seeing is where the security industry is investing all developments. If your SOC is not leveraging threat intelligence today, it soon will be. If you are using threat

intelligence, you will find in this chapter that there are many different ways to use threat intelligence, including looking beyond obtaining lists of things that are considered a high risk.

- **Chapter 8, "Threat Hunting and Incident Response":** One core service many organizations expect from the SOC is responding to incidents. This chapter covers how to develop a robust incident response service. Topics include how to plan a response based on the incident, how to contain, eradicate, and recover from an incident, when to use digital forensics, and what post-incident response activities should occur before you close out a case. Incident response is made up of many different services ranging from skillsets in analyzing malware to how to properly investigate an artifact without modifying it if legal action could occur. I include many tools and techniques so you can build a lab and eventually go live with a proper incident response capability.
- **Chapter 9, "Vulnerability Management":** Any SOC that is responsible for incident response should also include services for incident recovery and vulnerability management. By doing this, the SOC is able to reduce the risk of future security events based on the principle that an attack needs a vulnerability in order to succeed. Removing the vulnerability means the risk is reduced. This chapter focuses on vulnerability management services both from a proactive standpoint and a reactive standpoint. I believe if your SOC dedicates time to these topics, your organization will experience less attack behavior, leading to more time to focus on proactive services versus continuously reacting to security events.
- **Chapter 10, "Data Orchestration":** Many SOCs around the world have great data repositories and services but are finding their staff is being overwhelmed with tedious work. As a response to this problem, many organizations are investing in automation and orchestration with the goal of reducing mundane tasks and establishing a formalized and repeatable response to how they deliver SOC services. This chapter focuses on these topics, taking many of the concepts from previous chapters and looking at ways to apply orchestration and automation. Topics include tools, techniques, and programming, including an introduction to DevOps.
- **Chapter 11, "Future of the SOC":** This final chapter forecasts the future

of the security operations center. I present a few industry trends that I believe will change the SOC of the future and explain how they look today as well as predict what they will look like in the future. Topics include Secure Access Service Edge (SASE), software-defined wide-area network (SD-WAN) technologies, general cloud trends, IT services, training, and the future of automation. I close this chapter and book with a focus on the future of your own SOC, a synopsis of how to take everything covered in this book and apply it to your SOC's journey to maturity and success.

We Want to Hear from You!

As the reader of this book, *you* are our most important critic and commentator. We value your opinion and want to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our way.

We welcome your comments. You can email or write to let us know what you did or didn't like about this book—as well as what we can do to make our books better.

Please note that we cannot help you with technical problems related to the topic of this book.

When you write, please be sure to include this book's title and author as well as your name and email address. We will carefully review your comments and share them with the author and editors who worked on the book.

Email: community@informit.com

Reader Services

Register your copy of *The Modern Security Operations Center* at informit.com/aw for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account*. Enter the product ISBN 9780135619858 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Acknowledgments

I had a ton of help validating the content in this book and would like to recognize the reviewers, Anthony Giandomecino, Kevin Tigges, and Willow Young, for their hard work and valuable feedback. This book was a two-year journey to develop and these three kept me honest and the content on point. You will find this fine group of engineers have all spoken at industry events, have published various types of security content, and contribute to the security industry in many different ways. Hats off to them for helping me make this book project happen.

I would also like to thank James Manly and the rest of the Pearson team for their support during this long two-year journey from concept to publication. They are a very professional group and a pleasure to work with.

I would like to give a huge “thank you” to my friends and family for supporting me throughout this and my other projects. Also “thank you” to my managers at work and real boss, Anjelica Ruda, for allowing me to have time to work on this book in between all of the madness.

About the Author

Joseph Muniz is an architect and security researcher in the Cisco Security Sales and Engineering Organization. He is driven by making the world a safer place through education and adversary research. Joseph has extensive experience in designing security solutions and architectures as a trusted advisor for top Fortune 500 corporations and the U.S. government.

Joseph is a researcher and industry thought leader. He speaks regularly at international conferences, writes for technical magazines, and is involved with developing training for various industry certifications. He invented the fictitious character of Emily Williams to create awareness around social engineering. Joseph runs The Security Blogger website, a popular resource for security and product implementation. He is the author and contributor of several publications including titles ranging from security best practices to

exploitation tactics.

When Joseph is not using technology, you can find him on the fútbol (soccer) field or raising the next generation of hackers, also known as his children. Follow Joseph at https://www.thesecurityblogger.com and @SecureBlogger.

Figure Credits

Cover image Sdecoret/Shutterstock

Figure 1-1, video posted by Anonymous on YouTube © 2020 Scripps Media, Inc

Figure 1-18, part of the MITRE ATT&CK matrix for Enterprise © 2015–2021, The MITRE Corporation

Figure 1-19, chaining together attack behavior using ATT&CK modeling © 2015–2021, The MITRE Corporation

Figure 2-5, Rapid7's Nexpose Vulnerability Scanner © Rapid7

Figure 2-6, Cisco Firepower passive vulnerability data © Cisco Systems

Figure 2-7, Cisco reputation block page © Cisco Systems

Figure 2-8, Google's reputation warning banner ©2020 Google

Figure 2-13, raised floor tile courtesy of Alibaba.com

Figure 2-14, sample SOC layout courtesy of Cisco Systems

Figure 2-26, Figure 5-16, SPLUNK Dashboard example © 2005–2020 Splunk, Inc

Figure 2-27, QRadar Dashboard example © IBM Corporation 1994, 2020

Figure 3-8, Tenable.sc vulnerability tracking © 2020 Tenable, Inc.

Figure 3-9, OpenVAS GUI example © Greenbone Networks 2020

Figure 3-11, MITRE ATT&CK Framework © 2015–2020, The MITRE Corporation

Figure 3-12, Atomic Red Ream website © 2014–2020 Red Canary

Figure 3-13, Atomic Red Team example for Windows © 2015–2020, The MITRE Corporation

Figure 3-14, Kali Linux Tool Categories © OffSec Services Limited 2020

Figure 3-15, searching Metasploit for Adobe vulnerabilities © Rapid7

Figure 3-20, Incident Response Consortium Playbooks © 2019 Incident Response Consortium

Figure 3-21, Malware Outbreak Playbook © 2019 Incident Response Consortium

Figure 3-22, diagram about hidden extensions © Microsoft 2020

Figure 3-24, Peframe analyzing a packed file © Microsoft 2020

Figure 3-26, using SET to clone Gmail © 2020 by TrustedSec

Figure 3-27, cloned Gmail website © 2020 Google

Figure 5-2, poorly parsed log within Splunk © 2005–2020 Splunk, Inc

Figure 5-3, Windows Event Log © Microsoft 2020

Figure 5-5, Splunk customized dashboard example © 2005–2020 Splunk, Inc

Figure 5-7, Splunk data input options © 2005–2020 Splunk, Inc

Figure 5-8, Cisco Stealthwatch configured to syslog data to Splunk © Cisco Systems

Figure 5-9, using default parsing template within SMC © Cisco Systems

Figure 5-10, results from poorly formatted Syslog © 2005–2020 Splunk, Inc

Figure 5-11, using custom syslog template in SMC © Cisco Systems

Figure 5-12, converting syslog data into reports and widgets © Cisco Systems

Figure 5-13, IBM QRadar search screen © IBM Corporation 1994, 2020

Figure 5-14, searching for data exfiltration in IBM QRadar © IBM Corporation 1994, 2020

Figure 5-15, IBM QRadar dashboard example © IBM Corporation 1994, 2020

Figure 5-17, Stealthwatch application within Splunk © 2005–2020 Splunk, Inc

Figure 5-18, IBM QRadar asset and vulnerability management example © IBM Corporation 1994, 2020

Figure 5-21, Firepower leveraging Rapid7 vulnerability data © Cisco Systems

Figure 5-25, Cisco Umbrella big data results example © Cisco Systems

Figure 6-1, Incident Response Consortium Playbook options © 2019 Incident Response Consortium

Figure 6-7, various Apache Struts exploits in Metasploit © 2020 The Apache Software Foundation

Figure 6-8, using Armitage to exploit a struts vulnerability © 2020 The Apache Software Foundation

Figure 7-2, Sophos threat prevalence usage example © 1997–2020 Sophos Ltd

Figure 7-3, example of Splunk not correctly processing threat data © 2005–2020 Splunk, Inc

Figure 7-5, Google Alerts example © 2020 Google

Figure 7-6, Google Chrome Scraper collecting hashes from the Cisco Talos blog © 2020 Cisco Systems, Inc

Figure 7-7, Twitter threat data behavior examples © 2020 Trend Micro Incorporated

Figure 8-2, Incident Response Consortium malware outbreak playbook example © 2019 Incident Response Consortium

Figure 8-4, opening a new incident ticket in Cisco SecureX © 2020 Cisco Systems, Inc

Figure 8-5, example of new case © 2020 Cisco Systems, Inc

Figure 8-7, Cisco AMP indicators of compromise example © 2020 Cisco Systems, Inc

Figure 8-8, hidden extension example © Microsoft 2020

Figure 8-10, TrIDNET Free File Analysis tool © Marco Pontello

Figure 8-17, Ghidra Disassembler viewing WannaCry ransomware kill switch © Lazarus Group

Figure 8-18, example of using Joe Sandbox for malware analysis © Lazarus

Group

Figure 8-19, Cisco Stealthwatch identifying threats based on NetFlow © 2020 Cisco Systems, Inc

Figure 8-20, threat hunting maturity model © 2016–20 Sqrrl Fintech Private Limited

Figure 8-21, Cisco threat response analyzing success20.hopto.org © 2020 Cisco Systems, Inc

Figure 8-25, IRC Malware Outbreak playbook eradicate step © 2019 Incident Response Consortium

Figure 8-26, IRC Data Theft playbook eradicate step © 2019 Incident Response Consortium

Figure 8-27, IRC Malware Outbreak playbook recover step © 2019 Incident Response Consortium

Figure 8-29, Chain of Custody documentation bag example Courtesy of Cisco Systems

Figure 8-30, Autopsy main page © 2003–2020 Brian Carrier

Figure 8-31, Autopsy usage example © 2003–2020 Brian Carrier

Figure 8-40, Lessons Learned Meeting Agenda template example © Template.net

Figure 8-42, template for documenting parties involved © 2002–2020 Blackboard, Inc.

Figure 9-4, screenshot of Struts vulnerability example © 2020 Cisco Systems, Inc

Figure 9-5, screenshot of CVSS v2 base score calculator © National Institute of Standards and Technology

Figure 9-6, screenshot of CVSS temporal and environmental calculators © National Institute of Standards and Technology

Figure 9-7, screenshot of Struts CVSSv2 example © National Institute of Standards and Technology

Figure 9-8, screenshot of CVSS v3 base score metrics © National Institute of

Standards and Technology

Figure 9-9, screenshot of Struts CVSS v3 example © National Institute of Standards and Technology

Figure 9-10, screenshot of Struts CVE-2017-9793 resource example © National Institute of Standards and Technology

Figure 9-11, screenshot of Struts vulnerability shown in Rapid7's Nexpose © Rapid7

Figure 9-12, screenshot of Rapid7 Nexpose dashboard © Rapid7

Figure 9-13, screenshot of passive vulnerability scanning example © 2020 Cisco Systems, Inc

Figure 9-14, screenshot of OpenVAS example © blackMORE Ops

Figure 9-17, screenshot of Certero dashboard example © 2007–2020 Certero

Figure 9-18, screenshot of network access control asset list example © 2020 Cisco Systems, Inc

Figure 9-19, screenshot of Zenmap © nmap.org

Figure 9-20, screenshot of Cisco Firepower tuning with vulnerability data © 2020 Cisco Systems, Inc

Figure 9-25, screenshot of Rapid7 Nexpose automated actions configuration example © Rapid7

Figure 9-26, screenshot of Nexpose Asset dashboard © Rapid7

Figure 9-28, screenshot of Cisco Firepower Apache Struts rules © 2020 Cisco Systems, Inc

Figure 10-2, screenshot of Splunk Phantom main dashboard example © 2005–2020 Splunk, Inc

Figure 10-3, screenshot of Splunk Phantom case management dashboard example © 2005–2020 Splunk, Inc

Figure 10-4, screenshot of Splunk Phantom Playbook template list example © 2005–2020 Splunk, Inc

Figure 10-5, screenshot of high-level Splunk Phantom Playbook example © 2005–2020 Splunk, Inc

Figure 10-6, screenshot of zoomed-in Phantom Playbook example © 2005–2020 Splunk, Inc

Figure 10-7, screenshot of phantom example of DevOps coding © 2005–2020 Splunk, Inc

Figure 10-8, screenshot of CrowdStrike Falcon dashboard example © 2020 CrowdStrike

Figure 10-9, screenshot of Falcon event graph example © 2020 CrowdStrike

Figure 10-10, screenshot of CrowdStrike Falcon example of event details © 2020 CrowdStrike

Figure 10-13, screenshot of IRC's Prepare playbook for malware outbreak © 2019 Incident Response Consortium

Figure 10-14, screenshot of IRC's Analyze playbook for malware outbreak © 2019 Incident Response Consortium

Figure 10-16, screenshot of Cisco ISE configured with Rapid7 Nexpose example © 2020 Cisco Systems, Inc

Figure 10-17, screenshot of Cisco Firepower configuration rule example © 2020 Cisco Systems, Inc

Figure 10-18, screenshot of Cisco SecureX orchestration example © 2020 Cisco Systems, Inc

Figure 10-19, screenshot of example workflow validation and run options © 2020 Cisco Systems, Inc

Figure 10-20, screenshot of Splunk Phantom workflow execution example © 2005–2020 Splunk, Inc

Figure 10-24, screenshot of new installation of MediaWiki © 2020 Cisco Systems, Inc

Figure 10-31, screenshot of Cisco Engineering and Software certification programs © 2020 Cisco Systems, Inc

Figure 10-32, screenshot of Cisco DevNet Sandbox Lab catalog of free labs © 2020 Cisco Systems, Inc

Figure 10-33, screenshot of Postman dashboard © 2020 Postman, Inc

Figure 10-34, screenshot of configuring Postman to communicate with a Cisco router © 2020 Postman, Inc

Figure 10-35, screenshot of configuration pulled into Postman example © 2020 Postman, Inc

Figure 11-5, Cisco SD-WAN dashboard example © 2020 Cisco Systems, Inc

Black and white portrait of fortune-teller with crystal ball © Aniriana/Shutterstock

Figure 11-9, article about the “Emily Williams” penetration test © 2020 Reed Exhibitions Ltd

Fortune Teller with Crystal ball © Pete Saloutos/Shutterstock

Figure 11-10, lab guide converted to Moodle © 2020 Cisco Systems, Inc

Figure 11-11, Khan Academy dashboard © 2020 Khan Academy

Woman fortuneteller with crystal ball in darkness © Konstantin Shevtsov/123rf.com

Seer working over glowing crystal ball © Phil McDonald/Shutterstock

Figure 11-22, Cisco DevNet Sandbox catalog of free DevOps labs © 2020 Cisco Systems, Inc

Chapter 1

Introducing Security Operations and the SOC

The journey of a thousand miles begins with one step.

—Lao Tzu

Security is a simple concept: protect something from threats. Although this sounds easy, many organizations, from small government agencies to Fortune 500 businesses, do not know how to transform their current efforts into a formal security operations center (SOC). As a result, the security teams within these organizations have trouble obtaining the proper support and funding to improve their capabilities. Having static SOC capacities leads to failures in how the SOC functions because too much time is spent on reactive and manual efforts with no clear path for improvement of any SOC service. The combination of these challenges causes organizations to experience breaches of security, loss of talent, large fines, and other negative outcomes, possibly including the complete failure of the business.

In this chapter, you learn about fundamental security and SOC concepts. I cover why it is important to build a mature security operations center that combines people, processes, and technology. Security topics include how to develop a defense-in-depth security architecture using industry recommendations found within standards, guidelines, and frameworks. I show you how to better understand potential threats using threat models and vulnerability assessments. I introduce the eight fundamental services I find in mature SOCs around the world and show you how to assess yourself against those services. This chapter is the foundation for everything covered in this book.

Introducing the SOC

The security operations center, more commonly called “the SOC,” is a centralized unit that deals with security issues on both an organizational level and a technical level. This occurs through the use of people, process, and capabilities to deliver one or more services. Services could include identifying and reducing risk, addressing vulnerabilities, adhering to compliance requirements, responding to incidents, collecting forensic evidence, and performing other tasks deemed essential to the security posture of the organization. Which services any particular SOC offers depends on the nature of the business the SOC is protecting. Some SOCs might outsource services using on-demand or external service providers to fill the need for a capability. Other SOCs might accept the risk and ignore services or pass on responsibilities to other groups within the organization. An example is handing off vulnerability management to the desktop support team. Chapter 3, “SOC Services,” covers all of the services typically offered in mature SOCs around the world.

Every organization can have a SOC, regardless of the services and capabilities that it offers. The size or type of the business shouldn’t matter, because every organization has one or more business goals that are threatened by various elements ranging from cybercriminals to poor IT practices. There are some exceptions to this rule, such as a small firm or organization with a small IT footprint; however, many such organizations will leverage more technology in the future, leading to the need for a SOC. The massive growth in the Internet of Things (IoT) represents the impact from many nontechnical markets leveraging technology.

Whether you are the sole person responsible for protecting the security of your organization or you are part of a large team with the same responsibility, you essentially are operating as a SOC, even though you might not be labeled as such. How you are viewed by your organization depends on how you (and your group, if applicable) are organized and present your job responsibilities to the organization. For example, imagine that a team of two IT administrators who also are responsible for the security of their organization grows into a dedicated security team of ten administrators with the quality technology and authority to enforce proper practices even if the person violating a process is the CEO. I have assisted with this type of change by

helping security professionals mature their security job roles into a documented practice that is backed by leadership and given the authority to make strategic decisions and obtain budget for growth. This is the foundation of a mature SOC practice.

Factors Leading to a Dysfunctional SOC

Several factors can lead to a SOC becoming dysfunctional. The first problem is a lack of educated security professionals to meet demand. Many organizations have trouble finding and retaining the right people for SOC-related work. There are a lot of job opportunities in the cybersecurity marketspace and not enough skilled professionals to meet the demand. That has become even more acute during the pandemic, as reported in a September 2020 article published on CNBC.com by Kate Rogers and Betsy Spring titled "We are outnumbered – cybersecurity pros face a huge staffing shortage as attacks surge during the pandemic" (https://www.cnbc.com/2020/09/05/cyber-security-workers-in-demand.html). Citing a report by $(ISC)^2$, the article states "2.8 million professionals work in cybersecurity jobs globally, but the industry would need another 4 million trained workers in order to properly defend organizations and close the skills gap. That includes about half a million workers needed in the U.S. to meet demand." The skills gap is also a moving target, as technology shifts toward the need for skills in programming and development versus traditional management of security tools.

The second factor that may lead to a dysfunctional SOC is the cybersecurity industry's hyperfocus on preventing compromises. Preventing a compromise from happening is ideal, but the more realistic approach is to prepare for when a compromise does occur. Organizations should defend against all parts of the attack process rather than assuming the SOC will prevent exploitation 100% of the time. Lacking capabilities to detect adversaries that have compromised a network will lead to nefarious actions taking place within an organization that go unnoticed.

A third issue that may cause a SOC to become dysfunctional is that it cannot be scaled to meet the current demand, resulting in poor reporting,

dysfunctional tools, and analyst burnout. Many drivers such as cloud computing, data transfer and storage, and IoT increase bandwidth requirements, necessitating that security tools increase in size and power to accommodate the increase in data that must be monitored. As data increases, the backlog of events requiring an analyst to sift through becomes 12 to 18 months' worth of continuous review, leading to analyst burnout. The combination of underestimating technology and overwhelming workload demand can quickly bring a SOC to a grinding halt.

Finally, a SOC may become dysfunctional if the organization moves to cloud services without consideration for proper security. Cloud computing offers new challenges to securing data, including leveraging cloud resources that the organization is not permitted to manage or have visibility into. Also, some traditional security vendors might not offer cloud options of their technology that the organization is already familiar with, forcing the organization to acquire new technology that needs new skillsets and has additional costs. Recent technology trends, such as software defended networking (SDN) and work-from-anywhere strategies, are further driving the need for cloud technology.

Any of these challenges can cause the breakdown of how an organization runs its SOC. These challenges can lead to uneducated decisions, gaps in security capabilities, and ineffective procedures. Uneducated decisions include seeing the impact of a problem but not correctly addressing the issue due to a lack of understanding regarding what should be done. There are hundreds of sources claiming "best practices," ranging from vendor publications to industry guidelines, yet different teams will have different missions they are trying to accomplish and, hence, different views of any particular problem. This in turn leads to running the SOC as various siloed groups responding to events with no plan to improve how the security practice operates. A dysfunctional SOC puts an organization at high risk of being compromised by adversaries, and being compromised can lead to the end of an organization.

Cyberthreats

A SOC that is dysfunctional for any or all of the reasons outlined in the

previous section will eventually fail to secure the organization, giving cyber adversaries an opportunity to abuse an exposed vulnerability. An effective SOC, on the other hand, is aware of the wide range of cyberthreats and knows how to protect the organization from them. This section outlines the various cyberthreats as a foundation for subsequent discussion of how to defend against them.

One type of cyberthreat is a malicious actor attempting to compromise a network. An essential element of defending against such attacks is to understand who would do this and what motivates them. According to "Data Thieves: The Motivations of Cyber Threat Actors and Their Use and Monetization of Stolen Data" (a RAND Corp. publication documenting the March 15, 2018 testimony of RAND associate Lillian Ablon before the U.S. House of Representatives Committee on Financial Services Subcommittee on Terrorism and Illicit Finance), there are four categories of cyberthreat actors: cybercriminals, hacktivists, state-sponsored actors, and cyberterrorists.

- **Cybercriminals:** These threat actors focus on making money. They typically are members of organized crime groups or small-time criminals trying to capitalize on using technology to steal data and then sell it to make money. The key to preventing cybercriminals is to make their actions more costly than profitable so that they move on to another target. The key points to consider for this category of cyberthreat actors are as follows:
  - Cybercriminals are driven by profit. Reducing potential profit reduces this adversary's interest in investing time into an attack strategy.
  - Many cybercriminals operate as independent contractors taking work on a for-hire term. Cybercriminals can be recruited by nation states, organizations, or other parties that need to outsource their criminal activity.
  - Larger cybercriminal organizations are essentially the mafia. Rather than robbing banks physically, organized crime groups have shifted their focus to cybercrime as it involves less risk of being caught and very high profit possibilities.
  - Some cybercriminal organizations set up call centers dedicated to specific attacks including making phone calls to deliver social engineering attacks, packing malware to bypass host-based security,

and sending phishing emails. These call centers function similarly to any legitimate business, providing full- or part-time jobs for employees and providing the benefits an employee would expect from any employer.

- **Hacktivists:** These threat actors are people driven by belief in a cause. Anonymous is an infamous hacktivist group that targets people or organizations they feel have violated human rights or other political agendas and need to be punished. Figure 1-1 shows a screen capture from a video featuring a person in the typical mask worn by Anonymous members. Defending against hacktivists is different from defending against typical cybercriminals because hacktivists are not driven by making money. The key points to consider for this category of cyberthreat actors are as follows:
  - Because hacktivists are motivated by a cause rather than profit, they are likely to target a specific entity much more persistently than would a financially driven adversary.
  - It is common for hacktivists to be associated with conspiracy theories, including those involving anti-government concepts
  - Hacktivists have caused major breaches, including the takedown of the PlayStation network and the takedown of HBGary and its CEO Aaron Barr by publishing 68,000 private emails when Barr announced he would reveal the names of some “leaders” of Anonymous.
  - Hacktivists can contract cybercriminals to help with their mission as well as for burst support based on the issue they are addressing.

FIGURE 1-1 Video Posted by Anonymous on YouTube

- **State-sponsored actors:** These threat actors are similar to hacktivists in that they are driven by a cause based on the state that sponsors them. It isn't a secret that most governments are investing in cyberwarfare. Any large-scale war will include disruption of technology using cyber-exploitation tactics. This means if you are responsible for your government's critical infrastructure or other key services, you are a potential target for this adversary. The key points to consider for this category of cyberthreat actors are as follows:
  - State-sponsored cybercrime tends to be very well funded and elite.
  - Many organizations do not have the capabilities to prevent a state-sponsored attack.
  - Most technologically advanced countries are continuously growing their cyberattack capabilities in secret. No country really knows

precisely what other countries have in regard to cyber-offense capabilities, creating a cyber cold war based on an ongoing silent military race.

- It is extremely difficult to track, document evidence of, and enforce laws against international-based crime.
- State-sponsored cybercrime typically represents very targeted attacks commonly referred to as advanced persistent threats (APTs)

- **Cyberterrorists:** These threat actors can be anybody who is motivated to intimidate, coerce, or influence an audience, cause fear, or physically harm. Basically, these are terrorists using technology. Some cyberterrorists are very skilled and are responsible for developing malware never seen until it is used, known as *zero-day threats* since all known detection signatures will not be effective. Other times cyberterrorists are leveraging pre-built scripts to launch attacks making it easy to perform largescale damage to systems. They just point the tool at a target and execute the attack. The key points to consider for this category of cyberthreat actors are as follows:
  - Cyberterrorists can be contracted the same way cybercriminals are obtained, meaning cyberterrorists can operate as independent contractors with skills specializing in causing destruction.
  - Cyberterrorists are not the only adversary that can use a zero-day exploit.
  - The impact of cyberterrorists has changed the cybersecurity industry, prompting requirements for multifactor authentication, improved password policies, and the use of digital certificates to reduce the risk of global events caused by cyberterror.
  - The compromised systems of some unwitting victims of cyberterrorists become part of an attack. Examples include spreading malware through a compromised system, leveraging a compromised system as a gateway into a network, and pushing emails through a compromised system during a phishing attack.

Note

The specific motivation for a threat actor can vary from passion about supporting a cause to being involved only if the pay is right. The threat actor marketplace functions similarly to other marketplaces, with hackers for hire, hackers that treat their work as a 9 to 5 job, and hackers willing to spend months or even years to execute an attack without any pay based on whatever is driving them to be involved with cybercrime.

In addition to the categories of threat actors in the previous list, another category is insider threats. An insider threat could be someone with malicious intent, such as an employee about to leave the company with sensitive data or a security, or it could be an accident, such as an administrator making an honest mistake that exposes the organization to additional risk or being compromised. As an example of the latter scenario, I have seen security administrators accidentally misconfigure security tools such as honeypots and sandboxes, turning these tools into gateways for malware to infect the environment.

One final threat that your organization must be prepared for is *change*. The industry has seen continuous change in how threats operate over the last few decades. In the late 1980s and early 1990s, threat actors primarily attacked computer operating systems. As the operating systems became more secure, threat actors turned their attention to attacking the Internet browsers installed on operating systems. As Internet browsers became more secure, threat actors began attacking browser plugins such as Java and Flash. Looking at security tools, when defenders invented the sandbox to detect malware, malware writers purchased the sandbox, learned how it functioned, and configured malware to bypass the sandbox. When bitcoin technology made it effective for adversaries to remain anonymous while requesting payment from users infected with ransomware, bitcoin became the method to obtain payment from victims of ransomware. Some ransomware creators found it more lucrative to infect systems with crypto-mining software rather than ransomware, so those ransomware writers gravitated toward creating crypto-mining tools.

The cyber battlefield is a constantly changing environment, which means you need to expect constantly changing variations of exploitation against your

organization from different types of threat actors. If you do not continue to invest in your security program, it will quickly become obsolete. If you focus all of your energy at defending the attack of the month, the next change will bypass your security. The following is a great axiom to keep in mind: "Security is a journey, not a destination." You don't become secure; you continue your security journey as you run the security operations center.

This quick overview of cyberthreats should help you to understand what is out there waiting for your SOC to slip up in defending its people, technology, and data. This leads us to the next topic, which is how you can defend against cyberthreats. Let's next look at the concept of security.

Investing in Security

What is the proper investment to improve security within an organization? Some people believe security is all about having the latest or "best of breed" technology and that obtaining such technology should be the highest budget priority. Others think success depends on the quality of the people within the security team and therefore money is best spent on highly skilled IT personnel. A third idea is that the best security comes from well-defined and executed policies that include how to restrict risky behavior as well as respond to threats. The truth is that best practice is a combination of these concepts representing investments in people, process, and technology. There are many industry models and certification programs that reference the ingredients to security using these or very similar terms. For example, the U.S. National Security Agency (NSA) substitutes "operations" for "process" in its information assurance and defense-in-depth strategy, as shown in Figure 1-2.

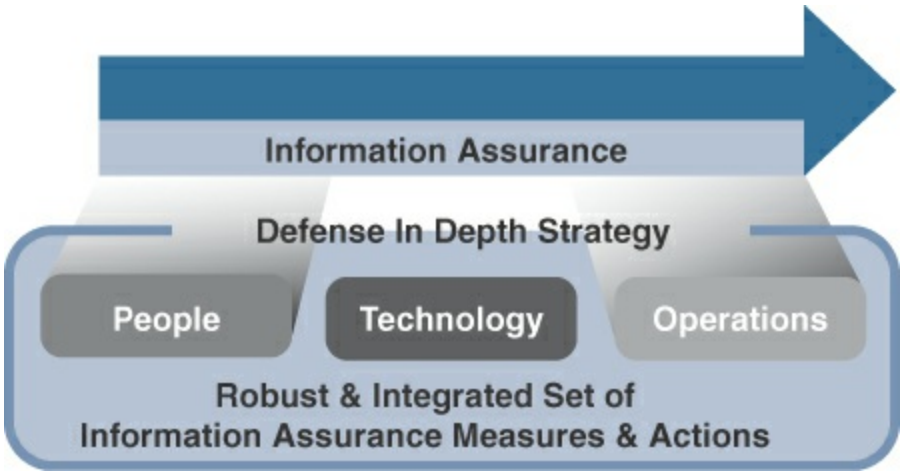

FIGURE 1-2 NSA Information Assurance and Defense-in-Depth Strategy

Figure 1-2 includes some key terms to explore, as they will be important underlying themes for many of the concepts in this book:

- **Information assurance:** The practice of assuring the confidentiality, integrity and availability (CIA) of information and managing risks related to the use, processing, storage, and transmission of information and the systems and processes used for those purposes. Essentially, information assurance means protecting data. Data represents the crown jewels of an organization and is very valuable both to the organization and on the black market.
- **Defense in depth strategy:** To protect data, a SOC uses a combination of people, processes, and technology to create different layers of defenses, which ensures that when one layer of defense fails, another layer steps in, making it harder for a cyberthreat to accomplish its goal. For example, when a firewall fails to prevent an open port from being exposed, an intrusion protection system (IPS) can monitor the traffic through the open port for exploitation behavior. If the IPS fails to see the exploitation behavior, a host-based anti-malware tool can evaluate which files successfully made it through the open firewall port and were not blocked by the IPS.

A SOC creates and enforces a defense-in-depth strategy to protecting data by following high-level polices supported by detailed procedures, which all make up the instructions that guide the success of securing data, hence

providing information insurance. Pay attention to the language being used throughout this book, as it aligns with how the industry speaks about cybersecurity concepts.

The Impact of a Breach

Why does developing a mature security operations center even matter to an organization? The answer is apparent by measuring the impact of incurring a cyberbreach. The impact includes a wide variety of pain to an organization, starting with potentially huge financial losses. According to the IBM Security *Cost of a Data Breach Report 2020*, based on independent research conducted by the Ponemon Institute, the average total cost of a data breach is $3.86 million and the average cost per lost or stolen record is $148. Many organizations would not be able to recover from this level of cost, and even larger Fortune 500 organizations would go out of business if a few incidents of this level of cost would occur, assuming they have not invested in cybersecurity insurance. This cost can span across a long period of time based on what is required to do post incident in terms of services, discounts, new tools, and other damages.

Another cost of a breach besides the direct financial impact is the loss of trust in the organization. This could result, for example, in customers not wanting to buy certain technology from the organization or even being afraid to pay for a service through its website due to a fear that their credit card information may get stolen. Depending on the location of where a breach occurs, there could be breach notification laws that require an organization to inform the public of potential losses. For example, some laws require an organization to alert all record owners about a breach based not on whether the adversary *actually* accessed their records but on the length of exposure and existing forensic evidence indicating the *possibility* that their records were accessed. An example is the Target breach that Kevin McCoy stated "the data breach that affected 41 million customers." It's hard to say that 41 million customers were directly impacted by this breach; however, because the potential existed for those records to be accessed, Target was forced to release that number to the public and alert all of those 41 million customers of the incident.

A third impact of a cyberbreach is potential fines and loss of staff. Some services such as leveraging credit card information (industry regulation) or people's private data (government regulation) include requirements that must be addressed to protect such data. If a violation is found within a government regulation, fines will be issued and parties responsible for the violation could serve time in jail. Outside of required punishment, many organizations respond internally to a cyberbreach by assigning blame, which may lead to the termination of one or more employees or their reassignment to a different role within the organization. This tends to compound the stressfulness of a security event because the organization then has to replace critical staff in an industry that is limited in qualified talent.

One final potential impact of a cyberbreach is the loss of data. The loss of data could lead to hefty fines and loss of trust, as covered earlier, but there are also other negative outcomes from losing data. The loss of company proprietary data could give competitors an advantage. For example, a breach of sales contacts or documentation on future technology would be devastating to some businesses if their competitors had access to such information. The loss of data can expose employees to identity theft and cause a loss of staff or loss of partnerships with other businesses (for example, according to a 2014 article in the *Wall Street Journal*, the Sony Pictures hack exposed personal data, including Social Security numbers, of 47,000 Sony employees and Hollywood stars including Sylvester Stallone, Judd Apatow, and Rebel Wilson."

Any of the previously discussed outcomes will cause a tremendous negative impact on any organization, from small businesses to Fortune 500 companies. Smaller businesses that run on a tighter margin have less resources and can be driven to bankruptcy. Larger organizations might not go out of business; however, they will experience losses leading to stock devaluation and huge costs to return business back to an operational state. It is for these and other reasons that running a mature SOC is critical to the safety and sustainability of any business. Small business can't live by the concept "we are not important enough to be a target" and larger organizations can't believe "we have enough resources to handle the blowback from a security incident." Actions must be taken to reduce the risk of having to deal with a breach. Those actions are the responsibility of the SOC.

Establishing a Baseline

Before you can make any improvements to your security practice, you need to assess the maturity of your current practice. This is an evaluation of everything from how your practice aligns with the goals of your business to your specific capabilities and processes. Consider this a baseline of your existing capabilities and services, which enables you to determine when and where improvements are made or lost. Having a baseline permits the SOC to develop goals for future capabilities and services as well as establish milestones leading to those goals. If you don't know how to establish your baseline, you can consult frameworks such as the Cybersecurity Framework from NIST and self-assessment strategies, covered later in this chapter, to help develop your baseline. SOC leadership can give rewards and recognition to SOC staff for meeting milestones toward accomplishing development goals as a way to encourage improvements that enhance the SOC atmosphere. SOC members can align requests for resources to specific goals to help justify those requests to non-SOC parties, typically members of nontechnical leadership teams that control the budget. Chapter 2, "Developing a Security Operations Center," walks you through how to align your SOC to the business.

Note

Details on the NIST Cybersecurity Framework version 1 can be found at https://www.nist.gov/cyberframework/framework. I will reference this and other guidelines throughout this book, as they are great resources for establishing a SOC baseline.

The Impact of Change

Improving the maturity of your security practice comes from making changes. Know that it is close to impossible to receive 100% benefit from any change, because changes also introduce some form of complexity. For example, a new tool might provide value, but there will be complexity

involved with setting it up and operationalizing it to obtain any benefit. Sometimes, the tools are so complex that they require many additional steps to use properly, pulling resources and time away from other areas of security, essentially hurting your security capabilities rather than helping them. If this is likely to occur, you should recognize that the complexity outweighs the benefit and avoid the change.

An example of a decision that requires comparing the impact of change is the choice between a free, open-source option and an enterprise option for a specific tool. The open-source tool might not have a cost to acquire it, but it will have a cost to install, configure, learn, and maintain. The enterprise option has an upfront acquisition cost, but it could offer simpler deployment and configuration options as well as include features and support from a vendor that would not be provided with the open-source option. It is important to weigh all of these factors before determining the true cost of a change. You will learn more about comparing investments in building your own tools, using open-source tools, and purchasing enterprise options in more detail in Chapter 10, "Data Orchestration."

A good change is one where the capability outweighs the complexity. An example of this is improving the SOC's capability to identify and remediate vulnerabilities, a practice that many organizations find extremely difficult due to a lack of visibility regarding what is on the network and what types of vulnerabilities the known and unknown devices introduce. One method to simplify a SOC's vulnerability management practice is to leverage network access control and vulnerability management solutions. Network access control (NAC) solutions are designed to automate control of what can and cannot connect to a network. Vulnerability management solutions are designed to identify any known vulnerability, including details based on a Common Vulnerability Scoring System (CVSS) score. A CVSS score provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. The numerical score can then be translated into a qualitative representation (such as low, medium, high, and critical) to help organizations properly assess and prioritize their vulnerability management processes. Chapter 9, "Vulnerability Management," covers vulnerability management in much greater detail.

Integrating NAC and vulnerability scanning technologies can automate identification of what is connecting to the network and scanning any device

for vulnerabilities upon connection. Many NAC technologies can even limit access to systems that are found to have a critical vulnerability, such as a CVSS of 8.0 or higher. Figure 1-3 is a SANS Institute concept model for vulnerability management best practices that reflects this concept.

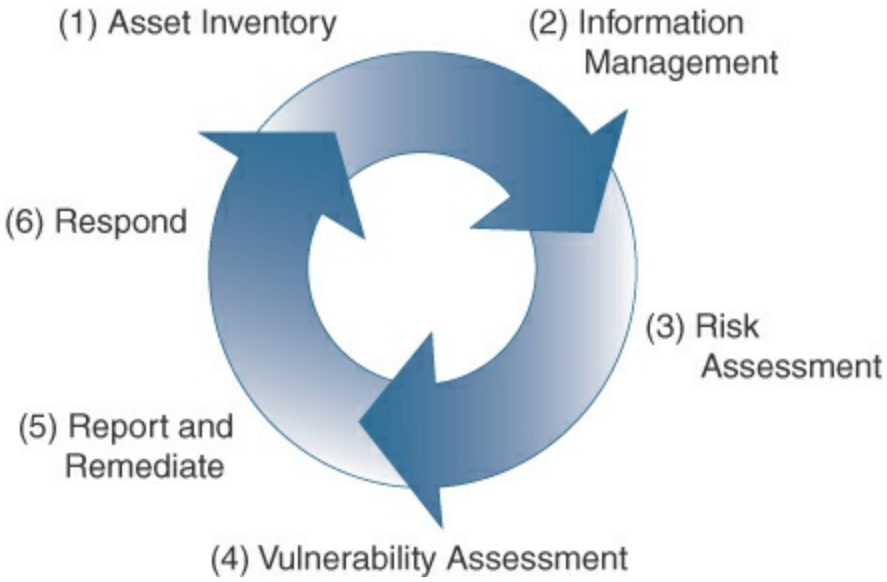

FIGURE 1-3 SANS Vulnerability Management Best Practices Concept Model

Automating vulnerability management using these tools and processes is something every organization should consider; however, remember that every change introduces its own level of complexity. Technologies such as NAC and integration with vulnerability management solutions can be complex to deploy. The question that should be asked is whether the benefit of the capability outweighs the complexity. For many medium to large organizations, the answer is yes, due to the existing risk of not having an effective automated vulnerability management program as well as existing efforts used to function in a manual reactive manner. Smaller organizations may not need this level of automation since it would be overcomplicating something that could be controlled at the desktop level by a small IT support team. This is why the capabilities and complexity associated with change will always be specific to the organization to which it is being applied.

It might be hard to recognize the complexity that comes with a capability or service. Many vendors love to promote how easy their technology is to

implement and use and they tend to oversell the effectiveness (after all, they are in the business of selling products and services). It will be up to the SOC or an external audit to judge the effectiveness of your capabilities and what could be done to increase your organization's security effectiveness. I will cover how to audit your security capabilities and services later in this chapter. Before looking at how to audit capabilities, we need to first review the concept of capabilities.

Fundamental Security Capabilities

What are capabilities in regard to security? They are your ability to identify and respond to a threat. Managing risk caused by threats is a key service most SOCs address, which is a topic covered in Chapter 3. What is important in understanding risk is the likelihood that a threat could exploit a vulnerability. Simply put, threats exploit vulnerabilities, and the SOC's job is to attempt to detect and prevent this from happening. Detection and prevention involve different security capabilities. The security industry is notorious for claiming that security tools can offer complex capabilities such as using human-like behavior (artificial intelligence) to make decisions about whether something is a risk, leveraging cloud resources to further evaluate potential threats, or having layers of different "checks" to catch the stealthiest malware. Some of these claims are true; however, the truth is that there are three fundamental detection capabilities that are used by security tools. Security tools are designed to leverage one or more of three fundamental capabilities: the capability to detect known attacks, the capability to detect known bad behavior, and the capability to detect anomalies (such as unusual behavior by a new type of attack). Industry capabilities can be boiled down to these fundamental detection concepts. Figure 1-4 represents this concept of the three core security capabilities used by security tools to detect and prevent threats.

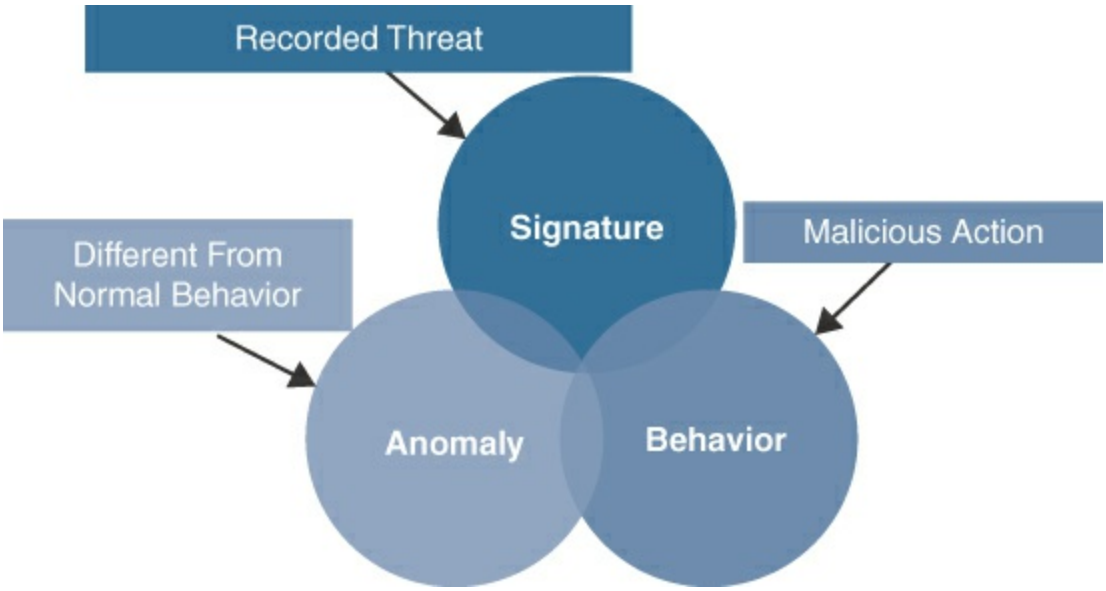

FIGURE 1-4 Security Detection Capabilities

Signature Detection

Let’s start with the capability to detect known attacks, which relies on signature-based technology. One way to look for known threats is by using signatures that represent detection of specific threats that have been identified before in the security industry. For example, antivirus solutions have many signatures for known malware and will compare files against a list that is continuously updated with signatures of recent threats. Detecting malware could be based on various characteristics, such as a hash of file, but it isn’t as simple as it may seem. There are challenges associated with creating and managing signatures. The more specific a signature is, the easier it is for an attacker to modify an attack to the point of avoiding detection. The more general a signature is, the more likely it could generate false positives, meaning triggering against things that are not actual threats. It is common for adversaries to modify existing malware in a way that changes the look of the file, known as *encoding the file*. They do this so that the malicious object is seen by an antivirus solution as something different from what is found within its list of known bad files. There are other methods of avoiding detection, including encrypting files, adding useless lines of code, or adding no operations (NOPs). Adversaries can test malware against public signature

lists such as VirusTotal to see if malware could possibly be triggered by popular security tools. (If you have never heard of VirusTotal, check it out at https://www.virustotal.com.)

> **Note**
>
> Open-source penetration testing tools such as Metasploit by Rapid7 offer various methods to hide files. For example, you can use Metasploit to encode test payloads and see if your security tools can detect your encoded threat. The creators of Kali Linux, Offensive Security, posted a great article explaining this concept at https://www.offensive-security.com/metasploit-unleashed/generating-payloads/.

Behavior Detection

The concept of detecting known bad behavior is based on actions seen rather than scanning for specific things (pattern matching). Malicious behavior could be anything from a computer scanning the network to a file attempting to gain root access to a system. Returning to my malicious file example, assume the adversary encoded the file so that it is not detected by an antivirus signature. Antivirus software can monitor the file's behavior to identify it as malicious. Let's say for this example the file is ransomware, which means it would attempt to encrypt the hard drive of the system it has infected. If ransomware is configured to use asymmetric encryption, it would need to reach out to an external source on the Internet to perform the key exchange before the encryption process can be completed. This means the file will beacon out a web source owned by the threat actor to complete the encryption process. Any of these actions caused by a file should trigger a security tool to prevent the file from proceeding with these actions. This makes behavior rules ideal for validating that threats are not bypassing signature rules.

Anomaly Detection

What happens when a threat is unknown by signature and behavior detection capabilities? This is where anomaly detection can be extremely beneficial. Although detecting anomalies might seem similar to detecting known bad behavior, they are different. Anomaly detection is based on baselining a network and flagging anything that is unusual. For example, some organizations might permit employees to send email from their corporate email accounts to their personal email accounts. Although this activity may occur all of the time in an organization with hundreds of employees, the SOC could configure an anomaly rule that flags when a corporate email account sends an unusual amount of emails to a personal email address. Why would an organization want to implement this rule? This activity could indicate that an employee is about to quit the organization and is emailing a load of sensitive internal material to their personal mailbox before they resign and turn in their computer. For this use case, the rule could be either behavior-based (meaning a set number of emails during a specific time would trigger the alarm) or anomaly based (look at the user's average email activity and flag an unusual spike based on that specific user). Some users might have a higher average of email activity based on their role, making anomaly detection capable to adjust to real-world activity. One key point that is shown in this example is how anomaly detection isn't as accurate as other methods. For this example, the email rule could trigger if a legit reason is occurring to send an unusual amount of email from a corporate email account to a personal account. Legit reasons could include a user sending out tons of email for an authorized email campaign or performing a backup of their email.

Note

I experienced the email anomaly detection alarm in my career when I was upgrading my computer. I spent a few minutes sending files over email from my corporate email account to my personal account because it was quicker than connecting an external hard drive or using a cloud hard drive to move files over. My goal was to quickly move files from my old computer to my new computer, which both were connected to the Internet, because I needed to clean my old computer so I could return it to my employer. After I sent around 15 to 20 emails to my personal Gmail account, I received a phone call from

human resources asking if I was satisfied with my current position. Based on my unusual email activity, my employer had concerns that I could be leaving the organization. I simply explained I was migrating to a new computer, and that was the end of that conversation. Also note that this rule was not put in place for data loss prevention purposes. If I had sent only a few emails, I wouldn't have received the call from HR or the data security team since the actual data wasn't being evaluated. That means the anomaly rule was only designed for a spike in data being sent targeting potential employee flight risk.

Another example of an ideal use of anomaly detection capability is to monitor home appliances that are being connected to the IoT. Imagine a thermostat connected to the network that periodically downloads small updates from the manufacturer's website. An anomaly alarm should go off if this thermostat starts exporting large amounts of data from the corporate network. Anomaly detection could be combined with behavior detection, meaning the risk in the thermostat example could be increased if the thermostat also starts performing port scanning for the first time, indicating this IoT device is potentially being leveraged by an external party to survey the inside network. Once again, it is important to point out that anomaly detection isn't always accurate. For this example, the IoT device might be downloading a very large firmware update or might have a new feature enabled that has it send large amounts of data to the vendor.

Note

To avoid false positives, you need to configure security tools to ignore devices that will trigger your security capabilities. For example, you will want to add a vulnerability scanner to an ignore list; otherwise, every time it performs a vulnerability scan, it will trigger reconnaissance alarms in your detection tools.

It is common at a security event to hear a vendor claim their security tool can do a lot more than the three detection capabilities I just covered. Hopefully,

knowing how to boil capabilities down to their core concepts will help you understand the true value of tools you are considering for your environment. I point this out because it is absolutely critical to understand which capability you are evaluating so that you can understand what it protects, what it doesn't protect, and how to maintain it for maximum return on investment.

Best of Breed vs. Defense in Depth

What is the best approach for using the available security capabilities? Is it better to have one very strong capability, known as *best-of-breed capability*, or to layer different capabilities, known as a *defense-in-depth strategy*? Best practice is to layer different capabilities rather than using the same type of detection. For example, a firewall permits or denies traffic based on rules. Having multiple firewalls wouldn't provide any additional defense against an exploit targeting a system over port 80 if all firewalls have port 80 open. What is more ideal is to have an IPS and some form of breach detection layered so that if the firewall permits the traffic, the IPS can analyze the traffic for exploitation. If the attack goes unnoticed, bypassing the firewall and IPS, the breach detection technology, such as an anomaly-based tool, can identify the unusual change to the target and flag it as being exploited. By following this strategy, an attacker has to beat different forms of detection in order to accomplish his or her goal. The more layers with different capabilities, the less likely an attacker will be successful. Figure 1-5 represents the concept of layering capabilities.

FIGURE 1-5 Best of Breed vs. Defense in Depth

Note

It is important to point out that purchasing multiple firewalls is not a bad investment. Different parts of your network will need their own firewall, hence would be their own string of technology capabilities that include a firewall. Each string of capabilities would be graded independently from the rest of the network's capabilities, including other firewalls. An example of this is comparing the edge of a company's network against the security within its private cloud or internal datacenter. It is ideal for organizations to have separate firewall, IPS, and breach solutions for the datacenter and the network edge.

Evaluating Technology

The security market can be confusing with regard to understanding whether the technology you are considering is offering a best-of-breed option or a defense-in-depth option. Vendors often claim their stack of defense-in-depth technology is the best of breed based on how the different capabilities work together. There are resources such as Gartner Magic Quadrants that attempt to categorize a security product and compare multiple vendors within that category based on who they believe is the closest to best of breed. The results of some resources are not technically accurate either because the ranking is influenced by nontechnical data such as customer feedback or vendor financial endorsements or because the tools are tested subjectively. An example of the latter case is a vendor who sponsors a test and rigs the testing criteria in favor of its own product. I once saw testing results for the "IPS Category" show that one vendor had a 100% detection rate while others did not. After further reviewing the testing criteria, I found the test was based on searching for a customized signature that only the tool of the vendor sponsoring the test had enabled! This type of test is obviously not a representation of a real-world use case and is only done as a way to develop a report used to generate sales. My recommendation is to question any report that ranks technology based on the following items:

- How is the testing performed?

- Who set up the vendor technology?
- Who conducted the test?
- Is the latest version of each vendor technology being evaluated?
- Has the vendor confirmed its tools were properly configured?
- Is the test based on real-world situations and based on vendor-neutral concepts?
- Are there nontechnical factors such as vendor sponsorship or other potential bias in the results?
- Do the associated vendors agree with the results?
- Were some vendors given access to the testing criteria before the tests were performed?

Basically, you are looking to see whether all tools were tested in a fair manner. I recommend to always question the results of these types of tests rather than trust them at face value. While I was employed for one security vendor that was involved with all of the leading third-party evaluation reports, I was shocked to find how much effort was required to participate with groups such as Gartner and NSS Labs (no longer in business). Effort included hundreds of hours of top engineering time, travel between testing facilities, providing hundreds of thousands of dollars of free equipment, countless hours involved with analyzing alpha/beta results, and responding to questions from the third-party evaluation team. There are smaller vendors with great technology that can't endure the financial burden associated with some third-party tests, leading to poor results that are not truly accurate as to how the tool actually performs against competitive technology.

Researching Technology

My recommendation regarding determining what is the best solution for your organization is to use a combination of external resources, align what you find to your business goals, and test. Testing can be challenging based on the capabilities involved as well as general performance concepts such as providing the promised traffic throughput when capabilities are enabled. One example of a common challenge security tools deal with is encrypted traffic.

One workaround is having the capability of decrypting traffic so it can be evaluated by a security tool and encrypted before sending it on to its destination. This decryption/encryption process will impact performance and is important to consider as a tool is evaluated. I have seen performance losses as high as 60 to 80% when certain security capabilities are enabled. Make sure to start your testing criteria with how the tool should be sized for your needs and if your desired capabilities will still deliver the expected performance.

When testing capabilities, it is common to use testing tools to speed up the process as well as to provide a third-party, vendor-agnostic view of the testing process being delivered. BreakingPoint is an example of a tool that can apply stress to a tool to evaluate how it performs under real-world conditions. Tools such as BreakingPoint also can provide templates for common attacks, sometimes called a *strike pack*, which can be used to test how a security tool detects and prevents attacks. To avoid obtaining inaccurate results, it is absolutely critical to consider how close to real-world conditions a strike pack is designed to represent. Using an outdated strike pack would mean testing security tools for threats that might not exist anymore or have been patched; hence, the security tools may no longer be blocking certain attacks launched by the strike pack, resulting in a reported miss even though the threat really doesn't exist since a patch would mean the vulnerability has been mitigated. Consultants such as technology resellers that have access to different vendors can help with building a lab to test different vendor capabilities as well as provide their experience with different tools.

One common question I often receive is, "I understand how to test security tools when I know what I'm looking for; however, how do I know what security capability my organization needs based on our existing investments and, more importantly, which tool should I buy first?" Essentially, the question is asking what materials should be referenced prior to researching what is best of breed for a specific security capability. My answer to this question is to not rely solely on the advice of a specific vendor or even a third-party consultant. Both parties can evaluate your organization and provide recommendations; however, both parties could also be financially motivated to recommend specific technology, hence offering biased advice. To truly obtain a vendor-neutral view of what security tools you need to

consider for your organization, you should leverage industry standards, guidelines, and frameworks.

Standards, Guidelines, and Frameworks

Many organizations look to industry standards, guidelines, and frameworks for help with developing security architectures for their environment. With the exception of industry standards such as PCI DSS, standards, guidelines, and frameworks provide recommendations and guidance that organizations can choose to follow, not mandatory practices that they are required to follow for compliance reasons. Many organizations will turn industry recommendations into corporate policies, which could have both advantages and disadvantages. Benefits of standards, guidelines, and frameworks come from how they are typically developed by industry experts and tested against common threats such as using threat modeling, which is a topic I cover shortly. Recommendations from standards, guidelines, and frameworks are usually vendor-agnostic and focus on capabilities and services generic enough to provide value to any type of organization.

The downside of these resources is that they take time to develop and update, during which time threats continue to rapidly change. I have seen situations where technology is found to be vulnerable and the manufacturer has developed a fix, but the customer will not install the fix until the version of software meets an industry standard, guideline, or framework. This exposes the customer to unnecessary risk during the time it takes for industry recommendations to catch up with the change in the threat landscape. Recommendations from these sources can also be too generic to address threats that are specific to an organization. As an example, suppose an organization has an HR employee who will open any resume file without considering security. The risk of the employee exposing the organization to malware infection could be reduced by following generic framework recommendations for implementing segmentation and anti-malware; however, to fully address this situation effectively, the organization also needs to directly educate the user of the malware risk and develop a specific process to sanitize files before they are sent to the user. The key point of this example is that you should use standards, guidelines, and frameworks only as

a baseline for your security architecture rather than as an all-encompassing blueprint for implementing security in your specific environment. Organizations need to develop their own maturity-grading structure and work on improving security based on what matters to their specific organization, which will extend well beyond the average standard, guideline, and framework.

> **Note**
>
> Many of these external security resources now include self-assessment procedures to help an organization more closely align the recommendations with what the organization actually needs to implement. As a result, fewer adjustments are needed to be made outside of what is being recommended by the specific guidance the organization needs to incorporate. I believe that this trend of including self-assessment capabilities within standards, guidelines, and frameworks will continue and the capabilities will get even better.

Some examples of commonly used standards, guidelines, and frameworks are NIST, ISO, and frameworks from FIRST.org. Chapter 6, "Reducing Risk and Exceeding Compliance," will cover each of these in greater detail. For now, let's take a quick look at each of these resources, starting with what NIST has to offer your SOC.

NIST Cybersecurity Framework

The National Institute of Standards and Technology (NIST) developed the NIST Cybersecurity Framework (CSF), one of the most popular frameworks consisting of standards, guidelines, and best practices related to dealing with cybersecurity-related risk. The latest version, CSF Version 1.1 (https://www.nist.gov/cyberframework), presents a Framework Core based on the following five Functions, as depicted in Figure 1-6 (from NIST):

- Identify applies to managing systems, people, assets, data, and capabilities.

- Protect focuses on defending services.
- Detect covers how to identify that a specific event has occurred.
- Respond is what actions are taken when an incident is detected.
- Recover applies to how an organization attempts to be resilient during the attack as well as how to restore services impacted by the event.

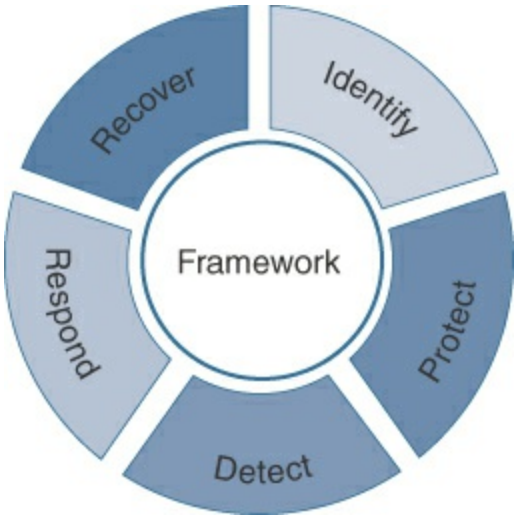

FIGURE 1-6 NIST Cybersecurity Framework Core

Within this Core Framework are Categories, Subcategories, and Informative References, which all are methods to further divide the five functions into more focused topics. For example, the Protect Function includes a Category named Identity Management and Access Control (PR.AC), which has a Subcategory labeled PR.AC-7 that specifies that users, devices, and other assets must be authenticated commensurate with the risk of the transaction. The industry references 802.1X as a leading method for providing access control with multi-factor authentication, meaning it isn't required but highly suggested by the industry. Figure 1-7 shows the structure of the 2018 NIST Framework Core as presented by NIST in CSF Version 1.1.

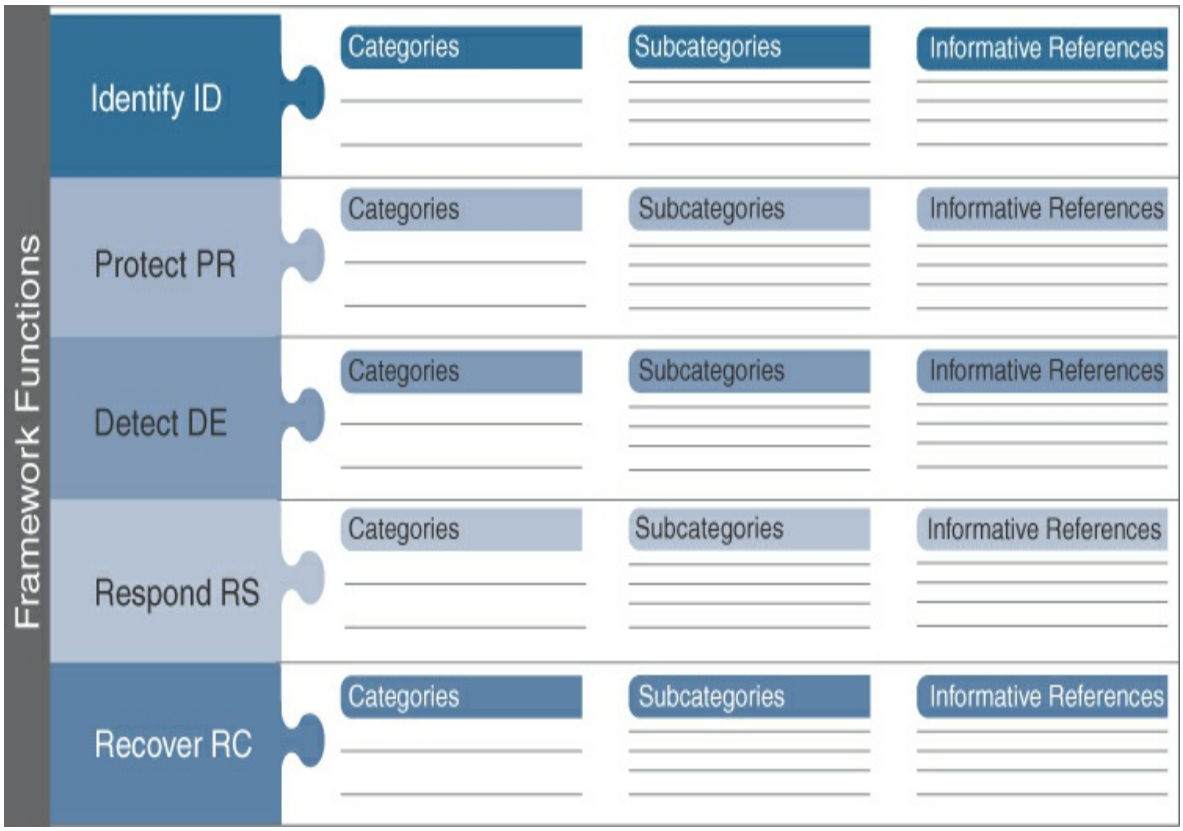

FIGURE 1-7 NIST Framework Core Structure

Using NIST

Many SOCs will review NIST recommendations and develop requests for capabilities they feel will improve their security based on how NIST grades the maturity of that specific Category or Subcategory within the Framework Core. Returning to the example of the PR.AC Category within the Protect Function, it includes best practices for controlling physical access to assets, handling identities and credentials, provisioning remote access, and so forth. Best practices found within this category include controlling physical access to assets, how identities and credentials are handled, and how remote access is provisioned. The NIST CSF also includes references to other industry guidelines to back up its recommendations. The Informative References section of the PR.AC Category includes references to ISO/IEC 27001:2013, COBIT 5, ISA 62443, and CIS CSC, including specific section references, confirming how the NIST CSF document aligns directly with other industry guidelines.

Recommend practice dictates that you download the latest version of the NIST CSF and validate your existing capabilities against each Category NIST covers. Doing so not only can help you identify areas of improvement but also can give you a list of Informative References to back up why you are requesting a people, process, or technology change. I have seen customers achieve success using the NIST CSF in this manner. An example is a member of the SOC team presenting to executives within the organization the risk of not having access control before requesting budget to purchase technology, services, and training for the access control capability. The SOC can back up the request by citing the recommendations of the NIST CSF and the other authoritative resources it supplies. Validating your budget request by citing best practices issued by well-regarded third parties can go a long way toward convincing decision makers who don't understand the technology!

Some organizations treat NIST recommendations as mandatory policy, making budget request conversations much easier. For example, U.S. military organizations use NIST documentation for many of their policies. I see that organizations which follow NIST in this fashion will identify any published NIST requirement, and if they don't have whatever technology is being suggested by NIST, they buy it. No questions asked. It is also important to point out that NIST has many other publications for areas outside IT that can provide value to your organization.

Note

If you have not explored the vast array of NIST documentation, visit https://www.nist.gov/publications

ISO 3100:2018

Another source of popular guidelines is the International Organization for Standardization, more commonly known as ISO. Like NIST, ISO is made up of vendor-agnostic industry experts that provide industry best practices. ISO is a worldwide federation and is well respected in the IT industry. Some organizations will even talk about being ISO certified even though ISO does

not certify organizations meaning the certification part is developed by third parties. The most that an organization can legally claim is that its product or system has been certified to a specific ISO standard by an accredited certification body, such as marketing a product as "ISO 9001:2015 certified." People and organizations talk about being ISO certified as a means to show they take cybersecurity seriously, so they can win over customer trust.

ISO 3100:2018, *Risk management – Guidelines*, helps organizations to deal with risk. The 2018 version replaces the 2009 standard and, like any guideline, is voluntary. There are three risk management focus areas for ISO, which are based on Principles, Frameworks, and Processes. Risk Management Principles targets how to develop an approach that is structured and comprehensive. This will take into consideration many factors such as what is valuable to the company, culture elements, how to ensure improvement, and so on. The end goal is to ensure the risk management approach is effective, dynamic, and customized to your organization's needs.

The risk management framework attempts to identify business goals and establish a formal framework that is sponsored by leadership. Having leadership's buy-in is critical and is a concept you will find not only in ISO but also many other guidelines covering SOC mission statements and business objectives. The ISO framework has five parts that are designed to be repeated, with the last step as Improvement, emphasizing that security needs to continue to improve. I will cover ISO in much more detail in Chapter 6.

Note

Learn more about ISO standards at https://www.iso.org/standards.html

FIRST Service Frameworks

FIRST.org is a nonprofit that brings together incident response and security teams from every country across the world to develop best practices for cybersecurity. One very useful guideline provided by FIRST is the FIRST Services Framework. The FIRST Computer Security Incident Response

Team (CSIRT) Services Framework (version 2.1.0 at the time of writing) provides recommendations for areas of services used by CSIRT teams around the world. Those service areas include Information Security Event Management, Information Security Incident Management, Vulnerability Management, Situational Awareness, Communication, and Knowledge Transfer. Recommendations are broken down into service area, service, and function. For example, the service of Information Security Management provides the service of monitoring and detection, which has multiple functions including log and sensor management, detection use case management, and contextual data management. The CSIRT Services Framework explains what is expected from each function to help guide other CSIRT operations on what expectations would exist within these services.

FIRST also offers a similar framework for product incident response teams (PSIRT). The PSIRT Services Framework defines the scope and operational activities of a PSIRT without the change actions an organization needs to take with regard to the specific products impacted at the organization. This is critical to provide value to any organization regardless of which products that organization is responsible to protect. Activities are specific to the PSIRT rather than what an entire organization would do. Figure 1-8 is a high-level diagram of what general PSIRT activities would entail.

FIGURE 1-8 General PSIRT Activities

Applying Frameworks

As previously discussed, you can use a NIST, ISO, or FIRST framework to validate your security capabilities and services based on industry best practices and reference that framework to request budget for change in people, process, and technology. Other industry standards, guidelines, and frameworks provide similar value, such as what is offered by SANS, ISACA's COBIT, and the Center for Internet Security (CIS) control, each of which gives another take on what experts consider are best practices. Just be mindful of the concepts pointed out earlier regarding how any of these reference materials have limitations based on how often they are released, how they must be generic enough to apply to most organizations, and many other factors that could lead to hurting your security posture if not leveraged properly. Frameworks are a core focus of Chapter 6.

The best way to use the recommendations from standards, guidelines, and frameworks is to apply them based on the specific areas of risk and threats your organization wants to be prepared for, identified from the results of using threat modeling and tabletop exercises. This will make the results of using standards, guidelines, and frameworks more accurate to what needs to be protected with your organization and why it matters. For example, let's look at comparing the security requirements for two schools. One school may permit teachers to access the network only while on campus, while the other school may permit teachers to use remote-access technology to work from home. This means the school permitting remote access would need to consider the associated risks of attacks over VPN, malware that could be introduced by systems connecting over VPN, and the whole process behind implementing and enforcing security for VPN users. The other school would not have to worry about this threat vector. For this example, the school looking for recommendations to secure the VPN service could seek guidance from NIST or ISO for best practices for people, process, and technology in the area of VPN security.

There is a difference between security concepts and the reality of security you need to provide to your organization. Industry standards, guidelines, and frameworks are great isolated use cases for referencing how you should design your security, but as shown in the previous example of comparing two schools, every business is going to function differently, leading to different

security needs. Your security needs today will be different tomorrow, meaning you need to accommodate ongoing changes as you manage your security practice. Part of the focus of Chapter 6, "Reducing Risk and Exceeding Compliance," is how to build policies, which provide very high-level guidance and don't change often, and procedures, which are more specific than policies and constantly change. The bigger challenge is to understand why change is needed and when it should occur. The answer to this challenge is not only understanding what is considered security best practice, but also understanding your potential threats. One popular method used to evaluate how a threat could attack different parts of a network is the use of threat modeling. Threat modeling can also help with developing your criteria for the security technology you plan to acquire. Next, I will look deeper into the concept of threat models.

Industry Threat Models

The security industry uses threat models to represent attack and defend concepts. The purpose of these models is to help organizations understand the type of capabilities they need as they develop a defense-in-depth architecture. For example, it is common for gateway or edge technologies, such as firewall/IPS and host-based firewall/IPS, to be heavily focused on preventing exploitation by using signature-based capabilities. The reason is that these are the first line of defense technology, which will see the most malicious traffic. When the first line of defense fails, the next phase of the attack is to establish a foothold and do things with the newly compromised system. The goals for the capabilities to counter this stage of the attack are different because these capabilities assume the gateway tools have failed. It is common for breach detection tools to be more behavior- and anomaly-based because the failed gateway tools tend to be more signature-based. In the real world, the gateway and breach detection tools can have a combination of all three capabilities, but what is key to understand is that different steps of an attack will have different types of objectives for the associated defense. The same concept applies based on different areas or types of devices, meaning email defense is different from network defense, and web application defense is different from network defense.

As organizations pile on all the possibilities for the types of tools potentially needed, they become overwhelmed and need industry threat models to help them understand what tools and technology apply to their business needs based on the types of threats they expect to encounter.

The Cyber Kill Chain Model

One of the most popular threat models used in the industry is the Cyber Kill Chain created by Lockheed Martin. Figure 1-9 is an example of the Cyber Kill Chain showcasing the lifecycle of a common cyberattack, which is an external party exploiting and gaining keyboard access to a victim's system.

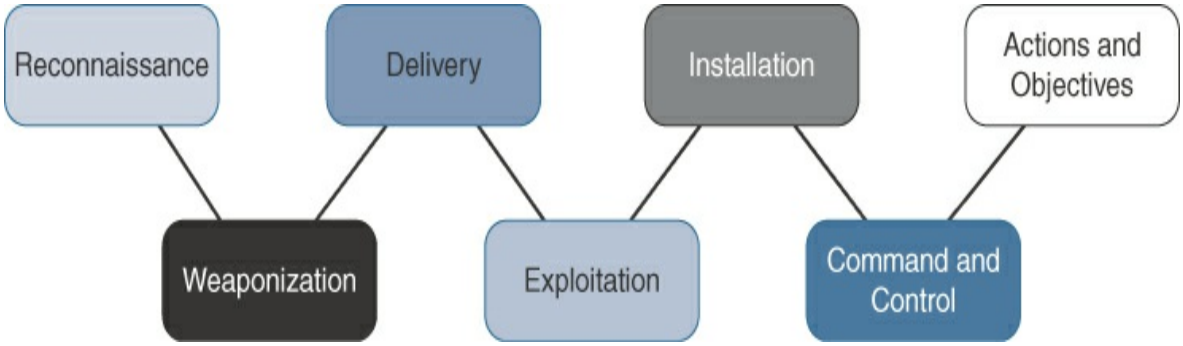

FIGURE 1-9 Generic Cyber Kill Chain Threat Model

The best way to understand the Cyber Kill Chain is to analyze each step of the attack lifecycle, starting with preparing the attack and ending with full-blown keyboard access to the exploited system.

- **Reconnaissance:** The attacker researches the target by probing and assessing publicly available content. This can also include harvesting login credentials or scanning for open ports on Internet-accessible systems. This step is critical for the attacker to learn what is the easiest and more effective method to compromise the target.
- **Weaponization:** Using data found during the reconnaissance phase, the attacker develops an attack technique or tool based on the easiest and most effective route to compromise the target network or system. This could be wrapping software with malware, building an exploit using a tool

like Metasploit, creating a phishing email that asks for data, or linking a file to a malicious website. This step could also include testing the attack against known security tools like VirusTotal.

- **Delivery:** The attacker makes contact with the target and delivers the exploit that was built in the weaponization step. This is the first step in which the attacker interacts with the victim. This could be the result of a user clicking the wrong link, for example, exposing the system to the attacker's malicious tool.
- **Exploitation:** The cyberweapon is delivered and abuses a vulnerability within the target. This causes unwanted behavior such as opening a backdoor on the system or taking down security defenses so that the attacker can install a payload. This could harm the system, but the real damage is what follows this step.
- **Installation:** Once the victim's system is compromised, the attacker can use the exploit to install malware on the target. Installation is the result of a successful exploitation of a system. Malware can be anything from ransomware to crypto mining to a remote access tool (RAT).
- **Command and Control:** One common step that follows the successful installation of malware is beaconing back to the attacker to inform them that the victim's system is available to control. Once the attacker knows the system is available to access, he or she can remotely connect and take control of the compromised asset. This is common for attacks that are not targeted, meaning they exploit any victim that can be attacked and wait to see which victim is successfully exploited through the call back from the compromised system.
- **Actions and Objectives:** The final step could be anything from stealing data to taking down the victim. In the real world, many cyberattacks are a combination of multiple attacks, meaning multiple kill chains are executed. For example, if an attacker was targeting the datacenter of an organization, they would need to breach systems on the network edge and pivot between internal systems to eventually make their way to the environment that contains the datacenter servers.

Know that all attacks do not have to follow this particular attack flow, meaning sometimes the attack situation does not apply to this model. A user

logging into a fake website and having their password stolen would be a different attack model. The threat caused by user error that takes down the network would be a completely different model. The Cyber Kill Chain model is specific to a threat actor attempting to compromise a network by gaining direct control of the compromised system, typically using a tool that provides keyboard access, which can be any of a variety of types of exploitation, such as browser injection or software abuse, or any type of endpoint compromise, such as installing a RAT or dropper.

Using the Cyber Kill Chain

When would you use the Cyber Kill Chain model, and why consider a model based on only one type of attack? It can be beneficial to understand how your capabilities and services match up to the steps associated with the Cyber Kill Chain threat model based on the range of exploitation and malware that could be used in this fashion. Keeping the model generic allows for change in attack behavior, such as ransomware moving to crypto mining, or flash exploitation changing to using an EternalBlue exploit (Microsoft vulnerability). The specifics don't matter since you are measuring layered capabilities and services against the entire lifecycle of the attack, not just a specific step of the attack. Preventing the attack at any point is a win for the defender. The earlier the prevention occurs, the bigger the win it is for the defender.

The first step of the Cyber Kill Chain model represents how malicious actors research and prepare an attack based on what they find using various forms of reconnaissance. Defense strategies include methods to limit how a system is exposed to outside parties and preventing access to high-risk external resources. Think about how to prevent an attack before it happens by reducing the exposure of being attacked. If an outsider can scan your systems for vulnerabilities, attackers will do that as a way to find your weaknesses. If your employees can access any website regardless of its potential risk, employees are going to connect to websites that will attempt to exploit their systems (commonly referred to as exploit kits). If your physical network ports are not performing any form of access control, you are at a high risk of the wrong person plugging something into your network. In summary, if you use technology and best practices for limiting exposure, attackers will have a

harder time identifying your weaknesses and, hopefully, will either attack somebody else or attempt to exploit you where you are monitoring and better prepared for attacks. One common saying in the industry is that your weakest link is your highest level of security. This translates to attackers will find where you are most vulnerable and hit you there.

The middle part of the Cyber Kill Chain looks at how the attacker abuses a vulnerability to gain access to the system. Common exploitation includes abusing out-of-date Java or Flash software or tricking somebody into installing malicious software. Security strategies should prevent the exploitation by identifying the attack and blocking it or quarantining the malicious software before it can install. Capability examples include intrusion prevention, antivirus, and other signature-based detection technology.

Note

As you will read in Chapter 3, many signature-based detection tools leveraging known threat signature lists are filled with enabled signatures that do not apply to what they are supposed to be protecting. Huh, that doesn't sound right. Think about the different types of customers that use a vendor's technology. How could a vendor automatically know what to protect for a large retail store, oil company, school, and casino using a default signature category shared by these different organizations? The truth is, it's impossible! Instead, vendors provide a best guess at what all customers will need protection from. This means that a significant number of signatures enabled by default on many vendor solutions are looking for things that don't exist on your network. This also means there are things on only your network that signatures are not enabled for using default signature settings. This is why tuning security solutions is so important. Tuning can only happen if you understand what you are trying to protect through the use of vulnerability management, understanding how the tool operates, understanding what capabilities the tool leverages, and threat modeling.

Many customers I speak with have some form of signature-based security,

and many employees in those organizations are hyper-focused on monitoring detection-based alarms. This can be a bad thing if all security capabilities and services are designed to prevent the exploitation part of the Cyber Kill Chain model. Every organization needs to prepare for a threat breaching their defenses based on the likelihood that either the organization will miss securing a vulnerability or malware will use a method that will go undetected by the organization's existing security defenses.

The final steps of the Cyber Kill Chain look at the result of a successful attack. Attackers can use an exploit to push malware to systems, disable security features, take down systems, and perform many other nefarious actions that will negatively impact your business. An example of the results of an attack is the attacker overloading a switch to cause it to become a hub, essentially opening up the ability for an attacker to hop between networks. Another example of the result of a successful exploitation of a system is exploiting a server and installing a backdoor, permitting the attacker to gain remote control of that system. With control, the attacker can accomplish his or her goal, which may be repeating the kill chain within the target's network to compromise internal systems, remove data, or shut down the system.

It is critical to use some form of breach detection and continuous monitoring capabilities to validate that previously covered capabilities and services are effective. An example of a breach detection technology is baselining the network and looking for anomalies. Another example is placing vulnerable decoy systems, or *honeypots*, on the network that will alarm the SOC when attacked. These tools could also be bypassed if the attacker knows how to beat them, but the goal of these capabilities is to be different than other capabilities to provide another layer of detection beyond what is at the perimeter of a network or first layer of defense on a host system.

Different Kill Chain Models

One key concept to consider is where you apply a threat model. For example, before using the Cyber Kill Chain threat model, you need to determine what part of the network you want to evaluate. A host laptop has a different kill chain than a datacenter. Both environments may have antivirus and an IPS, but a laptop's traffic is based on a single user and would require different tactics to attack than attempting to exploit a datacenter monitored by a SOC

within a company's network. It is important to evaluate different parts of your organization using the Cyber Kill Chain principles while also considering the specifics to the environment. Your goal as the defender is to "break the kill chain" as early as possible; hence, this threat model helps you to prepare for hypothetical cyberattack behavior so that you can evaluate your defenses against each step of the attack. (I will touch more on how to perform a capabilities assessment later in the chapter in the "SOC Capabilities Assessment" section, which looks at what capabilities and services you could use to break the kill chain.) Figure 1-10 represents different parts of the network that could be targeted by an attacker. Each part of the network should be assessed by its own version of the Cyber Kill Chain threat model.

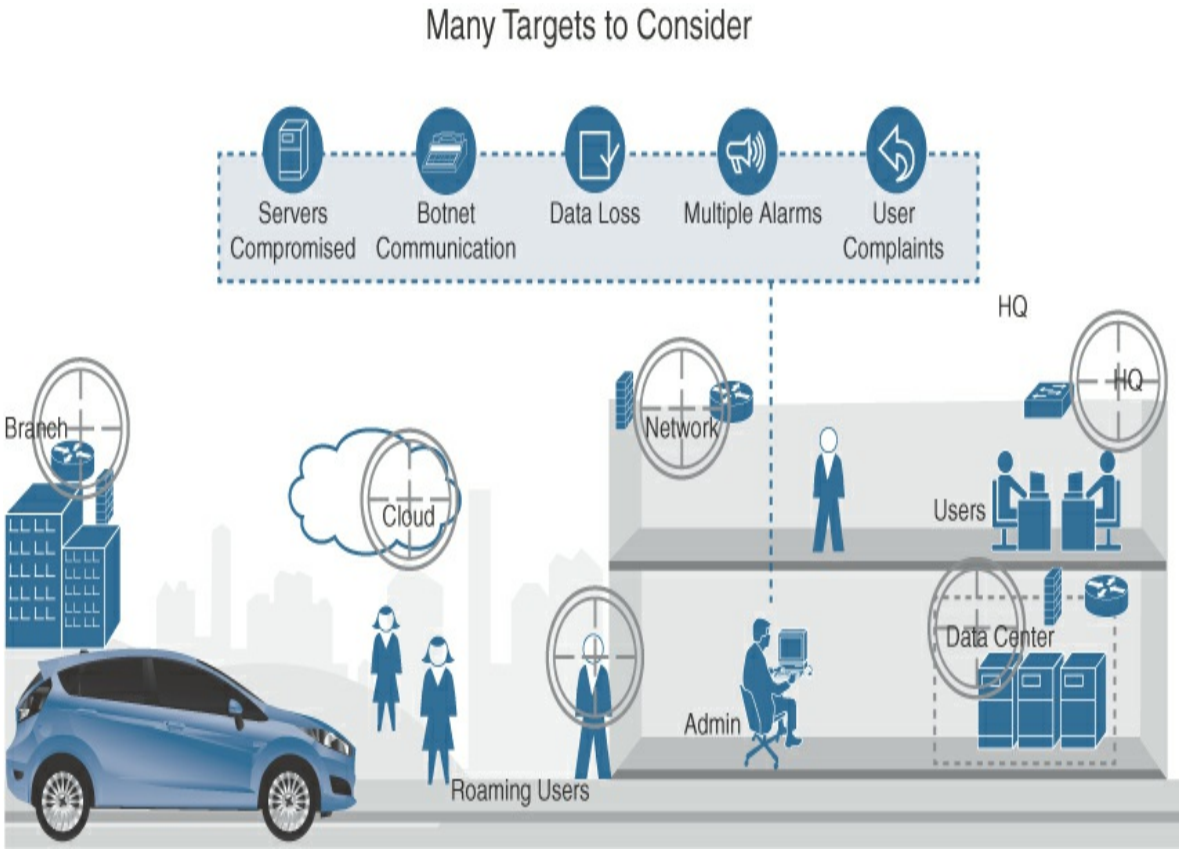

FIGURE 1-10 Various Network Targets

The Diamond Model

The Cyber Kill Chain threat model is a very effective method to evaluate one specific type of attack behavior; however, there are many other ways an adversary could attack your organization. I'm not downplaying leveraging the Cyber Kill Chain, as it is extremely valuable to assess your capabilities against the threats represented during each step of the Cyber Kill Chain. The challenge is considering scenarios in which the attacker approaches your defenses in a different method than represented by the Cyber Kill Chain. Rather than collecting multiple threat models that play through all of the different potential attacks, other threat models review attack behavior from a more holistic viewpoint, providing a way to accommodate any type of attack. One popular threat model that uses more of a broad look at potential threats is the Diamond Model of Intrusion Analysis, commonly known simply as the Diamond Model.

The Diamond Model is designed to represent a security incident made up of four parts, as shown in Figure 1-11. Active intrusions start with an adversary who is targeting a victim. The adversary uses various capabilities along some form of infrastructure to launch an attack against the victim. Capabilities used by the attacker are various forms of tools, techniques, and procedures (TTPs), while the infrastructure is what connects the adversary and victim. The lines connecting each part of the model depict a mapping of how one point reaches another. For example, a SOC analyst could see how a capability such as a phishing attack is being used over an infrastructure such as email and then relate the capabilities back to the adversary. All concepts represented in the Diamond Model are high level by design to accommodate different types of threats, making this model much more general in its approach than the Cyber Kill Chain.

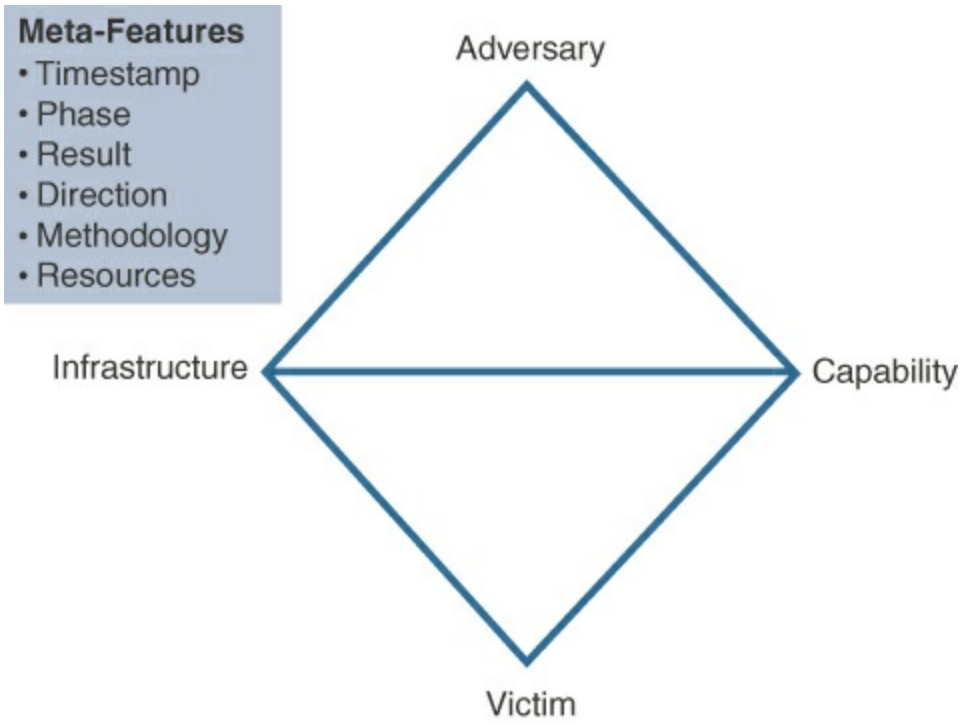

FIGURE 1-11 The Diamond Model of Intrusion

Moving between each part of an attack is called *analytic pivoting* and is key for modeling the event. The Diamond Model also includes additional meta-features of an event (see Figure 1-11), such as a timestamp, kill chain phase, result of the attack, direction of the attack, attack method, and resources used. An example of a meta-features list might show a timestamp of 1:05 p.m., a kill chain phase of exploitation, a result of success, a direction of adversary to victim, an attack method of spear phishing, and resources related to a specific vulnerability on the victim's host system. Meta-features provide useful context but are not core to the model, so they can be disregarded and augmented as necessary.

Extended Diamond Model

The Diamond Model can be further expanded by adding two additional meta-features that establish connections between relations. The *technology meta-feature* connects capabilities and infrastructure by describing the technology used between these two parts of the model. An example of a technology meta-feature could be the Domain Name System (DNS) if it is used by

malware to determine its command-and-control point. The *social-political meta-feature* represents the relationship between the adversary and victim. This is critical to determine the intent behind the attack so that the analyst can understand the reason the victim was selected and the value the adversary sees in the victim, as well as sometimes identify a shared threat space, meaning a situation where multiple victims link back to the same adversaries. A shared threat space is similar to threat intelligence insofar as it is a way of understanding threat actors in a specific space to potentially forecast and react to future malicious activity. An example might be threat actors identified for launching an attack campaign against schools. Figure 1-12 represents the extended version of the Diamond Model.

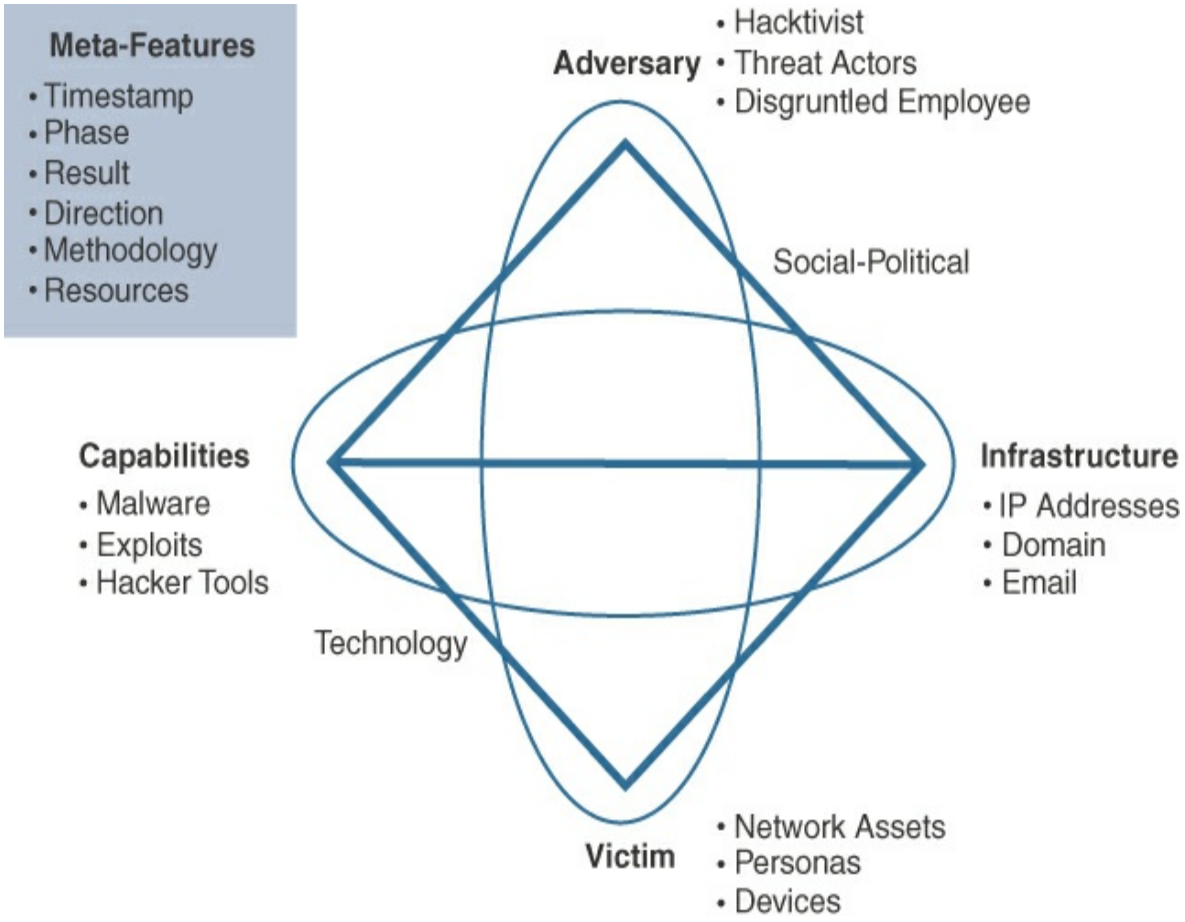

FIGURE 1-12 The Extended Diamond Model of Intrusion

Diamond Model for Incident Management

Each incident is considered a diamond using this threat modeling approach. An incident management practice should use the Diamond Model as the basis for grouping and organizing incidents. The goal would be to review multiple diamonds and identify a common adversary. For example, let's consider an attack where the adversary is delivering ransomware to a victim. The first part of the attack could involve the adversary using a malicious email message to trick the victim into accessing a website. The goal is to have the website scan the victim for vulnerabilities and deliver ransomware by exploiting one of those weaknesses. The first stage of the attack could be represented as one diamond, as shown in Figure 1-13.

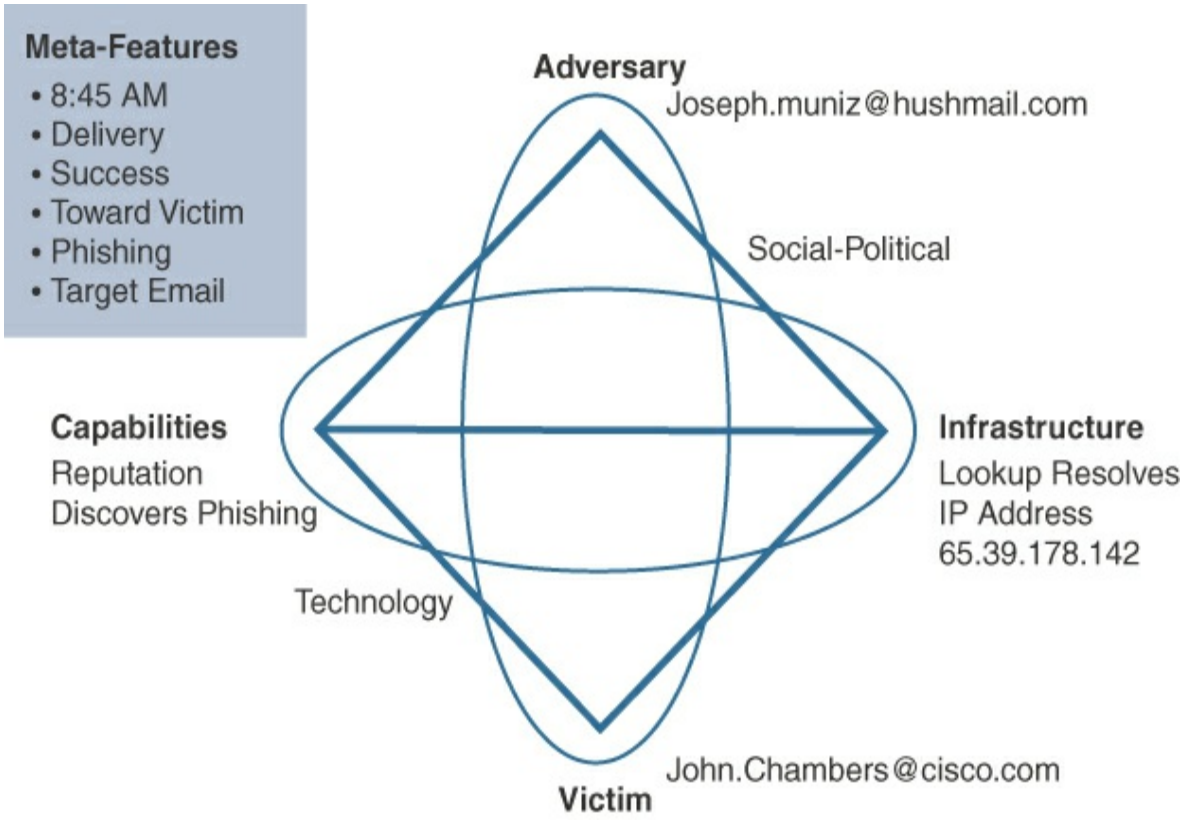

FIGURE 1-13 Diamond Model for Stage 1 of Ransomware Attack

Stage 2 of the attack follows the phishing email that redirected the victim's system to the malicious website. Now that the victim's system has accessed the website, the malicious website will push down the ransomware by exploiting a vulnerability. The adversary is still the same attacker; however, the capabilities and infrastructure involved with the second part of the

security incident have changed, which is common when identifying all stages of an attack according to the kill chain concept. Figure 1-14 showcases a diamond for stage 2 of this attack.

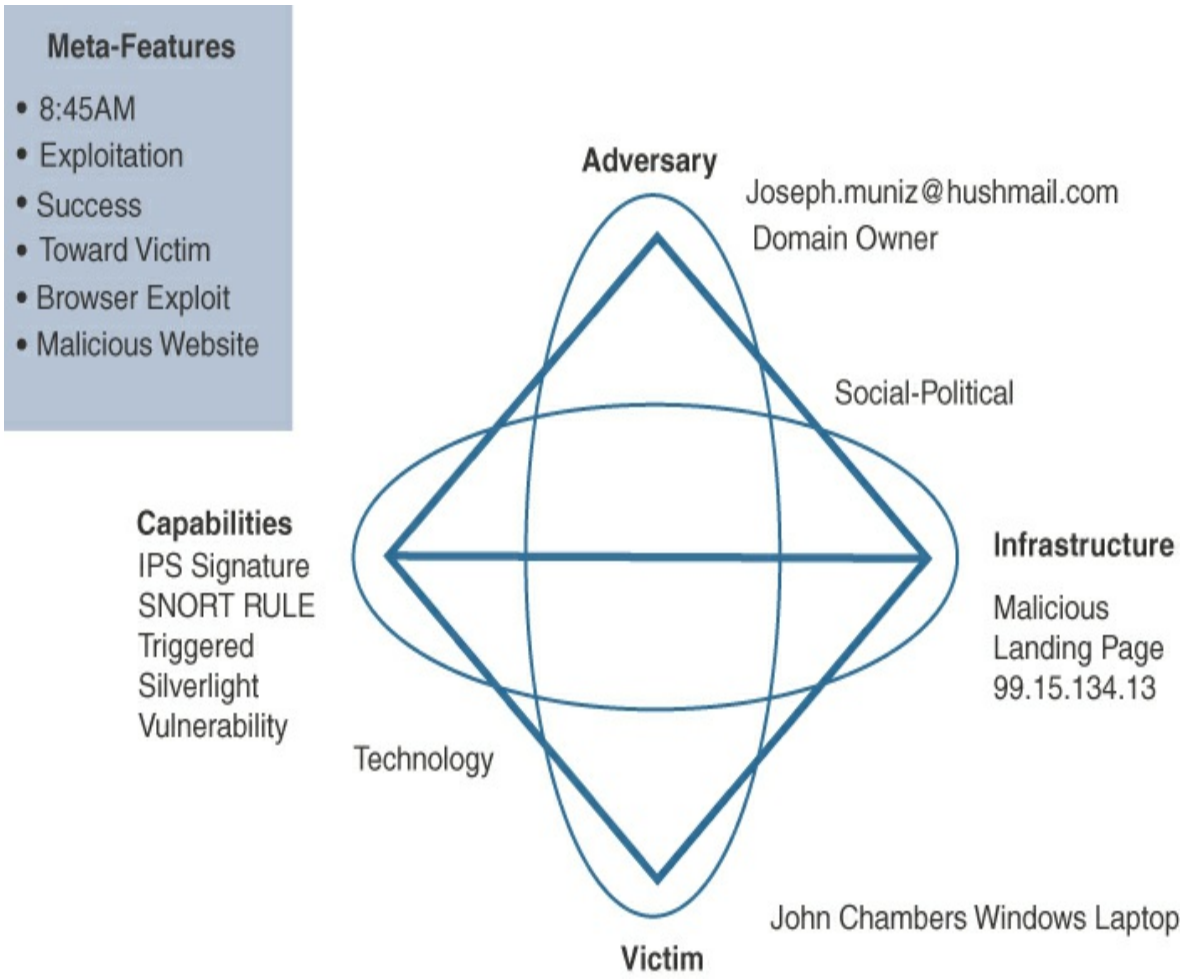

FIGURE 1-14 Diamond Model for Stage 2 of Ransomware Attack

Instances of the same event occurring over the course of a few weeks could be linked together through multiple diamonds and then linked back to the same adversary. Linking the spear-phishing attack to the delivery of ransomware can give an analyst a method to diagram the attack and all associated adversaries. The incident response team can create an activity group based on the various connected diamonds and attempt to define what combinations of elements are criteria for grouping diamonds together. As new diamonds appear, activity groups can grow as diamonds are grouped

together based on newly available data. The relationships between diamonds are known as *activity threads*, which can spread across the same attack as well as connect other attacks, depending on intelligence gathered that meets activity group requirements. Figure 1-15 provides an example of building an activity thread based on the previous sample attack data.

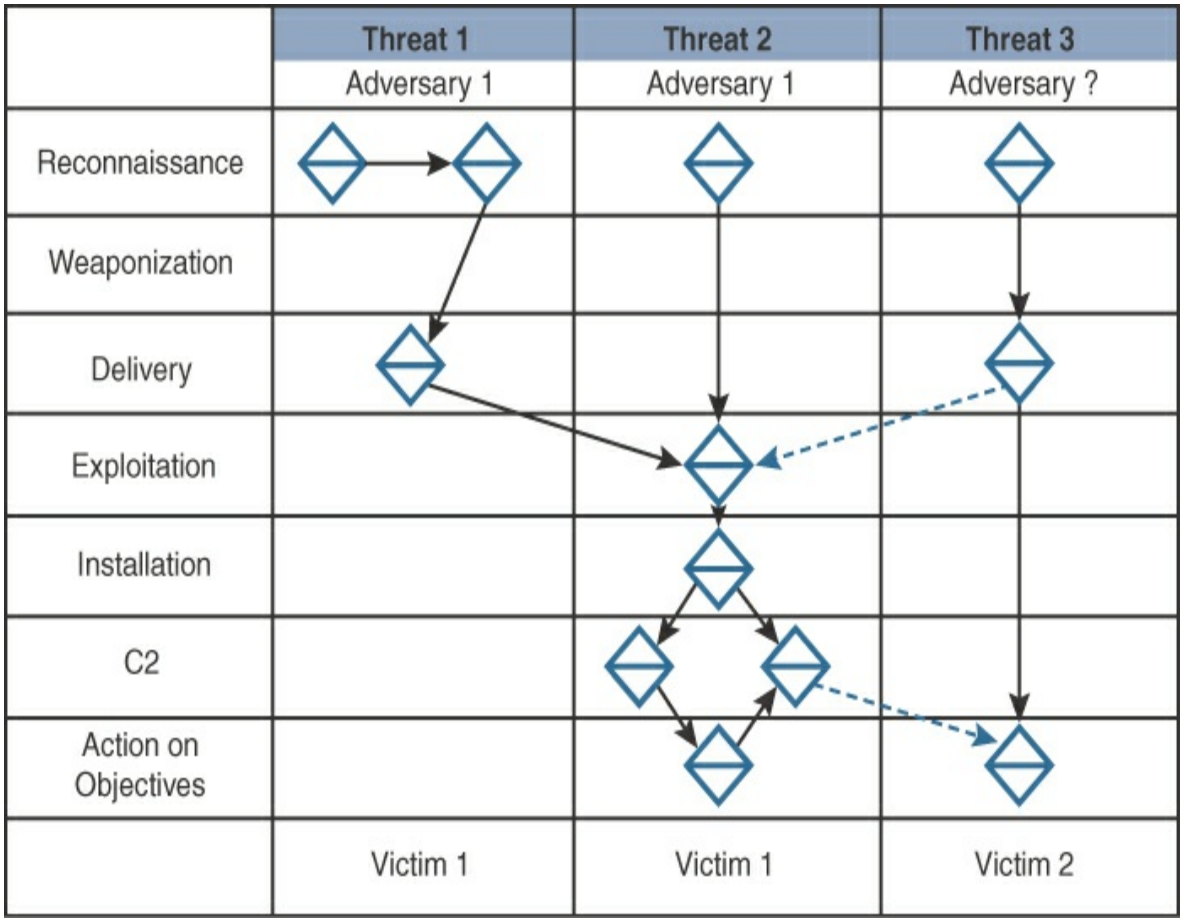

FIGURE 1-15 Developing an Activity Thread

Figure 1-15 shows an adversary is linked to two different attacks against the same victim as well as possibly another victim, represented with the dashed line. There is also another possible adversary attacking a similar victim as the previously identified adversary. This visibility into the attack data gives analysts the ability to integrate any hypotheses that can be tested as additional evidence is gathered. The activity thread process displays the current research status, which can help an analyst identify knowledge gaps and adversary campaigns through documentation and testing proposed attack hypotheses.

Diamond Model Attack Graph

Once the incident management team builds a decent-sized activity group mapping out multiple incidents, the team can better analyze the data to fill in missing knowledge gaps and potentially start to predict future attack paths. This threat intelligence data can be built into a graph, known as an *attack graph*, representing the paths an adversary could take against the victim. Within the attack graph are activity threads, which are paths the adversary has already taken. Combining the attack and activity data gives the team an *activity-attack graph*, which is useful for highlighting the attacker's preferences for attacking the victim as well as alternative paths that could be used. This gives the incident response team a way to focus efforts on defending against the adversary, by knowing where to likely expect the attack as well as being aware of other possible risks to the victim. Figure 1-16 is an example of an activity-attack graph for my ransomware example.

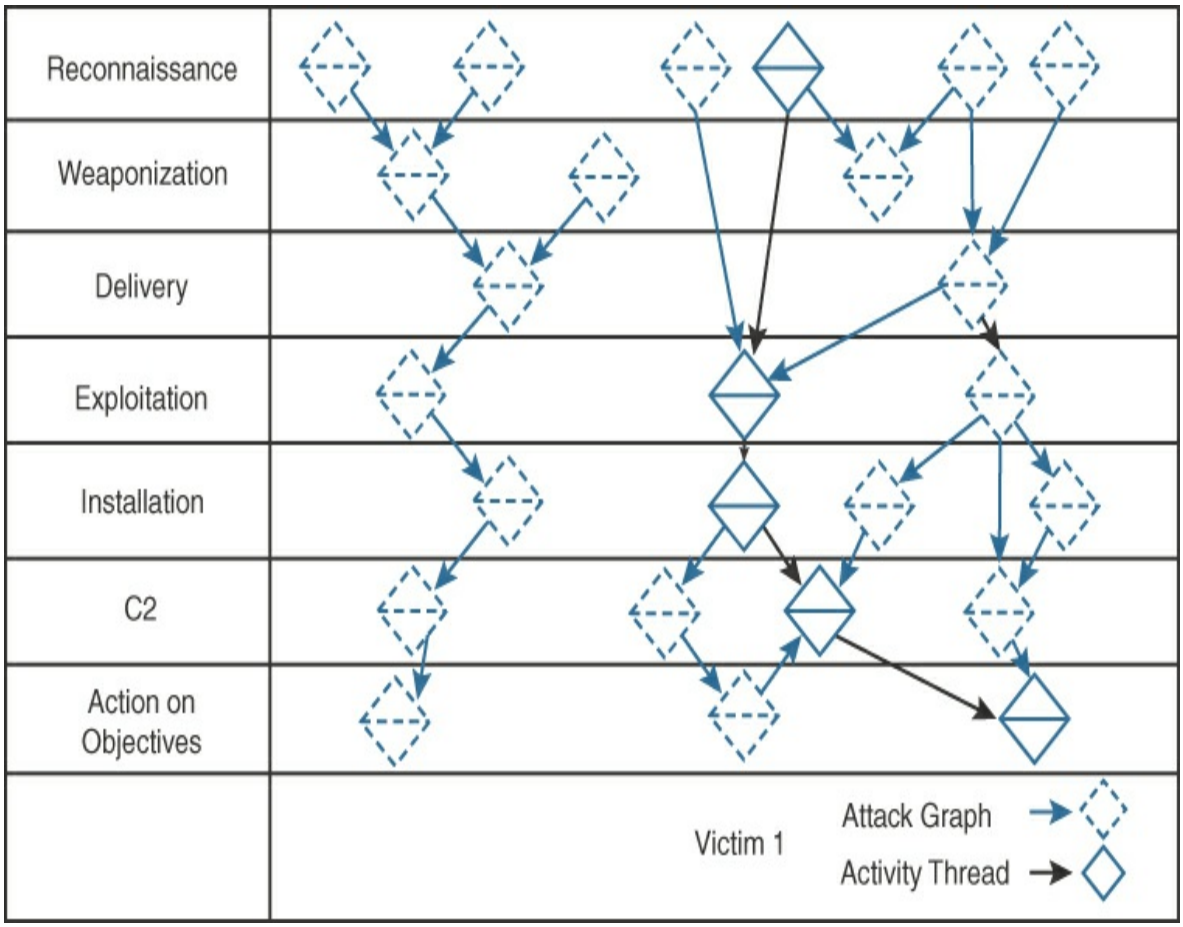

FIGURE 1-16 Activity-Attack Graph Example

If the analyst was concerned that this was a persistent attack, using the activity-attack graph could show not only where defenses should be considered for the identified active attack but also additional areas that could be used by the adversary and therefore should be secured proactively. By grouping common malicious events, adversary processes, and threads, the analyst can create activity groups. Figure 1-16 would help the analyst determine which combination of events makes up an activity group based on similar characteristics. Activity groups can then be grouped into activity group families used to model the organizations behind the various incidents, such as identifying a particular organized crime syndicate. The end result could be the identification of a particular group out of Ukraine attempting to plant ransomware at a specific U.S.-based hospital through the analyst grouping together various events against multiple hosts linked to the hospital.

The Diamond Model is a broader view of attack modeling that allows you to accommodate different attack types as well as identify association between attacks that can represent a larger attack campaign. The Diamond Model also includes the flexibility to add details about the attacker to better understand intent, leading to better decisions based on predicting behavior. This approach to threat modeling lacks some details regarding how an attack is carried out but offers a lot of value regarding general planning against attack behavior.

One final model to consider is a hybrid approach between the value seen from the Cyber Kill Chain and Diamond threat models: the MITRE ATT&CK model.

MITRE ATT&CK Model

Another globally accessible resource for modeling adversary tactics and techniques based on real-world observations is the MITRE ATT&CK knowledge base. This can be used for development of specific threat models and methodologies based on common adversary behavior for emulation and intrusion detection research. Customized threat models based on continuously updated real-world data can be more accurate than the Cyber Kill Chain and

Diamond models, which can lead to a better view of detection of post-compromise cyber-adversary behavior.

ATT&CK organizes the ecosystem an adversary operates within as *technology domains*. Adversaries must circumvent or take advantage of the ecosystem in order to accomplish a set of objectives. The two ATT&CK domains are enterprise networks and mobile devices. Within these domains are the platforms representing the systems the adversary operates within. Adversaries apply techniques to one or more platforms, which is how the adversary attempts to accomplish its goal. This approach is similar to the broad view of an attack offered by the Diamond Model but offers tons of details, such as describing the technique, which platforms apply to the technique, system requirements for the technique, permission requirements, data sources, examples, detection strategies, and even mitigation recommendations. Figure 1-17 shows a high-level view of how ATT&CK represents an attack model relationship. Notice that the goal or end result of an ATT&CK model is labeled as a tactic, which explains why the adversary is performing the previous actions.

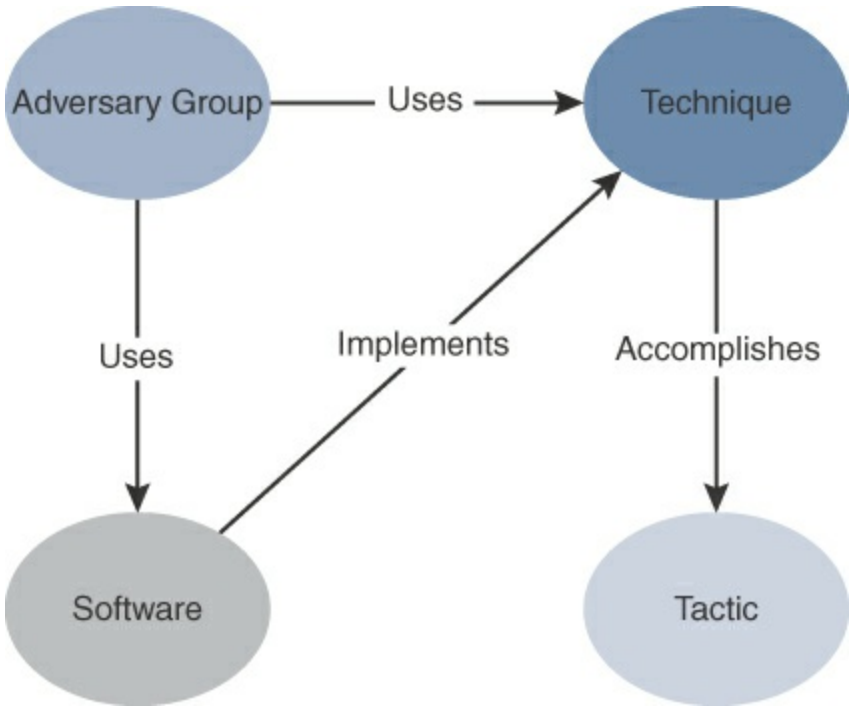

FIGURE 1-17 ATT&CK Model Relationships

PRE-ATT&CK Research

ATT&CK also includes behavior beyond what occurs during an attack within the PRE-ATT&CK research. PRE-ATT&CK covers documentation of adversarial behavior during requirements gathering, reconnaissance, and weaponization, before exploitation leading to access to an unauthorized network is identified. This is similar to the first few stages of the Cyber Kill Chain, but the results from PRE-ATT&CK are much more specific and based on recent real-world data. Figure 1-18 shows a matrix representing many of the steps of a potential attack against an enterprise.

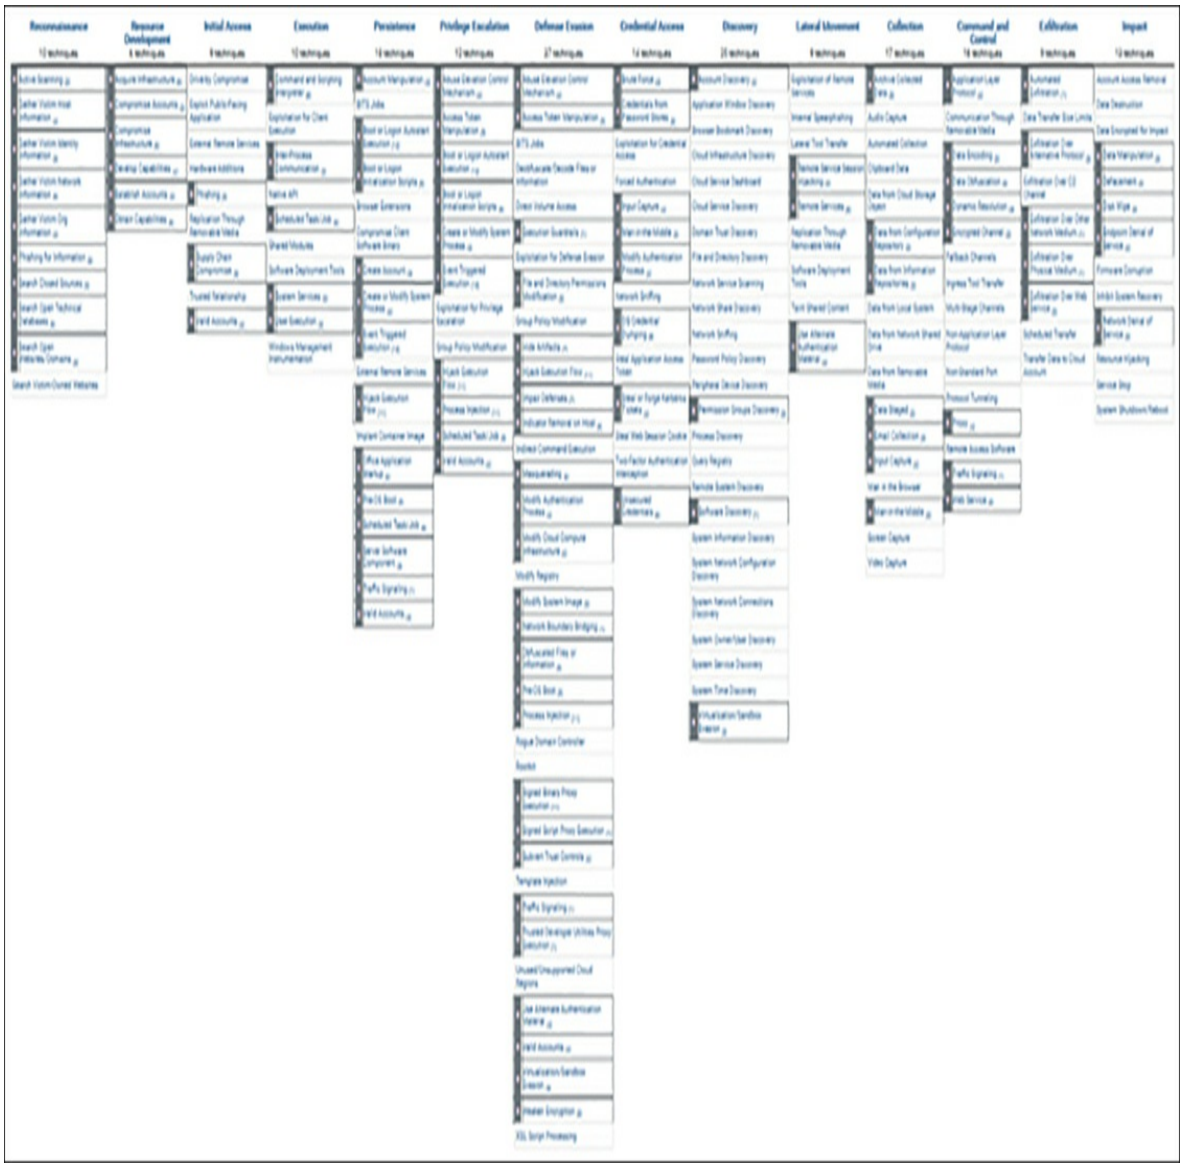

FIGURE 1-18 Part of the MITRE ATT&CK Matrix for Enterprise

Using MITRE ATT&CK

Common use cases for using ATT&CK details shown in Figure 1-18 include improving an organization's ability for detection and analytics based on threat modeling, providing a form of threat intelligence, emulating adversary behavior, and assessing existing security capabilities. This can occur based on working through the ATT&CK matrix and chaining together tactics, leading to a very powerful visual of the process that was taken by an adversary. Remember that a single attack will have multiple steps; the ATT&CK threat model offers a way to collect all of those steps and "chain" them into one larger attack. Figure 1-19 provides an example of how this chaining of attack steps could look as the ATT&CK model is leveraged to better understand the attack behavior of an adversary.

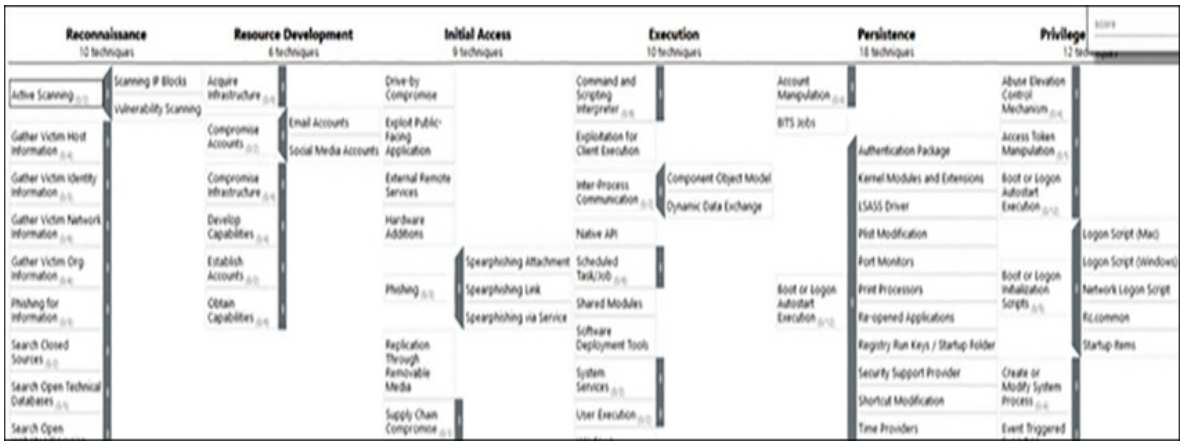

FIGURE 1-19 Chaining Together Attack Behavior Using ATT&CK Modeling

Note

Learn more about the MITRE ATT&CK model at https://attack.mitre.org/.

Choosing a Threat Model

The ATT&CK approach to threat modeling can seem overwhelming based on the level of provided detail, which is why simpler threat models such as the Cyber Kill Chain and Diamond models are still being used today. The best way to choose a threat model is to be outcome-focused and match the best approach based on obtaining your desired outcome. There are other industry models that can also be used to represent hypothetical attack and defend concepts. I highlighted the Cyber Kill Chain, Diamond, and ATT&CK models because they are very popular threat models used by many industry professionals, and each model offers a different approach to threat modeling leading to different types of value.

One challenge with these threat models is that they might be too hypothetical for some use cases, such as deciding which specific capability or service should be selected to respond to a potential threat or risk. Performing *tabletop exercises* is a more common practice used by many SOCs to hypothetically test security capabilities. These are meetings that walk through various attack situations without actually performing the attack and defend behavior. For example, a company could gather the head of desktop support, management, human resources, and the SOC team to go over what should occur if a cyberbreach such as one from the Cyber Kill Chain example is detected. Each team member could act out their role with the goal of testing if a process is in place and how it could work for that situation. I will cover how to perform a tabletop exercise in Chapter 4, "People and Process."

Threat models are useful for understanding what threats and risks could impact your organization. One key point to take away is how threats abuse vulnerabilities; without vulnerabilities, the threats are no longer a risk. This means that in order to deliver strong security, you must understand not only what are best practices for security and the potential threats, but also where you are vulnerable. The next topic to introduce is vulnerabilities and risk, which will also be the focus of Chapter 9.

Vulnerabilities and Risk

Vulnerabilities are weaknesses that can be exploited by an attacker. For

example, a vulnerability could be a door that is unlocked, opening the possibility of an intruder walking through it. A vulnerability could be a system missing a security patch, exposing a weakness that an attacker could digitally leverage to cause unwanted behavior. A vulnerability could even be an oversight in a business policy, such as setting a password policy that requires passwords to be a minimum of six characters and a maximum of eight characters and have a special character at the end. Why could this password policy be considered a vulnerability? If an attacker discovers this policy is being used by a specific organization, the attacker can adjust brute-force tools to search only for passwords that are six to eight characters and include a special character as the last character, dramatically reducing the complexity of guessing the password. Essentially, the brute-force attack does not have to attempt any passwords shorter than six characters or longer than eight characters and can assume that the last character is one of only a handful of special characters.

Looking back at the Cyber Kill Chain model, the hypothetical attacker in that threat model is exploiting a vulnerability to deliver a payload. This means that if a vulnerability doesn’t exist, the attacker can’t deliver the exploit. An example of removing a vulnerability is applying a patch that fixes the known weakness in code. The attacker can also be prevented from exploiting the vulnerability by using a defense tactic. That means the vulnerability continues to exist but the attacker can’t exploit it to accomplish his or her goal. An example of this is using a security tool such as an IPS to block exploitation behavior against a vulnerable system.

Endless Vulnerabilities

It is important to realize that all organizations have vulnerabilities and that there will never be a point when all vulnerabilities can be identified and patched. One reason for this is that networks and systems are constantly changing, which continues to introduce new vulnerabilities. Another reason is that technology changes cause errors to occur as new versions or capabilities are introduced. An even more common cause of vulnerabilities is how technology is used by people. The technology might not be vulnerable if used a certain way; however, people could misconfigure or misuse technology

outside of how it was intended to be used, causing a vulnerability. An example of a misuse of technology is placing a honeypot within the network and configuring an external connection. This is a bad idea because a honeypot is designed to be so vulnerable that it would attract a malicious element. Including an external connection would mean sources outside the network could use the honeypot to access the internal network, essentially turning the honeypot into a gateway for threat actors to breach the network.

Security tools are ideal for identifying and preventing threats from exploiting vulnerabilities. It would be ideal to fix all vulnerabilities; however, that will not happen for the previously described reasons. This means security tools can identify and prevent exploitation of a vulnerability until the organization is able to fix the vulnerability. In some cases, such as the IoT examples previously covered, a fix may not exist or be possible.

In many organizations, the SOC is responsible for managing vulnerabilities or partners with system support to oversee vulnerability management to ensure risk of exposure is reduced and ensure security technology and services are protecting where the most critical vulnerabilities exist. The key focus here is that vulnerability management is a risk reduction effort, meaning that it falls under the risk management services because vulnerabilities are a subset of the many things the make up what a risk is to an organization. Chapter 6 covers risk management concepts in more detail, while Chapter 9 provides a more focused conversation around vulnerability management best practices.

Technical vulnerabilities are not the only challenges a SOC is responsible to deal with in regard to risk management priorities. There are many business challenges and threats that introduce risk into an organization, which should be included under a SOC's risk management practice. Let's look at some of those business-related risks.

Business Challenges

Earlier in this chapter you learned about the challenges that organizations face that lead to a future compromise regardless of having a security program and security tools in place. Those challenges include the rapidly changing

landscape, lack of experienced security professionals, difficulty understanding which tools to choose, lack of security capabilities with some types of devices like IoT, and the list goes on. There are even more challenges to consider that relate to the business the organization is involved in. Examples of such challenges include adapting to changes in technology, complying with regulation requirements, and finding the right people to fill security roles.

A very common example of challenges with adapting to technology is how to leverage cloud services in a secure manner. There is a challenge to secure users accessing the cloud, which may require a secure Internet gateway. There is a challenge for software as a service (SaaS) cloud application, which may require a cloud access security broker (CASB). There are public cloud offerings such as Amazon Web Services (AWS) and Microsoft Azure, commonly referred to as infrastructure as a service (IaaS), that should be treated like datacenters within your environment, meaning they need to include their own layered security capabilities and services. Many organizations I meet with have a cloud-first business objective, yet they lack an understanding of how to move forward using cloud services in a secure manner. Technologies such as software-defined networks (SDNs) are pushing cloud business to new levels of interest across all organizations.

An example of a regulation challenge is meeting compliance requirements. Many organizations must comply with industry-specific regulations or risk incurring large fines and potential legal action if they are found out of acceptable compliance levels. For example, any organization leveraging credit card data must comply with the requirements of the Payment Card Industry Data Security Standard (PCI DSS), and any organization within the United States with access to healthcare records must comply with the requirements of the Health Insurance Portability and Accountability Act (HIPAA). Certain countries have compliance requirements such as only the government can own phone lines, which means it is against the law to have a private phone line. This requirement also means technologies such as voice over IP (VoIP) would be considered a violation of the government's ownership of communication. Chapter 6 covers compliance in much more detail.

People will always be a business challenge for every organization. Today's market has more security jobs available than people with the right skills that

can fill them. This means filling job requirements and employee retention is a major business challenge for organizations. It is becoming uncommon for security professionals to stay with an organization more than a few years, causing concerns for data privacy and increasing the need for job rotation to ensure there is always somebody ready to fill in when a key member leaves the organization. Chapter 4 dives into people challenges and best practices in more detail.

There are other business challenges that are specific business sectors. Organizations responsible for utility services and other critical infrastructure have to deal with being targeted by other countries' cyber militaries. Banks are responsible for keeping transactions secure and extremely fast, as milliseconds lost could cost hundreds of thousands of dollars to customers. Smaller churches have budget challenges for investing in proper security solutions since the majority of their profits are based on donation and volunteer work. You will not be able to cover every business challenge your organization will face; however, it is best practice to assign a team to deal with identifying and managing risk. That team should be the SOC. Risk management, however, should just be one of the handful of services a SOC can offer.

At this point, I have introduced the security operations center and covered at a high level the security capabilities designed to reduce the risk of threats. The rest of this book delves much deeper into all of these topics. To help you understand how a SOC can improve security, this book takes the approach of focusing on outcome, which means a focus on the services provided by the SOC. The next topic to address is what I find are the security capabilities offered by mature SOCs around the world. Each of these services will be the focus for the chapters ahead.

In-House vs. Outsourcing

While I have talked about threats and security capabilities, how does this all relate back to a security operations center? I defined a SOC as a centralized unit that deals with security issues on an organizational and technical level. This is accomplished through various types of services, which are directly handled by the SOC or outsourced. There are advantages and disadvantages

both to using in-house SOC services and outsourcing SOC services. The best decision for going either way will depend on your business needs. Some organizations may outsource using an on-demand or ad hoc approach, meaning they don't have the service internally but may add it if a situation ever demands it. Other times, services are ongoing, such as hiring an external company to handle any calls related to potential security incidents. A common example of using an on-demand approach would be an immediate need for digital forensics following a major security event. It is important to point out that most ad hoc approaches are the least effective and most expensive option versus the cost and value from dedicated or preplanned services. That decision will depend on your organization's need for each SOC service.

Table 1-1 shows a comparison of the advantages of using in-house SOC services and the advantages of outsourcing SOC services.

TABLE 1-1 Advantages of In-House SOC Services and Outsourcing SOC Services

In-House Advantages	Outsourcing Advantages
Knowledge of business	OPEX costs that can be spread out
Data stored internally	No conflict of interest
Cross-department correlation	Scalability and flexibility
Tailored requirements	Leverage other customer trends

Services Advantages

To summarize the advantages of in-house SOC services, it is all about having more control over and ability to customize the service being performed by the SOC. Members performing the service know the environment as well as all people involved with the incident. Having this knowledge can be extremely helpful for responding quickly and dealing with any internal politics. Using in-house services also builds a roadmap for growing the SOC by offering training and career development as the organization makes investments in security. In-house services can simplify compliance requirements for

protecting sensitive data and are flexible to adjustments in service goals and procedures. The reason for this flexibility is that in-house services are more familiar with the data, data owners, and business purpose of the data compared to an outsourced service that is just responsible for monitoring for security incidents. If changes need to be made to the outsourced service that are outside of the agreed-upon contract, meetings and changes in cost will be necessary before anything can be accomplished. In-house resources don't have these limitations regarding adapting to change.

While offering in-house SOC services is ideal for many organizations, sometimes the business model makes it cost-prohibitive to achieve. Advantages of outsourcing services might include a reduction in cost because the organization doesn't have to hire and pay employees, provide benefits such as healthcare insurance, or provide workspace and equipment. People are the most expensive asset for a business, and finding and retaining the right people that specialize in specific SOC services can be challenging in today's market. In my experience, many organizations that outsource SOC services have chosen to do so based on comparing the cost to build a team with the right skills and/or train internal employees to be part of a SOC against the cost of plugging in outside services that can provide impact much faster as well as have additional value like experience with other customer events. Other customer data can be a form of threat intelligence, meaning the service provider sees incidents with other organizations that help it to proactively prepare for future attacks against your organization.

Some organizations are required to outsource some SOC services due to conflict-of-interest situations. For example, an organization could have in-house digital forensics experts, but because those experts know the parties involved in any internal investigation, they may not be permitted to testify in court if the judge determines their relationship with the impacted parties might have influenced their investigation or might influence their testimony. In this situation, the organization should hand off the investigation to an outsourced digitals forensic team, so the results are considered an unbiased opinion.

Services Disadvantages

As indicated in the previous discussion of the advantages of outsourcing, the primary disadvantage of in-house SOC services is the cost to properly create and maintain them and the corresponding challenge of obtaining the required budget. Determining the cost to create an in-house SOC service is challenging because many uncertainties exist, such as the cost to find the right people, unforeseen changes in business, required training, and time needed to stand up the practice. This is why having an executive sponsoring the SOC is so important. Without this level of support, any of these cost hurdles may not be addressed, causing a SOC's service to become dysfunctional. I will provide recommendations for obtaining leadership support for the SOC in Chapter 2.

In situations where the total expected cost to stand up a SOC team is difficult to translate into dollar values, selling the idea to leadership is very difficult because you essentially are asking for a blank check. There are some formulas you can use to overcome this concern. One formula I will cover in Chapter 6 is used to calculate the likelihood of an incident and identify the cost to the organization if an event occurred. The formula will look at how to spread that cost over a period of time. The end result will be a hypothetical dollar amount of the cost of an incident and how its cost can be spread over time, giving leadership a value to compare against the risk based on how likely it will occur compared as well as how often it will occur. If that value is very high, such as millions of dollars for every security incident that could occur, it will make sense to price out SOC services to reduce the chance of a devastating incident occurring. If the per-incident cost is low, other options such as contracted or ad hoc services may be the best option for the business. These computations are typically hypothetical, whereas most contracted and on-demand services have much clearer associated costs, which makes outsourcing a much more attractive approach to parties responsible for funding the reduction of risk. If you plan to attain CISSP or CompTIA CySA+ certification, you will be required to know the formulas for calculating security costs.

As previously indicated, the largest disadvantages of outsourcing services are the limitations in the service provider's knowledge about the organization and the lack of flexibility because outsourced services are tied to a set contract. Changes to contracts have a cost, and it is possible the provider doesn't have resources available to support the changes being requested.

Another disadvantage of outsourcing is that it is common for external services to use tools that sit outside the network or place technology within the network in order to monitor the environment, limiting the organization's visibility of where those tools are installed. Further, many outsourced services also swap around personnel, meaning the resource assigned to a contract is not dedicated and not very knowledgeable of the environment he or she is assigned to protect. Finally, outsourced services commonly offer different tiers of coverage, the result of which is that if a top-tier organization experiences an incident, it will consume the top talent from the service provider, leaving limited support for other customers with lower-tier coverage. Table 1-2 outlines the disadvantages of in-house SOC services and outsourcing SOC services.

TABLE 1-2 Disadvantages of In-House SOC Services and Outsourcing SOC Services

In-House Disadvantages	Outsourcing Disadvantages
Cost	Limited business knowledge
People (hire/maintain)	External tools and data flow
Potential conflict of interest	Lack of communication
ROI concerns	Usually not dedicated people
	Limited customization
	Services are limited based on cost (e.g., tiered Gold, Silver, and Bronze services)

Hybrid Services

It is common for organizations to use a hybrid approach to obtain the maximum benefit of both in-house and outsourced SOC service approaches. I've encountered a lot of companies that outsource tier one support and train specialists within the organization to handle anything that is escalated above

tier one's capabilities. The value of this process is having generic requests outsourced, reducing the workload for the higher-cost internal assets. Internal assets know the environment and internal politics and have access to all internal tools. Tier one has a limited view of the organization, but that is sufficient for handling many of the first-level calls for support. This hybrid example also gives an organization the opportunity to grow its internal assets so eventually those people can handle the tier one support requirements and the organization can dissolve future outsourcing needs. The same approach can be used for specific services, such as first outsource incident response services until internal members are required or trained to cover this responsibility. Services can also be outsourced during demanding times such as during a major incident potentially requiring specific expertise.

SOC Services

Mature SOCs around the world tend to have in common a core set of security services. Those services might be in-house, outsourced, or even on demand, enabling the SOC to pull desired services when needed. On demand could be a contract that retains the services if they are ever needed or as simple as a saved services quote that can be executed upon at some future point. Regardless of the approach of delivery, there are services that every SOC needs to offer. To summarize those common SOC services, they can be defined as the following offerings:

- **Risk management:** Identifying and making decisions to deal with organizational risk. This pertains to managing any type of risk, from physically securing assets to patching digital vulnerabilities that exist within software. This can also apply to remediating weak policies and lack of education regarding security awareness within members of an organization.
- **Vulnerability management:** Identifying and managing risk from technical vulnerabilities. This commonly involves targeting vulnerabilities within software found on servers, laptops, and IoT devices. Most SOCs use vulnerability scanners and outside threat intelligence to identify vulnerabilities.

- **Incident management:** Responding to security-related events. This covers what actions the SOC takes when certain events occur, such as isolating systems, alerting team members, and implementing remediation steps to resolve the issue. Other subcategories that fall under incident management include incident response, incident investigation, and other incident-related topics. Technologies such as orchestration tools, artificial intelligence, and playbooks are becoming extremely popular to help assist SOCs with incident response services.
- **Analysis:** Analyzing various types of artefacts. This includes identifying characteristics, reverse engineering, vulnerability/exploitation analysis, root-cause analysis, remediation, and mitigation analysis. What separates an analyst focusing on analysis versus incident response is the type of required skills. Analysis uses tools such as IDA Pro to disassemble malware and understand how it functions. An analysis engineer can answer the question "Is this file malicious?" by running it in a sandbox to learn about its behavior. These skills are different from those of a SOC analyst responding to a potential breach.
- **Compliance:** Assessing and maintaining organizational compliance requirements. This can include legally obligated requirements such as HIPAA and PCI DSS compliance as well as organization-driven goals such as meeting a NIST or ISO standard, which are not required by law but could be seen as a required policy by the organization or its customers. The compliance service also prepares the organization for assessments and assists with gathering required information for outside parties validating an organization's compliance.
- **Digital forensics:** Gathering evidence post incident to determine the cause of the incident and prepare for legal action. There is some overlap in digital forensics, incident response, and analysis skillsets since all three include some form of understanding what malware or a malicious party has done. What separates digital forensics is the legal aspect regarding how evidence is collected. For example, if you manipulate a file during your investigation, you ruin any chance of using that evidence in a court of law (based on the concept of evidence contamination). Chapter 8, "Threat Hunting and Incident Response," addresses digital forensic concepts in more detail.

- **Situational and security awareness:** Providing the organization with awareness of its operational environment and potential threats. This includes education about critical elements that could impact the organization's goals, potential threats, and actions to reduce risk against operational risk and threats.
- **Research and development:** Researching the ever-evolving threat landscape, developing new tools and techniques, and modifying existing tools to improve effectiveness.

You might have some form of all of these services within your SOC, whether outsourced or covered using in-house employees. If your SOC lacks one of these services, I recommend at least obtaining quotes for on-demand services so that you have an option if that service is ever required. I already gave an example of how some organizations do not have digital forensics covered in-house and are sometimes forced to use an ad hoc approach following a major incident. It would be wise to know who to call before this service is needed versus trying to figure out who to contact while also dealing with the major incident.

Note

Regarding digital forensics, time is critical to success, meaning evidence will quickly be lost or contaminated if a proper investigation isn't launched. Imagine the time that would be lost if you had to start researching forensics services after normal business hours or on the weekend when most businesses are closed. According to the Ponemon Institute *2018 Cost of Data Breach Study* report, sponsored by IBM, in a consolidated sample of companies in various countries and regions, the mean time to identify (MTTI) a breach was 197 days and the mean time to contain (MTTC) a breach was 69 days. These numbers represent identifying and containing breaches, not the time to perform forensics post-incident response!

One important exercise is determining which of these SOC services your organization currently has and doesn't have. It is also important to figure out

how effective are the services that exist within the SOC. If those services are not effective, your organization can decide whether to improve those services based on the cost of the change compared against the business need. Remember that simply having a service doesn't mean it provides value to the organization. Just because you purchase a vulnerability scanner doesn't mean you have a formidable vulnerability practice. The best way to not only understand what services you currently have but also evaluate how effective they are is by using maturity models.

SOC Maturity Models

One common question I receive when speaking with SOCs is how effective they are compared to the industry. This is hard to answer since the goal of the SOC should be supporting the business objectives of its own organization rather than measuring itself based on the type of technology or processes SOCs in other organizations are using. My answer is that a better approach is to validate how effective the SOC is in supporting the goals of the business rather than how other businesses are running their SOCs. This leads to a conversation about SOC maturity and business relevance.

Assessing SOC maturity is the process of determining characteristics and features of the SOC, such as specific technologies and processes. The goal of creating SOC maturity models is to establish an understanding of the quality of SOC services as they currently exist and develop a roadmap with milestones for improvement. The results of a SOC maturity model can be used to predict and request budget for SOC services, establish criteria for rewarding improvement, and tracking the success of each SOC service as change is made. Assessing the maturity of the SOC is a critical step to running a successful SOC program and is a process all SOCs need to do on a regular basis.

SOC Maturity Assessment

How does assessing a SOC's maturity work? A basic maturity assessment evaluates how each SOC service is functioning. For example, consider a SOC

service that is used only when needed. Suppose a particular SOC has an incident response service that only a few people know how to do. They don't share knowledge about how they do it; they just do it when an incident occurs. Those incident responders don't perform the same steps for each investigation, and if they were to leave the company, that service would be lost. In this example, this SOC service would be considered acting in a low maturity state. To improve the maturity of this service, they first would need to formalize the steps they perform in their incident response service, so that it becomes a repeatable service with expected results. As the service matures, the process can be documented so others can follow the same steps to produce similar results. Eventually, certain steps could be automated and optimized with tools and processes to improve the speed and effectiveness of the incident response service. Each improvement milestone can be seen as an increase in maturity based on the maturity model being used. Figure 1-20 is CMMI Institute's view of modeling maturity. The CMMI Institute example uses five levels of maturity for assessing people, process, and technology.

	Level 1 Performed	Level 2 Managed	Level 3 Defined	Level 4 Quantitatively Managed	Level 5 Optimized
People	General personnel capabilities may be performed by an individual but are not well defined	Personnel capabilities achieved consistently within subsets of the organization, but inconsistent across the entire organization	Roles and responsibilities are identified, assigned, and trained across the organization	Achievement and performance of personnel practices are predicted, measured, and evaluated	Proactive performance improvement and resourcing based on organizational changes and lessons learned (internal & external)
Process	General process capabilities may be performed by an individual but are not well defined	Adequate procedures documented within a subset of the organization	Organizational policies and procedures are defined and standardized; policies and procedures support the organizational strategy	Policy compliance is measured and enforced Procedures are monitored for effectiveness	Policies and procedures are updated based on organizational changes and lessons learned (internal & external)
Technology	General technical mechanisms are in place and may be used by an individual	Technical mechanisms are formally identified and defined by a subset of the organization; technical requirements in place	Purpose and intent is defined (right technology, adequately deployed); proper technology is implemented in each subset of the organization	Effectiveness of technical mechanisms are predicted, measured, and evaluated	Technical mechanisms are proactively improved based on organizational changes and lessons learned (internal & external)

FIGURE 1-20 CMMI Institute Standardized Definitions of Maturity

SOC-CMM Model

Another SOC maturity model is the SOC-CMM model created by the master program of Lulea University of Technology (LTU). This model consists of five domains. The first three domains, Business, People, and Process, are all evaluated for maturity, while the remaining two domains, Technology and Services, are evaluated for both maturity and capabilities. The SOC-CMM model ranks maturity based on similar categories as NIST. The categories are non-existent, initial, defined, managed, quantitatively managed, and optimizing and will vary depending on the type of domain being evaluated. Figure 1-21 shows each of the five SOC-CCM domains and the 25 associated aspects.

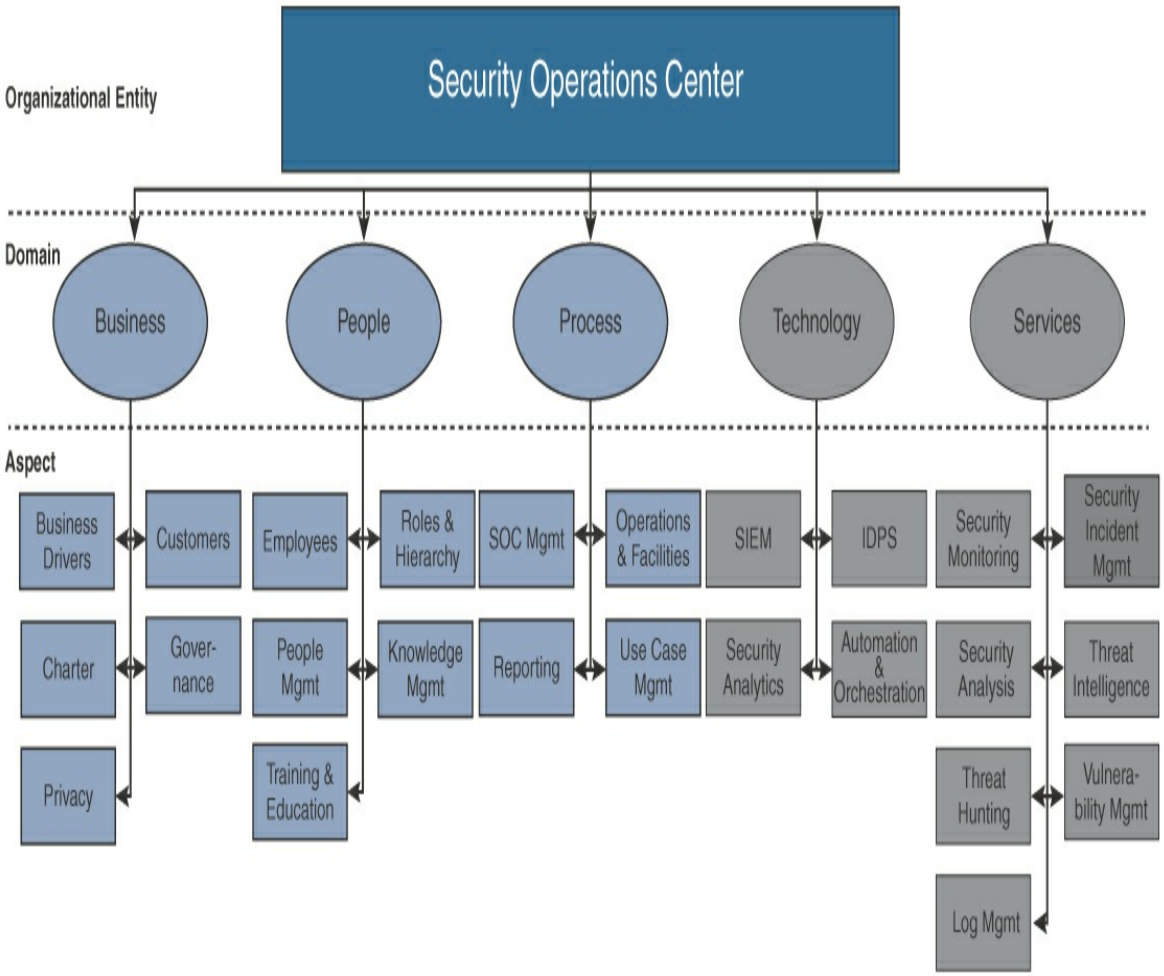

FIGURE 1-21 SOC-CCM Model Diagram

What is really beneficial about the SOC-CCM model is that it includes references to other standards, guidelines, and frameworks for its recommendations to enhance maturity of a SOC service or capability. For example, the subcategory for ID.AM-1: Physical devices and systems within the organization are inventoried, includes references to CCS CSC1, COBIT 5, ISA, and NIST. The SOC-CCM model is free to use and can be downloaded from https://www.soc-cmm.com.

ISACA COBIT 5 Process Assessment Model

Another popular industry maturity model by ISACA is the COBIT 5 Process Assessment Model (PAM) based on the ISO/IEC 15504 standard for

performing a process assessment (COBIT). According to COBIT, there is a six-point system of scoring, 0 through 5. Level 0 means a service doesn't exist or is incomplete. A level 1 service represents an ad hoc capability, meaning that it at least achieves its purpose even though it is not an effective approach for the business. Once that capability is repeatable, it moves to level 2. This requires the capability to be managed with expected results. Once the repeatable capability is fully documented and defined as a formal process, it moves to a level 4. This means the process can be used throughout the organization rather than just by a specific group or individual. As the process is executed, it can be measured and managed for improvement. Once checkpoints for success and failure are established and maintained, the maturity becomes level 4. Finally, the process can start to be optimized for ongoing improvement, making it a level 5.

This model should look similar to the CMMI Institute model. Figure 1-22 breaks down the COBIT scoring model. The details of how to perform a COBIT assessment can be found on the ISACA.org website.

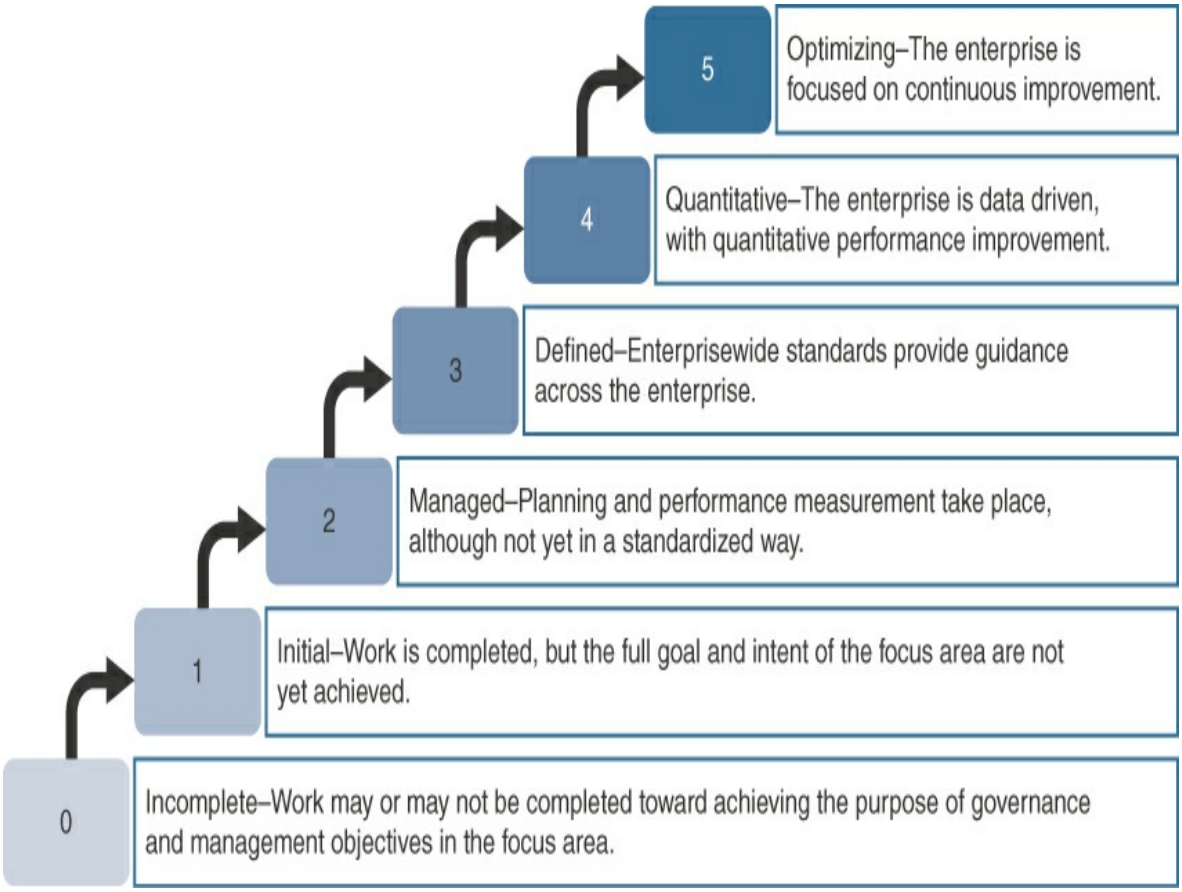

FIGURE 1-22 COBIT Capability Scoring

SOC Program Maturity

The NIST, SOC-CMM, and COBIT models are great resources for evaluating the maturity of specific features of a SOC service. A more general approach to modeling the maturity of the SOC is to evaluate the entire SOC program. Using this approach consolidates the associated people, process, and technology for the entire SOC rather than breaking up each SOC service as is done in the NIST, SOC-CCM, and COBIT models. This approach starts with defining the most basic SOC maturity level:

- **First-generation SOC:** This level is a SOC that just monitors device logs, which means it has limited coverage based on the data that is monitored. A basic SOC has limited data retention capabilities and is not effective at responding to security incidents. The basic SOC has a few security tools sending event logs to a centralized tool such as a security information and event management (SIEM) system, which is what the SOC uses for all security awareness. The services expected from a first-generation SOC are some form of risk management and limited continuous monitoring for security incidents.
- **Second-generation SOC:** A second-generation SOC leverages data correlation and consolidation to turn log data into security events. This simplifies monitoring, dramatically improving incident response. This can also lead to developing a tracking system to manage events and eventually playbooks that represent the proper response to a specific type of incident. A second-generation SOC offers more advanced risk management and a more mature incident response service. This level of SOC can be effective but is still very reactive based.
- **Third-generation SOC:** A third-generation SOC has more experience with SOC capabilities and is able to offer more services, such as vulnerability management and compliance. This level of SOC has assessment services looking for potential weaknesses and areas that violate policy as well as required compliance. A key point is that a third-generation SOC has moved from reactive to proactive security practices, because it has services that are designed to prepare for attacks before they

happen by reducing potential risk. This also means the SOC can develop better playbooks and perform lessons learned exercises to better prepare for future attacks following a security incident.

- **Fourth-generation SOC:** A fourth-generation SOC leverages the latest SOC technologies and services. This level of SOC further tunes tools and expands visibility to other networks through threat intelligence, reputation security, and cloud services. It enhances data correlation by using artificial intelligence, not only improving decision making but also supporting development of new security rules and playbooks based on live data. A fourth-generation SOC uses data sources such as NetFlow and packet captures to deliver network forensics services. It not only is proactive but also continuously measures results and sets growth and maturity goals. I provide recommendations and guidance throughout this book to help you increase your overall SOC maturity to become a fourth-generation SOC.

Figure 1-23 represents the overall SOC maturity breakdown.

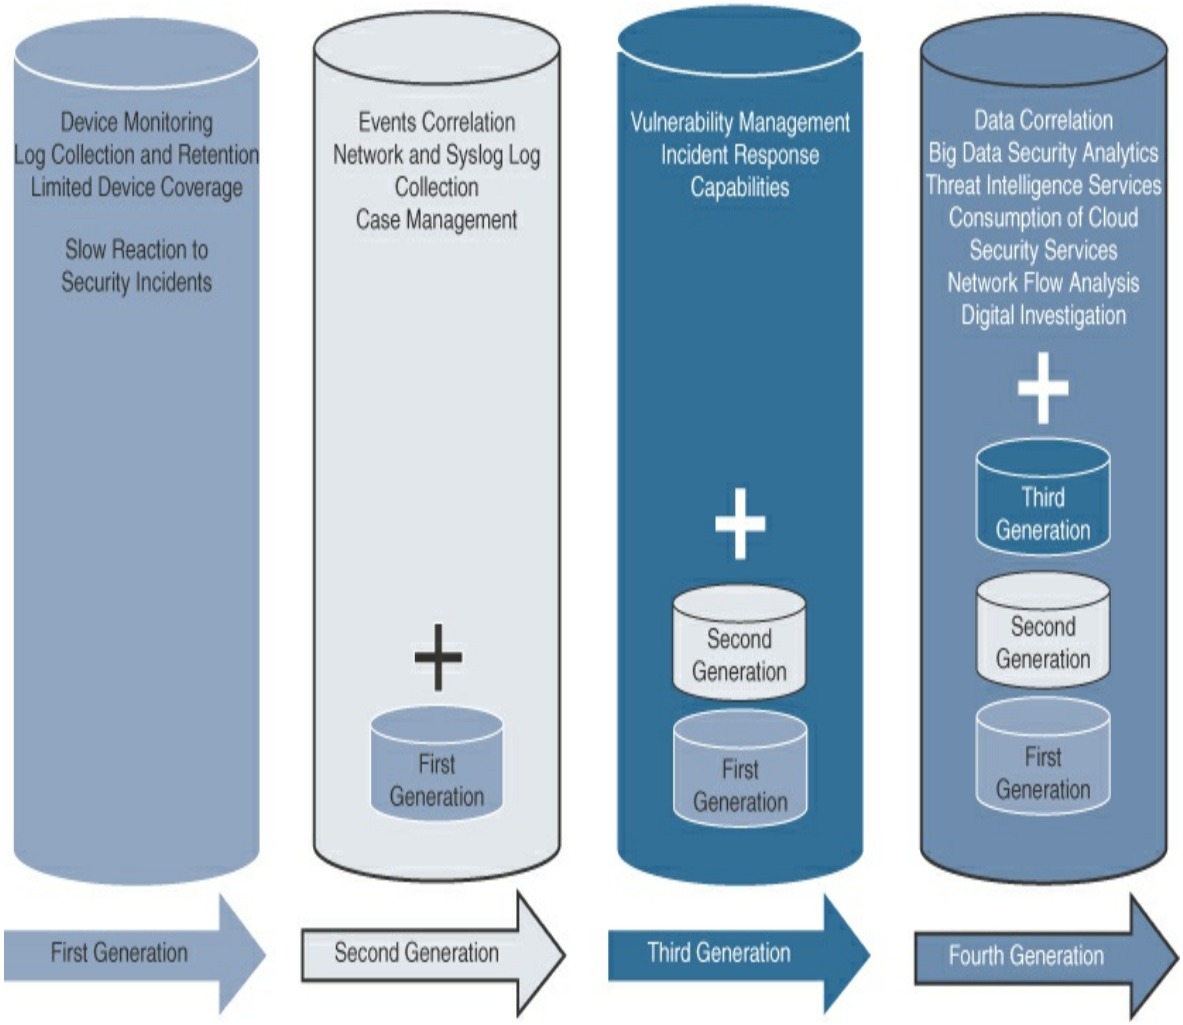

FIGURE 1-23 SOC Maturing Models

Within each SOC generation are services that can be broken down and evaluated using the CSSI Institute, SOC-CCM, or COBIT models or other standards, guidelines, or frameworks. You may find that your SOC has a basic forensics service that you would grade as a first-generation SOC service because it is ad hoc, while the incident response program is much more mature based on how your SOC has invested in that service. You may find that your SOC works with a different vulnerability management team rather than running the vulnerability management program within the SOC. This approach to vulnerability management means the first-generation SOC service is combining its resources with the desktop support team to move its capabilities and services to the third-generation performance arena. I recommend grading your SOC as an operation as well as evaluating each of

the common SOC services (described earlier in this chapter) in this fashion using a standard, framework, or guideline such as NIST, SOC-CCM, or COBIT.

Establishing the maturity of your SOC services and developing milestones for improvement are critical steps to formalizing your SOC program. By understanding the status of maturity, your SOC program can assess what level of improvement is required for the business. This provides a foundation for requests for resources and developing a reward behavior to encourage a healthy SOC environment. Roadmaps lead to change, which is how a group of people responsible for security can develop their practice into a responsible SOC.

One challenge regarding assessing the maturity of a SOC is assessing the specific security capabilities that exist within the network of the organization protected by the SOC. Network standards, guidelines, and frameworks can be helpful, but in the next section I provide a methodology that I have used to develop customized capability assessment maps and goal ranking. This approach helps with making decisions about what is the best security capability and/or service to invest effort for improvement of the organization security posture and SOC's services.

SOC Goals Assessment

As you have read, there are very useful standards, guidelines, and frameworks available in the industry that provide recommendations for best practice for security. I highlighted some of the most popular options in this chapter as well as identified some limitations to using those resources. Those limitations include the fact that they are not updated at the same pace as technology changes, they are generic by design to accommodate various types of industries, and they don't provide a method to prioritize which changes are most important to your specific business. This leaves decisions such as prioritizing which investment in security capabilities and services to make first and where consolidation of existing security capabilities could occur up to the organization. Imagine an organization that has a budget for one security tool but needs both a next-generation firewall (NGFW) and a web application firewall (WAF). Which one should the organization invest in

first?

I have collaborated with some fantastic security architects at Cisco to develop a methodology that complements the value of standards, guidelines, and frameworks by providing a customized list of goals with capability assessment diagrams showcasing gaps in security as well as potential areas for consolidation. Having a list of goals simplifies where investments could be made if those goals align with the focus of the business. Ranking the goal list against how important each goal is to the business provides a clear view of how important each goal is so decisions can be focused on top-priority goals. Developing capability maps provides a clear viewpoint of where a gap or overlapping capability exists so technical and nontechnical people can understand why a change is needed. You may be wondering, why not just jump into assessing your security capabilities? Why first review the goals of an organization and SOC? Consider again the scenario where an organization with a limited budget has a need for multiple security tools but does not know the order in which to invest its money and time. By first establishing and ranking the goals of the business and SOC, the organization can align security tools to that list to provide a roadmap for how each security need should be addressed and in what order.

Defining Goals

The first step to the SOC assessment methodology is to define the goals of the business. This is critical to ensure that all other goals for the SOC align directly to the goal of the business. If a business is focused on delivering online video games, for example, the goal for the SOC should target supporting delivering games in a secure manner as well as protecting all systems associated with delivering the online gaming service. If the business is a school whose primary focus is to offer the most modern online learning environment, then the goal of the SOC should be to provide a reliable network and digital learning resources. Determining the goals for the organization should be done at the executive level. The SOC then builds its IT goals based on the goals of the organization. It is important to have an executive sponsor validate that the business goals are correct, to ensure all goal planning for the SOC is relevant for the organization. The business goals

can also be seen in the SOC's mission and scope statements, which I will cover in Chapter 2.

> **Note**
>
> I highly recommend confirming the SOC team's understanding of the business goals with their SOC executive sponsor or other executive to ensure the SOC's goals are properly aligned with the business.

Once the business goal(s) are established and validated by leadership, you can develop the goals for the SOC based on meeting the business goals. These goals will be based on the people, process, and technology offered by the SOC and part of one or more SOC services. I call these the *IT goals* for the SOC. IT goals align with business goals, meaning the technology helps accomplish the business goals. Revisiting the previous examples, an obvious IT goal for the online gaming company would be to protect the online gaming service from external threats, because the online gaming service is needed for the business to be profitable. An IT goal for the school offering online classes would be to provide a reliable network and identify any systems that could impact the network's performance. Once again, this should make sense, because a business goal for the school is to provide a modern learning environment, which would heavily depend on a reliable network. IT goals can be documented and transformed into company policy, making them a requirement to follow. IT goals should be very high level and explain the vision for the goal without explaining details on how to accomplish the goal. This helps keep the goal relevant regardless of specific changes to people, process, and technology. The aim of IT goals is to provide the vision for the SOC.

Because IT goals are high level, the details for delivering each goal are defined in one or more IT processes. An IT process establishes the step-by-step approach to delivering a goal. A process for the online gaming company would be to monitor for spikes in online requests with the intent to identify potential denial-of-service (DoS) attacks. All of the steps and tools used to accomplish the IT goal to protect the online service, as well as who is responsible to ensure this occurs, would be explained within one or more IT

processes. If somebody asks why a process exists, it could be explained by showing how the process directly aligns with an IT goal that also aligns with a business goal. An IT process for the school would be the steps used to monitor for network outages. This could include what tools are used to monitor network traffic usage, what to do when an outage occurs, and which people are responsible for providing this service. In both examples, these could be one or more IT processes aligned to a specific IT goal or more than one IT goal. Figure 1-24 represents how business goals, IT goals, and processes should align. This model is similar to the CompTIA CySA+ view of policies, standards, and procedures alignment.

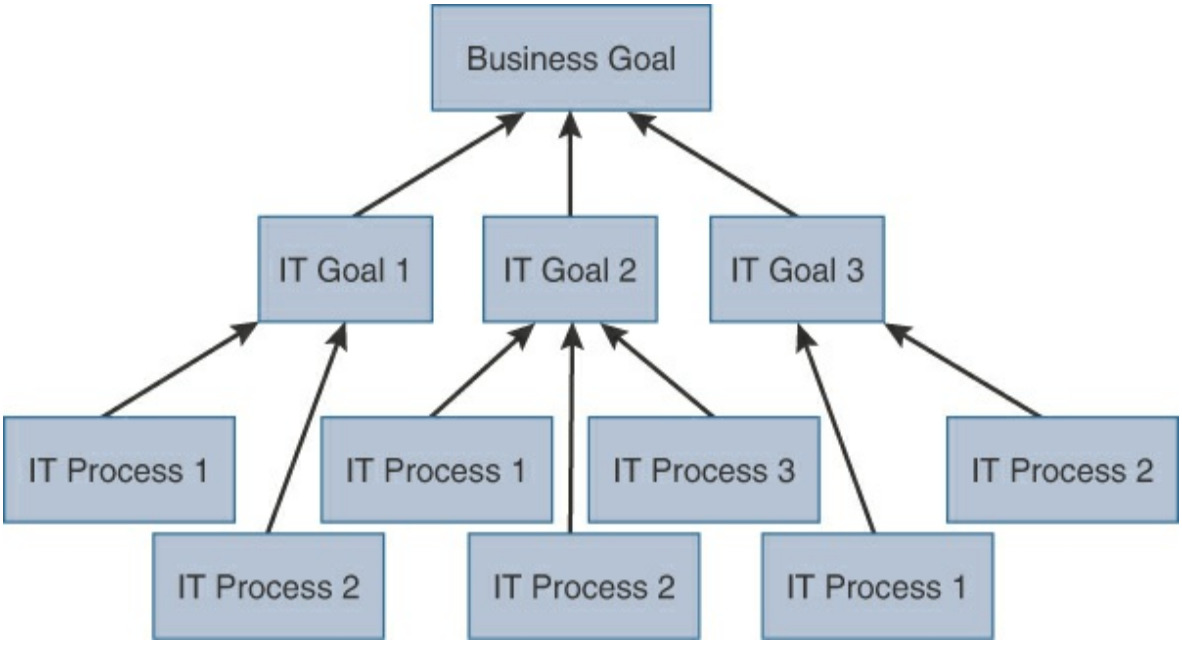

FIGURE 1-24 Business Goals, IT Goals, and Process Alignment

The following is a short summary of the SOC goal assessment process:

1. Meet with the SOC executive sponsor or business leadership to confirm business goals.
2. Develop SOC goals, also called IT goals, that support the mission for the business goals. IT goals must align to the business goals.
3. Create IT processes representing the more detailed documentation of how to execute an IT goal properly.

4. Identify any missing people, process, or technology within an IT process.

SOC Goals Ranking

Looking back at my examples, I mentioned the IT goal of the online gaming company would include protecting online resources, which will require DoS and web exploitation defense technology capabilities and services. I mentioned the goal of the online school would be to protect the performance of the network, which would require IT goals for vulnerability management, incident response, and continuous monitoring of the network. Assessing a SOC in this manner helps justify the specific tools and processes that are put in place, which later can be evaluated for maturity using any of the standards, guidelines, or frameworks previously reviewed. Table 1-3 and Table 1-4 outline examples of the high-level business and technology goal mappings of the online school and online gaming company, respectively.

> **Note**
>
> The following examples are similar to results from customers I have performed this work for in the past. Every organization will have different goals and priority ranking.

TABLE 1-3 Online School Goal Mapping

Policy Goal	Priority
Data privacy	1
Support learning environment	1
Services for students and faculty	2
Reduce time to detect and correct events	3
Learning outside of walls	3
Reputation of school	4

Road warrior support	5

TABLE 1-4 Online Gaming Company Goal Mapping

Policy Goal	Priority
Providing 99.9% uptime	1
Company reputation	2
Quick software release and updates	3
Protect customer data	3
Remote worker support	4
Employee retention	5
Vendor partnerships	6

Notice in Tables 1-3 and 1-4 that some of the policy goals are ranked with the same priority. This is ok as long as all goals are not ranked as a top priority. The reason behind the ranking is to determine how to prioritize goals. It is recommended to first list out all business goals and rank them based on input from various parties, which could include desktop support, legal, finance, facilities, security, HR, and leadership, so everybody is on the same page regarding how important each goal is to the organization. I recommend validating the results of how business goals are ranked with the executive sponsor of the SOC.

Table 1-5 and Table 1-6 demonstrate performing a technology goal assessment for the online school and online gaming company, respectively. Notice that the technology goals align with how the business goals are ranked. If the school ranks data privacy as a top concern, data loss prevention technology should also be a top technology goal. Technology goals are just one example of an IT goal. Other examples could focus on people and process elements of the SOC.

TABLE 1-5 Online School Technology Mapping

Technology Goal	Priority

Data loss prevention	1
Network uptime	1
High availability	2
Segmentation	2
Endpoint security enforcement	3
Authorized vs. unauthorized cloud services	4
User training	4
Least privilege enforcement	5
Vulnerability management	6
Configuration management	7

TABLE 1-6 Online Gaming Company Technology Mapping

Technology Goal	Priority
Denial of service defense	1
High availability	1
Network monitoring	1
Application and WAF needs	2
Configuration validation	3
Data loss prevention	3
Least privilege enforcement	4
VPN and remote routing	5
Endpoint security enforcement	6
Internal segmentation	7

Developing the technology goals and ranking them should be much easier than performing the same process for the business goals, since the previously created business goals influence and narrow the scope of the technology goal conversation. It is critical that the business goals are developed and ranked first for this purpose. The order of how goals are assessed matters! I have

performed these assessments for hundreds of customers and many times find, for example, that the organization believes they need a new firewall before the assessment but discover through the assessment that their business goals highlight a larger need for investment in areas where they completely lack capabilities and services. In the example of the online gaming system, a WAF and DoS technology should be a higher priority than a new NGFW or segmentation technology. I opened this section with posing the question of whether a company should choose a WAF or an NGFW. For the online gaming company example, the answer would be a WAF, based on the results of the business goal and IT goal ranking. If another organization went through this process, they may find the NGFW is the better investment based on the results they come up with.

Note

You might be wondering why the SOC would be involved with purchasing security tools (some customers I speak with think of a SOC as only being responsible for responding to security incidents). Among the SOC services I commonly find in mature SOCs around the world, one key service is research and development. This service focuses on researching and evaluating security tools so that the expected security experts within the organization, the SOC, are choosing the best technology match for the organization rather than another team within the organization. I highly recommend the SOC's involvement in any evaluation of a tool that will be used by the SOC. I have seen hundreds of times an organization procurement office holding full responsibility for acquiring tools and selecting a tool only based on price. Trust me when I say that approach tends to lead to a poor decision. Chapter 10 covers evaluating whether to create your own tools, take advantage of open source tools, or step up and pay for enterprise options.

IT goals should be high level, while IT processes should be specific to which people, process, or technology is being offered. An IT process should include step-by-step instructions for how something is performed, the details about

the tools involved, and who is responsible for doing the service. One or more IT processes can align to a single IT goal. For example, one IT process may be the specifics around how the WAF is configured, while another IT process may cover how the WAF is monitored. In both IT processes, details about who does the work, when the work is done, and how the work is done are just some of the information that should be included. IT processes can also be ranked similarly to how IT goals were ranked in the previous examples.

Threats Ranking

A third aspect that can be evaluated outside of goals is looking at top threats to the organization. This can complement the threat modeling exercise I covered earlier in this chapter. Ranking threats can justify how the technology goals align to the business goals as well as validate if capabilities and processes exist to combat a situation of high concern. The types of threats and expected action of threats can be pulled from threat modeling exercises. Looking at the online gaming company example, it would make sense to consider threats against their goal of 99.9% uptime to be a top priority for the business. This means threats that can take down their service would be ranked at a 1. Table 1-7 and Table 1-8 show examples of doing this for my example online school and online gaming company.

TABLE 1-7 Online School Threat Mapping

Threat Concern	Priority
Denial of service	1
Data compromise	1
Lack of visibility	2
Exploitation	3
Stolen accounts	3
Lateral movement	3
Malware	4
User error/Layer 8	4

Unauthorized devices (hubs/routers)	5
Process violations	6

TABLE 1-8 Online Gaming Company Threat Mapping

Threat Concern	Priority
Denial of service	1
Stolen user accounts	1
Website exploitation	2
Software compromise	2
Stolen accounts	3
Configuration errors	4
Endpoint malware	5
Internal threats	6
Unauthorized devices (hubs/routers)	7
Partner risk	7

Having these ranking and alignment results provides a solid foundation for understanding what goals and processes are important to the SOC based on how they align with the overall business. Keep in mind that if certain groups are not involved with the process of creating and ranking goals, you may have to repeat the assessment with the missing parties to get a true balanced opinion of the goals and ranking. I have seen situations where desktop support was not involved with these decisions and later pulled away their part of budget due to not supporting what was being proposed as the next step for the organization’s security investments. I highly recommend including at least a director or higher to represent business goals and managers from each key service within the organization to obtain maximum impact from this exercise. Typically, this person is the executive sponsor of the SOC.

SOC Goals Assessment Summarized

Once this work is complete, your organization will have a good idea of what parts of the organization are top targets for investing for SOC services and capabilities in relation to accomplishing goals. What is still missing is identifying specific gaps in capabilities as well as areas where consolidation could occur to create new budget. In the next section, I will focus on a methodology used to evaluate security capabilities based on capability models.

The following steps summarize the ranking of goals:

1. Establish the business goals and align SOC/IT goals and IT processes.
2. Rank the business goals and IT goals according to their importance.
3. Debate ranking and validate with different groups within the organization to ensure all voices are heard.
4. Apply gap analysis against each item ranked, starting with the most critical.
5. Turn gap analysis into a three- to five-year plan to enhance SOC capabilities.

The two assessment programs I just covered will result in a list of goals that align with business. The next step of assessing the SOC and developing a roadmap for improvement is identifying gaps or overlapping capabilities within an IT goal. For example, if an IT goal is to reduce the risk of vulnerabilities, how do you know which capabilities exist to accomplish that goal? To assess people and processes, you will need to do a tabletop exercise, covered in Chapter 4. To assess technology capabilities, I recommend using a SOC capabilities assessment.

SOC Capabilities Assessment

The previous section discussed assessing SOC business goals, IT goals, and IT processes. You learned how to rank those goals to get an idea of how to prioritize focus for investments into the SOC services and capabilities. What I haven't looked at is how to validate what SOC capabilities exist so specific areas of investment can be identified. For example, what capabilities does the

online gaming company have to accomplish its goal of 99.9% uptime for its games and services? What capabilities could the online school use to ensure the student network is performing properly? Should the focus be on the edge, within the network, or on the hosts? What are the best investments to make, assuming there is limited budget to use for the SOC?

In this section, I will look at developing capability maps that can be used for consolidation of similar capabilities as well as identifying gaps. Consolidation leads to having more budget for future tools, people, and services. Identifying gaps helps SOC and non-SOC members understand what types of services and capabilities need to exist to improve maturity of a SOC service. Industry standards, guidelines, and frameworks can also provide recommendations for improvements, but, as previously discussed, they will not be specific to your organization. Capability maps are customized to your environment, but they can include missing capabilities that are found within industry standards, guidelines, and frameworks if desired.

Capability Maps

Creating capability maps involves identifying the different parts of your network and mapping how users and systems interact within that environment. To better understand this, let’s walk through mapping an average organization’s branch network. I will start with the end user and attempt to define what security capabilities or services exist to protect that user’s laptop. Does the organization standardize on the same hardware and software or do employees bring their own technology? Do the employees have rights to install or modify software? You may or may not care about the level of detail depending on your business requirements and what you want to include in your capability map. Looking at the online school example, let’s say the school does not care which operating system is used by students, meaning students are permitted to bring their own system. This would be different from an organization such as the online gaming system that requires all employees to only use their corporate-issued device on the company network.

Capability Map Example: Endpoint Security

For this capability map example, I am going to disregard system and host details so that I can purely focus on security technology capabilities. Capability maps could include other concepts, such as services running on a host, but to keep this example simple and focused, I will look only at security capabilities. Starting with a focus on end-user systems, those computers will have antivirus software to prevent known malicious files. What about anti-malware that is used for more advanced unknown threats, which has the ability to detect threats based on behavior that bypasses antivirus? Is there any content filtering or reputation security installed on the host for protecting users from accessing external known malicious sources? Do host systems have firewalls enabled or host-based intrusion prevention (HIPS)? You can use industry standards, guidelines, or frameworks covered earlier in this chapter as a method to identify industry best practices for security capabilities for endpoint protection and compare those recommendations against existing capabilities. For example, NIST has a set of recommendations for what security capabilities should exist on a host system. Which specific standard, guideline, or framework you use should be based on which you find the most useful for your organization.

The key to this example is to list capabilities that are different or at least provide additional checkpoints for threats. This leads to a defense-in-depth approach to security. Redundant capabilities such as two antivirus programs can be identified as being a repetitive capability, represented by the same icon twice in the capability map. Figure 1-25 is an example of mapping out host-based security capabilities. Remember that for this example, I am only focusing on the capabilities that exist within the host. SOC services such as endpoint vulnerability scanning and managing user privileges are not covered but are just as important as IT capabilities.

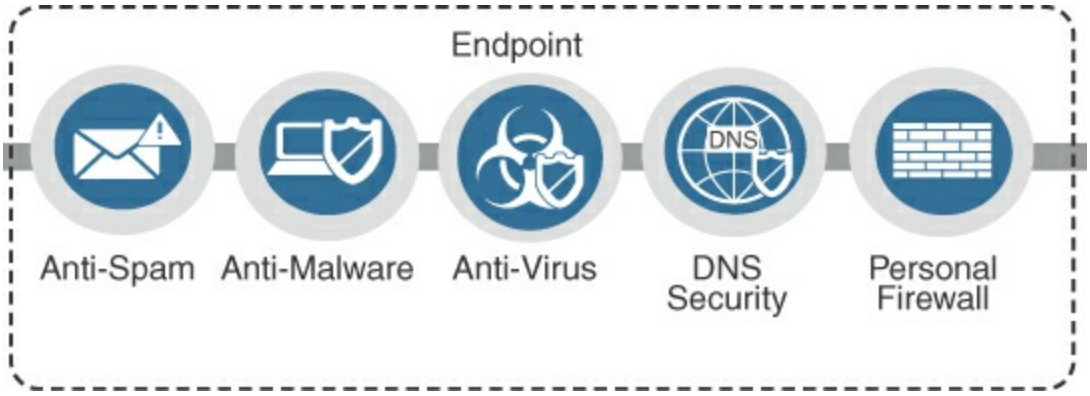

FIGURE 1-25 Host Capabilities

Capability Map Example: Endpoint OS Security

You might want to dive deeper into assessing capabilities within an endpoint because the last capability map shown in Figure 1-25 glossed over it. For the next example, a desktop support member might be concerned about security for the operating system and want to assess the host OS for capabilities and services. You could develop a separate capability map for the host OS to serve this purpose. Once the host OS capability map is complete, it could be referenced by other capability maps, so you don't have to repeat the process every time a host OS is listed in a capability map. Figure 1-26 is an example of creating a host-specific defense security diagram. These host OS capability suggestions come from NIST SP 800-30, *Guide for Conducting Risk Assessments*.

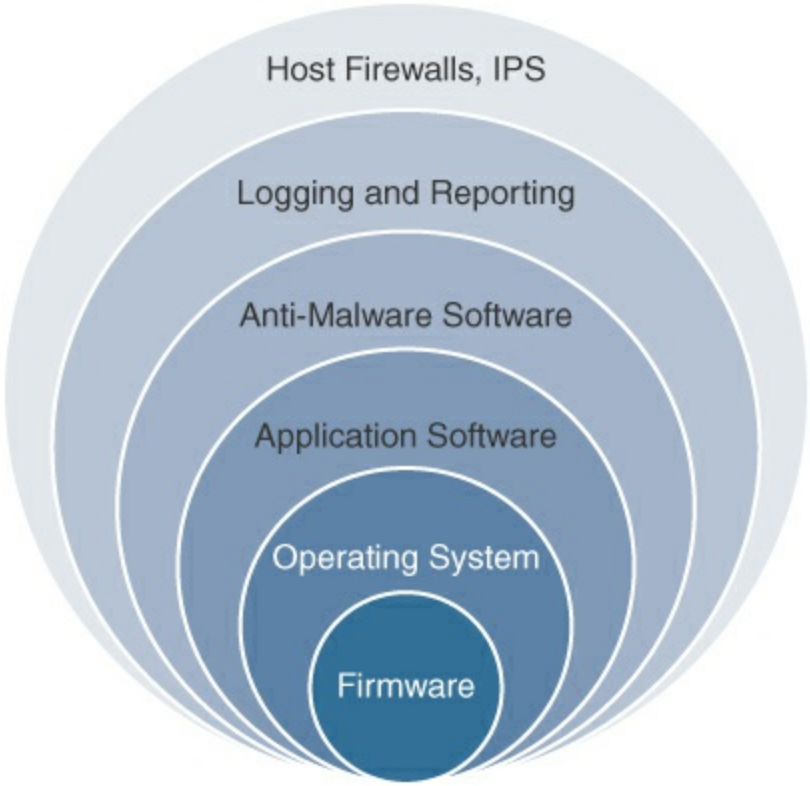

FIGURE 1-26 Host-Specific Defense Diagram

Note

I recommend starting the capability mapping at a high level to establish a general understanding of information flow and how it interacts with existing capabilities. As you identify areas of interest, you can create more detailed capability maps for that specific area. For example, you can first group "gateway security capabilities" as a single checkpoint and later map out all of the capabilities that would be within a secured gateway such as firewall, IPS, WAF, etc.

Capability Map Example: Network Security

For the next part of this capability mapping exercise, I will look at how users connect to the network. I need to think about how any user connects to the network and what the network does to reduce the risk of compromise. I

should also think about any associated processes and policies, such as requiring any device that connects to be scanned for vulnerabilities before being permitted access. Here are some questions to consider for this part of the assessment exercise:

- Is there a form of access control in place?
- How is authentication handled?
- What is the policy for permitting certain levels of access to people and systems?
- How are systems monitored that can connect to the network?
- Is the hardware profiled and are certain hardware types granted more access than others?
- Is segmentation enforced based on device and user types?
- What happens once the system is granted online services?
- Is user behavior monitored for malicious behavior?
- Who is responsible for these services?
- What security is in place as traffic leaves the user's system and goes out to the Internet?

What is key is understanding the flow of traffic and different capabilities that are in place. It is recommended to not turn this process into a product conversation, but rather focus on the capabilities and services regardless of how they are provided. For example, one product such as a "Next-Generation Firewall" could have multiple capabilities. Many "Next-Generation Firewalls" found in the industry are providing multiple capabilities including firewall, application firewall, intrusion prevention, anti-malware, and so on. Figure 1-27 is an example of a NGFW broken down into a capability flow. Notice that traffic is first going to hit the firewall capability and then be filtered through the application controls if traffic is permitted through. That traffic will next be scanned by the IPS for threat behavior. Finally, any files that are permitted through will be evaluated for malware. The flow is important, so you understand at what part of the kill chain the defense should be taking effect. Once again for this example I am just focusing on capabilities. Factors such as who should manage this, what policies should be

enforced such as what content should be filtered, and many other items are not included for simplicity purposes.

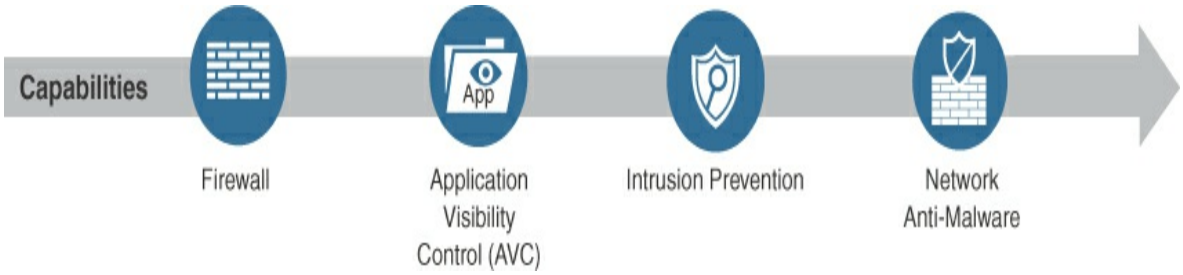

FIGURE 1-27 Next-Generation Firewall Capabilities Flow

> **Note**
>
> You can create boxes around capabilities that are included in a single existing product to improve the understanding of what is being documented. I find confusion can occur regarding the number of products needed or existing if the reader believes each capability represents a product. The reality is that most of the capabilities will be grouped into different products, meaning a few multifeatured products will make up all of the capabilities represented in a capabilities map.

Capability Map Example: Branch Network

I recommend that you keep mapping of capabilities vendor agnostic, but you could label them for documentation purposes if desired. The key is to develop a diagram of what capabilities and services exist even if more than one technology is providing the same capability. An example could be an NGFW with application-layer firewall capabilities and web proxy both having the capability to filter traffic. Figure 1-28 is an example of what an organization's branch network could look like after mapping out capabilities from a user to how that user's traffic goes out to the Internet. This example isn't focused on only what actually exists but looks at all of the possible capabilities, including what doesn't exist but should be part of the map, based on industry standards, guidelines, and frameworks like NIST and ISO. In this

example, I am only focusing on technical capabilities for simplicity purposes.

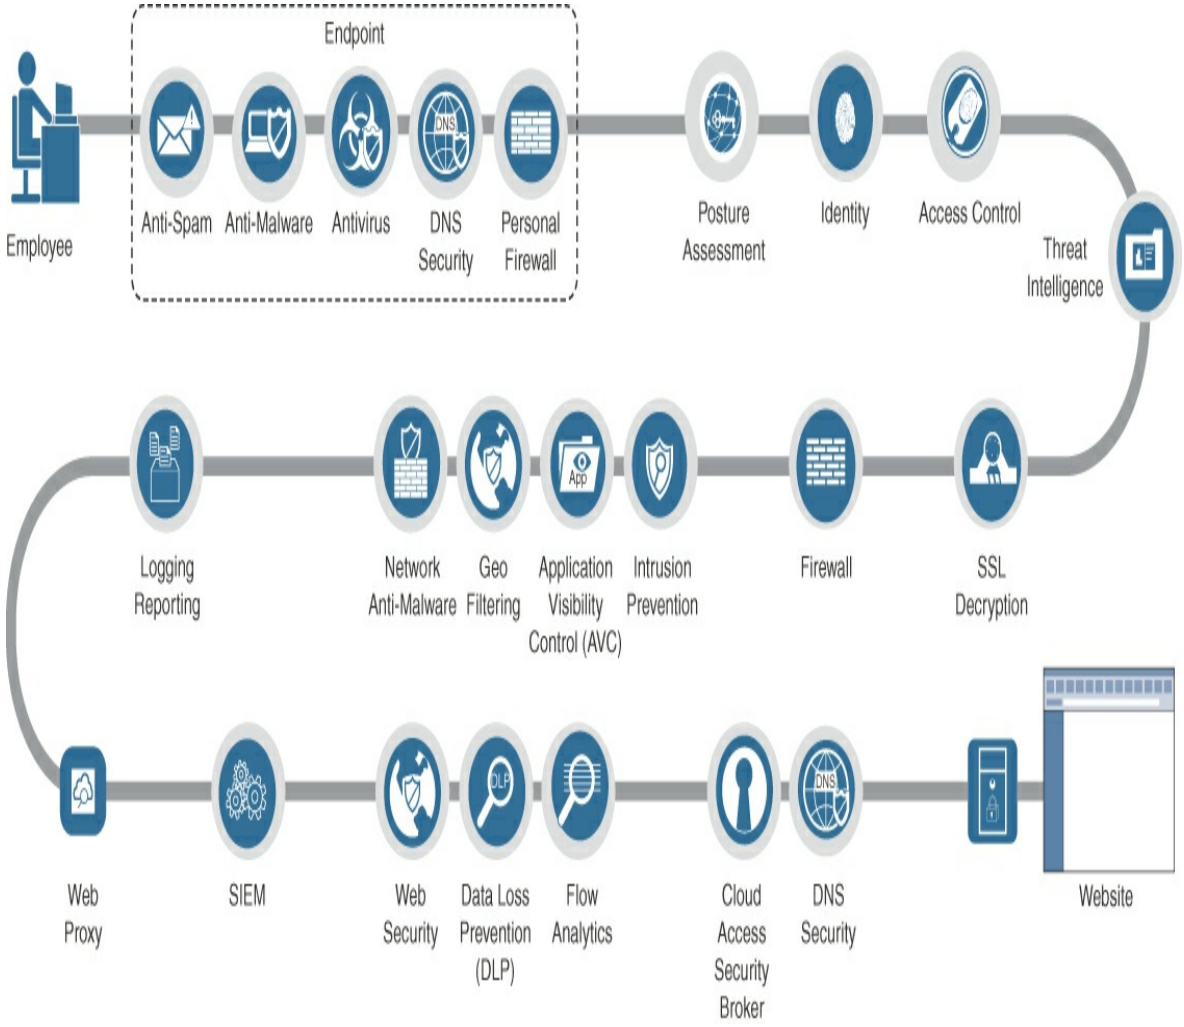

FIGURE 1-28 Branch Office Capabilities Example

SOC Capabilities Gaps Analysis

Once you create a capability map, you evaluate it based on gaps and overlapping capabilities and services. I suggest marking areas of competent capabilities in green shading or with a check mark and identifying unsatisfactory capabilities in yellow shading or with a slash. An example of a weak capability could be an organization using manual port security to enforce access control versus using automated network access control.

Manual port security means to configure a network switch to only permit one or more devices access to the network. When any device not within the approved list attempts to connect to the network, the network switch will disable the port. Although this provides port security, it can be tedious to manage the approved list of devices and respond when network ports become disabled. Manual port security is also known to be vulnerable to spoofing attacks since adversaries can spoof an authorized system's MAC address if they can determine what type of system should be plugged into a port. An example could be spoofing the MAC address of a printer that is connected to a specific point, unplugging the printer, and plugging into that port with the printer's spoofed MAC address.

I suggest marking capabilities and services that do not exist in red shading or with an X. The lack of a capability or service doesn't mean it must be added. There may be business or technical reasons for not having a capability or service. You may find that industry guidelines, standards, and frameworks suggest more capabilities and services than you can invest in or need for your particular business and IT goals. Figure 1-29 shows the previous capability map marked up with how it applies to an example customer. For this example, a check mark represents an existing effective capability, a slash represents an existing capability that is not effective, and an X represents a capability that is missing.

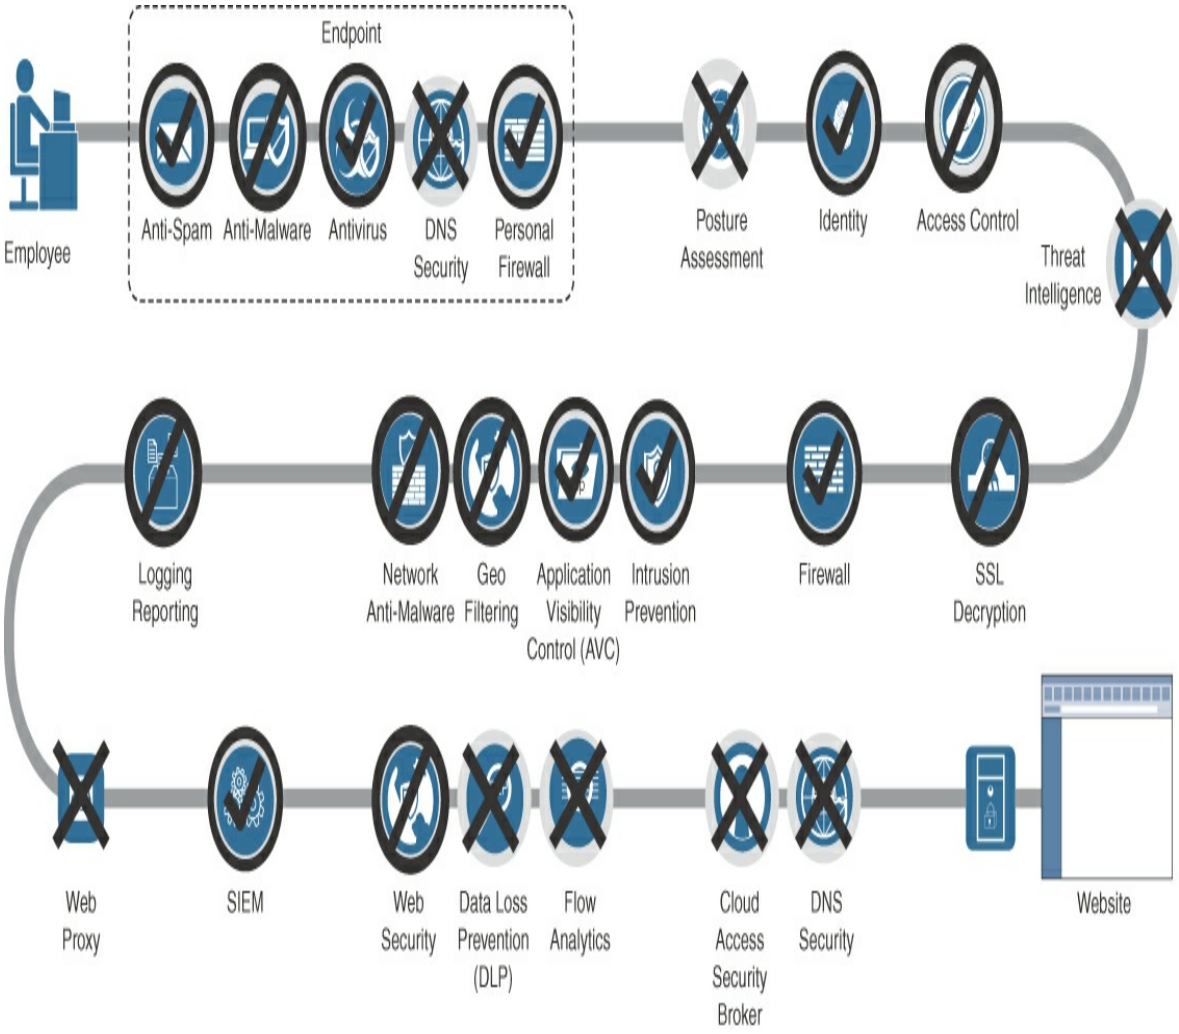

FIGURE 1-29 Capabilities Evaluation Example

This approach to performing a capabilities gap analysis does not show how capabilities will prevent specific threats. The purpose of this approach is to develop an understanding of the organization's defense-in-depth architecture regardless of the type of threats it will encounter. Remember that threats that are a major problem today will be replaced by another threat tomorrow. Using a defense-in-depth approach accommodates for such change by layering different capabilities. Maybe a threat today will exploit a vulnerability that your IPS isn't aware of using signature-based detection. If that attack is successful at exploiting a target, now the threat must deal with the next layer of defense, which will analyze the files being installed on the target. If the attacker can plant malware on the target, the next layer of defense will evaluate if the file is functioning in an unusual behavior based

on how all files of the same type are supposed to function. As you layer more defense capabilities, you reduce the risk of an attacker having the ability to bypass all of the layers of security that need to be beat in order for the attacker to accomplish his or her goal. If you want to align a test of capabilities to this approach of assessing the organization's security capabilities, you can apply a threat model such as MITRE ATT&CK.

Note

At some point, the conversation of which vendor should be chosen will come up during these different assessment meetings. It is ideal to postpone vendor-specific conversations until after all assessment work has been done, to keep the focus purely on capabilities and business goals.

Capability Map Next Steps

There are a couple of next steps that could follow developing a capability map. First, you might find that you have overlapping capabilities and/or services, which could be consolidated to simplify data streams, simplify management of tools, and reduce cost to the organization to maintain vendor contracts. Second, when presenting to nontechnical parties, you can identify gaps in capabilities and services by using the visual diagram(s) you just created. This data can be applied against business goals, technology goals, and threat ranking to determine which area of concern should be addressed first. For example, if your capability map shows a lack of insider threat capabilities and business priorities point to this being an area of concern, that would be the first place to invest future people, process, and technology improvements. Looking back at the online gaming company example, its capability map would show a missing capability for next-generation firewall and web application firewall. The company's business goal and IT goal ranking can help with the decision of what to invest in if only one tool could be purchased due to budget constraints. For this example, the WAF would be the best one based on web-based attacks being a high priority.

Note

Having overlapping capabilities is not always a bad thing. It may represent an opportunity to consolidate and acquire new capabilities to further expand the defense-in-depth architecture. For example, why have a proxy and application-layer firewall both performing content filtering when either approach can accomplish that goal for the organization? Maybe there are political or operational reasons to maintain this approach or maybe by consolidating that responsibility to one tool, an additional license, product, or job function could be repurposed to a capability that doesn't exit.

This exercise can be performed in a similar manner for other parts of the network, such as the datacenter, wireless network, headquarters, remote offices, and so on. Figure 1-30 is an example of mapping out a datacenter focusing only on IT capabilities. In this example, I first looked at clients accessing the network, which will access an inside firewall segmenting off the network from the datacenter network. Inside that datacenter network are suggested insider threat or east/west-based security. I find many organizations put most of the security on the edge of a datacenter even though the majority of the traffic exists within the datacenter. This means that if the datacenter is compromised, it may go unnoticed until malicious actions occur across the edge, such as stolen data going out the door! If that was the case, this diagram would show X's representing a lack of capabilities beyond the datacenter edge firewall.

FIGURE 1-30 Datacenter Capabilities Example

Performing a SOC capabilities and services assessment will take time and

require somebody with a decent understanding of security as well as industry best practices to properly capture the right data points. If you are not sure what capabilities should exist, look to industry standards, guidelines and frameworks for recommendations.

SOC Development Milestones

Developing a mature SOC is a journey with key milestones. The first milestone is to identify one or more executive sponsors. The purpose of a sponsor is to have an authoritative figure help the SOC enforce its policies as they are created. The sponsor can also validate the true business goal(s) of the organization and help align any SOC IT goals to the business. Together, the SOC and the sponsor can work toward the second milestone, which is establishing the SOC program. Establishing the program includes various development milestones such as developing what budget would be needed for the SOC, what talent should be recruited, and what types of technology should be acquired. These milestones are dependent on each other. The budget needs to be established before talent can be targeted. Higher-quality talent will require a higher cost than entry-level positions, but the latter will require more training. Technology shouldn't be acquired until the staff is in place and can be assessed for skills as well as asked which type of technology they prefer to use. Once the people and technology are acquired based on the program goals, processes can be developed that consist of the steps that should be taken to accomplish each goal. Processes might start off in an ad hoc fashion; however, over time they can be improved based on the maturity models covered in this chapter. As maturity is established, collaboration can occur between other teams. An example is the SOC taking on some responsibilities for vulnerability management while the desktop support team is responsible for desktop vulnerability remediation. Figure 1-31 represents a high-level diagram showcasing each of the SOC development milestones.

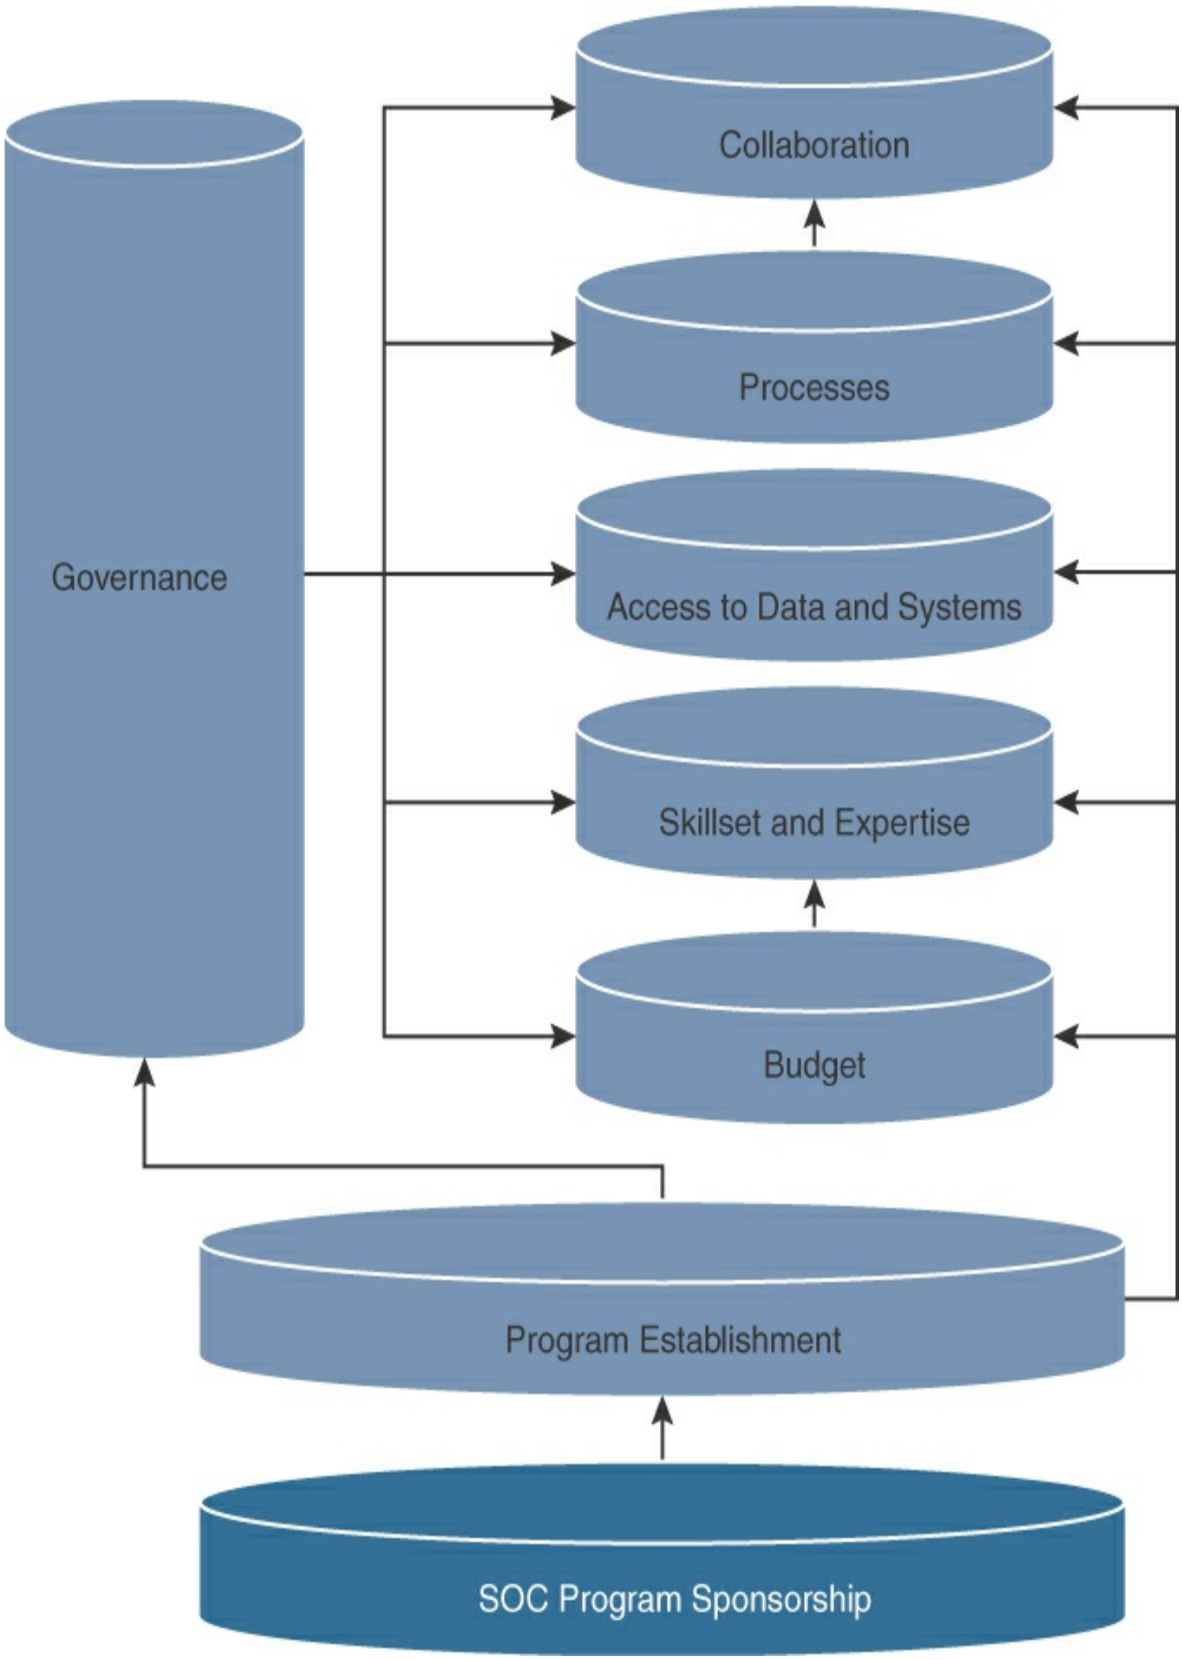
Collaboration
Processes
Governance
Access to Data and Systems
Skillset and Expertise
Budget
Program Establishment
SOC Program Sponsorship

FIGURE 1-31 Developing a Successful SOC

The strategy shown in Figure 1-31 should seem straightforward when evaluating how to develop a SOC from a high level. The key to success, however, will be how to execute the steps within each milestone. I will dive into the details of each of these milestones in the following chapters. You will learn how to approach executives and obtain proper sponsorship for the SOC. You will learn about best practices for recruiting and retaining people. I will also cover security tools, capabilities, and services in much more detail, and I will explain the processes seen within effective fourth-generation SOCs. Your journey to an effective SOC starts with the next chapter, where I will look deeper into building a SOC and how to choose what services you plan to offer.

Summary

This chapter kicked off with a review of the basics behind cyberthreats and vulnerabilities. Next, you learned the purpose of the SOC and what capabilities could be used by a SOC to identify and respond to cyberthreats. Next, the chapter dove deeper into understanding cyberthreat behavior by reviewing threat models. You also learned about industry best practices for defending against cyberthreats using standards, guidelines, and procedures. You looked at business and technical challenges as well as common SOC services. The chapter wrapped up with a review of how to assess a SOC using maturity models, goal assessments, and capabilities assessment exercises.

Many of the concepts in this chapter were high level and designed to establish a fundamental understanding of how to develop a SOC. The following chapters will provide the details for the concepts covered, so you improve the capabilities and services of your SOC. Maturing a SOC will lead to less vulnerabilities, effective SOC services and more support from the organization.

References

Ablon, L. (2018, March 15). Data Thieves. Rand Corporation. https://www.rand.org/content/dam/rand/pubs/testimonies/CT400/CT490/RAND_CT490.pdf

Caltagirone, S., Pendergast, A., & Betz, C. (2013, July 5). The Diamond Model of Intrusion Analysis. U.S. Department of Defense. http://www.activeresponse.org/wp-content/uploads/2013/07/diamond.pdf

Cimpanu, C. (2019, December 12). A Decade of Hacking: The Most Notable Cyber-security Events of the 2010s. ZDNet. https://www.zdnet.com/article/a-decade-of-hacking-the-most-notable-cyber-security-events-of-the-2010s/

Cisco. (2018). Figure 37: Patching Behavior Before and After WannaCry Campaign. Cisco 2018 Annual Cybersecurity Report (pp. 41). Cisco. https://www.cisco.com/c/dam/m/hu_hu/campaigns/security-hub/pdf/acr-2018.pdf

Fritz, B., & Yadron D. (2014, December 5). Sony Hack Exposed Personal Data of Hollywood Stars. Wall Street Journal. https://www.wsj.com/articles/sony-pictures-hack-reveals-more-data-than-previously-believed-1417734425

IBM Security. (2020). Cost of a Data Breach Report 2020. IBM Security. https://www.ibm.com/security/data-breach

International Organization for Standardization (2018, February). ISO 31000:2018: Risk management – Guidelines. ISO. https://www.iso.org/standard/65694.html

ISACA. (2019). COBIT 2019 Framework (various publications). ISACA. https://www.isaca.org/resources/cobit

Joint Task Force Transformation Initiative. (2012, September). SP 800-30 Rev. 1: Guide for Conducting Risk Assessments. NIST. https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final

Kaspersky. (2018, June 27). Ransomware and Malicious Crypto Miners in 2016–2018. Securelist. https://securelist.com/ransomware-and-malicious-crypto-miners-in-2016-2018/86238/

McKay, B., & McKay, K. (2014, September 15). The Tao of Boyd: How to Master the OODA Loop. The Art of Manliness. https://www.artofmanliness.com/articles/ooda-loop/

National Institute of Standards and Technology. (2018, August 10). The Five Functions. NIST. https://www.nist.gov/cyberframework/online-learning/five-functions

National Security Agency. (2010). Defense in Depth. NSA. https://citadel-information.com/wp-content/uploads/2010/12/nsa-defense-in-depth.pdf

Ponemon Institute. (2018, October). 2018 Cost of a Data Breach Study: Impact of Business Continuity Management. IBM. https://www.ibm.com/downloads/cas/AEJYBPWA

Proffitt, T. (2009, March). Achievements and Pitfalls of Creating and Maintaining Vulnerability Assessment Programs. SlidePlayer. https://slideplayer.com/slide/5705807/

Quote Investigator. (2013, February 10). I Rob Banks Because That's Where the Money Is. Quote Investigator. https://quoteinvestigator.com/2013/02/10/where-money-is/

Robertson, A. (2018, September 28). California Just Became the First State with an Internet of Things Cybersecurity Law. The Verge. https://www.theverge.com/2018/9/28/17874768/california-iot-smart-device-cybersecurity-bill-sb-327-signed-law

Rogers, K., & Spring, B. (2020, September 6). 'We Are Outnumbered' – Cybersecurity Pros Face a Huge Staffing Shortage as Attacks Surge During the Pandemic. CNBC. https://www.cnbc.com/2020/09/05/cyber-security-workers-in-demand.html

van Os, R. (n.d.). SOC-CCM Model. SOC-CCM. https://www.soc-cmm.com/

Chapter 2

Developing a Security Operations Center

One great building does not make a great city.

—Thomas Heatherwick

Chapter 1 opened with an introduction to security concepts and covered how to rate the quality of existing capabilities and services against industry best practices. The starting point for your security operations center (SOC) journey may be anywhere between creating a new practice to enhancing a functioning operation. Regardless of the maturity of your operation, it is important to review everything from the high-level mission to the details for each process of each service your SOC is responsible for. This chapter begins with a review of high-level SOC development concepts and works down into the technical details found within highly effective SOCs showcased in later chapters. The absolute first thing any successful organization must define is its *purpose*, which is captured in the mission statement and scope. Even if you have an existing mission statement or defined purpose, it is worth your time to review it periodically to see if any improvements can be made.

Mission Statement and Scope Statement

It is important for any employee of an organization to know why they come to work every day. If employees don't know the purpose of the business, they won't know how to make the most impact with their daily contribution, leading to a lack of career-driven workers. Employees performing job duties without any guidance from leadership regarding why their position exists and

how it benefits the mission of the organization become dissatisfied and a flight risk. An example of an employee performing their daily duties without considering the impact to the business can be seen with poor customer service or support. Imagine a situation where a customer calls for support on a product that is no longer covered. If the support representative doesn't know the business goal, he or she will simply inform the customer that the product is no longer supported, since that is what the rep's specific job instructs them to do regardless of how it impacts customers. The long-term impact of poor customer service is fewer repeat customers and damage to the brand of the business, leading to diminishing sales. If the goal of the business is to put the customer first and that goal is communicated clearly to customer support representatives, the support reps will prioritize satisfying the customer and provide support based on meeting the business goal for support.

The mission and scope statements of an organization inform employees what the purpose of the organization is. A *mission statement* is essentially a formal summary of the organization's core purpose and values. The statement is high level and rarely changes over time, if at all. The mission statement expresses what is important to the business, giving direction to what the organization is looking to accomplish. Any employee should be able to see how their role within the organization impacts the goals found within the mission statement. This accessibility makes the mission statement the ideal starting point for business planning. A *scope statement* is similar to the mission statement but contains more details. Whereas the mission statement remains the same unless major change to the organization occurs, the purpose of the scope statement is to be more flexible to change.

Developing Mission and Scope Statements

If a SOC is being developed within an organization, its mission and scope statements must relate to the business the SOC is responsible for protecting. The key to creating a successful SOC mission statement is to be crystal clear about what the responsibilities of the SOC are. The mission statement must be a concise explanation of the SOC's purpose and overall intention. Statements must be long-term goals rather than specifics of what type of services are being offered, because that is covered in the scope statement.

Note

It is common for a SOC mission statement to contain goals similar to those of the organization it protects; however, there are exceptions. Some SOCs are made of a group of SOCs, known as a global SOC (GSOC). Individual SOCs might have their own mission statements or share the overall GSOC mission statement. Some SOCs are outsourced or are isolated from the customer being serviced for compliance or other purposes. This type of SOC has a mission statement that is focused only on its service, rather than considering the mission of its customer.

If you are creating or contributing to a mission statement for a SOC, the key is to ensure that it answers the following questions. Remember to keep mission statement language at a high level representing the vision for the SOC. The specifics are described in the scope statement covered shortly.

- **Purpose:** What is the purpose of the SOC? Why does it exist?
- **Customers:** Who are the SOC's customers? All employees within the organization or only those in a specific department? Are external customers included?
- **Alignment to customer's mission:** How does the SOC's mission align with the organization's mission? Does the SOC help the organization accomplish its goals along with protecting its assets and data?
- **Services:** What services does the SOC provide? Is it only monitoring or is it also responding to incidents? The services should not be specific, meaning monitoring the network could apply to multiple services such as incident response and vulnerability management. Those details are ideal for the scope statement.
- **Service availability:** When are SOC services available? Is it 24/7 or only during business hours? The mission statement should not provide details that may change for purposes of availability. It should only be a general statement about coverage.

Sample Mission Statements

Here are two sample SOC mission statements. As you read each mission statement, try to identify the answers to the questions provided in the preceding list. (The answers are given in the sublist following each statement.)

- **Mission Statement 1:** The security operations center (SOC) is responsible for monitoring, detecting, and responding to incidents and the management of the organization's security tools, network technology, end-user devices, and other systems as assigned. The goal of the SOC is to reduce risk from cyberthreats targeting the organization. The function of the SOC is performed seven days a week, 24 hours per day. The SOC is the primary location of the staff and the systems dedicated for this function.
  - **Purpose:** Monitoring, detecting, and responding to incidents and managing security tools.
  - **Customers:** All end users employed by the organization and guest users.
  - **Alignment to customer's mission:** Reducing the risk from cyberthreats targeting the organization.
  - **Services:** Incident management and management of security tools, network technology, end-user devices, and other systems assigned.
  - **Service availability:** Seven days a week, 24 hours per day.
- **Mission Statement 2:** The security operations center (SOC) monitors the overall security posture of the IT network. The goal of the SOC is to help the organization identify and respond to security incidents and return impacted systems back to a normal operational state in a timely manner. The security operations center responds and tracks security incidents with the objective of maintaining the overall security posture while reducing risk when possible. The functions are performed around the clock in support of the organization's operation model.
  - **Purpose:** Maintaining the overall security posture of the organization while reducing risk when possible.
  - **Customers:** The organization.

- **Alignment to customer's mission:** Respond to security incidents and return impacted systems back to normal operational state in a timely manner.
- **Services:** Respond and track security incidents.
- **Service availability:** Around the clock in support of the organization's operation model.

Each of these two sample mission statements represents a single SOC taking on most of the security responsibilities at the organization that it functions within. Statement 1 has more specific details, such as the days of operation and location, while Statement 2 is broader in how it explains its purpose. Both statements are fine; however, relegating items that could change to the scope statement is recommended so that the mission statement can endure relatively unchanged and represent the high-level goals. Statement 1 isn't bad, but if the SOC changed its hours of operation or location, the mission statement would need to be adjusted. Statement 2 provides broader language for the hours of operation that would be ideal so specific changes in service availability won't impact the mission statement. Both statements could also have more details regarding how they align to the organization, but are good starting points.

Note

If you do not have a dedicated SOC mission statement or have been using the organization's mission statement, it is highly recommended to develop a SOC mission statement.

SOC Scope Statement

A SOC scope statement is similar to a SOC mission statement but contains a lot more details about the SOC. The scope statement includes where the SOC is located, which services it provides and does not provide, details about the hours of operation, and what technologies it uses. The scope statement needs to complement the mission statement, but unlike the mission statement, the

scope statement can contain items that are subject to more frequent change. Some SOCs combine the scope and mission statements; however, in those situations, changes to the SOC require an update to the combined mission and scope statement rather than just updating the scope statement. For this reason, I recommend keeping the scope and mission statements separate so the scope statement is more flexible to change while the mission statement remains the unchanged vision of the SOC. The scope statement can also accommodate responsibilities for partners and subsidiaries.

Looking back at the mission statement examples, some important topics were not covered. These topics will fall under the scope statement and are subject to change. The ideal topics for a scope statement include the following.

- **Locations and networks:** What parts of the organization will the SOC be responsible for? What physical locations, remote offices, and cloud environments would be in and out of scope for support? What network types are being considered? Will the guest network, datacenter, or wide-area network be part of the SOC's responsibility?
- **Ownership:** What level of ownership will the SOC have regarding systems and information? Does the SOC have ownership on securing endpoints, or is that a responsibility of desktop support? Does the SOC monitor the datacenter and have responsibility for provisioning access to data, or are access rights managed by a datacenter custodian? It is very likely that the SOC will not be responsible for managing systems or data outside of data within specific security tools. This ownership needs to be clearly listed in the scope statement.
- **SOC objectives:** The SOC's objectives need to be outlined so that the rest of the organization can understand the SOC's purpose. This can be explained within the mission statement and referenced within the scope statement or be another mission statement from the organization that aligns with the overall organization's or GSOC's mission statement. If a mission statement already exists, the scope statement should be more specific about the objectives, including timelines for accomplishing specific goals. An example could be increasing the maturity of a capability such as responding to incidents or managing vulnerabilities.
- **Technologies and services:** Scope statements can specifically call out what will be used by the SOC. For example, the scope statement can

specify “All intrusion prevention appliances,” which would mean other network gear used by the organization would not be a responsibility of the SOC. The scope can also identify services that it offers, such as forensics or incident response. If some services are outsourced, then the scope statement could identify only what the SOC would cover. For example, a scope statement could state that the SOC is responsible for Tier 2 or higher incidents, which would mean another resource would handle Tier 1 support.

- **Specifics regarding when SOC services are available:** This includes the exact start time and ending time each day, if holidays are included, and how non-supported hours are covered.

Let’s look at a few SOC scope statement examples.

- **Scope Statement 1:** The SOC is responsible for every company location across the country that hosts systems and permits systems to connect to the general network. The services that are offered by the SOC include around-the-clock security monitoring of systems, applications, and networks with the objectives of detecting and reacting to all external attacks as well as insider malicious behaviors. The SOC services are also responsible for handling incident response; collecting and correlating the various system events; capturing and analyzing raw packet data; discovering and tracking vulnerabilities; and consuming threat intelligence information received from external sources.
  - **Location and networks:** Every company location across the country that hosts systems and permits systems to connect to the general network.
  - **Ownership:** Monitoring systems, applications, and networks.
  - **SOC objectives:** Detecting and reacting to all external attacks as well as insider malicious behaviors.
  - **Technologies and services:** Incident response; collecting and correlating the various system events; capturing and analyzing raw packet data; discovering and tracking vulnerabilities; and consuming threat intelligence information from received external sources.
  - **Specifics regarding services:** Around-the-clock monitoring of systems, applications, and networks.

- **Scope Statement 2:** The SOC scope covers all systems that are managed and operated by IT, including those located in national and international offices. The SOC services are offered around the clock and include the collection and correlation of security event messages; detecting internal and external malicious activities; responding to security incidents; and conducting awareness training when required.
    - **Location and networks:** National and international offices.
    - **Ownership:** All systems that are managed and operated by IT.
    - **SOC objectives:** Covering all systems that are managed and operated by IT.
    - **Technologies and services:** The collection and correlation of security event messages; detecting internal and external malicious activities; responding to security incidents; and conducting awareness training when required
    - **Specifics regarding services:** Around-the-clock services.

What should stand out is how a scope statement contains a lot more details than a mission statement about what the SOC will provide. Every SOC should have a separate mission statement and scope statement (or at least a combination of these statements), even a smaller SOC within a larger GSOC, so that key details such as the services, time of operation, and purpose of the SOC are documented for anybody to reference.

Scope Statement Challenges

A challenge that is common when creating the scope statement is committing to specific services. There might be questions regarding whether the SOC should be responsible for a particular service or another department should handle that service. For example, if the SOC scope statement claims responsibility for identifying vulnerabilities, does that mean desktop support is not responsible for identifying vulnerabilities? Who would handle *remediating* vulnerabilities? Would that be desktop support and/or the SOC? How would this handoff occur if the SOC does not handle remediating vulnerabilities and hands off systems identified as vulnerable to desktop support? Looking back at the language used in the first example scope

statement, the statement points out that the SOC is responsible for only "discovery and tracking vulnerabilities." The scope statement does not include the responsibility for remediating vulnerabilities, which means the SOC is not responsible for this service.

For this example, another team such as the vulnerability management team can have the responsibility of remediating vulnerabilities. The vulnerability management team can also have a scope statement claiming responsibility for identifying vulnerabilities, which would be an overlap in services offered by the SOC since it also claims responsibility for identifying vulnerabilities. If this situation occurs, both teams need to recognize this overlap in responsibility and develop a plan for how responsibilities are handed off and/or how these teams collaborate. The best approach to this specific example would depend on various factors within the organization. Some organizations might decide to remove vulnerability identification from the SOC's scope, while other organizations might decide to remove vulnerability identification from the vulnerability management team's scope, which would mean the SOC would identify vulnerabilities and hand off those findings to the vulnerability management team so the vulnerability management team can handle the remediation.

Governance and Risk Management References

Scope statements can also include concepts related to security governance and/or risk management. The purpose is to provide clarity regarding why specific services are being offered. For example, a scope statement may specify that the reason for providing data control services is to meet HIPAA or PCI DSS compliance. A scope statement can also relate to a business goal, such as to ensure an organization's product is safe or service is always available to customers. Including clarity regarding the purpose of a SOC service helps leadership better understand the purpose of the service, leading to better support for the SOC. Another benefit of clarity is that those using or impacted by the SOC service will understand and respect why the service exists. Respecting a SOC service leads to less pushback when the SOC has to take actions.

Developing a SOC

Developing a SOC takes more effort than picking the right location and purchasing a bunch of equipment. You must have a process to plan, design, build, and operate the SOC, which can lead to transferring the SOC to another team once the SOC is up and running or help the SOC mature from a small team to a dedicated group within the organization that has its own leadership and funding. SOC planning starts with topics covered in Chapter 1 regarding assessing the SOC's current state as well as the development of a mission and scope statement.

How you plan and design your SOC will impact it in a positive or negative way, including the cost and time required to develop and improve SOC services. Knowing the expectations of the SOC based on its scope is useful for selecting the type of people to hire and the type of technology to acquire. Figure 2-1 provides a high-level summary of the steps to designing a SOC.

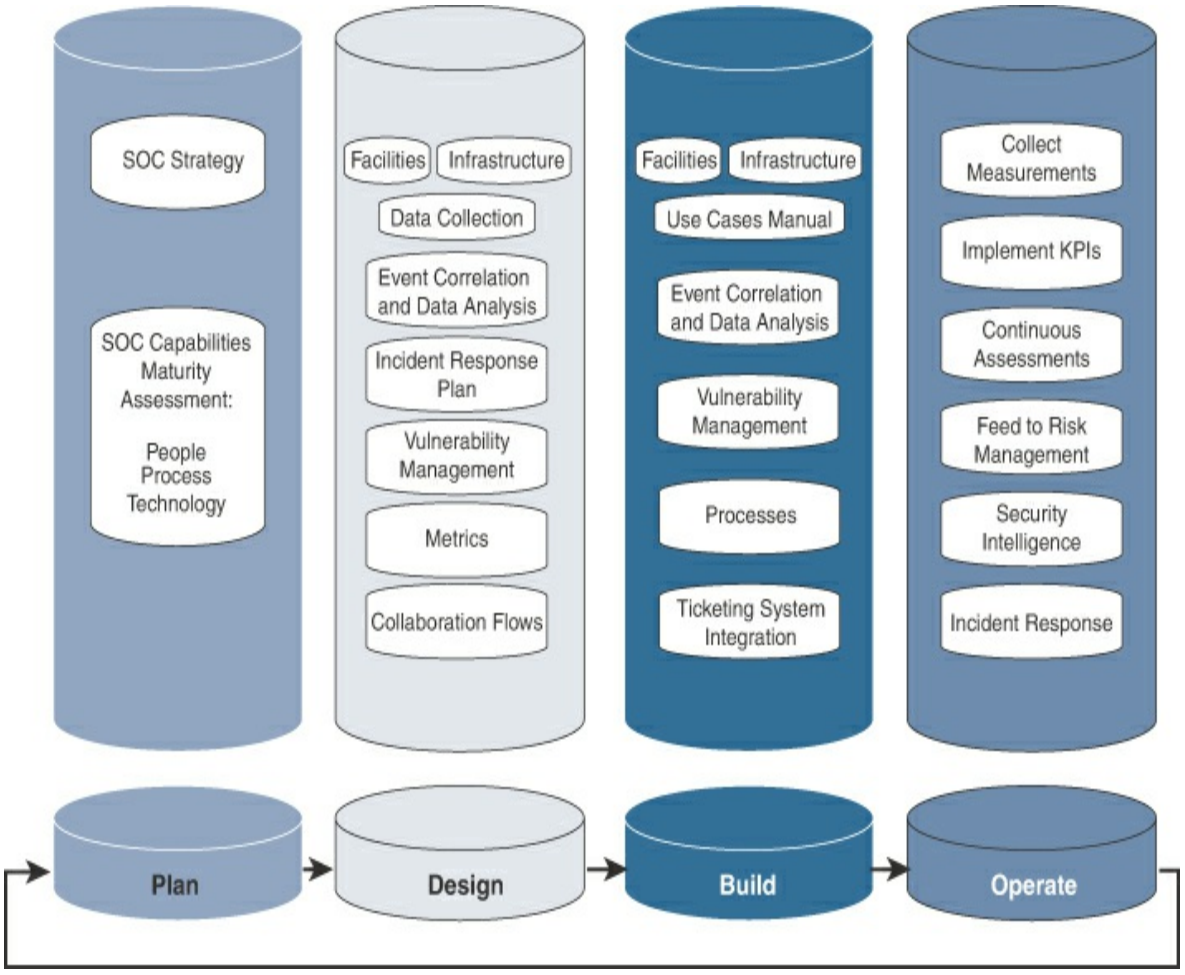

FIGURE 2-1 Phases of Developing a SOC

In Figure 2-1 the planning phase identifies the creation of the mission and scope statements specified as SOC Strategy. These statements are used to develop the SOC strategy made up of policies and processes. Using the assessment tactics from Chapter 1, SOC capability and maturity models are developed in the planning phase. A SOC can compare the existing capability and maturity models against the SOC strategy to develop a design plan that will require people, process, and technology.

In the design and build phase of Figure 2-1, technology is acquired, which involves considering factors such as bandwidth, high availability, and other requirements. Those acquired capabilities lead to services specified in the planning phase that will require processes to be developed and documented. Eventually, those services will go online into the operate phase.

Notice in Figure 2-1 that there is an arrow from the operate phase back to the plan phase. This represents the fact that additional requirements for new services and capabilities will follow the operate phase. As an example, after the SOC is operating, leadership may decide to assign a new responsibility for research and development to the SOC. Even though many of the SOC services are in the operate phase, the new R&D service would start at the plan phase and follow the steps shown in Figure 2-1 before it would be considered operational.

Maintaining the specific people, process, and technology within the SOC might not follow the development process outlined in Figure 2-1. Individual people, process, and technology can have their own quality testing and upgrading processes as shown in Figure 2-2, representing different SOC aspects that are independently developed or acquired. For example, training can be used to onboard a new employee would not follow all four stages shown in Figure 2-2.

FIGURE 2-2 People, Process, and Technology Core Elements

Consider Figures 2-1 and 2-2 as high-level models of the process for developing a SOC. As these steps are completed, procedures are developed by the SOC to represent how the SOC will execute each service it is responsible for. Next we will look at developing SOC procedures.

SOC Procedures

Chapter 1 touched upon common services offered by mature SOCs, which include risk management, vulnerability management, incident management, analysis, compliance, digital forensics, situational and security awareness, and research and development. All of these services are covered in greater

detail in Chapter 3, "SOC Services," and throughout the book. If your SOC is going to offer any of these services, you need to decide what should be outsourced and what will be covered using internal services. Although Chapter 1 covered the benefits and disadvantages of each approach, it did not delve into developing procedures for SOC services.

Procedures are how services are executed by the SOC. Procedures include step-by-step instructions to execute policies, such as the steps to operate a vulnerability scanning tool, how often the tool is run, and what parts of the network or which hosts are scanned by the tool. Each of these procedures is used to enforce a policy for identifying vulnerabilities within the network. Remember that procedures are the details regarding what is done to execute a policy.

Designing Procedures

When designing procedures, there are some important items to consider. The questions in the following list need to be addressed by a SOC procedure:

- What is the purpose of the procedure and what policy does it align with?
- How should the SOC be involved with this procedure?
- How long is the SOC responsible for this procedure?
- What other groups or outside elements impact the procedure?
- What threat does this procedure deal with?
- What resources are required for this procedure?
- Are logging or reporting required for this procedure?
- What notifications should be included within this procedure?
- What is the notification escalation process?
- How are notifications delivered (email, mobile, home, chat, etc.)?
- Are there any compliance elements involved with this procedure?

The SOC's involvement with these procedures will vary based on whether the SOC will be responsible for all or part of these actions. For example, a

SOC can have another team handle the monitoring and alerting of threats while the SOC is responsible for anything that is escalated. Other times, the complete opposite can occur, having the SOC handle basic incidents and hand off anything escalated to external investigation services.

Here is a list of what are considered common SOC procedures that are used to deliver various types of SOC services:

- **Monitoring:** Surveillance of specific system(s) and network(s)
- **Alerting:** Notification of a threat, problem, or event
- **Escalation:** Passing responsibilities for responding to an event to the next level of support
- **Investigation:** Examining, studying, or inquiring into an event
- **Incident logging:** Case management efforts, which include tracking events that occur and the process of how the events are handled
- **Compliance:** Relating activities and events to meeting compliance requirements
- **Reporting:** Developing and delivering metrics about events that have and have not occurred
- **Remediation:** Returning systems back to operational state and closing vulnerability to reduce future attack

It is critical that any procedure clearly defines the actions required from the SOC so it can develop proper support for people, process, and technology. As procedures mature, recommended practice dictates developing templates so that new team members can easily follow what they are expected to perform. Over time, templates can be automated to improve the efficiency and effectiveness of the SOC services, leading to mature SOC capabilities.

Procedures must define any required communication as well as action. Communication responsibilities include how communication should occur between teams as well as what communication method should be used. This is especially important as SOC tiers are developed.

Procedures Examples

Let’s look at an example of procedures associated with responding to a security incident. Procedures should specify which group is responsible to perform the first response to a security incident. Procedures should include the hours of operation for this team, steps to escalate the problem if it can’t be solved, steps to solve problems that fall within scope of the first responding team, what are the backup options when the team isn’t available, and what the first responding team is authorized to perform. If a new person is recruited to perform the first responder role for an incident, procedures will provide directions for how each of these tasks is performed.

Figure 2-3 is an example of creating a mapping of how a SOC supports a procedure. In this example, the SOC activities include monitoring alerts from three different groups. Tier 1 support handles monitoring, alerting, and escalating to tier 2 when an issue requires more skilled individuals who are authorized to perform such work. Tier 2 can also escalate to a more advanced tier 3 group when such skills are needed. Tier 3 support also includes additional SOC services such as forensics and malware investigation.

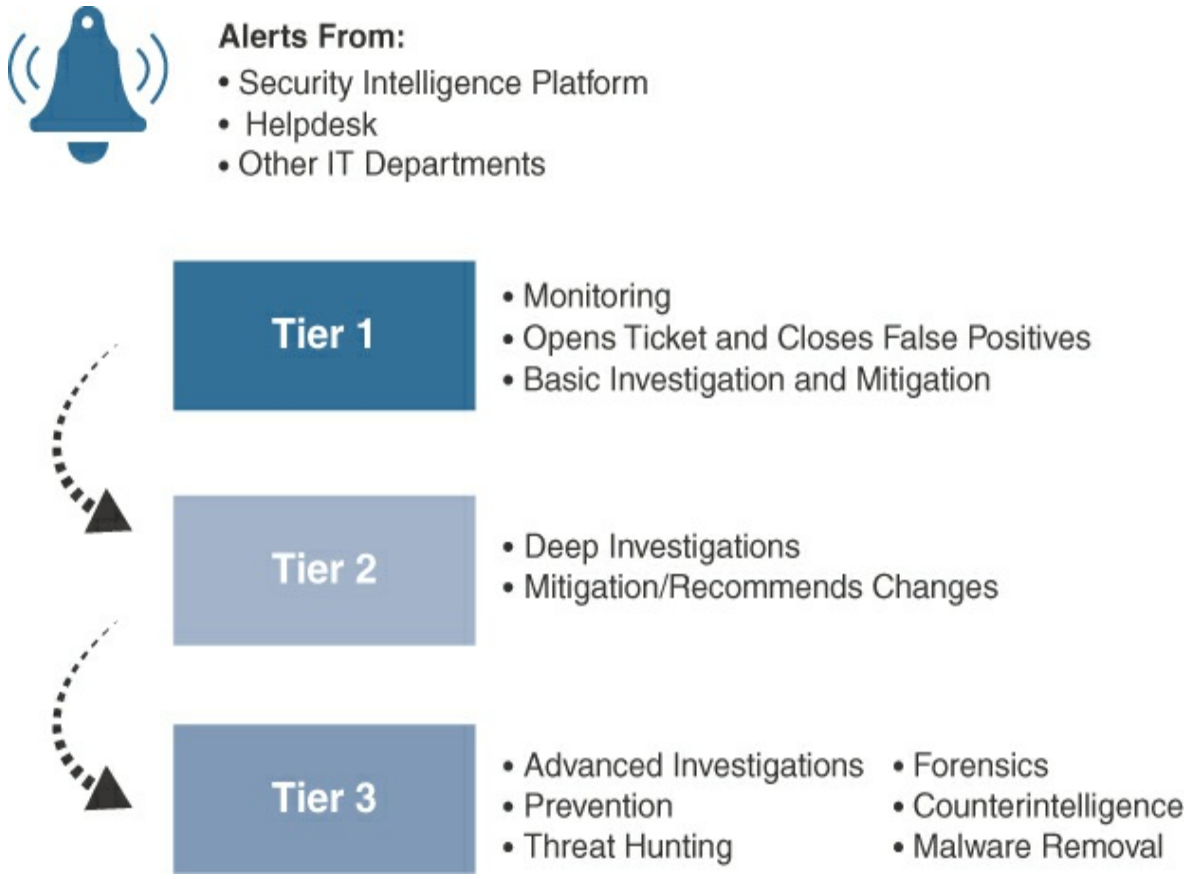

FIGURE 2-3 SOC Procedure Example

Figure 2-3 is great for a high-level view of SOC tiers; however, the communication between these teams isn't defined. For example, how is tier 1 contacted? Who is authorized to escalate an event to a higher tier? Who handles logging and reporting? What if a tier is unavailable when contacted? Figure 2-4 provides a more detailed approach to documenting a procedure. For this example, the procedure diagram covers how customer inbound requests are handled. The helpdesk receives alerts through a phone call or email. Incidents are handed to a SOC tier 1 analyst, who can attempt to close the case or escalate the case through the SOC tier 1 supervisor or team lead. The SOC tier 2 attempts to handle an incident or pull in or hand off an incident to a tier 3 support member. Tier 3 members include external services they can leverage if they are unable to cover the incident with in-house resources. Know that this diagram can include even more details such as the names and locations of parties involved, specific tools available and hours of operation. Many of these details were omitted from the diagram for the purpose of providing a generic example.

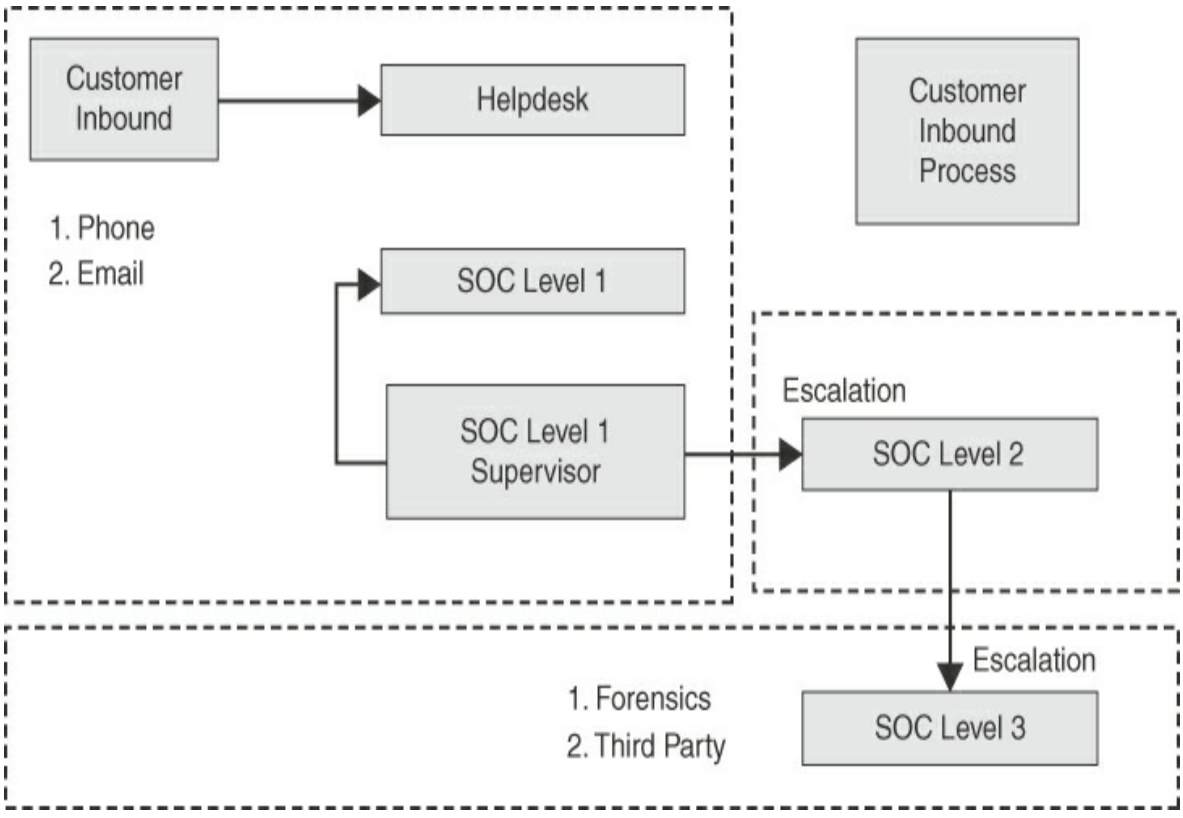

FIGURE 2-4 SOC Contact Flow Example

Security Tools

Chapter 1 pointed out that the best practice for developing a security strategy is to develop capabilities against all of the steps an attacker would take to breach a network. You also learned that those steps will be different for different parts of the network. Let's now start looking at the different types of tools that a SOC can use to provide a defense-in-depth approach to alerting and responding to security events. All of these tools should have the ability to export logs to the SOC's security information and event management (SIEM) solution. Chapter 5, "Centralizing Data," will cover the processes for how that is done in more detail. While many of these tools were touched upon already in the Chapter 1 discussion of protecting the SOC, this section delves into the specific capabilities associated with security tools used by SOCs around the world to deliver various types of services.

Evaluating Vulnerabilities

Chapter 1 pointed out how adversaries exploit vulnerabilities in order to exploit systems. This means as a SOC, it is ideal to have a mature vulnerability management practice. Chapter 9, "Vulnerability Management," dives deeper into how to build a vulnerability management practice; however, there are some tools you should consider dedicating to your vulnerability management practice. The most obvious tool is something that can assess the network and endpoints for vulnerabilities. Vulnerability scanners serve this purpose by comparing attributes about systems and software being analyzed against known weaknesses. A known vulnerability is classified based on risk impact, which is summarized in its Common Vulnerability Scoring System (CVSS) score (see https://www.first.org/cvss/). A CVSS score provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. The numerical score can then be translated into a qualitative representation to help organizations properly assess and prioritize their vulnerability management processes.

Active Vulnerability Scanning

In order for a vulnerability scanner to detect and classify system weaknesses in computers, networks, and communications, it must have access to the target being scanned. Access can be from the network or directly on the host. The level of access can be full read-level access, known as *authenticated scanning*, or no ability to log into the system, known as *unauthenticated scanning*. Having full read access provides more details about how systems, networks, and communication are vulnerable. Results of authenticated scanning provide a more accurate report, which leads to a better remediation response. Although authenticated scanning provides better results than unauthenticated scanning, it does not realistically represent what a potential adversary will see. Unauthenticated scanning is considered the attacker's point of view, meaning tools like firewalls will prevent some scanning use cases from providing any useful data. A SOC should use both authenticated and unauthenticated scanning to identify all potential vulnerabilities as well as understand how adversaries would view computers, networks, and communication they could target using the unauthenticated scanning approach.

Figure 2-5 shows an example using Rapid7's Nexpose vulnerability scanner platform. The dashboard shows a summary of the environment being continuously evaluated for vulnerabilities, which can be broken down into sites or network segments. An example could be scans set up for different parts of the customer and SOC segments. In this example, a summary of the vulnerabilities found within each network segment is presented in two pie charts. The first pie chart rates vulnerabilities based on their CVSS score. The second pie chart ranks vulnerabilities based on the skill level required to exploit them. Below the pie charts are specific vulnerabilities identified by Nexpose. An analyst can click any part of either pie chart or the specific vulnerabilities listed to learn more about what Nexpose discovered.

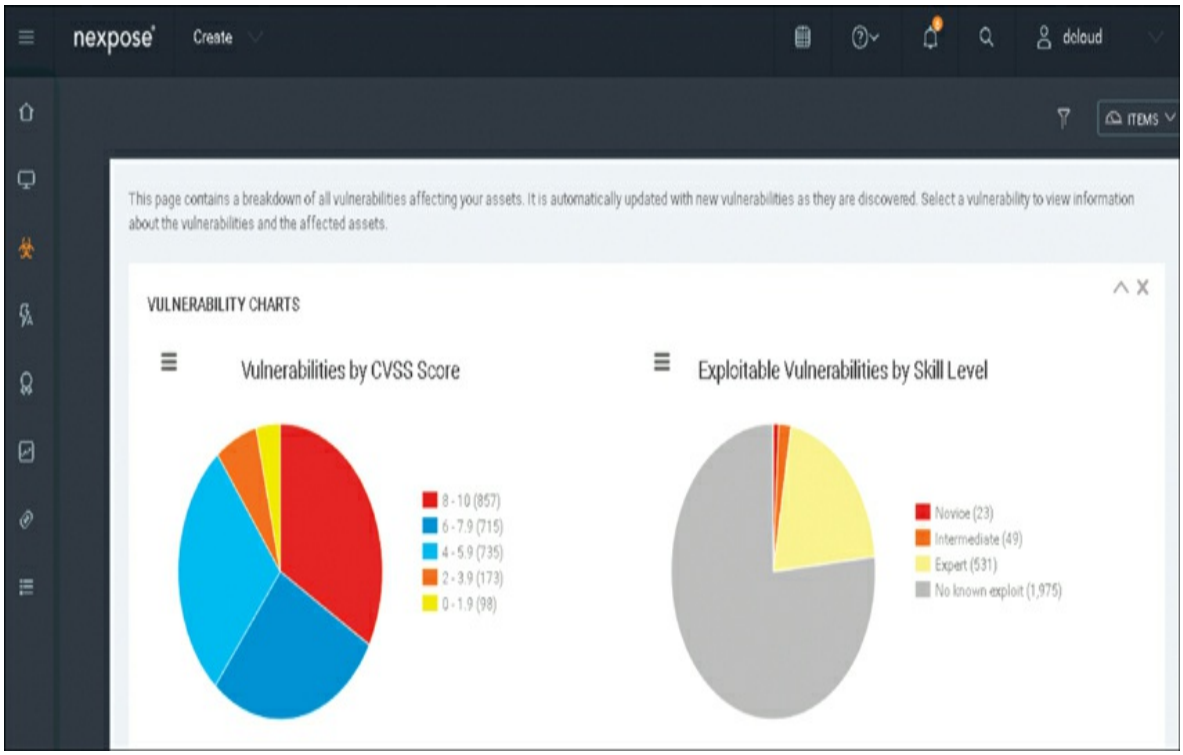

FIGURE 2-5 Rapid7's Nexpose Vulnerability Scanner

Passive Vulnerability Scanning

Tools such as application-layer firewalls can gather data about endpoints and send that data to vulnerability scanners to perform passive scanning. I am finding that more security solutions are leveraging passive security technology with the intent of adjusting defense capabilities to systems found vulnerable. Figure 2-6 is an example of the Cisco Firepower solution showcasing passive vulnerability scanning information. This approach to passive vulnerability scanning works by leveraging the next-generation firewall's capability of seeing all layers of traffic. This permits host profiles to be created, which are compared against a database of known vulnerabilities similar to an active vulnerability scanner. The difference between an active vulnerability scanner such as Rapid7's Nexpose and a passive vulnerability scanner like the Cisco Firepower solution is that a passive scanner only sees what data can be profiled, while an active scanner can be configured to scan any network or host it can access.

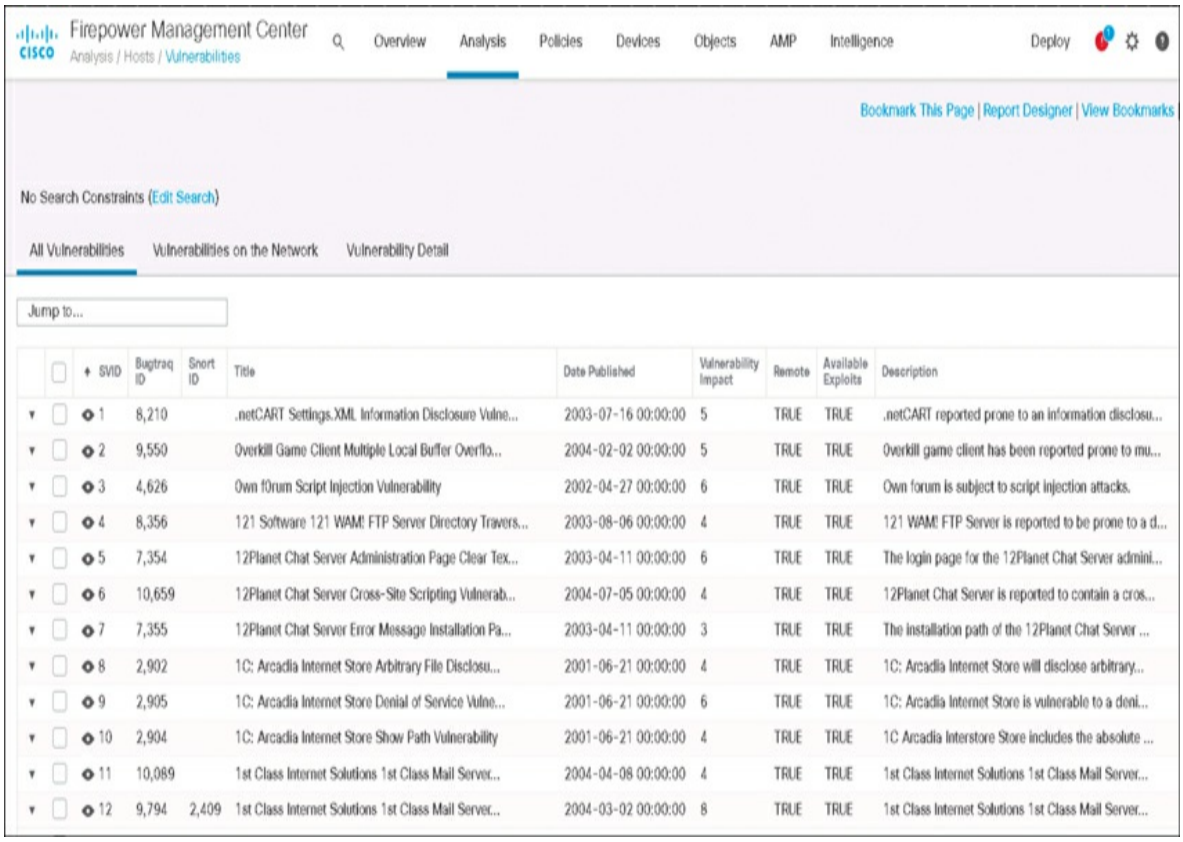

FIGURE 2-6 Cisco Firepower Passive Vulnerability Data

The results of a vulnerability scan (regardless of which approach is used) will produce a handful of vulnerabilities found within computers, networks, and communication on your environment. Analysts will want to identify all systems that are vulnerable and create tickets to assign an engineer to patch these weaknesses. While systems are exposed, the SOC will work to prevent an adversary from exploiting these weaknesses using security tools such as an IDS/IPS and antivirus. Security tools can be tuned using results from vulnerability scanners to help protect computers, networks, and communication that have been identified as being vulnerable. Tuning security tools helps to protect vulnerable computers, networks, and communication from attack while the asset owner is notified to fix the vulnerability.

Note

If you search for a recent threat and don’t see matching signatures within your security product, contact the vendor and ask if you are

protected. The vendor might have hidden the rules so that adversaries can’t reverse engineer their defense strategy.

Preventive Technologies

If a system, network, or communication is found to be vulnerable, removing the vulnerability through patching, upgrading, or other method will prevent exploitation from occurring. Unfortunately, remediation cannot always be performed, for a variety of reasons. Some technology such as IoT devices do not offer the ability to fix an identified vulnerability due to limitations in how the technology functions. If a vulnerability can’t be fixed, the SOC needs to find other methods to prevent an adversary from exploiting the vulnerability. One approach is to hide the vulnerability from any potential threat. Another approach is to set up security tools to monitor the vulnerable computer, network, or communication and alert as well as prevent exploitation that is attempted against the vulnerability. Preventive technologies can offer both of these approaches to protecting vulnerable computers, networks, and communication.

There are different types of preventive technologies that the SOC should use to protect vulnerabilities found within the network and system it is responsible for.

Preventive Technology: Firewalls

Firewalls and proxies can perform network segmentation and limit what ports, protocols, and applications are permitted to cross a network segment. Best practice is to enforce least privilege access through each network segment to reduce the risk of exposing vulnerabilities to adversaries outside of network segments. For example, if an organization doesn’t have a business need for permitting Tor traffic, this traffic should be prevented from crossing network segments. Blocking Tor reduces the risk of adversaries and malware communicating between a network segment and Tor network, which is common behavior for threats such as ransomware.

Preventive Technology: Reputation Security

Another preventive technology the SOC should use is reputation security. Reputation security reviews various factors of an external source and judges how trustworthy that source should be seen as. One factor is how long the source has been online. If a source claims to be an online bank but has been online for only a few hours, the source is not a bank. Another factor is the reputation of the domain owner. The current source being evaluated might be seen as safe, but if the domain owner has been flagged as being responsible for other malicious domains, that factor will impact the trustworthiness of the current source being evaluated. Another factor is the type of content that is being presented from the domain. If files that have been downloaded from a domain have been flagged as malicious, that will impact the trustworthiness of the source being evaluated. The list goes on for what factors could be evaluated, and different vendors will have a different approach to how reputation security is evaluated and enforced. An example of reputation security's value is blocking an adversary from creating a malicious website that looks similar to a trusted site, which is a common tactic associated with phishing attacks. If an adversary attempts to create a fake website representing a trusted bank, factors such as the website only being online for a short time, being hosted from GoDaddy, and the domain owner being associated with other malicious websites would all impact this source's reputation score. Figure 2-7 is an example of a Cisco block page triggered when accessing a web source with a negative reputation score. Security solutions leveraging reputation security will block the fake bank described in the previous example and display a page such as what is shown in Figure 2-7 based on how the vendor informs users of a source being blocked.

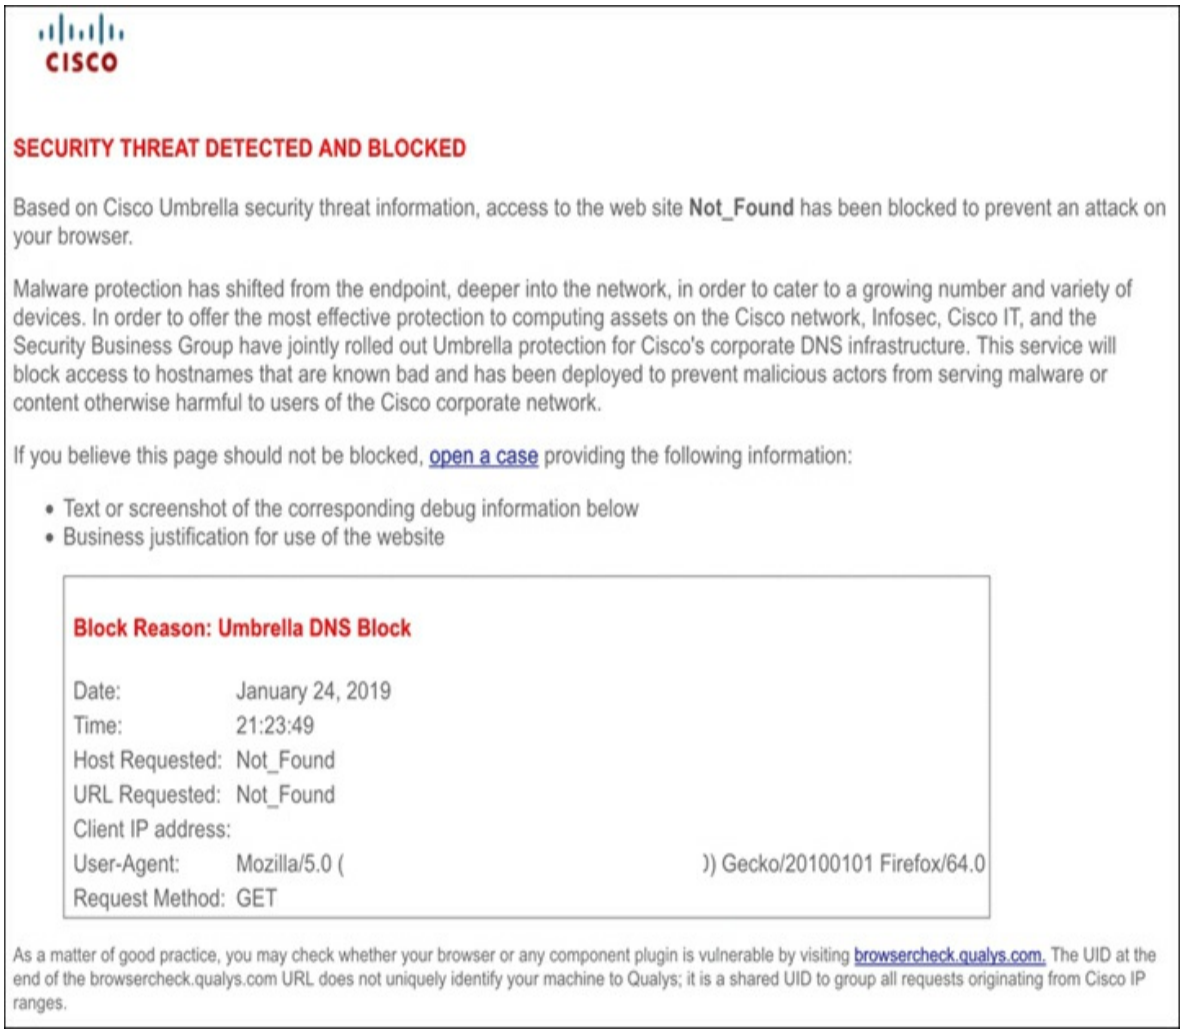

FIGURE 2-7 Cisco Reputation Block Page

Note

You can test if your organization has reputation security by going to www.ihaveabadreputation.com. This website will not harm your computer but has a bad reputation for testing purposes. If you see an image of a ghost rather than a vendor-generated block screen, you do not have reputation security.

Similar reputation security technology can be used on other parts of the network such as host systems and email. Google Gmail is one of the many

email services that have adapted reputation security to reduce spam and phishing attacks. Figure 2-8 is an example of a phishing email being flagged by Google as being dangerous based on various factors Goggle uses to flag malicious email. Factors can include the content of the email, the sender of the email, the audience that is receiving the email, and negative results associated with similar emails.

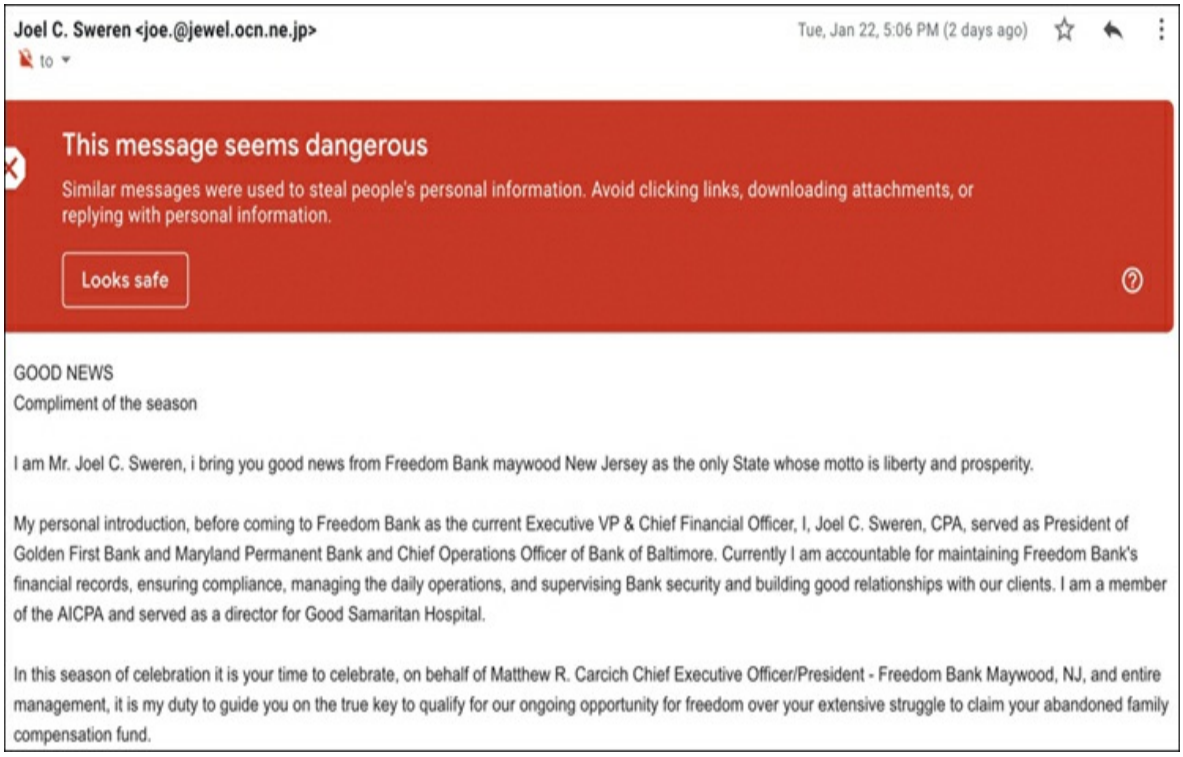

FIGURE 2-8 Google's Reputation Warning Banner

Note

Reputation security is not 100% effective. That doesn't mean it isn't great at reducing the number of attacks due to the increase in effort needed by an adversary to launch an attack. For example, adversaries can compromise a small business or school and pivot from that source using the compromise source's reputation to bypass this reputation security defense.

Preventive Technology: VPN

Virtual private networking (VPN) is another preventive technology. VPN reduces the risk of man-in-the-middle attacks because all traffic is encrypted. There are different VPN options a SOC can consider. For remote devices requiring access to secured network segments, the SOC can use different options for remote-access VPN such as Secure Sockets Layer (SSL)/TLS or IP Security (IPsec). A SOC can also use VPN to encrypt traffic between different segments such as a SOC headquarters and a branch office. Table 2-1 is a quick summary comparing the differences between remote-access and site-to-site VPN technology.

TABLE 2-1 Comparing Site-to-Site VPN and Remote-Access VPN Technology

Parameter	Site-to-Site VPN	Remote-Access VPN
Concept	Uses IPsec to build an encrypted tunnel from one customer network (generally HQ or DC) to the customer's remote site.	Connects individual remote users to private networks (usually HQ or DC).
VPN Client on End Devices	Not required to be set up on each client.	Every user may (client VPN) or may not (clientless) be required to have own VPN client.
Tunnel Creation	Each user is not required to initiate to set up VPN tunnel.	Each remote-access user needs to initiate to form VPN tunnel.
Target User	Office LAN users of branch office need to connect to servers in HQ.	Roaming users who want to access corporate office resources/servers security. Employees who travel frequently.

Encryption/ Decryption	The VPN gateway is responsible for encapsulating and encrypting outbound traffic, sending it through a VPN tunnel over the Internet to a peer VPN gateway at the target site.	The VPN client software encapsulates and encrypts that traffic before sending it over the Internet to the VPN gateway at the edge of the target network.
Technologies Supported	IPsec	IPsec and SSL
Multiple User/ VLAN Traffic Flow	Allows traffic from multiple users/VLANs to flow through each VPN tunnel.	Does not allow multiple user traffic to pass through each VPN tunnel.

Preventive Technology: Network Access Control

Network access control (NAC) is a preventive technology that prevents an unauthorized device from connecting to the network. Because potential attackers don't have an authorized device, they can't connect to the network to attempt an attack. NAC can also enforce the principle of least privilege through provisioning predefined access to devices based on various characteristics. Networks for guest users can be limited to only what guest users need, such as only Internet access with limited bandwidth, while trusted employees can be placed on a network with more access capabilities. Controlling access is fundamental to enforcing proper network security.

Preventive Technology: Data at Rest/In Motion

Data at rest and data in motion technology can protect data from unauthorized access. Data at rest technology encrypts files containing sensitive data, preventing adversaries from having access to such data. Data in motion solutions, commonly called data loss prevention (DLP) technology, can

prevent sensitive data from leaving the organization using pattern matching, such as identifying data with credit card numbers or other sensitive data. Cloud access security broker (CASB) technology can also perform data in motion capabilities by controlling who can access data within a cloud service to prevent it from being exposed to adversaries.

In summary, *preventive* means blocking the attack before it happens. Preventive technology can provide value at all stages of the Cyber Kill Chain and needs to be included in your SOC's defense-in-depth strategy.

Detection Technologies

Preventing any attempt to exploit a computer, network, or communication is ideal, but prevention technologies are not 100% effective due to many factors, including gaps in preventive capabilities as well as how adversaries are continuously changing their tactics. Detection technologies are designed to detect malicious behavior within networks and systems. A SOC needs to include detection technologies for a few reasons:

- Detection technologies provide validation that preventive technologies are or are not preventing adversaries from exploiting systems, networks, and communication.
- Detection technologies provide another layer of defense, reducing the risk that an adversary can successfully complete all the steps required to exploit an intended target.

Another value of some detection technology is having the ability to learn about exploitation techniques to improve preventive capabilities. For example, if a detection capability identifies a file as being malicious based on its behavior, that security tool can share the hash value of the file with other preventive tools so other security tools can block the file based on the identified hash value.

Detection technologies typically use one or more core capabilities, such as signature-based detection of known threats, malicious behavior detection, and anomaly identification. Antivirus and IDS/IPS are examples of capabilities that attempt to block exploitation behavior. The method used to detect malicious behavior will vary depending on the type of security tool and

vendor offering it. For example, some antivirus offerings only scan files for threats. Any threat that doesn't use files, such as PowerShell exploitation that occurs within memory, would not be detected by a file-based antivirus.

Detection Technology: NetFlow and NBAR

NetFlow and Network-Based Application Recognition (NBAR) can be used to detect threats based on how devices behave on the network as well as deviations from a baseline of those devices' normal behavior. Both of these approaches offer a lot of value but have limitations in the details they provide about a potential threat. Packet capturing technology can provide more details than NetFlow, such as the specific file that was lost during a potential data loss incident, where NetFlow would only be able to provide records of the parties involved with the event. Records can include who sent the file, what protocols were used, the type of data, and other metrics depending on the technology being used and type of flow being collected.

Detection Technology: Baselines

User behavior can be monitored using detection capabilities. This technology baselines user login trends, what systems are accessing different parts of the network, where the user is located, and what level of access they should have. For example, if a user logs into the network from the United States and moments later logs into the same system from Russia, two different users likely are using the same account, meaning an account login has been compromised. Similar technology is becoming popular for cloud environments, particularly the previously mentioned cloud access security broker (CASB) that monitors access to services such as Dropbox and Salesforce.

Detection Technology: Honeypots

As introduced in Chapter 1, other popular breach detection tools are honeypots, systems designed to attract adversaries and malware. Like bears are attracted to honey, the concept of a honeypot is to attract attackers to a decoy system, diverting their attention and effort from valuable production

systems. Malware and adversaries tend to attack what is perceived as the easiest target, and the honeypot system is configured with tons of vulnerabilities to catch their attention. The honeypot is accessible via the network but is firewalled from the rest of the network systems. This tactic can be very effective; however, some malware is designed to target specific behavior rather than exploit the weakest system on a network to avoid compromising a honeypot that would alarm the target of the presence of the adversary within the network. Honeypots also pose a risk to the network if they are not configured properly, such as configuring a honeypot with a public IP address while placing it within a trusted network. Any adversary outside of the trusted network can use this misconfigured honeypot as a gateway into the trusted network.

Chapter 1 covered security capabilities and how a SOC needs to layer different capabilities to develop a defense-in-depth approach to security. Defense in depth includes layering different types of detection and prevention technologies.

Mobile Device Security Concerns

Today, the average employee has more than one device. Many of these devices are mobile and yet contain the same level of sensitive data that an authorized desktop has access to, creating a need for similar levels of security to be enforced between trusted mobile devices and desktops. One popular method to secure mobile devices is to use mobile device management (MDM) technology. SANS Security Awareness, a division of the SANS Institute, provides its recommendations for securing mobile devices in the blog post “Mobile Device Security (https://www.sans.org/security-awareness-training/blog/mobile-device-security). The following list is a shortened version of the security recommendations included in the blog post and features you will want from the MDM technology your SOC chooses to leverage:

- Enforce passwords
- Screen lock timeout
- Remote lock and remote wipe

- Password reset
- Deployment of settings, policies, and applications
- Remote monitoring
- Logging for compliance

For the most part, many MDM vendors offer these commonly requested features. Sources such as PCMag that provide evaluation of MDM vendors base their decisions on how users and devices can self-register, how policies and settings can be pushed out to devices, how lost devices are located, and data security. Choosing the right MDM vendor will come down to how their cost and capabilities meet your business requirements.

An alternative to installing clients such as an MDM agent on mobile devices is to create portals that any device can access. This limits exposure of an untrusted mobile device from infecting the network with malware or other threats since the portal is limited to what it has access to. An alternative to a portal is to provide a remote system only keyboard access to a system within a secured network segment. By only providing keyboard access, threats such as malware that have infected the remote host can only see keystrokes to the system within the secured environment. The remote system will never have access to the secured environment, denying the malware the ability to access the secured system and environment.

Planning a SOC

Developing your SOC objectives, defined through a mission statement and scope statement (or a combination mission and scope statement), sets the bar for what your SOC will accomplish. Aligning security tools, policies, and procedures builds a foundation for how your SOC services will be delivered. The next challenge to address is how to support the daily operations of your SOC. Questions such as where analysts will work, how much compute power is needed, and how the SOC will be secured need to be addressed before your SOC can go live. These topics also need to be continuously evaluated to ensure the plumbing and foundation of your SOC keeps up with the increase in demand for various resources.

Capacity Planning

The first SOC foundation planning concept to tackle is understanding the capacity requirements you will need today as well as in the future.

Capacity planning involves developing the physical and technical requirements for building a SOC:

- Physical requirements include the location(s) to physically host the SOC employees and technology, power requirements, space to hold equipment, seating arrangement, how trash is handled, physical security considerations, and other items that are used within a SOC.
- Technical requirements include network throughput, types of technology, monitors, and other technical SOC needs.

These requirements will be based on supporting the services outlined in the SOC mission and scope statements. For example, if monitoring a local datacenter is part of the scope statement, then physical and technical requirements will include what is required for a SOC analyst to perform that work. If capacity is not properly assessed, you might encounter complications such as lack of network bandwidth or lack of rack space to secure and power the tools needed to perform monitoring. If capacity planning does not take into consideration a healthy SOC layout, problems ranging from a noisy work environment to lack of visibility to monitors displaying critical information will lead to a dysfunctional SOC. For these and many other reasons, it is critical to properly plan the capacity requirements for your SOC.

SOC Goal Alignment

Designing a SOC depends on a handful of factors. First off, the mission statement and scope statement set the boundaries for the location(s) where the SOC will operate. These statements also identify the different services expected from the SOC. Chapter 3 covers SOC services in more detail; however, understand that before a service can be offered, there must be people, process, and technology available to perform the service. That means the people and technology must reside somewhere and there will be expected resources available. Capacity planning ensures that those requirements are met.

Chapter 1 covered how to assess existing capabilities and available technology based on best practices from industry standards, guidelines, and frameworks. Assessing existing capabilities and available technology provides a blueprint of the current state of the SOC. Chapter 1 also covered how to use maturity modeling to develop a plan with milestones to improve SOC services. Those maturity goals include how the future desired SOC will look regarding people, process, and technology:

- People goals reflect the number of people for required roles as well as the expected skillsets. Some skillsets might be taught on the job while other skills will be required the day the new hire takes the role.
- Process goals impact what is required to deliver services. This includes what policies need to be enforced and how those policies are enforced.
- Technology goals complement the people and process goals, acting as an enabler to accomplish goals for SOC services. For example, a goal to monitor the network for threats requires technology that permits monitoring network traffic as well as tools for identifying when malicious activity occurs.

Growth Planning

Maturity goals influence the SOC's capacity planning requirements. Planning must account for required hardware, number of employees, how teams will collaborate, seating arrangements, and many other factors that are needed to properly deliver a specific SOC service at the expected level of maturity. When considering the physical location for the SOC, one key principle is that choosing a location that the SOC can grow into usually is less expensive than having to move to a larger location at a later date because the SOC has outgrown its initial space. If the SOC capacity planning team is unsure about future requirements for the physical location, it is ideal to predict, at a minimum, a three- to five-year growth percentage and develop the location around those numbers. An example could be viewing the maturity goals for the SOC and target each SOC service to be at a specific state three to five years from the current date. Capacity planning must take into consideration what changes are needed in people, process, and technology to be at the three- to five-year maturity state. Using this approach leads to a SOC design

that has a percentage of open seats, growth room for increased bandwidth usage, and available rack space for future technology.

Key Point

Maturity models must be part of capacity planning. More advanced phases of a maturity model will require certain resources in order to be executed properly. This is why before you perform a capacity plan, you first develop your maturity model(s). Then you can align your maturity goals with expected capacity requirements so your SOC sponsor understands your capacity requirements today as well as what they could be in the future.

If your SOC budget does not allow for including future growth in a capacity plan, there are other options that can be used as the SOC outgrows its physical location. Options include permitting teleworking, opening a new branch facility that is connected to the SOC headquarters, and/or outsourcing some services. Any of these options is more ideal than having to migrate an entire SOC and its people and technology to a larger facility while maintaining its services.

Technology Planning

Another capacity planning factor is related to the technologies used by the SOC. As technology capacity planning is performed, it is recommended to set up meetings with vendors of existing equipment within your organization so they can help future-proof their equipment. Vendors will also know the best options for configuration and hardware that meet your SOC maturity goals. It is best to start with existing vendors because they are less likely to have to "sell" against anybody else and they might already understand your environment. Even by doing so, you still might encounter a vendor attempting to upsell or displace other vendors within your organization. For example, if you have products from two different firewall vendors in your environment and invite each vendor to help you with capacity planning, each vendor will offer a future design that replaces the competitive firewall.

Another option for technology capacity planning that captures a vendor-neutral viewpoint is to seek consulting from technology service partners, also called resellers, that broker sales between your organization and vendors. Many enterprise vendors do not sell directly to customers but instead rely on resellers, which can provide the SOC with a comparison between industry leaders and advise on what other SOCs are doing in the specific technology areas of interest. Knowledge is power and there are likely many experts willing to speak with you at no cost regarding the best ways to design different parts of your network.

Before starting conversations with vendors or specialists, you will want to verify all hardware and connectivity requirements based on your SOC maturity goals. Remember that the requirements you develop are a starting point for the design; be prepared for outside specialists to suggest change. For example, some products can include proprietary components or have special requirements such as additional power. It is best practice to lead with technologies that are flexible by supporting application programming interfaces (APIs) and open standards; however, you will want to ensure capacity planning accommodates all requirements, including those systems that are not flexible regarding APIs and open standards. Sometimes you will find that some proprietary technology may be much cheaper and already used within the organization, while other times investing in the effort to convert to a new system will make more sense for your SOC. Chapter 1 provided details on how to assess for capabilities, plan for maturity, and rank which investment would provide the most impact to the organization. The results from these exercises will provide guidance for which technology options are best for your SOC.

SOC Resource Planning

It is important to include all teams that will work with the SOC during the SOC's capacity planning. There might be areas of cost savings that could be accomplished by leveraging other teams' resources, including using other teams as designated backup options. For example, another team might agree that its office space could serve as the overflow or emergency workspace for the SOC rather than having to build or rent a redundant location. Having conversations with other teams will involve budget, which can lead to sharing

the cost of hardware and service support and consolidation of existing hardware. Working with teams outside of the SOC can help overcome challenges to certain SOC goals. An example is the SOC using the organization's datacenter to host and manage its technical equipment rather than having the SOC responsible for datacenter management tasks.

Some capacity planning decisions will be influenced by whether in-house or external services are used. Once again, vendors and resellers can advise on the benefits from either approach, but keep in mind that vendors want to sell you something. Rely on your maturity assessments (discussed in Chapter 1) to help keep the conversation nonbiased and in the best interests of the SOC. Reasons for using external services in response to capacity planning are reductions and potential savings in hardware, personnel, required process, and compliance.

Redundancy Planning

Another capacity topic that you must address is the requirement for redundancy. Redundancy decisions are based on the SOC's risk appetite (for example, backup systems can be configured as active, passive, or on demand) and redundancy plans are developed using in-house services, outsourced services, or both. Chapter 1 covered how to rank the importance of capabilities and services against the goals of the organization. Anything ranked as high importance requires active and more costly redundancy, while lower-ranked services and technology require passive, less expensive redundancy options or just an on-demand option if a backup of a service or technology is needed due to a failure in the primary offering. The section "Disaster Recover," later in this chapter, looks at disaster planning concepts in greater detail.

Developing a Capacity Plan

To properly perform capacity planning, you should document your capacity plan. Begin by including the time, date, and version of the capacity planning report so readers know how relevant the data is that they are viewing. Next, you need to provide executive statements and/or purpose statements, so

readers know why you performed the planning. For example, you want to explain whether the capacity planning is for a new SOC or to future-proof an existing SOC. Another example is explaining that the capacity planning report will define what is needed to improve an existing SOC service to a desired maturity goal. The quality of the executive statement for this report can be the determining factor if executives approve or deny the resources being requested.

When stating capacity specifics in your capacity planning report, you should list the capacity type, current analysis, and any expected growth or changes. You also might want to include other subsections depending on the item being analyzed, but at a minimum, these three capacity specifics must be in the capacity planning report for any items being requested. Once all items are logged, you need to provide details about each item being analyzed. This can include where it is located; what the specific requirements are to operate; what changes are needed to meet the future requirements; any threshold limitations based on hardware, software, or services; and a summary of the strategy that can be implemented to increase the capacity to the desired state. Other details to consider are budget, compliance planning, and an appendix to define key terms or other reference items.

The document in Figure 2-9 is a sample template for capacity planning. You can adjust these questions to meet your SOC's goals.

Version Number: 1.0

Version Date: mm/dd/yy

VERSION HISTORY

Version Number	Implemented By	Revision Date	Approved By	Approval Date	Description of Change
1.0	<Author name>	<mm/dd/yy>	<name>	<mm/dd/yy>	<description of change>

TABLE OF CONTENTS

Content

EXECUTIVE SUMMARY

Content

PURPOSE OF CAPACITY PLAN

Content

ANALYSIS OF CAPACITY

Capacity Type	Current Capacity Analysis	Planned/Expected Growth and Recommendations

FINDINGS SUMMARY

Area/Item Monitored	Capacity Requirement(s)	% Increase Needed Per <time period>	Capacity Threshold(s)	Threshold Response Strategy (Action to Be Take Upon Reaching Threshold(s))
<Hard Drive Storage>	<enter capacity requirements and measures>	<enter projected increases over intervals of time>	<enter acceptable capacity threshold(s)>	<enter response strategies to varying threshold limits. Threshold is defined as the level at which an event or change occurs>
<Meeting Room Tables>				
<Number of Project Staff>				
<Ratio of Quality Assurance Staff to Development Staff>				

BUDGET MANAGEMENT

Content

COMPLIANCE REQUIREMENTS

Content

APPENDIX

Content

FIGURE 2-9 Capacity Planning Template

There are additional items that you could include with your capacity planning documentation, such as success criteria and a ranking system that is ordered starting with what is most important to the SOC down to what would be a "nice to have." This additional data will help answer questions such as how to invest limited budget into the SOC, meaning which capacity needs should be addressed now and which should be left for a future budget.

Designing a SOC Facility

This section looks at points of consideration as you develop a new facility or modify an existing facility in which the SOC functions. A great resource for recommendations in this area is the Whole Building Design Guide (WBDG) website (https://www.wbdg.org/) of the National Institute of Building Sciences. A 2017 WBDG article titled "Office Building" identifies the following (among others) as important design considerations (tailored here to SOC facility considerations):

- **Cost-effective:** Assess the performance versus cost of each design element and building component. In some cases, it will be ideal to invest more initially to save on long-term operations and maintenance.
- **Functional/operational tenant requirements:** Consider integrated requirements of the SOC staff. These include the desired image, degree of public access, hours of operation, growth demands, security issues and vulnerability assessment results, and other long-term items.
- **Flexibility:** The office building must easily accommodate frequent renovation and alteration. It must offer easy access for changes, such as using raised floors to permit access to cabling and power distribution.
- **Urban planning:** Consider the value the local neighborhood offers as well as the impact of the SOC to the neighborhood. How far will employees have to travel to purchase lunch? How close is a fire department or a police station? Will building the SOC and associated facilities such as the parking lot negatively impact the surrounding environment?

- **Productive:** Consider worker satisfaction, health, safety, and comfort. Will SOC employees be comfortable and satisfied with the workspace? How easy will it be for the SOC to have meetings with internal and external people?
- **Technical connectivity:** Ensure flexibility for IT infrastructure, allowing technological access and power for SOC equipment. Considerations include wired and wireless options, such as whether wall material permits wireless signals and whether wired network ports are available throughout the building.

The WBDG recommendations apply to an average business developing a workspace. There are some additional key points that should be included with the WBDG list of recommendations based on the needs of a SOC. Additional items include being able to facilitate effective collaboration between SOC members, maintain a secured SOC, adhere to compliance requirements, and accommodate the requirements from your capacity planning reports. Collaboration between SOC members includes how data is shared between individual team members as well as the large screens providing data to the entire SOC floor if a floor monitor is needed. Security for a SOC can depend on the goals of the business. I have seen many SOCs containing classified data require physical lights to be turned off when people without clearance are permitted on the SOC floor.

These examples are just some of the many factors to keep in mind as you consider the capacity planning requirements and building recommendations (such as those from WBDG) while designing the facility for the SOC. Some design features and functions will be mandatory and rank higher in importance than others; however, you should not overlook lower-ranked design features and functions. For example, a comfortable working environment might not be a mandatory requirement, but it will dramatically improve job retention, saving you down the road from having to respond to complaints and employees leaving due to an unsatisfactory work environment. I have seen people leave an organization even after receiving a promotion because they were not offered their own office space. Remember that employees will be spending a large part of their lives in the SOC facility. Investing a little extra into making employees happy can go a long way toward keeping them satisfied.

Physical SOC vs. Virtual SOC

A major factor regarding how much physical space a SOC facility requires is whether you plan to incorporate a virtual SOC design or external consulting services. A virtual SOC might mean that team members are working from various locations that are all connected through collaboration technology. In this example, a physical facility might be needed only for the SOC headquarters or an emergency workspace in the event of a major incident requiring virtual members to meet in a physical location. An example of a major incident could be a security incident that disrupts the organization's profits, renders employees unable to perform their work, or the loss of company classified data. These situations would demand an "all hands on deck" approach requiring field members to travel to the SOC headquarters to form an ad hoc group, commonly referred to as a "tiger team," to deal with the situation.

Outsourcing services can mean a virtual SOC design where the SOC analysts work from their own office and only physically show up during a major incident or quarterly meetings. Factors that contribute to whether remote work is allowed include security requirements, policies against data leaving the facility, and technology limitations. As an example, the U.S. government has specific requirements for handling certain levels of classified material. As content is deemed higher in classification, facility security requirements increase, including limiting access to such data as well as where the data is stored.

Chapter 1 covered the advantages and disadvantages of using in-house SOC services and external SOC services. Facility considerations are similar. Table 2-2 provides an overview of the benefits of using an internal facility (i.e., a single physical location for SOC employees to work from) or external facilities (leveraging smaller offices or permitting employees to work from home) for the SOC. This does not take into consideration the benefits and costs associated with outsourcing services, as explained in Chapter 1.

TABLE 2-2 Benefits of In-House and Outsourcing the SOC Facility

Internal	External
Face-to-face team interaction	OPEX savings (office space and

	resources)
Direct access to interact with employees	Further reach for recruitment
Local support for organization	Employment in different time zones to accommodate 24/7 support
Options for live events (team building/lunch and learn)	Job hour flexibility

SOC Location

Another decision to consider is where to locate the SOC facility or facilities. Some companies might find value by placing the SOC near the network operations center (NOC) or desktop support team. Smaller organizations might use a shared floorspace and leverage cloud resources for redundancy purposes. The previously referenced WBDG recommendations point out urban considerations, which include the surrounding neighborhood. I have seen organizations located where food, childcare, and other resources that employees need during a workweek were not accessible within an hour's drive. Employees would have to develop unappealing work schedules to accommodate childcare or take extremely long lunch breaks, which leads to an unhealthy work environment. If a fire department isn't located within a reasonable distance, the SOC needs to have a fire prevention strategy that doesn't rely on waiting for the local firefighters.

Outside of the local neighborhood are country-based restrictions to consider. Certain countries have data sovereignty requirements that restrict data from leaving the country. Other country-based restrictions include laws or culture aspects impacting local employees of the organization. Maybe certain procedures that are legally acceptable in one country would not be acceptable in another country and therefore need to be adjusted. For example, could a Dutch-based SOC monitor Chinese citizens? All of these factors must be evaluated as the locations for the SOC facilities are considered.

SOC Interior

There are interior aspects of the SOC facility that need to be considered against the capacity planning report and recommendations such as those from WBDG. The same decision-makers who decide how the SOC is designed need to be responsible for developing the SOC's mission and scope. This includes SOC managers that will run the SOC, executives that support the SOC, and facility specialists that understand how to create a facility that meets what is being requested for the SOC. Some decisions will concern aesthetics, while others will be based on politics or budgets. Examples of decisions regarding budget and politics are how large a manager's office should be and what type of workspace SOC analysts should have. Aesthetic considerations include the amount of natural light compared to artificial light. For example, security requirements could influence the decision to not use natural light based on the light ports being a potential vulnerability to physical intrusion.

SOC Interior Design Considerations

The following list highlights key aspects to consider for the SOC's internal design:

- **Lighting:** Lighting should promote a comfortable working environment permitting the analyst to not only see well in his or her workspace but also the video wall monitors, if applicable. Decisions regarding natural lighting and artificial lighting are made based on the building design, budget, and security requirements. Specialized lighting might also be required for certain SOC functions, such as viewing server rack space or cables under a raised floor.
- **Acoustics:** Poor acoustics can be very disruptive to daily operations. SOC analysts will be collaborating over the phone and within their teams, which will create noise. Noise levels need to be accounted for so that analysts can hear clearly during phone conversations. Also consider noise from equipment, which may necessitate placing it in a separate room such as the computer room rather than in the SOC analyst work area. Soundproof material for the walls could be used to reduce sound levels as well as protect conversations from being heard by unwanted parties.

Audio privacy rooms can be built as an escape from noise as well as to ensure privacy of conversations. ISO 17624:2004 also provides guidelines for leveraging acoustical screens as an option to reduce office noise. All of these options can be used to promote a healthy and productive work environment.

- **Security:** Every SOC will have specific requirements about who can and cannot access the operations room. Physical security must be part of the design of the SOC facility. Common options are door locks, common access card (CAC) or other smartcard door locks, video monitoring, and mantraps. Additional specialized lighting might be required for certain situations such as when visitors without clearance are in the SOC facility. I have been in SOCs that use police lights that are enabled when an uncleared person enters the SOC, notifying all analysts to turn off their monitors and the video wall screens while the person is present. External consultants can assist with meeting your SOC's security requirements and can advise on any constraints such as if CCTV is not permitted within a country.
- **Secure disposal:** People will produce trash that contains sensitive information. A SOC must have requirements for properly shredding and disposing of such trash. Dumpster diving is still a threat to modern organizations. Purchasing professional trash disposable services or tools such as shredders to accommodate securing trash are some options that are available.
- **Storage:** Certain types of data might require specific storage requirements. For example, classified documents or forensic evidence will have to be bagged, tagged, and stored securely. Secure storage needs can require a safe or other storage options, perhaps even a specialized part of the building with armed guards, known as a sensitive compartmented information facility (SCIF), discussed further in the next section.
- **Video wall:** Some SOCs use monitors that enable multiple analysts to view the same data simultaneously. This means having view of multiple dashboards, which will likely need to be placed on a wall so all analysts can see the data. For example, a SOC responsible for various services may display the dashboard for a SIEM or other tools that take into consideration the SOC's current event workload as well as external data

such as global events. Having a centralized monitor allows each SOC analyst to focus on their own responsibility while always having a view of data that applies to the entire organization. To maximize collaboration, analysts should have access to the same data on their personal computer as well as be able to share what they are looking at on a shared video wall. Make sure that the video wall is high enough that all analysts can view it, including if somebody in the front row is taller than an analyst in the back row. Also consider future systems and analysts that might be added according to the capacity planning report. Figure 2-10 shows an example of how to design for these items.

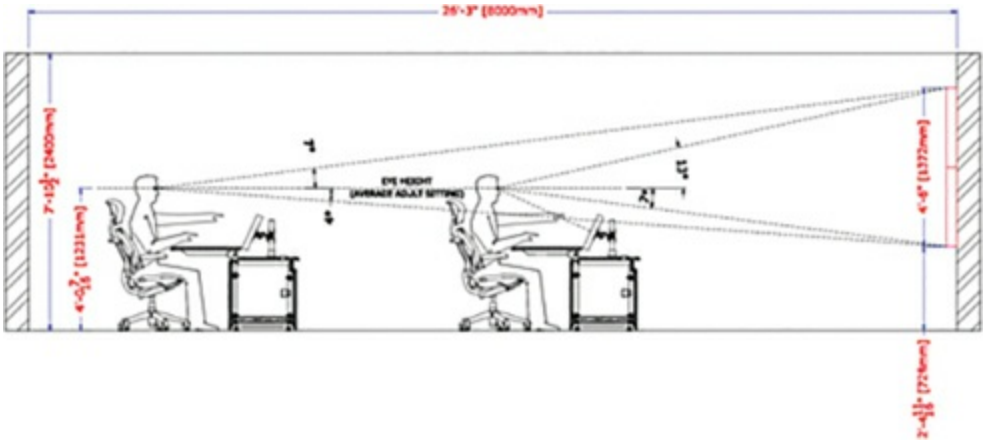

FIGURE 2-10 Video Wall Design Example

- **Workstations:** Consider the workspace for each SOC employee type. For example, should it include a phone, a desktop computer, and possibly a mounted flat-screen monitor? Some analysts will use a laptop only and not need a monitor, while other analysts will prefer having multiple monitors. Consider the number of cable holes that are available to help conceal power and network cables. The layout of how workstations will look can be divided into separate cubical work spaces, grouped into a large workspace, or possibly be a long table that permits open seating. Many office equipment manufactures offer options for configuring workspaces, which is ideal so that the SOC can adjust to the most effective layout (which fulfills the previously discussed WBDG recommendation for flexibility). Speak with an office furniture specialist for more information about workstation furniture options.

- **Lockers:** Providing lockers gives SOC personnel a secure means of storing personal items. There may even be compliance requirements, such as no mobile devices permitted in the SOC, that necessitate provision of lockers. Lockers can also be placed within personal workspaces for storing smaller personal items and documents.

A final factor to consider regarding the value impact of the SOC's interior design is the SOC's aesthetic appeal. Do not underestimate the value of a first impression when people walk through a functioning SOC. Many SOCs offer the option to provide tours, enabling non-SOC members to gain awareness of the SOC and its associated activities. Investing extra effort in aesthetics can capture additional funding and support from leadership, who may include the SOC in tours of the company to showcase the company's investment in security.

SOC Rooms

Another SOC facilities design consideration is the different types of rooms within the facility. These rooms could reside within the SOC headquarters, at a remote branch office, or at some other type of remote facility such as an employee's home office. The following are room types to consider for your SOC (all of which might not be required for your SOC):

- **SOC operations room:** This is the SOC floor occupied by analysts and other SOC members doing their daily tasks. This room can house various dashboards, incident boards, analyst workspaces, and other operational items. If a video board is being used, this room should offer unrestricted line of sight to the video board for all SOC members. Make sure to size up versus down or you might end up with a crowded floor that ruins the atmosphere of the SOC workspace.
- **War room/situation room:** This room is used for meetings. It should have collaboration technology such as videoconferencing and white boards. Most modern war rooms offer the capability to dial in remote users and share their desktop or the desktop of somebody in the room, so everybody is able to work from the same screen. One aspect to consider is the sensitivity of content that might be covered within the war room. Additional security measures might be required for the walls, door, or

collaboration technology based on sensitivity needs.

- **SOC manager's office:** It is common to design the manager's office with a clear view of the SOC floor and video wall as well as walls and a door to provide privacy when talking about sensitive topics.
- **Digital forensics lab/clean room:** A room that is dedicated to digital forensics work requires additional security, including documenting any access. This is critical for tracking chain of custody, which requires that all access to evidence be documented; otherwise, it can be considered contaminated from a legal perspective. Other aspects of a digital forensics lab include additional power sources to keep equipment that has been obtained with the system's power still on and plugged in so it doesn't power off, fume hoods, space for large lab tables, and higher cleaning standards. Resources for better understanding how to set up a digital forensics lab/clean room include ISO 31000 (risk management), the ISO 14000 family (environmental management systems), the ISO/IEC 27000 family (information security management systems), and OHSAS 18000 (occupational health and safety), depending on the type of lab you are looking to build.
- **Sensitive compartmented information facility (SCIF):** A SCIF (pronounced "skiff") is required if the SOC plans to work with sensitive compartmented information (SCI). SCI represents resources pertaining to sensitive information that is required by law to protect. Technical specifications for a SCIF take into consideration physical security; heat that radiates from computers, known as tempest; alarms, locks, and safes for containing sensitive information; and storage for all unauthorized devices such as phones, cameras, and computers. These requirements depend on whether the SCIF is for handling British classified data or U.S. classified data and what type of data will be handled. The physical construction, access control, and alarming of the facility has been defined by various government directives, which must be met and audited before sensitive data can be contained with the SCIF.
- **Computer room:** The SOC equipment such as servers and equipment for private networks used for testing need to be hosted somewhere. Sometimes the SOC's computer room is part of a shared space such as a datacenter, while other times the room is a space dedicated to SOC

equipment.

One specific room that is critical for the operation of the SOC is where its technology is housed. This brings us to the next topic of designing computer rooms for the SOC.

SOC Computer Rooms

Developing the computer room includes many requirements to consider. The SOC might or might not be responsible for building the computer room it will use, but it should at least be involved with planning it. There are specific regulatory compliance requirements for versions of hardware and software that are used within the datacenter. Regulatory compliance could also control how a SOC's data is architected for segmentation as well as require use of certain levels of encryption to protect traffic. Certain specialized equipment may need to be racked, powered, and secured, which must be accounted for.

Computer Room Considerations

The following key topics must be addressed as planning occurs for a computer room that will be used by the SOC:

- **Power requirements:** The first step for planning power requirements is to determine the number of services to install per rack. Some equipment might need special power, such as more current, that would require an electrician to set up. There might be requirements for power strips where large amounts of gear will be located. To avoid blowing fuses or equipment, I recommend gathering your power requirements based on capacity planning reports and speaking with an electrician about what can and cannot be supported. Power redundancy is also an important topic. Impacting factors include the type of redundancy that can be provided, the budget for redundancy options, the available power supplies in equipment, and risk tolerance. One option for power redundancy is a single redundant option, meaning if a primary system powerline goes down, a backup power source takes its place. Grid redundancy is another option, which means to have twice the number of power supplies on the same system, with each power supply using a different AC power source.

Power loads are dependent on the facility input plug. If the power load exceeds the rating on the input plug for a sufficient period of time, the input breaker will be triggered, causing the power to be disrupted. For this reason, power plugs need to be evaluated to ensure they can support expected power requirements based on capacity planning needs. This includes consideration for the different types of plugs used in different countries.

The Uptime Institute offers estimates for the expected power usage of different size computer rooms, which can be helpful for budgeting for power. For example, the Uptime Institute estimates that 1 kilowatt of server capacity costs US$25,000 for a Tier IV datacenter. See https://uptimeinstitute.com/ for a better understanding of estimated costs for powering a computer room and associated equipment.

- **Temperature and humidity:** The American Society of Heating, Refrigeration, and Air Conditioning Engineering (ASHRAE) Technical Committee 9.9 has created a widely accepted set of guidelines for optimal temperature and humidity set points in the datacenter. For example, the ASHRAE guidelines specify acceptable temperature ranges for the datacenter to prevent IT gear from overheating. Cooling options include server racks that have built-in cooling, cooling the server room, and HVAC filters that suck heat out of a computer room or individual server rack space. Warm rooms can produce mold, which is bad for equipment. You might need humidity units, which will require some method to be drained as they collect water. It is best to have a sink that humidifiers can drain in rather than manually having to empty a humidifier.

  Table 2-3 represents ASHRAE recommended temperature and humidity ranges for a computer room.

TABLE 2-3 ASHRAE Recommended Temperature and Humidity Ranges

Property	Recommended Value
Lower limit temperature	64.4°F [18°C]
Upper limit temperature	80.6°F [27°C]

Lower limit humidity	40% relative humidity and 41.9°F (5.5°C) dew point
Upper limit humidity	60% relative humidity and 59°F (15°C) dew point

Measuring temperature and humidity needs to be properly done to ensure that accurate readings are taken relating to all systems within the computer room. One side of the computer room may be hotter and moister than another side, so both sides need to be monitored. The design of the computer room can also impact airflow, permitting better control of temperature and humidity. Figure 2-11 is an example of air flowing through a computer room.

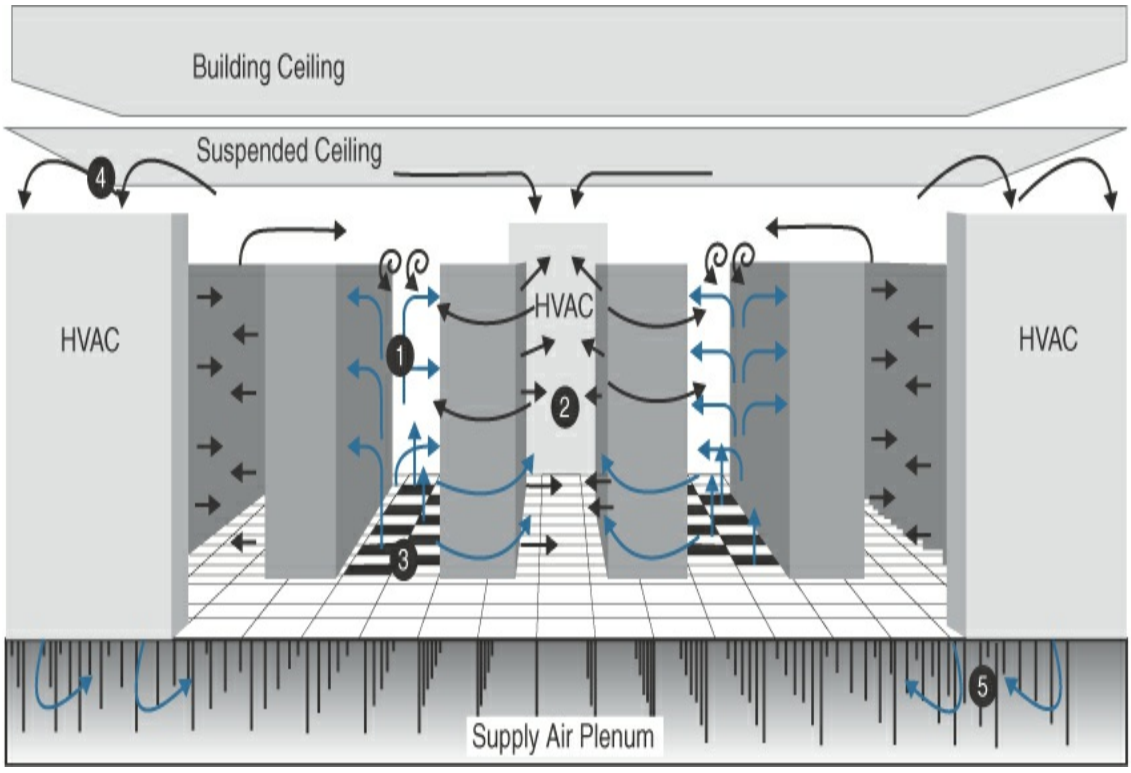

FIGURE 2-11 Example Airflow Within a Datacenter

Cisco published a site planning guide titled “Data Center Power and Cooling White Paper” for power and cooling that includes recommendations for the location and use of equipment racks in regard to their impact on temperature and humidity. One recommendation is to

consider the arrangement of hot and cold aisles of a server rack. Arranging racks into rows of hot and cold aisles minimizes the mixing of air, reducing temperature and humidity. If two hot aisles are permitted to mix air, cooler air-conditioning requirements will be needed to compensate for the additional heat. Figure 2-12 is an example of how a hot-aisle and cold-aisle layout could look.

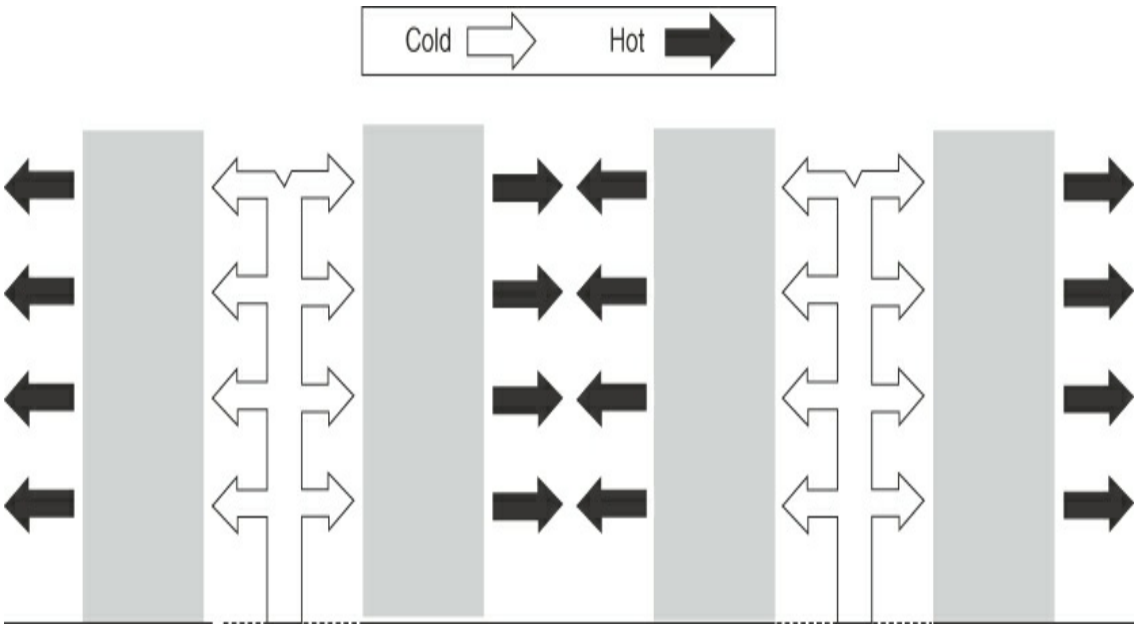

FIGURE 2-12 Sample Hot and Cold Aisle Design

The Cisco guide suggests powering the heaviest and most power-dense equipment at the bottom of the rack. This helps lower the rack's center of mass to reduce the risk of tipping. Power-dense equipment tends to draw more air and need the most cooling. Many datacenters have cooling near the bottom of the rack to give the hottest equipment the coolest air. Cisco also recommends, when possible, to leave unoccupied space in the rack between servers to permit airflow. Unoccupied rack space can be filled with blank panels to keep the airflow within the server space.

- **Equipment racks:** Different types of equipment have different rack requirements. Server racks are different from network gear racks. Make sure to consider what type of gear would need to be stored or racked. Racks that have perforated front and rear doors to maximize air flow are ideal. You might need cable plants that can accommodate different types

of connection types such as fiber and CAT6 with booted RJ45 connectors. Optional casters should exist to permit rack mobility if needed. Locks should be available to secure access to the equipment within the rack. Consider removable rack doors to simplify maintenance of systems within the rack as well as the rack itself. Finally, as previously discussed, consider the heat and humidity impact based on how racks are positioned within the computer room as well as the equipment within the racks.

- **Lighting:** Consider how the computer room is lit. Cabling can become extremely difficult to identify in a poorly lit computer room. Make sure your lighting accommodates people and the types of spaces they will have to work out of, meaning that if a tech is expected to go within a rack, there should be lighting that allows the tech to have full visibility of the equipment within that space. Emergency lighting should be available if the primary facility power becomes unavailable.
- **Redundancy and UPS:** The SOC deals with sensitive data that needs to be available even during a power outage. Redundancy can be accomplished through hardware, a backup location, routing, software, and backup power. For example, a dedicated uninterruptable power supply (UPS) or an existing UPS that supports the facility the SOC resides in can provide redundant power options. Backup routing options can be designed so that if a primary system goes out, traffic is routed to another location. Backup systems can sit at branch offices or within the cloud. Software can be configured to fail open or fail closed based on its business function. For example, a network access control solution could be configured to permit all connections (fail open) or deny all connections (fail closed) if a system failure occurs.

  Regarding expected costs for UPS support, a good guide is a whitepaper titled “Comparing UPS System Design Configurations” by Kevin McCarthy and Victor Avelar, which suggests potential costs associated with increasing the level of redundancy for a UPS architecture. As more scales of availability are added to the design, the cost per rack increases. McCarthy and Avelar use the following terms for UPS configurations:

  - **Capacity:** A single UPS system with no redundancy, meaning no backup system in place.
  - **Isolated redundant:** The traditional active standby two-system

design where the primary system takes on all the load and the standby system is not used unless an event occurs.

- **Parallel redundant:** Two systems activity sharing the load being used.
- **Distributed redundant:** A few UPS systems are all sharing the load in an active manner.
- **System plus system:** Systems not only are sharing the load but also have an additional local redundant system for backup if the local primary system goes down. An example of a system plus system design is having three facilities sharing the power load as well as each facility having a parallel redundant system.

  Table 2-4 represents the estimated costs associated with each UPS design according to the McCarthy and Avelar whitepaper.

TABLE 2-4 Estimated UPS Cost

Configurations	Scale of Availability	Tier Class	Datacenter Scale of Cost (US$)
Capacity (N)	1 = Lowest	Tier 1	$13,500–$18,000/rack
Isolated redundant	2	Tier II	$18,000–$24,000/rack
Parallel redundant (N+1)	3		
Distributed redundant	4	Tier III	$24,000–$30,000/rack
System plus system (2N, 2N+1)	5 = Highest	Tier IV	$36,500–$42,000/rack

- **Grounding:** All datacenter equipment needs to be grounded. This can be part of the equipment rack or require additional work. Ensure ground resistance is at least < 1 Ohm. There are many standards and guidelines that can be referenced for grounding best practices. NIST offers recommendations for grounding based on papers titled “Ground Connections for Electrical Systems” (Technologic Papers of the Bureau of Standards No. 108) and “A Consensus on Powering and

Grounding Sensitive Electronic Equipment."

- **Raised floors:** If the SOC plans to host a lot of equipment, that equipment will require cables, which can get really messy. Computer rooms can be built on raised floors to hide cables and power connectors. Figure 2-13 is an example of a raised floor tile. Another option is running cables and power connectors within the ceiling, which requires a ceiling that can support the associated weight. Speak with a facility specialist to confirm that the ceiling or floor option used to host cables will meet your capacity planning requirements.

FIGURE 2-13 Raised Floor Tile

- **Fire safety:** Different countries have different required fire safety codes. Sometimes, fire safety options can put equipment or people at risk when they are released, such as chemicals that can quickly put out fires but are poisonous to humans. Water may be safe to humans but is bad for equipment. You will have to consider all of these factors, including the different types of fires that could occur. For example, an electrical fire is different from a gas fire and requires a different form of fire safety to be put in place. You may need a professional to look at smoke detectors and carbon monoxide detectors to accommodate all of these factors. ISO/TC92, *Fire Safety*, is one of the many guidelines that you could

reference to address the different types of fire safety risk.

- **Flood protection:** The SOC facility could be at risk of flooding based on where it is located. Pumps and drains are countermeasures that can be put in place to handle a certain level of flooding. Water-resistant material and water barriers are also options to reduce the risk of flooding. The WBDG article “Flood Resistance of the Building Envelope” provides methods to evaluate if your facility is located in an area prone to flooding and describes various flooding countermeasures. Table 2-5 is WBDG’s overview of active and passive floodproofing methods.

TABLE 2-5 Examples of Active and Passive Floodproofing Methods for Buildings

	Dry	Wet
Active	Temporary flood shields or doors (on building openings)	Temporary relocation of vulnerable contents and equipment prior to a flood, in conjunction with use of flood-resistant materials for the building
	Temporary gates or panels (on levees and floodwalls)	
	Emergency sand bagging	
Passive	Waterproof sealants and coatings on walls and floors	Use of flood-resistant materials below design flood elevation (DFE)
	Permanently installed, automatic flood shields and doors	Installation of flood vents to permit automatic equalization of water levels
	Installation of backflow prevention valves and sump pumps	Elevation of vulnerable equipment above DFE

- **Monitoring:** All electronic equipment needs to be monitored to ensure that accessibility and acceptable performance levels are maintained. It is recommended to have a central monitoring console located outside of the computer room that is monitored 24/7 for this purpose. Some SOCs assign monitoring to the NOC or a third-party consulting company, while other SOCs monitor equipment using in-house resources. Monitoring can include ensuring power is being distributed, systems are reachable, and systems are functioning properly. An example of a network monitoring tool is WhatsUp Gold from Progress (formerly Ipswitch). This tool continuously pings network-based equipment and alarms administrators when a system in found to be unreachable. NetFlow tools can be used to alarm when spikes or drops in network bandwidth are identified. Cisco Secure Network Analytics (formerly Stealthwatch), SolarWinds' Network Performance Monitor, and Plixer Scrutinizer are examples of NetFlow-based network monitoring solutions.
- **Locks:** Systems in the computer room are critical to the operation of the SOC and need to be physically protected. Locks on the doors and server racks within the datacenter provide physical security and can be lock and key or based on various types of digital systems. Examples of digital systems include smartcards and biometric-based devices. Ensure that access to power and network cables is also secure. You want to avoid scenarios such as a cleaning person unplugging a critical system by accident during a standard cleaning routine. FIPS 140-3, *Security Requirements for Cryptographic Modules*, also provides guidelines to using locks to enforce proper physical security.
- **Video and access control:** A common requirement for a SOC computer room is the capability to control who accesses the room and monitor what they do while in the room. Access control can be based on physical controls such as mantraps and what type of locks are used on the doors within the SOC and its computer room. Various forms of video surveillance equipment can be used that are always active or triggered when motion is detected. Some SOCs require video surveillance to be recorded and stored for a specific period of time. The specifics of the video and access control requirements will depend on your SOC's business needs in regard to physical security for the computer room. ISO 22311:2012, *Societal security – Video surveillance – Export*

interoperability, is one of the many guidelines available that can provide recommendations for computer room physical security and monitoring options.

SOC Layouts

Figure 2-14 shows an example of what a SOC layout could look like. This example shows a SOC operation room, war room, and manager's office. In this example, the computer room is located somewhere else in the building.

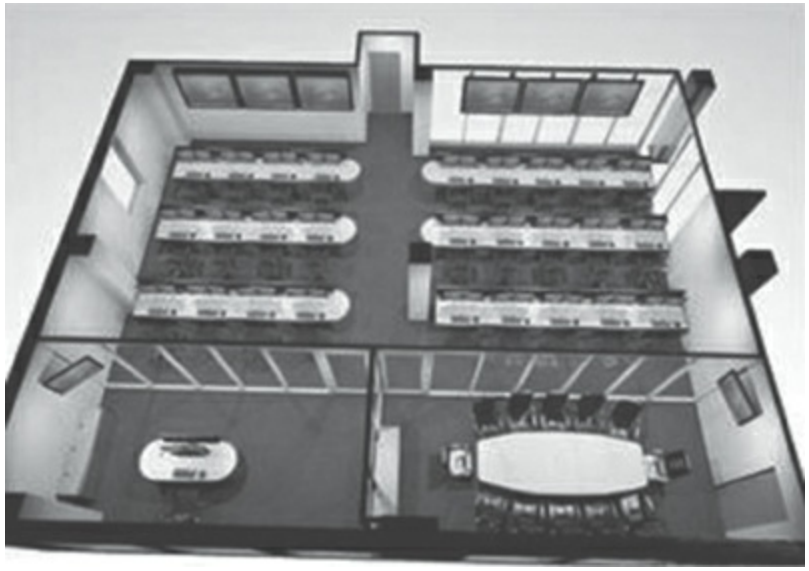

FIGURE 2-14 Sample SOC Layout

Figure 2-15 shows a draft for how the analyst area of the operation room could be designed. Notice how the center point for the design is visibility to the video wall. Some SOCs may not require a video wall, which means the design of the analyst area would be completely different from Figure 2-14. Your ideal design will be based on the results of your capacity planning compared against your SOC's business objectives.

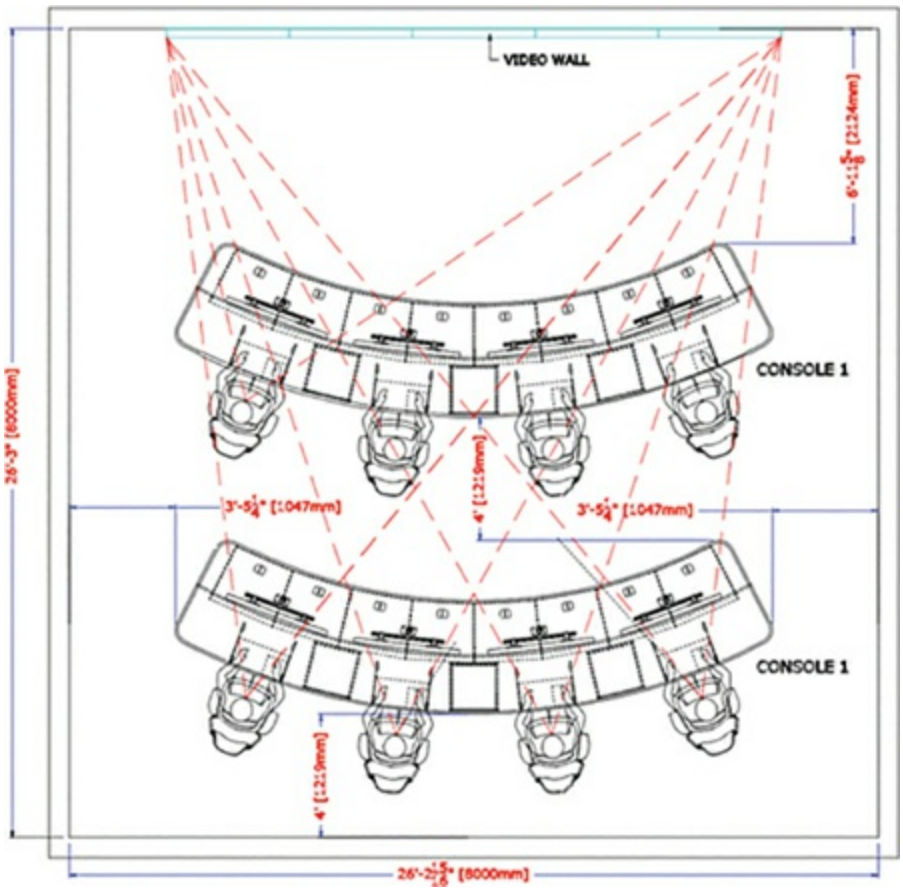

FIGURE 2-15 SOC Floor Layout

Note

For many organizations, other teams outside of the SOC will handle some or all of the topics covered in this section. There might be a dedicated facilities team that handles everything, or the previous tasks might be divided among multiple teams; for example, datacenter operations might handle anything server, power, and cooling related. I covered all of the core concepts in the event your SOC will be responsible for planning and/or building all of the resources it needs.

Network Considerations

The SOC needs a network. Many books are dedicated entirely to developing networks, but for the scope of this book, the following are some key concepts to consider when building a network for a SOC:

- Segmentation
- Throughput
- Connectivity and redundancy

The following sections focus on building a reliable and secure network for your SOC, starting with considerations for segmentation.

Segmentation

Segmentation by definition is to separate things into parts and sections. Regarding technology, segmentation is how systems and users are separated. Segmentation can be accomplished at many layers, starting from physically separating networks to logically separating traffic that is on the same hardware. If you are building a new network, you will want to plan segmentation proactively rather than deploy segmentation after the network is set up. Many organizations have trouble deploying deep segmentation or modifying existing networks once they are operational, causing unwanted disruptions and consuming lots of engineering hours. Being proactive about designing network segmentation will save you a ton of time and money. Figure 2-16 shows an example of the different layers of segmentation that are possible.

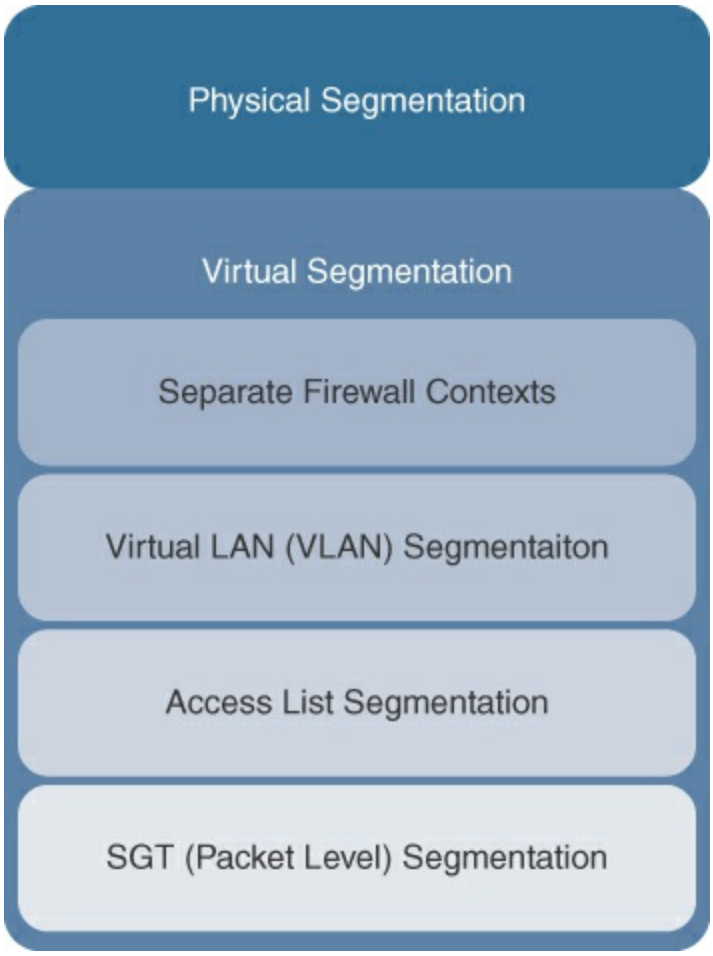

FIGURE 2-16 Different Segmentation Options

The top layer represents providing dedicated network equipment for each network segment, which would include its own routers, switches, and security equipment. This is the most secure and recommended approach, but also the costliest option. Some SOCs require a dedicated network, sometimes referred to as a *secure enclave*, in which case there is no choice but to invest in dedicated hardware. Reasons for using secure enclaves could be to protect certain types of data, such as a classified network, or to comply with regulatory requirements to keep traffic on a specific network, such as data sovereignty requirements. The secured enclave could also have more security requirements for encryption and other policies for how to control what is connected to the network and the type of hardware or system software that is permitted to run traffic within the network equipment. NIST 800-53 Rev. 5, *Security and Privacy Controls for Information Systems and Organizations*, is a great resource for learning more about secure enclaves and other

segmentation recommendations.

Note

Some SOCs are forced to meet specific regulatory compliance requirements for versions of hardware and software that are used within the network. Regulatory compliance could also control how a SOC's network is architected for segmentation as well as requirements to use certain levels of encryption to protect network traffic. For example, networks that are in contact with classified data will be required to have government approved technology and encryption. Upgrading to the latest code that isn't on the approved list can cause the SOC's network to be in violation of regulatory compliance requirements even though the update fixes known vulnerabilities!

Logical Segmentation

If you are permitted the flexibility of using virtual segmentation tactics, you can save money using the same network equipment by leveraging virtual LANs (VLANs) to logically separate SOC traffic from other traffic. Regarding concerns for VLAN security, there are tactics that can be used to protect traffic that are industry supported and proven to be secure. There is a low risk that an adversary could compromise switch security and tap into the SOC network, depending on whether vulnerabilities exist within the network equipment. An example could be overflowing the content-addressable memory (CAM) table or a switch-spoofing attack with the goal to hop between VLANs. Most modern switches have native security features to defend against these and other switch-based types of threats. This is why I highly recommend investing in proper networking equipment rather than relying on used or cheap equipment that lacks modern network and security capabilities.

The following are settings within most monitor switches that can be enabled to reduce the threat of Layer 2 exploitation in a way to prevent VLAN

hopping:

- Disable trunking, dynamic desirable, or dynamic auto when not used
- Prevent the use of Dynamic Trunking Protocol (DTP) on all trunk and access ports
- Prevent double tagging
- Shut down all interfaces that are not in use

Logically Segmenting Hardware

Another logical segmentation option is to configure one physical piece of hardware to be recognized logically as more than one solution. An example is configuring a physical firewall to act as multiple logical firewalls. Configuring a physical firewall to be viewed as multiple logical firewalls is known by some vendors as *multi-context mode*. This feature is great for maximizing an investment in technology because logical separate versions of the same technology can be dedicated for different purposes. Sometimes the segmentation can be at the administration level, meaning administrators can make changes and see everything while other users can only see or change part of the system based on their login rights. Other times, such as in the multi-context firewall example, the firewall hardware views the different logical firewall segments as isolated separate systems that are unaware of each other. Multi-context firewall deployments are popular for service providers that sell access to datacenter services and deploy virtual firewalls dedicated to each customer so that the customer can manage and protect their rented virtual space. The SOC can use a single firewall and virtually dedicate separate firewalls for the analyst network, the testing network, and the management network.

ACL Segmentation

You can use access control lists (ACLs) to filter what can and cannot pass on the same network. An example of using ACLs could be permitting analysts full access to the Internet from the SOC network while limiting network access to guests visiting the SOC based on ACLs denying specific ports and protocols. There are many ways to use ACLs, including dynamic ACLs that

can be pushed down by security tools if a policy is triggered. Regardless of your choice for VLANs, ACLs, or security group tags (SGTs), which are yet another way to perform segmentation, I advise not to overcomplicate your network design, or you will create a lot of unnecessary management work that will lead to confusion and unwanted disruption of services.

> **Note**
>
> I highly recommend considering a centralized network access control (NAC) technology for centralizing the management of VLANs, ACLs, and SGTs. NAC will be covered in more detail later in this chapter.

Choosing Segmentation

Which option you should choose for segmentation will depend on how critical the data is within the SOC. This choice essentially boils down to the SOC's risk appetite based on the risk and the likelihood of an event occurring versus the cost to deploy countermeasures such as implementing segmentation or other security technologies. Be aware that if the SOC shares network gear and that gear is compromised, that will put the SOC's operational state at risk as well.

Figure 2-17 shows a basic example of segmenting a SOC's network. The SOC Tools segment is dedicated to the management of various tools. The second network segment is dedicated for analyst tools. The third network segment is set aside for all analyst workstation traffic. The fourth network segment is dedicated to analyzing malware. The purpose of these network segments is to isolate the type of traffic within each network based on the different types of risk they introduce. As an example, if an analyst workstation is compromised, only other analyst workstations could be impacted.

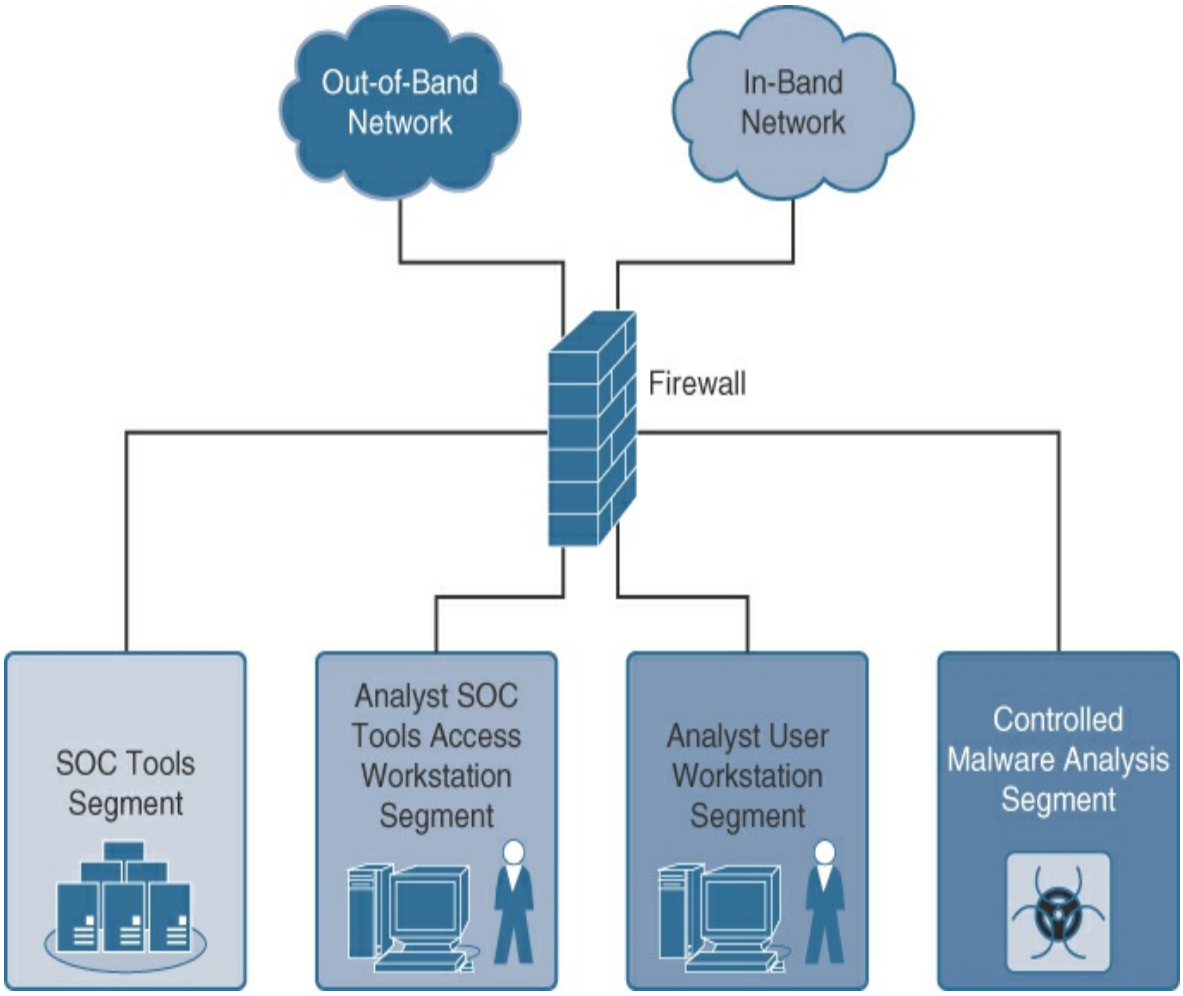

FIGURE 2-17 Logical SOC Network Segmentation

Note

The example shown in Figure 2-17 does not showcase any form of high availability.

Client/Server Segmentation

There are technology options that could run within a network segment that provide another layer of segmentation or access to other segments in a controlled manner, such as thin or fat client/server technology. A client/server

architecture is a distributed application structure that partitions tasks or workloads between the server and service requestors, known as clients. Clients can be thin or thick. Thin clients are designed to be small, pushing the bulk of the data processing on the server. Thick clients are much bigger and are capable of performing the bulk of the processing rather than depending on the server. It is common for servers and clients to communicate over a network, which also is used as a way to provision access to a system within a network segment. As an example, referring to Figure 2-17, a SOC network could have a tool located within the SOC Tools network segment that an analyst could not access directly from his or her workstation that is on the Analyst Workstation segment. A thin or thick client can be installed on the analyst's workstation that is able to access the tool within the SOC Tools network segment, which is a more secure method of providing access to the tool than allowing the analyst's workstation direct access. Risks such as malware installed on the analyst system would not be able to infect the tool since the workstation does not have direct access to the tool.

Table 2-6 compares the benefits and drawbacks of thin client devices compared to thick client (personal computers).

TABLE 2-6 Comparing Benefits and Drawbacks of Thin Client Devices and Thick Clients

Benefits	Drawbacks
Thin client endpoints are less expensive than thick clients.	Thick clients offer more optimal features than thin clients.
Thin clients require less power.	Thick clients require more compute power.
It is easier to scale a server, client-based technology versus deploying multiple thick clients.	Some applications will not run on a thin client based on power or other requirements.
Thin clients can be considered more secured based on being controlled by a centralized computer.	Thin client centralized management opens the potential for a single point of failure.
Thin clients are easier to manage based on the centralized management concept.	

There is a smaller risk of data loss based on a centralized view of the systems.

Active Directory Segmentation

Another form of logical segmentation is using an Active Directory (AD) administrative-tier model. This model divides three tiers that create buffer zones to separate administration of high-risk PCs and valuable assets such as domain controllers. Here is a short explanation of each tier:

- **Tier 0** is the highest level and includes administrative accounts and groups, domain controllers, and domains that have direct or indirect administrative control of the AD forest. Tier 0 administrators can manage and control assets in all tiers but only log in interactively to Tier 0 assets. In other words, a domain administrator should never interactively log in to a Tier 2 asset.
- **Tier 1** is for domain member servers and applications. Accounts that control these assets have access to sensitive business data. Tier 1 administrators can access Tier 1 or Tier 0 assets (network logon) but can only manage Tier 1 or Tier 2 assets. Tier 1 administrators can only log on interactively to Tier 1 assets.
- **Tier 2** is for end-user devices. For example, helpdesk staff would be part of this tier. Tier 2 administrators can access all tier assets (network logon) as necessary but can only manage Tier 2 assets. Tier 2 admins can log in interactively to Tier 2 assets.

Just like with traditional network segmentation, the goal of this approach is to add additional layers of segmentation that hackers would need to compromise before they could access a critical system. Also like other segmentation approaches, although it isn't impossible for a hacker to work from a Tier 2 system up to a Tier 0 system, doing so is much harder. Best practices for security within tier security includes other defenses such as monitoring for malicious activity between tiers. Figure 2-18 is a conceptual diagram showing the Active Directory three-tier model.

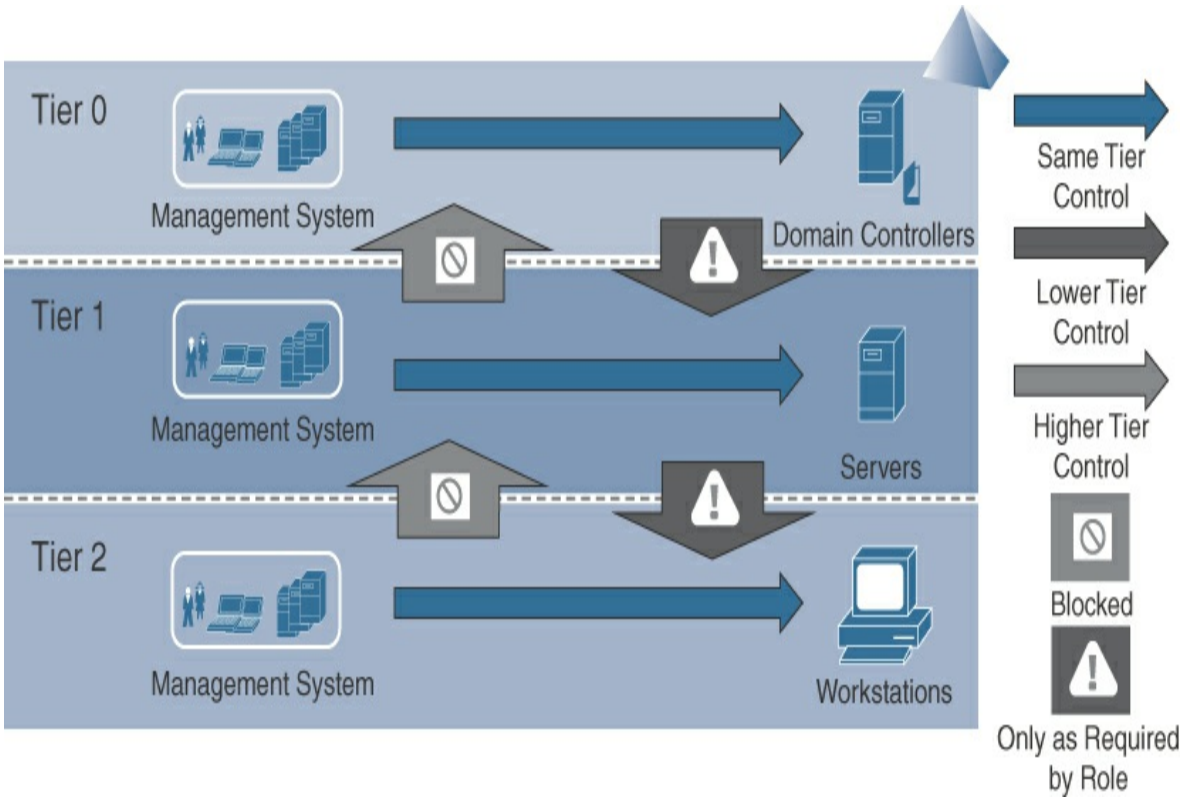

FIGURE 2-18 Active Directory Three-Tier Model

Throughput

Another consideration for the SOC network is throughput requirements. It is recommended to use dedicated circuits to provide services to the SOC. Sharing a network with other, non-SOC network traffic can cause services for all users to go down or be sluggish if the network experiences spikes in usage. There is a higher cost to having a dedicated circuit, which might not be obtainable for all SOCs due to budget constraints.

If your SOC must share a network, you can use configuration and technology to prioritize SOC traffic over other traffic and reduce the risk of a sluggish SOC network. The following items are some options that you can use to provide priority to SOC traffic as well as create an overall faster network for all systems:

- **Segmentation:** Isolate the SOC network logically.
- **Proper capacity planning:** Ensure enough bandwidth exists for the

expected workload.

- **Use wire-speed routing between VLANs:** Use the switch, not the router, to route inter-VLAN traffic. A switch's hardware can do that routing at wire speed, known as static or dynamic IP routing.
- **Prioritize applications and use traffic shaping:** Switches offer the capability to determine the importance of traffic. This can be accomplished using the following methods:
  - 802.1p/Q tagging to prioritize applications
  - DSCP/type of service (ToS)/Layer 3 switching to prioritize applications by header
  - Shaping traffic using bandwidth by throttling or rate limiting
- **Use automatic endpoint parameters:** Configure switch endpoint ports using storm control, number of devices, quality of service (QoS), and VLANs. If the smart ports option is available, use that to automate provisioning of endpoint parameter settings.
- **Leverage switch security:** Enable Address Resolution Protocol (ARP) inspection, IP Source Guard, and Dynamic Host Configuration Protocol (DHCP).
- **MAC layer prioritization:** Configure the network to give priority to devices within a specific list of MAC addresses.
- **Quality of services (QoS):** Configure the network to give priority to certain types of time-sensitive and mission-critical data protocols. For example, prioritize audio and video over nonurgent protocols such as SMTP or HTTP.

Note

The organization might have its own requirements for throughput that are different from the throughput needs for your SOC.

Determining Throughput Requirements

One important task in network capacity planning is determining the required amount of throughput for the SOC. Not properly sizing throughput can slow down or even render some SOC services and capabilities unavailable. An example is an interruption in connectivity between a system dedicated to monitoring various SOC tools for availability. If the monitoring tool believes systems are down, it can trigger a failover, which could cause wasted effort to reset all tools back to the primary state. Another example of the impact of not properly sizing the throughput for a SOC could be having the network become extremely slow when a large number of devices connect to the network around the same time. This could occur in the morning when everybody arrives at work or right after lunch when most people return from their break. Capacity planning needs to include a proper method to predict throughput to avoid these situations.

You can measure and predict needed throughput in a few ways. One way is to consider the number of users that will connect to the SOC's network. Each user will use some level of office Internet bandwidth. The amount of Internet bandwidth can vary from 10 × 1 Mbps to 300 × 20 Mbps depending on the number of users and devices being used. This can depend on whether expectations for each user are for light, moderate, or heavy usage. Light usage would include email and Internet, which may be what a manager would use. Moderate usage includes file downloads, streaming music, Voice over IP (VoIP), and cloud-based resources. Heavy usage includes large file downloads, intense Internet application use, multiple devices, and interactive web conferencing.

The following are some expected bandwidth values associated with common actions that occur on a network:

- **Opening a webpage:** 1 Mbps
- **Watching a live streaming video at 720p:** 5 MB/minute
- **Skype video calling:** 28 Mbps/second

You can use tools to test and validate throughput. Online sources such as Speedtest.net offer a free method to identify current throughput levels. Vendors such as SolarWinds offer network throughput tools that can also provide an estimated report of throughput usage. Whatever estimated throughput requirements you develop, you must also consider other factors beyond your current throughput readings. The following are additional

elements that you need to consider in your throughput capacity planning:

- **Network peak possibilities:** Estimate how much throughput would be required in a situation in which the maximum number of users need throughput. An example is accommodating for an event that causes all local and remote employees to connect to the network simultaneously.
- **Average bandwidth requirements:** What is the average usage? Tools such as Speedtest.net and SolarWinds Bandwidth Monitor (free tool) can capture these readings.
- **Expected growth:** What is the expected growth in the number of users and devices. Expected growth in the number of devices should include expectations for Internet of Things (IoT) and other types of devices that are not run by an end user.
- **System throughput dependencies:** All systems expected to connect to the SOC network need to be evaluated for throughput dependencies. The vendor's website or product documentation can provide this information.
- **High-availability/failover throughput considerations:** Some high-availability configurations will have specific throughput requirements to ensure primary and redundant systems can monitor the status of one another.
- **Possible traffic shaping options to reduce throughput needs:** Consider the configuration options within switches and routers to reduce throughput needs.

As you perform capacity planning for throughput, make sure to consider how traffic will change throughout the year as well as future growth rather than assuming current throughput numbers are good enough for the next three to five years. Getting a feel of what throughput would be needed the next three to five years might require sampling at different times of the year to get a more accurate view of actual bandwidth usage. You need to also consider what new services and capabilities the SOC plans to add based on maturity planning for SOC services. I recommend sizing at least 10–20% room for growth on top of the current throughput needs and expected demands from new SOC services and capabilities. If the SOC experiences bandwidth problems, you can free up bandwidth by using tools and tactics to adjust how the network is used. Best practice is to not depend on filtering tactics for

acceptable throughput.

Note

If performance criteria exceed what can be provided with local resources, outsourcing will be a better approach than expanding capacity. For example, if you find that current resources to investigate a single incident take 35 minutes to load and process a query within a security tool, that would lead to a poor SOC response time based on this bottleneck of wasted time waiting for a specific tool to finish its workload. One option to fix this problem is to add more capacity, but the cost to do so may be more expensive than simply removing the tool and using a third-party cloud service or migrating the existing platform to an IaaS or SaaS option. It is critical to weigh all options when capacity planning, including not adding capacity and switching to something else!

Connectivity and Redundancy

One final consideration as you plan the network for the SOC is how systems will connect to the network. Tools that monitor network traffic need to connect inline or out-of-band. Inline monitoring means to force traffic through the technology so that it can read the traffic. A common example is forcing all traffic that is leaving the SOC to travel through a proxy or gateway application-layer firewall to view what applications are used by the SOC. An out-of-band deployment means a mirror or SPAN port is set up so that traffic can be copied and passed over to a tool for evaluation. If a security tool is out-of-band, it will not be able to automatically block traffic because the traffic it is viewing is not live. In the example of using a proxy or gateway application-layer firewall, if either of those tools is configured to view traffic in an out-of-band fashion, it would be able to see if a user is attempting to access Facebook but would not be able to deny access to Facebook (which would require forcing the traffic to travel inline through the application-layer firewall or proxy). Make sure to evaluate whether tools the SOC intends to

use will need to evaluate traffic and whether they should be connected inline or configured as an out-of-band connection.

Inline Connectivity Risks

There are risks of having traffic flow inline through a security tool. The first risk is that the security device becomes a bottleneck for network throughput. For example, if a 10 Gigabit Ethernet network flows inline through a 1-gigabit firewall, then the users on the inside network can receive at most only 1 gigabit of traffic because the firewall can't deliver anything faster. Another concern is that having traffic flow inline through a security tool could create a single point of failure. To resolve this risk, an inline device can be configured in a fail open mode so that if the software fails, traffic can still pass through the device. Doing so essentially turns the device into a hub while the actual security aspects of the product are offline. Use cases might exist that warrant a configuration of a SOC tool to fail closed, meaning disabling the network connection while security is not functioning.

Risk Reduction with Redundancy

You should also consider redundant hardware and networks to reduce the risk caused by an inline device or other network bottlenecks from failing. Include redundancy for all critical systems as well as systems responsible for time, access control, event logging, and any other critical SOC service. Including redundant systems is also useful for performing upgrades that require a system to be rebooted, so that the secondary system can act as the primary is upgraded. You can accomplish redundancy by purchasing duplicate equipment or by using an alternative network path. If devices on one network are found to be unavailable, systems accessing those devices are instructed to go to another network to obtain the same resources.

The following are key terms for understanding redundancy options:

- **Failover:** When a primary system goes down and a second system immediately takes over the primary system role.
- **Hot standby:** When one system acts as the primary while another system waits in a standby inactive mode. If the primary system goes down, the

standby system is running and ready to become the primary.

- **Cold standby:** When one system acts as the primary while another system is powered off and assigned as the backup system. If the primary system goes down, the cold standby system must be powered on and set up to take on the role as the primary system.
- **Subscriber state failover:** When a failover occurs, the backup system maintains the subscriber state. An example is if a primary access control system goes down, and the secondary system is capable of taking over the load as well as maintaining awareness/logging of systems that were logged by the primary system.
- **Failure detect:** The method used to determine if a primary system is no longer available.

NIST 800-123, *Guide to General Server Security*, is one of the main standards, guidelines, and frameworks that can be leveraged to learn more about proper redundancy tactics. Redundancy options will vary depending on the technology you are using.

Note

It is critical not to assume the organization's existing redundancy meets the risk tolerance for the SOC. Your SOC may need to add additional redundancy on top of what currently exists.

Disaster Recovery

The SOC needs to prepare for the potential loss of systems and people that could be caused by various types of disasters. A disaster recovery plan (DRP) helps eliminate or reduce the potential for economic damage, loss of life, and destruction of property during a disaster. Every business needs to build a disaster recovery team, and it must involve members of the SOC team. There are many industry guidelines, standards, and frameworks that can be used as references to develop a DRP and assign roles for the disaster recovery team.

ISO 22301:2019, *Security and resilience – Business continuity management systems – Requirements*, provides guidance for protecting the organization, while ISO/IEC 27031:2011, *Information technology — Security techniques — Guidelines for information and communication technology readiness for business continuity*, provides guidance for the technical aspects of ISO 22301. ISO 27031 provides six main categories for business continuity planning:

- **Key competencies and knowledge:** What information is required to run critical systems and services?
- **Facilities:** Where is the facility located, and what concerns exist around how it can resist a potential disaster?
- **Technology:** Which technologies are business critical?
- **Data:** What data is required to restore business activities?
- **Processes:** What processes are in place to deal with an incident?
- **Suppliers:** Which suppliers and supplies are critical to the business?

The disaster recovery team should develop the following documents as part of the DRP:

- **Backup and recovery policy:** The team should document procedures for backup and recovery of any system deemed critical to the business. The details of a backup and recovery policy need to explain how often data is backed up, where it is backed up to, the type of data that is backed up, how to access backed-up data, and who is responsible for the systems and data being backed up. If a disaster occurs, the business will need to restore critical systems and populate those systems with the last known-good version of data. Having the backup and recovery policy documents for every critical system enables anybody within the disaster recovery team to restore systems and obtain critical data post disaster.
- **Business impact analysis (BIA):** The purpose of a BIA is to determine what people, process, and technology must be available to sustain the business. The disaster recovery team can associate values with the loss of people, process, and technology to determine the types of redundancy that should be used as a countermeasure to the risk of loss. For example, a SOC’s datacenter would be critical to the operation of the SOC. A BIA

would represent how critical the datacenter is based on the daily loss that would occur if the datacenter were to go down, how the SOC going offline would impact the business, and potential risk reduction items. Including BIA documents within the DRP will inform the team of the impact of losing any person, system, or data due to a disaster.

- **Step-by-step disaster recovery plan:** This plan informs the team where all other disaster recovery data is located as well as the location of the backup of the plan. The DRP also includes contact information for every member of the disaster recovery team, specifies where the team is supposed to meet, and specifies any backups to the meeting location and contacts in case the primary location or any team member is lost or unavailable during a disaster. It is recommended to store a backup of the DRP at a different facility from where the primary plan is stored. Another recommendation is to have a backup facility that is at another location that would likely not be impacted by a disaster that would compromise the primary disaster meeting place. The specifics of how far away the backup facility should be located would be based on your organization's business requirements and risk tolerance. The steps included in the DRP must accommodate all the items pointed out in ISO/IEC 27031. This includes documenting the knowledge needed to restore the business as well as the facilities, technology, data, and suppliers that are critical to the business.

Note

It is critical that all documentation is dated and tested to ensure the disaster recovery plan is effective and up to date. You do not want to find an out-of-date DRP during a real disaster!

Security Considerations

Another important topic to cover regarding planning the design of a SOC is what security technology and processes will be used to protect the SOC. If a key tool such as the time server or logging system is compromised by an

adversary, many of the SOC services will be rendered useless. This type of failure not only could undermine the value of the SOC but also could cause future investment in and trust of the SOC to diminish.

Remember, adversaries know they are up against the defense of their identified target. They know companies install antivirus measures, deploy security tools on the network, and have people responsible for monitoring those tools. Adversaries have labs and develop methods to beat security and hide from administrators.

Some security topics that need to be part of the security planning for the SOC are as follows:

- Policy and compliance
- Network access control
- Encryption
- Internal security tools

> **Note**
>
> Security requirements for the SOC will be different than what exist for the organization. The SOC will contain sensitive information about the organization and therefore should not just adopt the organization's security standards.

Policy and Compliance

The term *compliance* means to meet some goal. That goal will vary by industry and how an organization operates. Some compliance goals are requirements pushed by local government, which if not followed will lead to fines (or possibly even incarceration in extreme cases) for those responsible for the organization's compliance. Other compliance goals are mandated by the leadership of an organization, meaning there isn't any legal obligation; however, the organization has made meeting the compliance goal an obligation for the organization. Some leaders will convert industry

recommendations (known as guidelines) into mandated compliance goals with the intent of establishing a baseline of security practices within the organization. Chapter 6, "Reducing Risk and Exceeding Compliance," covers common forms of mandated and recommended compliance as well as how to enforce compliance goals with policies, procedures, and standards.

Regarding building a SOC, there might be compliance goals that exist or are desired by the organization that must be met before the SOC can move to a go-live status. The most common pre-SOC compliance requirements are those required by local and federal governments as well as those aligned with business operations as stated within an organization's policies. The Health Insurance Portability and Accountability Act (HIPAA) is an example of a U.S. law that mandates compliance by all organizations ("covered entities") in the United States that handle protected health information (PHI). The Payment Card Industry Data Security Standard (PCI DSS) is an example of a standard that mandates compliance by all organizations that accept, process, store, or transmit credit card information. It is critical that the SOC identify all required compliance and ensure requirements are met pre-launch. Chapter 6 will cover the most common legal and financial compliance requirements your SOC needs to be aware of.

Policies are high-level requirements the organization must follow and are set by leadership. *Procedures* contain the details for how to follow policies. Not only will the SOC need to adhere to all policies, but there might be adjustments to how the SOC will be developed and operate in order to be compliant with corporate policies. The SOC will need to verify any pre-launch compliance requirements with the SOC sponsor regardless of whether those requirements are legal or corporate driven. By doing so, the SOC will have executive support that all compliance has been met. I highly recommend ensuring that legal, financial, and other compliance identified by your SOC sponsor is addressed as early in the planning phase as possible to avoid disruption of later steps due to compliance violations. Compliance violations can prevent a SOC from going live!

Note

You will need a SOC sponsor that understands your organization's compliance requirements. If you are unsure if your SOC sponsor

would know all of those requirements, validate with other leadership. Do not assume you have met all compliance requirements without proper validation.

Network Access Control

Network capacity planning and security requirements provide the basic foundation for developing the SOC's network as well as how the SOC can secure the organization. A foundational security requirement for many modern SOCs is segmentation and controlling access to and between segments. Access control needs to be delivered in a least-privilege design, permitting only necessary services to access or leave a network segment. I recommend isolating device types based on their associated risk and dependencies, but very elaborate segmentation can be complicated to deploy and manage. My most basic device segmentation recommendation is to isolate guest users, employees, administrators, IoT devices, mobile devices, SOC tools, malware testing, and critical systems.

Profiling

A SOC needs a method of validating devices on network segments based on MAC address or more advanced profiling tactics. Without this capability, a SOC network administrator will not be able to determine what type of device has connected to the network. Profiling a device beyond its MAC address is a more ideal approach because it uses various network protocols to ensure the device is authenticated, rather than relying solely on a MAC address, which could be spoofed by an attacker. For example, profiling technologies can determine the difference between an iPhone and iPad even though both products are made by Apple and have similar hardware components. Determining the difference between two different Apple devices can be based on how the DHCP request is seen, what version of Safari is used, or other factors analyzed by the profiling technology. Specific characteristics can also be used to determine not only if the device is what it claims to be but also if it is authorized. A common example is determining if the same device type is

company issued versus an unauthorized personally owned asset. MAC addresses can be used for this purpose but, again, could be spoofed. Best practice is to use characteristics within the system such as placing a certificate on the asset that can be validated to ensure the device is issued by the organization.

NAC Value

Leveraging a NAC solution is recommended by many frameworks, including NIST's Cyber Security Framework and ISO/IEC 27001. One popular choice is to use a centralized policy monitoring and enforcement NAC solution based on the IEEE 802.1X protocol. NAC enables a network administrator to limit certain users or devices to specific access privileges, enforced automatically upon connecting to the network. For example, Anjelica, an administrator, can bring her laptop, printer, and VoIP phone to any office within the SOC's building and just plug them in. Based on the NAC profile for those devices, her printer would be placed on a dedicated printer network, her VoIP phone would be placed on the IP phone network, and her laptop would be placed on the administrator network. If a guest unplugs any of Anjelica's devices and plugs in his or her own computer to the network, the network would change that port to the guest network based on how the guest authenticates and the type of device being plugged in. In this example, the guest's device would not have the company certificate installed on the device connecting to the network and the guest user would not have an employee account within the accounting system. This would cause the guest to have his or her computer automatically placed on the guest network. This is just one of the many deployment options for a NAC solution—the default action the guest experiences could be complete denial of access or possibly provisional access, but an ACL would limit this guest user from administration-level access even though this user is on the same VLAN as Anjelica.

Regardless of how you deploy NAC technology, it is critical to ensure similar security is deployed across all areas and networks, including the LAN, wireless, and VPN networks. I have seen organizations deploy secure wireless while their LAN is open to anybody plugging in any device as long as they can get access to a LAN port. Best practice is to standardize on a network access policy across all access methods.

Additional value can be obtained from using a centralized network access control solution. Many NAC solutions offer the capability to share the database of network-connected devices, which is known as *context*. This database of connected devices can be used by other solutions, such as a SIEM solution, to match an IP address that the solution sees to what content the NAC solution knows about that IP address. An example is logging into a SOC's SIEM solution and identifying every device that is on the SOC network by username rather than by IP address. This can occur based on the NAC technology capturing the user's identity when the user accesses the network and sharing the database of users on the network with the SIEM technology.

Another value of a NAC solution is the capability to remove devices off the network that have violated a policy or are seen as a risk to the organization. Many NAC solutions scan a device as it attempts to access the network and verify that the device has required security features enabled before permitting access. For example, the SOC can use a NAC technology to validate whether a device connecting to the SOC network has antivirus installed and running and, if it doesn't, deny it network access. NAC technologies can also be integrated with vulnerability scanners. A vulnerability scanner will scan the network and identify any device that has vulnerabilities. Devices that have critical vulnerabilities are labeled with a very high vulnerability score. A NAC technology can be used to automatically remove devices that have very high vulnerability levels. Figure 2-19 is an example of SANS's recommendations for vulnerability management. Combining the value of a vulnerability scanner and NAC solution can automate the entire process recommended by SANS.

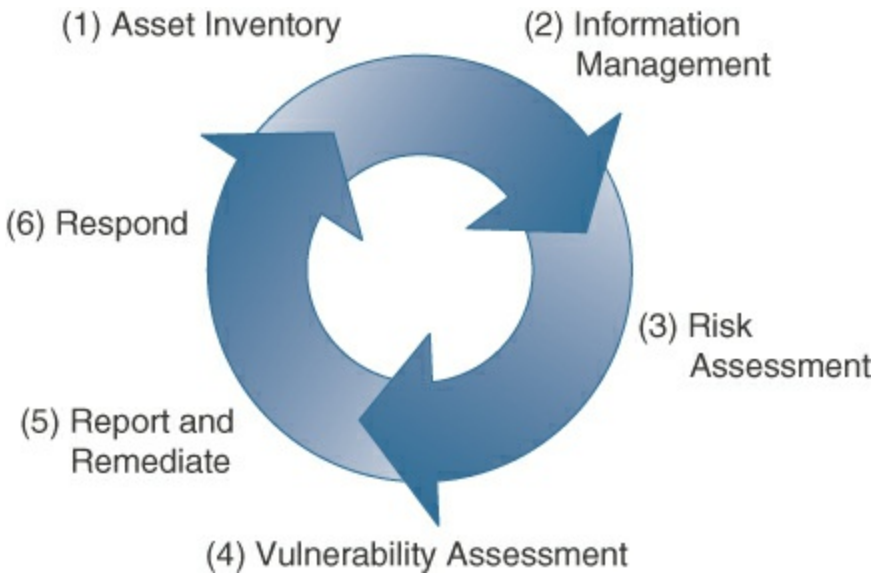

FIGURE 2-19 SANS Recommendations for Vulnerability Management

Note

The SOC must manage its own access control policies even if the NAC solution is a shared resource with the organization. There may be situations such as a forensic investigation led by the SOC that require certain data to be controlled. In the case of a forensic investigation, only the SOC's forensic team should have access to that data.

Encryption

Traffic needs to be safeguarded between the SOC and customer as well as within SOC segments. This is especially true for traffic sent over wireless networks that anybody within range of the SOC could potentially access as well as for traffic that remote users send over untrusted networks to the SOC network. There are different levels of encryption that can be used to protect the confidentiality and integrity of network traffic, some of which are more trustworthy than others. For example, using WEP to encrypt the SOC's wireless network is not recommended. WEP is very old and can be easily

beaten through known exploitation tactics. WPA2 and WPA3 are more secure options based on the features and security capabilities they offer. Some wireless equipment might not be able to support WPA3, so WPA2 would be the best option. Table 2-7 is a summary of the different types of wireless encryption options.

TABLE 2-7 Comparing Different Wireless Encryption Types

	WEP	WPA	WPA2	WPA3
Brief Description	Ensures wired-like privacy in wireless	Based on 802.11i without requirement for new hardware	All mandatory 802.11i features and new hardware	Provides more individualized encryption and WPA3-Enterprise boosting cryptophaic strength for networks transmitting sensitive data.
Encryption	RC4	TKIP + RC4	CCMP/AES	GCMP-256
Authentication	WEP-Open WEP-Shared	WPA-PSK WPA-Enterprise	WPA2-Personal WPA2-Enterprise	WPA3-Personal WPA3-Enterprise
Data Integrity	CRC-32	MIC algorithm	Cipher Block Chaining Message Authentication Code (based on AES)	256-bit Broadcast/Multicast Integrity Protocol Galois Message Authentication Code (BIP-GMAC-256)
Key Managem	None	4-way handshake	4-way handshake	Elliptic Curve Diffie-Hellman (ECDH) exchange and Elliptic Curve Digital Signature Algorithm (ECDSA)

LAN Encryption

The LAN can also use encryption, which could be accomplished using tunneling technology, such as IPsec or SSL encryption, between hosts and servers or specific hardware-encrypting devices that sit between different parts of the network. The difference between encryption options include the following:

- Encryption strength
- How encryption keys are managed and exchanged
- What interfaces, protocols, and ports are used
- What OSI layer they run on
- Ease of deployment
- Speed

A site-to-site VPN can be used to connect two locations together. A SOC can use a site-to-site VPN between a branch office and headquarters to protect traffic as it moves between locations over an untrusted network. A similar concept can be used between a host device and the SOC network, which is ideal for SOC employees that need to access the SOC network remotely. I recommend speaking with the vendor of your network equipment to learn more about what encryption options are available. Table 2-1 from earlier in this chapter provided a quick summary of a site-to-site VPN and remote-access VPN.

Internal Security Tools

SOC security tools must consider all steps an attacker will use to compromise the SOC and its users. Recall the discussion about different threat models in Chapter 1. Threat models are a common approach to understanding what steps an attacker could take as they attempt to exploit your vulnerabilities, deliver unwanted software, and execute unwanted actions. One example of a

threat model covered in Chapter 1 is the Cyber Kill Chain represented as Figure 2-20.

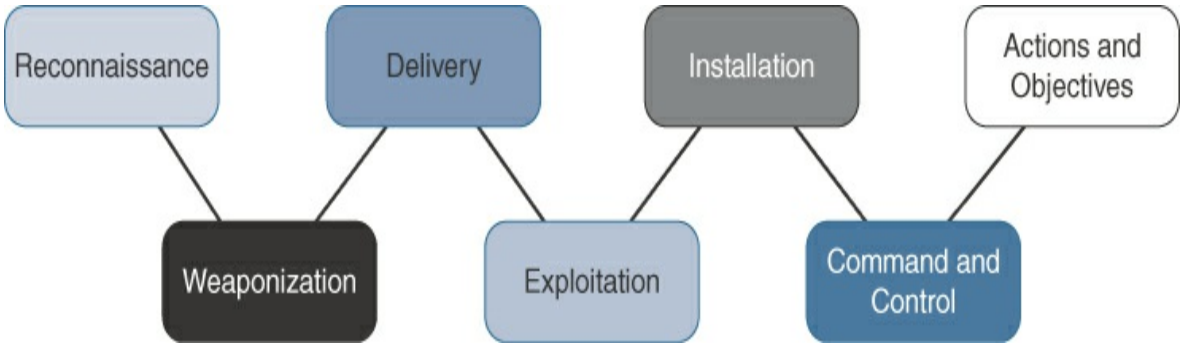

FIGURE 2-20 The Cyber Kill Chain

This particular threat model is used to evaluate different steps an adversary can take to establish a foothold on a network and showcases the need for different types of security capabilities to prevent the attacker's steps from successfully occurring. For example, the exploitation phase defense capabilities specifically look for attack behavior used for exploiting vulnerabilities and attempts to prevent the exploit from working. The installation step assumes that the attacker's exploit worked, making security capabilities for prevention in this phase of the attack different, commonly called breach detection capabilities. Breach detection capabilities are designed to detect when an insider is present, alerting the SOC and preventing further expansion into the network. Both exploitation and breach detection have detective and preventive capabilities but they are designed to function differently based on the types of attacks they are expected to defend against.

The goal for the defender is to prevent the attack as early in the Cyber Kill Chain as possible. By combining defenses used for different parts of the Cyber Attack Kill chain, the SOC creates layers of different security capabilities as part of its defense-in-depth strategy.

The rest of this section walks through common security capabilities that are common in mature SOCs around the world. I'll keep the language vendor agnostic for the most part and focus on capabilities common within a device type category.

Note

The SOC must control polices for all security tools that are used to protect the SOC even when technology is used that is shared by the organization.

Intrusion Detection and Prevention

One option for monitoring internal traffic is to use an intrusion detection system (IDS) deployed off of a monitoring port combination with an intrusion prevention system (IPS) monitoring traffic that travels between network segments. A recommended practice is to tune these tools to be aware of the potential vulnerabilities within assets they are protecting through integration with a vulnerability scanner as well as to perform periodic rule evaluation to maximize their effectiveness. For example, a SOC can run vulnerability scans often to identify and respond to known weaknesses. While those weaknesses are exposed, the IDS and IPS can be tuned to have defense signatures enabled until parties responsible to remediate the vulnerability are able to perform that work. Integrating a vulnerability assessment tool with security tools helps speed up this process. Think of this as defenses against the “exploitation” phase of the Cyber Kill Chain.

Network Flow and Capturing Packets

Baseline tools can be used within the SOC’s network that alarm when they detect an unusual deviation from normal activity. For example, an IoT device can be permitted access to the Internet, but if an unusual amount of traffic is seen from that device, the SOC should be made aware. The increase in traffic could indicate a large firmware update is being installed or it could indicate that an adversary has found a way to use the IoT device as a point of entry to the SOC network! Using NetFlow is one method to detect this type of activity. NetFlow is supported by many common network tools such as network switches. Essentially, NetFlow turns a network into multiple security

detection points continuously looking for unusual activity. This dramatically improves visibility compared to deploying specific security tools across a network.

Note

All NetFlow tools are not created equally! Many NetFlow tools only look at peak and valley traffic, meaning they would not be able to determine if activity is malicious.

NetFlow Technology

There are different types of network flow technology options depending on the vendor and version enabled. Examples include sampled flow (sFlow) and Juniper Networks' J-Flow, to name a few. Comparing the value of a sampled flow such as sFlow against a data-rich option like NetFlow 9 could mean the difference of knowing an event occurred over the last 24 hours (sFlow) versus knowing that a specific action occurred on a specific part of the network from a specific device at a specific point in time (NetFlow). The reason for this difference is that sFlow uses time-based sampling counters while NetFlow uses cached flow entries from the hardware producing it.

NetFlow tools can offer different capabilities based on what the tool can do with the flow that is collected. Many flow tools can alarm for unusual traffic patterns to enable network techs to determine routing problems or areas that are underperforming. This can be useful for detecting network congestion so that alternative routing or other adjustments can be made to improve the overall network performance. Using NetFlow for improving network performance is useful; however, many of these same tools do not have algorithms to detect a port scan or worm activity. Products such as Cisco Stealthwatch and Plixer Scrutinizer are able to convert flow data into security events. Examples of security events include port scanning, worm propagation, data exfiltration, and denial of service.

Note

To see a comprehensive list and descriptions of security events that Cisco Stealthwatch can detect, search the Web for “Stealthwatch Security Events and Alarm Categories” (the current version at the time of writing is version is 7.3).

There is a lot of value from NetFlow if an analyst is able to view all of the records. NetFlow is based on one-way communication, meaning the action of a user connecting to a server and the server responding will equate to multiple NetFlow sessions because each step of the connection will be a single NetFlow record. Multiply each single NetFlow log created as a user continues to use the network by the number of users on the network at any given time and you get a lot of NetFlow records to search through. If a NetFlow tool is only collecting NetFlow records, that tool will quickly fill up with more data than an analyst will be able to review. To deal with this issue, NetFlow tools with deduplication and stitching capabilities can remove duplicate records and stitch the entire session together to simplify mining large amounts of NetFlow data. One user going to a server may create multiple NetFlow records, but a single event can be recorded after all of those NetFlow records are stitched together.

Network-Based Application Recognition (NBAR) is another network protocol that provides intelligent network classification for network infrastructure. NBAR recognizes a wide variety of applications, including web-based applications and client/server applications that dynamically assign TCP and UDP port numbers. Once applications are classified, they can be monitored by security technologies as well as provisioned specific bandwidth through various network policy–based tools. Many NetFlow tools from vendors such as Cisco, SolarWinds, and Plixer also offer the capability to digest NBAR, the results of which complement what is found with NetFlow data.

There is a lot of value from using NetFlow and NBAR; however, these approaches to monitoring data cannot offer the same level of details as using a packet capture technology. Think of the value from the NetFlow approach as giving you the headlines of security events. This includes providing metadata timestamps, senders’ and receivers’ IP addresses, ports used, length of conversation, and amount of data transferred. SOC services such as

incident response need accurate details of an event, which NetFlow and NBAR cannot provide. When this need comes up, the SOC must use another source that provides more details.

Packet Capturing Technology

Packet capturing provides the full story, accurately reconstructing what exactly happened and when it happened. The difference between analyzing flow and analyzing a packet capture is analogous to the difference between reviewing a list of phone calls made (flow) and recording the actual calls and listening to them (packet capture).

Packet capturing must store the traffic and a tool must be used to analyze what was captured. This requires large storage systems, leading to a higher cost. Flow can help you to determine where to look, but many flow-based tools lack evidence of whether an event is genuine, meaning you know an IP address violated a policy, but you don't have the details of what was actually done, as you would if you have a packet-level capture of the activity. An example of details that packet capturing would include that a flow approach would not is the specific files that were exported off the network during a security incident. Having this level of detail would determine the value of the files that were lost, giving the SOC the specifics needed to apply the appropriate response to the incident. If the data was classified or protected by law, the SOC would be required to release to the public what was lost. It would be ideal to use the flow approach to view the network for threats and use the packet capturing approach to investigate a threat, but that is possible only if funding permits acquiring both of these types of tools. Either approach is ideal for monitoring the internal network for threats, but cost will likely cause you to initially invest in only one of these options. The good news is that vendors of either approach are starting to offer hybrid options, which will use flow to baseline the network and launch a packet capture upon seeing an event.

Note

There will be situations that require the details associated with packet capturing technology. Flow technology will not meet these

requirements. An example is proving specific types of data were lost during a security event for a forensic investigation.

Change Management

Traffic entering and leaving each SOC network as well as traffic that crosses SOC network segments needs to be controlled using a tool such as an application-layer firewall or proxy because these tools have visibility of all types of traffic. You can configure policies to block specific traffic types but also include a process by which a user can request that traffic be unblocked for a particular service the user needs to use. The process of making exceptions for blocked traffic needs to be part of the SOC's change management process, requiring at least the following items to be included in the change request:

- Requestor
- Purpose
- Impacted systems
- Creation and expiry dates
- Change management approval reference number

Best practice is to include change management tools that assist with tracking requests and automate reminders to respond to the expiration of an exception that was granted on a temporary basis. Many organizations fail to use reminders to remove exceptions to firewalls, causing temporary policies to get lost within the configuration, leading to a weak firewall and administrator uncertainty about which traffic they are required to permit or deny. Orchestration technology is becoming popular to assist change management tools with automated remediation of expiring requests using predesigned playbooks.

Note

Over time, poor exception management practices lead to failure in

security. An example is how organizations struggle with dealing with expiring access requests. Not addressing expiring access requests leads to exceptions that administrators are afraid to disable because the original purpose of the exception is not documented. As times goes on, the parties involved with the exception are no longer around to explain the need for the exception. Nothing is done and the exception becomes part of the permanent configuration.

Host Systems

Host systems and servers must include security products that provide system-level defense. Computers have different parts that require different types of security capabilities. Security capabilities must protect the firmware, operating system, Internet browsers, and even the plugins within the browsers. Standard signature-based antivirus is a good starting point but cannot be the only security capability installed on a host system. Current malware can operate in ways to avoid antivirus, such as using tactics that function in memory rather than as a file that can be scanned by a traditional antivirus solution.

Figure 2-21 provides an example of host layers and how security can be applied. For this example, the firewall and IPS provide edge defense; however, notice how anti-malware detection sits within the system to monitor activity on the host. Imagine a PDF file with malware that antivirus didn't flag as malicious and now a user is attempting to open. Ideally, the anti-malware technology is capable of noticing that the PDF is not a known threat, but its actions are not normal and are a concern. I recommend you ask any potential host security vendor about how their solution can look at defending the different host layers shown in Figure 2-21, including host firewalls and IPS options as well as anti-malware capabilities. Also, make sure to question how this software is supported by configuration management, monitoring, and other management tools to avoid being stuck with an isolated management system.

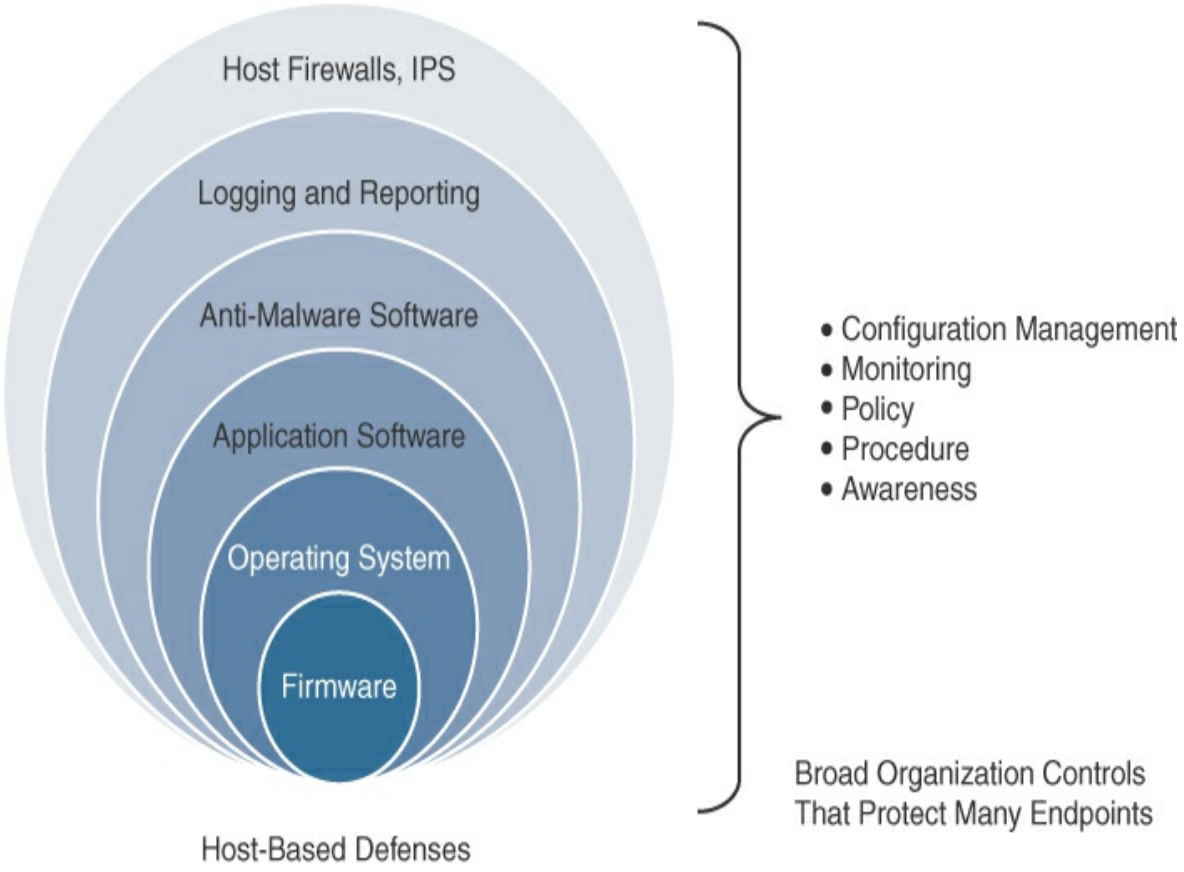

FIGURE 2-21 Defense in Depth for Endpoint Protection

Guidelines and Recommendations for Securing Your SOC Network

There are other security capabilities to consider that can provide more layers to a SOC's defense-in-depth approach. Honeypots can be used across the network with the intent of luring threats that have breached the SOC's network. Sandboxes can be used to analyze external files before permitting those files onto a secured SOC network segment or critical system. Vulnerability scanners can be configured to scan the network proactively and identify any potential weaknesses the SOC needs to patch and protect as specified by the earlier SANS recommendation for vulnerability management best practices. Hosts can have agents installed that complement the security technologies previously covered, such as a vulnerability scanner feature

within the antivirus or agent that provides host flow data to the network flow collector. As covered in Chapter 1 and earlier in this chapter, there are many sources such as standards, guidelines, procedures, and consultants that you can consult to obtain the best recommendations for securing your SOC network.

> **Note**
>
> There are risks associated with leveraging honeypots and sandboxes. Misconfiguring or misusing these technologies can lead to being compromised. Make sure to use a trusted source for guidance for configuration and proper use of these technologies.

Tool Collaboration

Regardless of the security technology you choose, all of these tools will need a place to send data so the SOC can monitor and analyze it. One popular tool many SOCs use for this purpose is a SIEM solution. Playbooks that are built through developing SOC processes can be automated using orchestration tools and combined with the SIEM solution's tracking system. This provides a method to enforce actions as instructed by the SOC's playbook to events identified by the SIEM solution, streamlining the incident response capability offered by the SOC. Many of these topics will be covered later in this book in more detail since they are used by SOCs to monitor the networks they are responsible to protect. The key is that you consider the same tools to protect the SOC network as the SOC would use to protect the network(s) it is responsible for.

There are general industry guidelines that show how many of these security tools can work together. This includes tying network security tools, host tools, and management tools across different parts of the organization. One example guideline is the U.S. Navy CND defense-in-depth strategy shown in Figure 2-22. This example lists different capabilities for each defense segment, which starts with the host layer and works its way out to the outermost network edge. A way to apply this strategy to a SOC's network

design is to replace the LAN in the Navy's diagram with the SOC's internal network and replace the outer network with the organization's network being protected by the SOC. The outer layer in this diagram can be seen as the external network containing the DMZ and other external-facing resources.

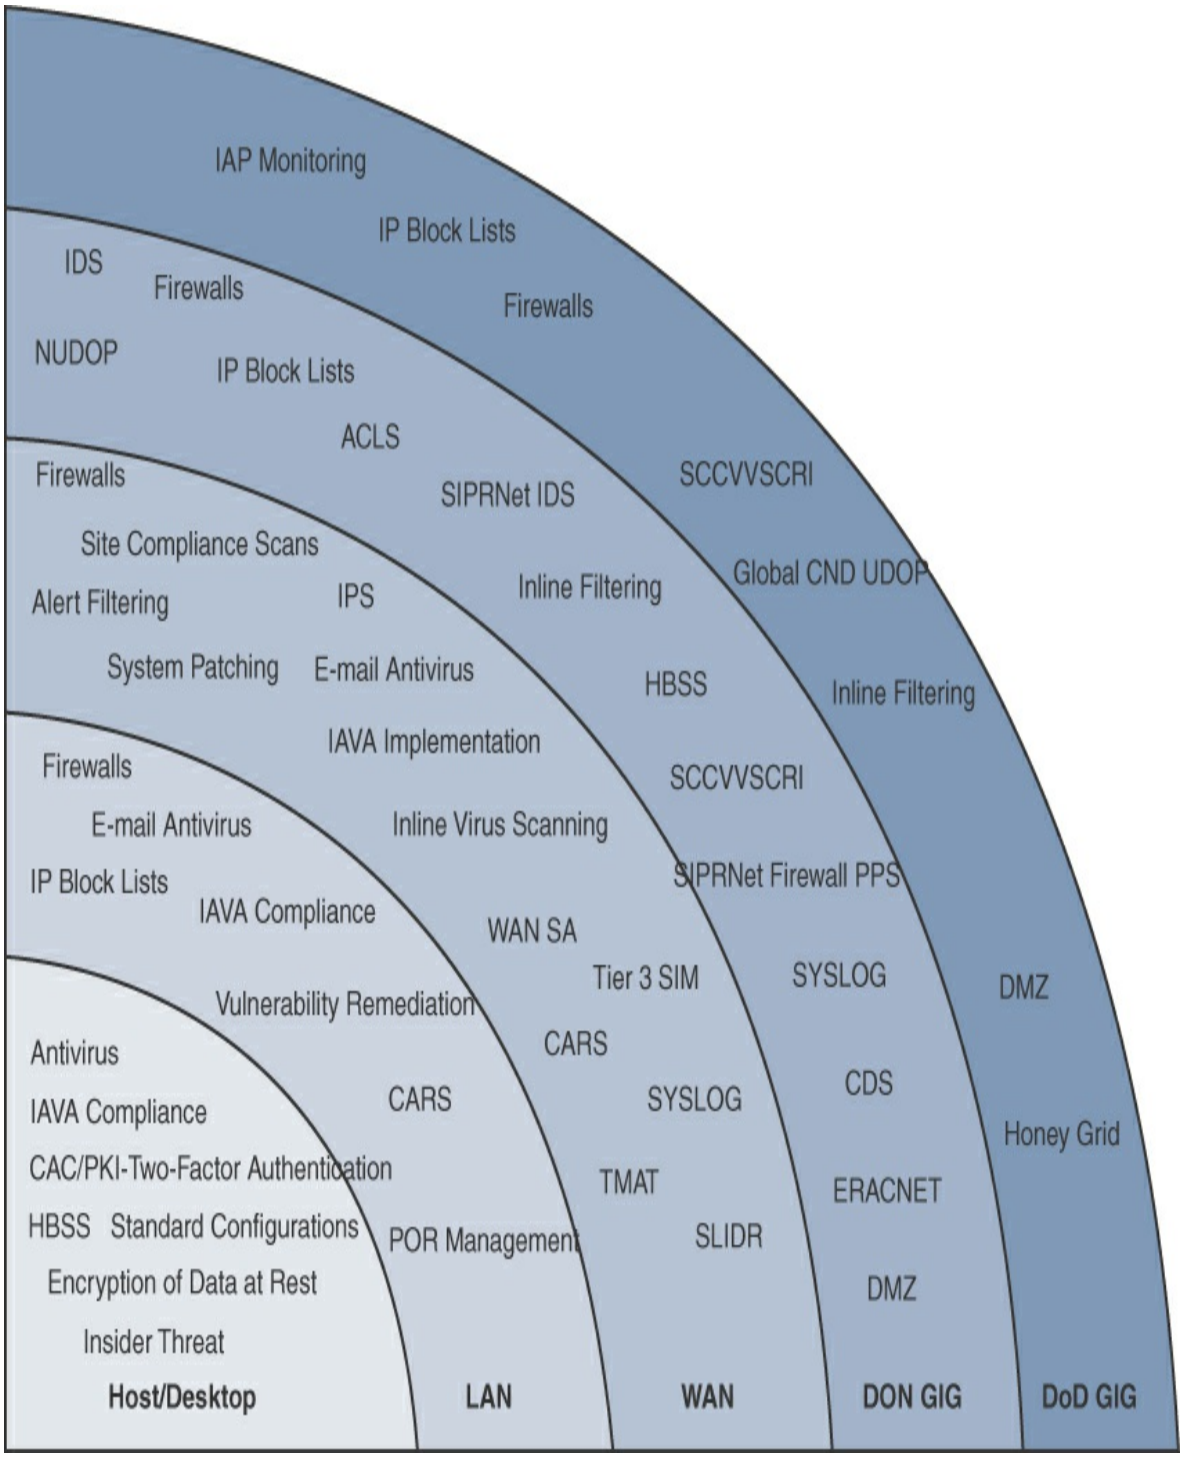

FIGURE 2-22 Department of the Navy Defense-in-Depth Strategy

> **Note**
>
> If you are unsure which technologies to choose, look back to the "SOC Capabilities Assessment" section in Chapter 1 as a method to develop a SOC-focused capabilities map. You can use industry guidelines to help fill in each capability that should exist within your SOC capabilities map. The same approach can be used when choosing security capabilities to secure the organization, which will be a different capabilities map than what is created to show how to protect the SOC.

Building on the Figure 2-17 diagram of an example SOC's network design, Figure 2-23 shows how a SOC can apply the various security tools covered in this section. Figure 2-23 adds NetFlow, IDS appliances, and packet capture tools to analyze traffic between network segments; adds application-layer firewalls and IPS appliances to secure traffic that travels between each segment; adds antivirus and anti-malware on host systems that support this technology; and adds a sandbox within the malware testing segment to assist with that type of work. This may seem like a lot of security, but as stated at the beginning of this section, the SOC's security must be a top priority.

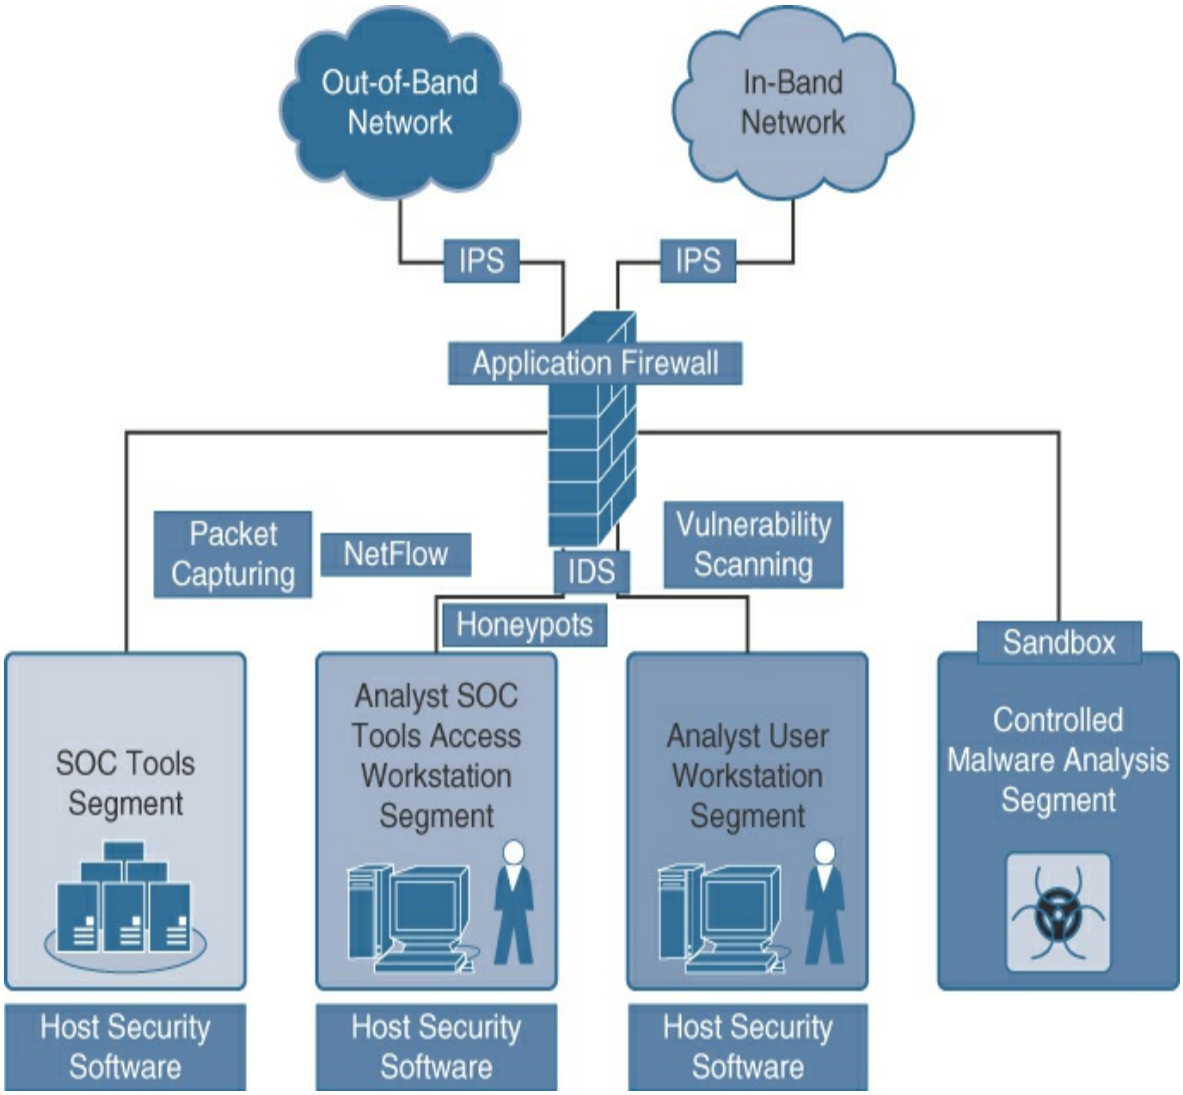

FIGURE 2-23 Logical SOC Network Segmentation with Security

SOC Tools

Many SOC procedures require the use of specific tools. I covered tools used to protect a SOC earlier in this chapter. The SOC needs those same tools to protect the SOC's customer and deliver SOC services, although the SOC may deploy the tools differently or have different expectations from them. Access to tools can be limited based on the role of the SOC member and the type of data within the tool. For example, members responsible for keeping the SOC secure should only see internal SOC data, while an analyst responsible for

customer behavior would not be able to see activity regarding fellow SOC analysts. Most security tools offer different views of data based on the type of account that is created for accessing the management interface of the tool. This feature is commonly known as *role-based access control* because each job role only has a view of the data required for that job role. Best practice is to limit the data seen by a user to least-privilege access using role-based access versus giving administrative access to any user.

Reporting and Dashboards

Reporting and dashboards are other important features available in most SOC tools. I recommend creating reports and dashboards for various parties, including executives and other key members, to provide them a clear value from the SOC. Analysts responsible for monitoring security events will receive reports and have dashboards designed specifically for an analyst role, including IPS logs, Windows security events, and other security tools that fall within the analyst's area of responsibility. Dashboards and reports for executives will be different than the dashboard for an analyst's, including how employees use the Internet, what types of devices are connected to the network, or maybe how certain people are using the network. A lot of this data can be captured from next-generation firewall, NAC, and other security tools, and reporting and dashboard options within these technologies can be customized for their intended audience targeting only desired data within the dashboard. SOC-specific dashboards can be created and displayed on the SOC wall, while specific analysts would have additional data available on their own customized dashboards and reports based on their job role. For example, tier 1 analysts will not need access to cases that are handed off to higher-tier support teams.

Note

Reports are a great way to align results with success criteria within SOC functions. Chapter 1 covered using maturity models as a way to measure how effective a SOC service is performing. Part of higher maturity is providing repeatable actions. To prove actions are

repeatable, dashboards and reports can be used to show a history of repeatable success.

Note

I highly recommend creating reports for key members of the organization, including the member(s) sponsoring the SOC. This will quickly earn respect and support for the SOC. Providing awareness to the SOC sponsor can lead to more budget, support for decisions made by the SOC, and support for security and policy awareness. Best practice is to interview key executive members and the SOC sponsor(s) regarding the type of data they would find useful and setting up systems to deliver that data on a weekly or monthly basis. An example is meeting with the compliance officer with the goal of assisting him or her with obtaining data needed for compliance reporting.

Throughput and Storage

Other factors that impact a device being used by the SOC are the amount of data the device must consume and store. In the security world, the amount of data thrown at a tool is known as *events per second (EPS)*. Event data could be many things, including alerts, log messages, and flows per second. Chapter 5, "Centralizing Data," will dive deeper into understanding and tuning tools to accommodate different types of EPS values, but for the purpose of sizing SOC tools, the factors to consider are the amount of storage required for saving events and how ESP impacts licenses. Chatty devices that generate a ton of EPS can quickly consume the storage of security systems. Figure 2-24 shows an example of the different types of data that could be sent to a SOC tool, which would be digested and converted into events an analyst could review.

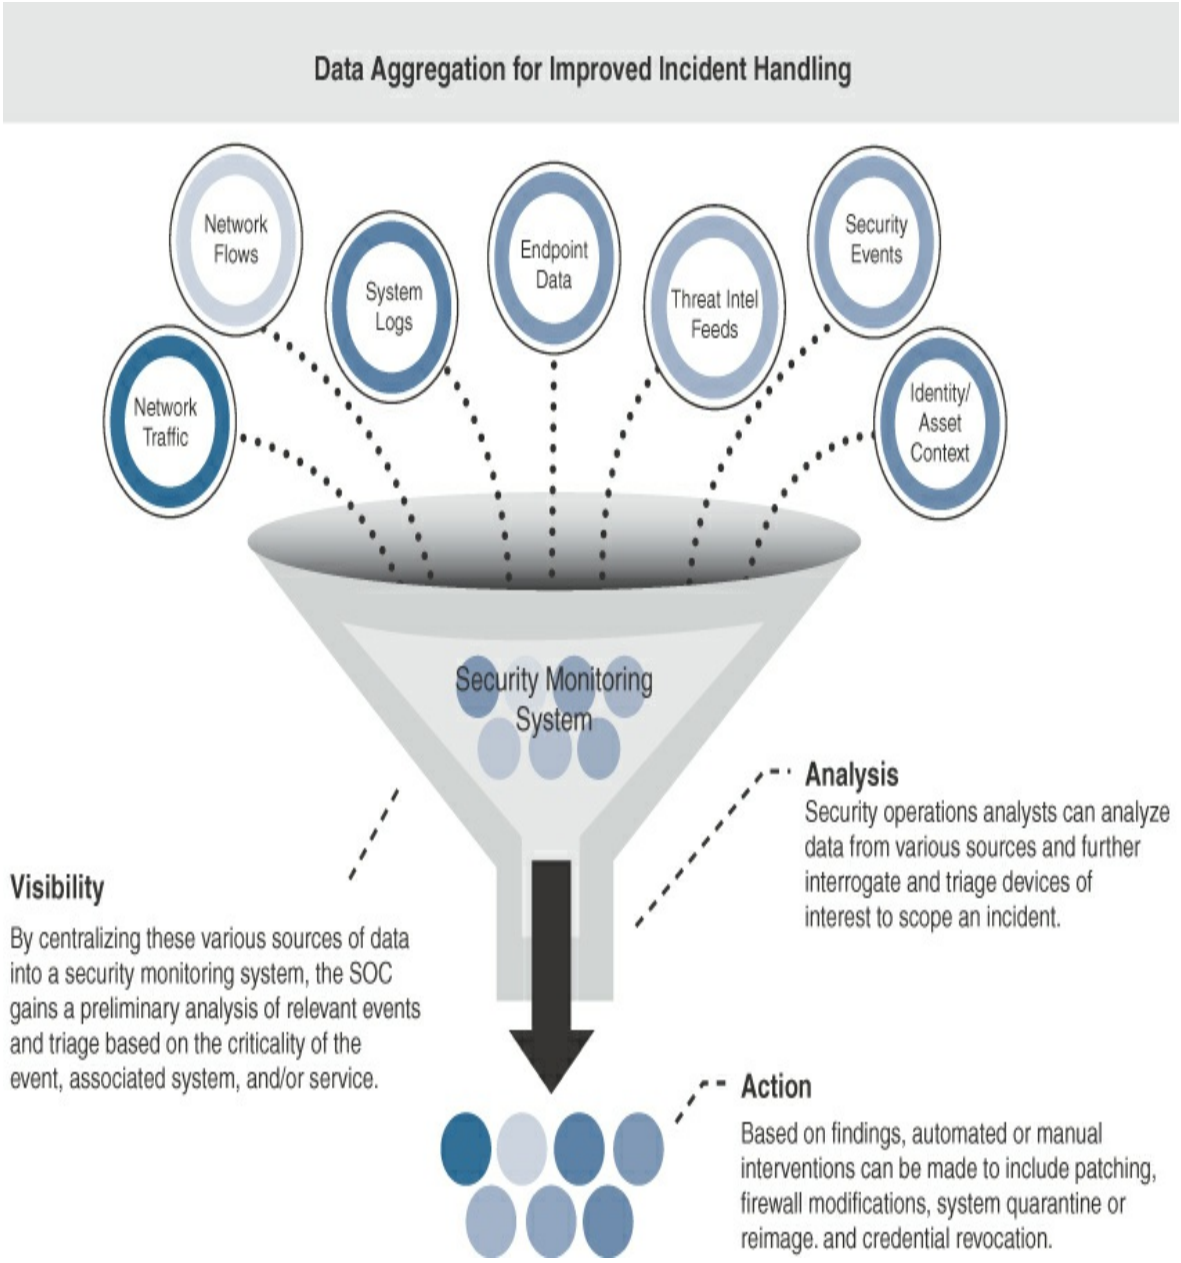

FIGURE 2-24 Events Digested by a Monitoring System

Note

Factoring the number of devices on a network does not necessarily provide an accurate method to size expected EPS. For example, 10 very chatty systems could generate the same data as 100 non-chatty systems requiring the same amount of storage. See Chapter 5 for more

details on understanding how to properly size EPS requirements. When in doubt on ESP numbers, size up to accommodate additional EPS and future growth.

Reducing Events Per Seconds

To reduce the impact from chatty devices, you can implement tuning and storage to reduce the ESP. You can tune devices to generate a log based on the type of event. When debugging a device, you can enable logging for every single event, and when monitoring a device, you can restrict logging to only when a critical action is encountered. Figure 2-25 is an example of alerting levels available in some Cisco products.

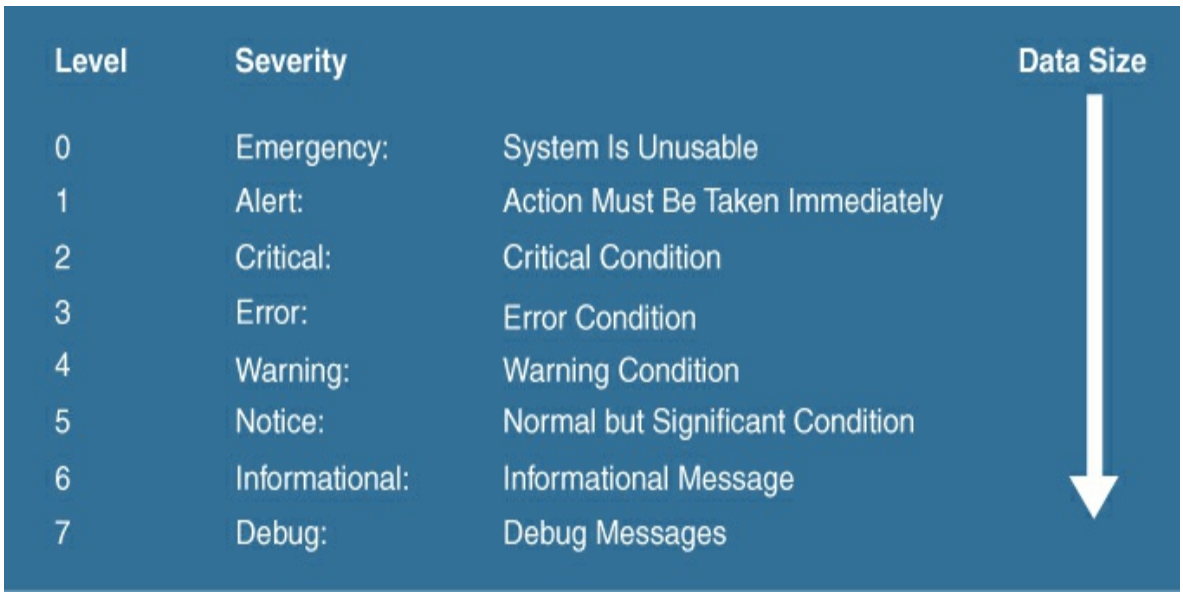

FIGURE 2-25 Cisco Alerting Levels Example

Another approach to reducing the impact of EPS is to adjust the amount of data included within an event. Including more details about an event helps the analyst understand why the event occurred but also increases the requirements for storing these types of events. It is recommended to reduce the level of event details to only what an analyst requires to understand the event. Metadata could also be used instead of including the full log message,

which would reduce the size of the log message. Some data fields may not provide value, in which case you can configure the tool generating the log to not include data from those fields. Consider all of these factors as you tune how tools will produce events. This will save you time on tuning, save money by enabling you to size monitoring products accurately, and improve the quality of the data being analyzed.

Note

You do not need to have all devices configured at the same severity and log detail level. For example, you can have your edge security tools trigger on less severe alarms while internal tools trigger on more severe logs or vice versa. It is best practice to make these decisions based on the value desired from each tool that is generating logs.

Storage and Data Retention

Another sizing concept to consider for SOC tools is sizing required storage and data retention requirements. For example, if three years' worth of log data must be stored before records are purged, you will need to look at the impact of EPS on available storage and size accordingly. Sometimes, data retention requirements have different on-box availability needs versus having the data stored and encrypted off box in the event the data is ever required for a future investigation. This could be based on capability goals or required by a regulation such as Sarbanes-Oxley (SOX) or North American Electric Reliability Corporation Critical Information Protection (NERC CIP), to name just a few U.S.-specific examples. It is common for security tools such as SIEMs to carry data for one to three years and export details for anything older to a separate storage system, in a compressed format to save space. If older data is required, there is a process used to identify the record, uncompress it, and upload it back into the SIEM solution. Some security tools such as NGFWs might store data locally only for a short period of time before an off-box storage option is needed. It is important to cover data retention options with the vendors of the products you choose to use and ensure they meet your SOC's requirements.

General cost can be another factor that determines what data is kept on box and when data is exported to an external storage system. It is common that an organization requires at least a year's worth of log data that is ready to access, while exporting anything older to an external system. Know that the cost of keeping data on box is higher than storing it. The good news is, the cost of storage continues to decrease as technology advances. According to a March 2017 article published by Computerworld, "in 1967, a 1-megabyte hard drive would have set you back by $1 million. Today, that same megabyte of capacity on a hard disk drive (HDD) costs about two cents." The specific amount of time before you export data to an external storage system will be based on your SOC's business objectives. The bad news is, more devices are generating larger amounts of data. increasing the need for more storage. The growth of IoT is causing a major impact on networks. including increasing EPS. Also, different storage options have different costs. Flash memory has a higher cost than magnetic tape or optical discs; however, flash may be worth the higher cost because it performs quicker and can provide stored data faster when requesting it.

Note

SOC data must be protected. I recommend keeping SOC data storage separate from the organization's data storage whenever possible. If a shared retention source is used, SOC data must be segmented from the organization's data.

Centralized Data Management

Most SOCs have a centralized tool for digesting events. The most commonly used tool for this purpose is a security information and event management solution. SIEMs specialize in storing events so that they can be mined and analyzed; however, all SIEMs are not created equally. In my experience, some SIEMs focus on SIM, meaning they offer various ways to manage logging and mine data. Splunk is a popular SIEM solution with very advanced log management and mining capabilities. Figure 2-26 is an example

of Splunk's main search field. You can pretty much find anything you want and create various types of dashboards and reports with the results.

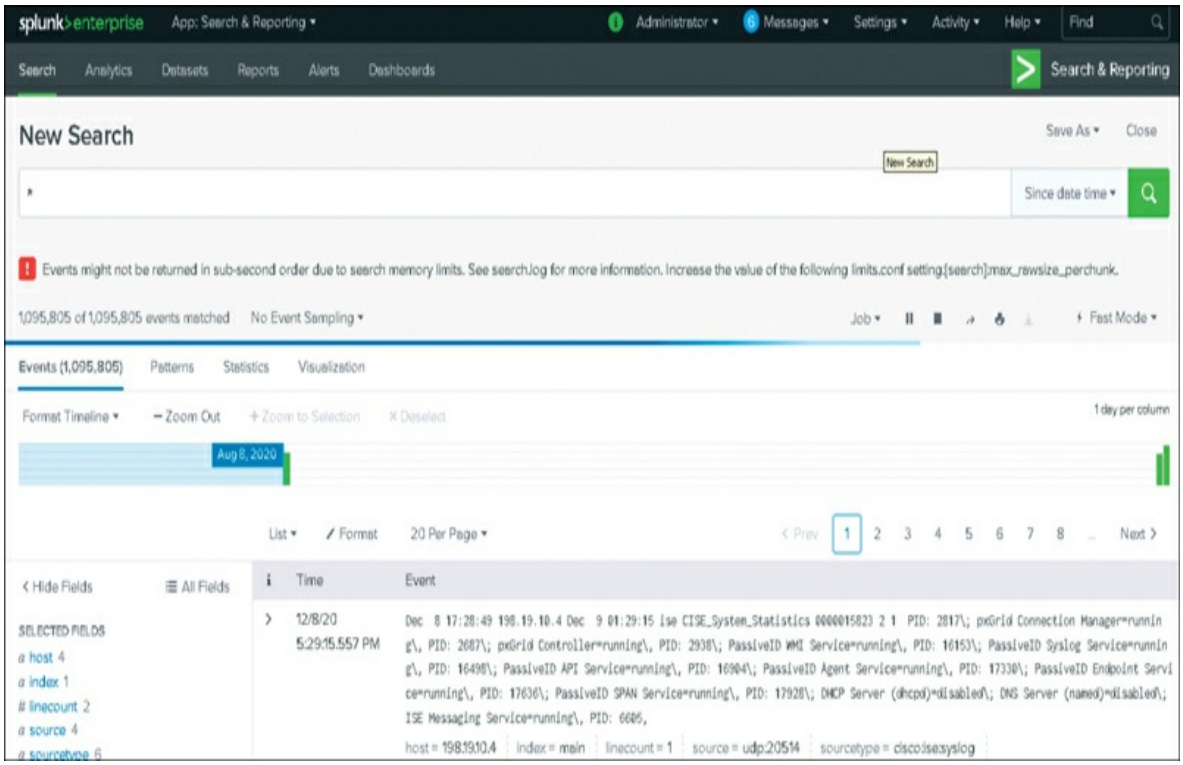

FIGURE 2-26 Splunk Dashboard Example

Other SIEMs specialize in SEM, meaning that they offer various event management and event analysis capabilities. IBM QRadar is an example of a SIEM with lots of SEM features. Examples of QRadar event features are the ability to develop various types of dashboards, identify every asset on the network, and even launch native or third-party vulnerability scanners directly from the dashboard. To be clear, Splunk has SEM capabilities and QRadar can be used to mine data as a SIM, but the experience is much different for those specific needs on both of these systems. Figure 2-27 is an example of a QRadar application dashboard known as Pulse. Chapter 5 covers SIEM and data management in more detail.

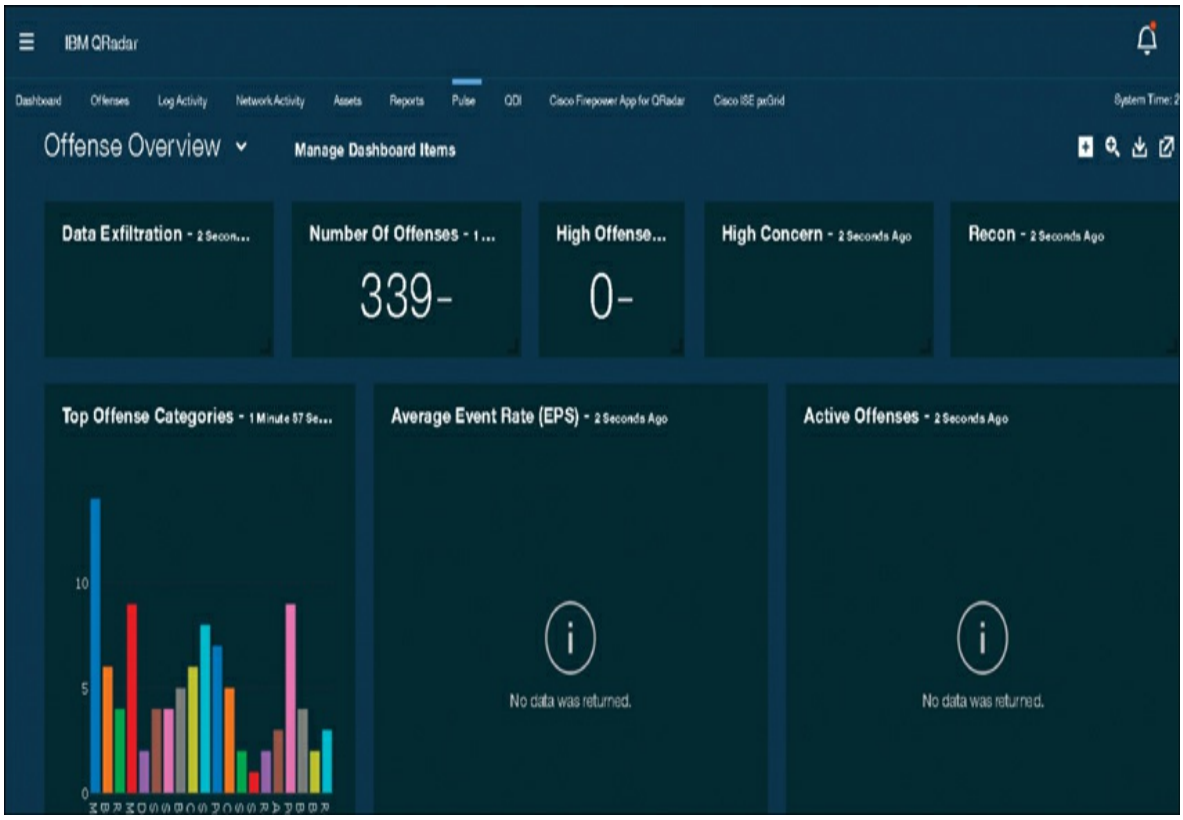

FIGURE 2-27 QRadar Dashboard Example

Some SIEMs do not offer the capability to take action on events. An example of an action can be blocking an event identified by the SIEM solution. If the ability to take action isn't available, the analyst would use the SIEM product for event awareness while using a different method to take action. Procedures for taking actions can be documented and converted into playbooks, which later can be automated. This is known as *orchestration*, which is becoming a feature SIEMs are adapting or partnering with other vendors to offer as a combined solution. We will look deeper at orchestration in Chapter 10, "Data Orchestration."

SIEMs can also include different approaches for tracking events. One popular approach is a ticketing system. Ticketing could be part of the SIEM or orchestration platform or something dedicated to this purpose. BMC's Helix ITSM and ServiceNow Security Operations are examples of incident tracking platforms that could service this purpose. Good ticketing tools include collaboration features such as wikis to search for similar cases that occurred in the past. Tools like Tiki Wiki CMS Groupware and MediaWiki are also

options for this. I recommend considering all of these needs as you evaluate which solution(s) you use for managing all of the data the SOC will be responsible to monitor and act upon.

> **Note**
>
> It is ideal for the SOC to have its own SIEM, SOAR, and/or XDR used to protect the SOC versus sharing what is used to protect the organization. If a shared system is used, administration between the SOC's network and organization's network must be segmented.

Summary

This chapter focused on how to develop a new SOC as well as increase the maturity of an existing SOC. The first concept covered was how to develop mission and scope statements. These form the building blocks for all SOC policies and procedures and are the foundation for why the SOC exists. The next concept covered was the importance of SOC policies and procedures, making up the building blocks for what services the SOC will deliver. This chapter then covered security tool concepts, which are the tools a SOC will use not only to defend the organization it is responsible for, but also to defend the SOC itself from attack.

The remaining chapter topics covered specific details that need to be addressed as you plan and build a security operations center. Topics included everything from how to customize the SOC's facility to what type of capacity and throughput will be needed to support the SOC services. Many of these planning concepts can also apply to building and protecting an organization, but the focus of this chapter is developing a SOC, which includes keeping the SOC-related data segmented from the organization and secured.

Up until this point, I have provided a high-level overview of building a SOC. In the upcoming chapters, we are going to look deeper at the elements that make a SOC successful. Chapter 3 surveys the variety of services that modern SOCs offer. This includes services offered using in-house resources

as well as outsourced services. I highly recommend that every SOC have at least a minimal version of each SOC service covered in the next chapter or at least have an outsourced resource available. The worst situation would be an organization needing a security service that is not offered by the SOC and being unaware of where to look to contract support for the missing service.

References

Bodwell, D. J. (1996, June 23). Team Charter. High Performance Teams. http://www.highperformanceteams.org/hpt_chtr.htm

Cisco. Data Center Power and Cooling White Paper. https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/white_paper_c11-680202.html

Conway, B. (2017, May 12). Office Building. Whole Building Design Guide. https://www.wbdg.org/building-types/office-building

Ferrill, P. (2018, January 8). The Best Mobile Device Management (MDM) Solutions for 2020. *PC Magazine*. https://www.pcmag.com/article/342695/the-best-mobile-device-management-mdm-software

International Organization for Standardization. (2004, May). ISO 17624:2004: Acoustics – Guidelines for noise control in offices and workrooms by means of acoustical screens. ISO. https://www.iso.org/standard/33148.html

International Organization for Standardization. (2011, March). ISO/IEC 27031:2011: Information technology – Security techniques – Guidelines for information and communication technology readiness for business continuity. ISO. https://www.iso.org/standard/44374.html

International Organization for Standardization. (2012, November). ISO 22311:2012: Societal security – Video surveillance – Export interoperability. ISO. https://www.iso.org/standard/53467.html

ISACA. (2019). COBIT 2019 Framework (various publications). ISACA. https://www.isaca.org/resources/cobit

Jones, C.P. (2017, June 9). Flood Resistance of the Building Envelope. Whole Building Design Guide. https://www.wbdg.org/resources/flood-resistance-building-envelope

Key, T.S., & and Martzloff, F. D. (1986, September). A Consensus on Powering and Grounding Sensitive Electronic Equipment. Conference Record, IEEE/IAS Annual Meeting. https://www.nist.gov/system/files/documents/pml/div684/Consensus.pdf

Kosutic, D. (2015, September 21). Understanding IT Disaster Recovery According to ISO 27031. Advisera. https://advisera.com/27001academy/blog/2015/09/21/understanding-it-disaster-recovery-according-to-iso-27031/

McCarthy, K., & Avelar, V. (n.d.). Comparing UPS System Design Configurations. Schneider Electric – Data Center Science Center. Retrieved from https://www.apc.com/salestools/SADE-5TPL8X/SADE-5TPL8X_R3_EN.pdf

Mearian, L. (2017, March 23). CW@50: Data Storage Goes from $1M to 2 Cents Per Gigabyte. Computerworld. https://www.computerworld.com/article/3182207/cw50-data-storage-goes-from-1m-to-2-cents-per-gigabyte.html

National Security Agency. (2010). Defense in Depth. NSA. https://citadel-information.com/wp-content/uploads/2010/12/nsa-defense-in-depth.pdf

Peters, O.S. (1918, June 20). Technologic Papers of the Bureau of Standards: No. 108. Ground Connections for Electrical Systems. Department of Commerce. https://nvlpubs.nist.gov/nistpubs/nbstechnologic/nbstechnologicpaperT108.pdf

Spitzner, L. (2016, December 14). Mobile Device Security. SANS Security Awareness. https://www.sans.org/security-awareness-training/blog/mobile-device-security

Think Tech Advisors. (2017). PCs vs. Thin Client Devices. Think Tech Advisors. https://thinktechadvisors.com/2020/05/pros-cons-of-thin-client-devices/

Chapter 3

SOC Services

Design is a funny word. Some people think design means how it looks. But of course, if you dig deeper, it's really how it works.

—Steve Jobs

What are SOC services? Simply put, *SOC services* is a broad term referring to the value the SOC provides to the organization as a whole. Some companies provide SOC-as-a-Service offerings. These service providers ingest your logs, monitor change management, and inform you of potential things that might be occurring in your organization. Some organizations use SOC service providers while other organizations develop their own SOC capabilities using internal resources. It does not matter how a SOC provides services but rather how effective the services are at achieving the goal(s) for the organization that the SOC services. Effectiveness can be expressed in the maturity of a SOC and its services, which will be the core topic for this chapter.

Fundamental SOC Services

This chapter focuses on SOC services found in organizations around the world. All mature SOCs offer some form of each of these fundamental SOC services, either active, on demand, or through a partnership. Core SOC services covered in this book include the following:

- **Risk management:** Identifying and making decisions to deal with organizational risk. This pertains to managing any type of risk, from physically securing assets to patching digital vulnerabilities that exist within software. This can also apply to remediating weak policies and

lack of education regarding security awareness within members of an organization.

- **Vulnerability management:** Identifying and managing risk from technical vulnerabilities. This commonly involves targeting vulnerabilities within software found on servers, laptops, and IoT devices. Most SOCs use vulnerability scanners and outside threat intelligence to identify vulnerabilities.
- **Incident management:** Responding to security-related events. This covers what actions the SOC takes when certain events occur, such as isolating systems, alerting team members, and implementing remediation steps to resolve the issue. Other subcategories that fall under incident management include incident response, incident investigation, and other incident-related topics. Technologies such as orchestration tools, artificial intelligence, and playbooks are becoming extremely popular to help assist SOCs with incident response services.
- **Analysis:** Analyzing various types of artifacts. This includes identifying characteristics, reverse engineering, vulnerability/exploitation analysis, root-cause analysis, remediation, and mitigation analysis. What separates an analyst focusing on analysis versus incident response is the type of required skills. Analysis uses tools such as IDA Pro to disassemble malware and understand how it functions. An analysis engineer can answer the question “Is this file malicious?” by running it in a sandbox to learn about its behavior. These skills are different from those of a SOC analyst responding to a potential breach.
- **Compliance:** Assessing and maintaining organizational compliance requirements. This can include legally obligated requirements such as HIPAA and PCI DSS compliance as well as organization-driven goals such as meeting a NIST or ISO standard, which are not required by law but could be seen as a required policy by the organization or its customers. The compliance service also prepares the organization for assessments and assists with gathering required information for outside parties validating an organization’s compliance.
- **Digital forensics:** Gathering evidence post incident to determine the cause of the incident and prepare for legal action. There is some overlap in digital forensics, incident response, and analysis skillsets since all three

include some form of understanding what malware or a malicious party has done. What separates digital forensics is the legal aspect regarding how evidence is collected. For example, if you manipulate a file during your investigation, you ruin any chance of using that evidence in a court of law (based on the concept of evidence contamination). Chapter 8, "Threat Hunting and Incident Response," addresses digital forensic concepts in more detail.

- **Situational and security awareness:** Providing the organization with awareness of its operational environment and potential threats. This includes education about critical elements that could impact the organization's goals, potential threats, and actions to reduce risk against operational risk and threats.
- **Research and development:** Researching the ever-evolving threat landscape, developing new tools and techniques, and modifying existing tools to improve effectiveness.

I look closer at each of these services later in this chapter under the section "Service Maturity: If You Build It, They Will Come," as well as hit fundamental concepts from all of these services within the remaining chapters of this book. There are many subsets of these services, and many of these services have synergy with each other, which I call SOC service areas—a topic I will cover shortly. Team members within the SOC can also deliver more than one service, and many SOCs do not label their services exactly as I have listed them. Every organization is different, of course, and the approach I am taking is a consolidation of what I see within mature SOCs mixed with industry guidelines for expected SOC services.

If you want to find how a SOC explains its services, look at the SOC's mission statement and scope statement. As explained in Chapter 2, "Developing a Security Operations Center," the mission statement defines the purpose of the SOC and the scope statement provides even more details, including what services are provided. The goal for your SOC is to grow your scope to include all of the fundamental services included in the previous list and mature those services. If another department handles one of these services, your SOC should coordinate with that department to have some involvement. For example, if another team handles vulnerability management, the SOC should at a minimum be aware of the status of active

vulnerabilities. Knowing this will be helpful as the SOC delivers other services such as incident response, part of which is knowing when a vulnerability has been exploited.

Even the smallest SOCs should work toward having some form of all fundamental SOC services, even if they are outsourced or quoted for on-demand usage only. One common SOC service that is not typically found in smaller organizations is digital forensics. Yes, it can be overkill to keep a digital forensic specialist on staff; however, it will be critical to be able to quickly recruit one when a cybercrime has occurred within your organization. Regarding digital forensics, timing and how the initial response is executed will make or break the entire service. If the evidence is tampered with, it cannot be used in most courts based on the potential of contamination, regardless of what the evidence shows. For this reason, you should proactively know who to call and their associated costs rather than wait until you need to find a digital forensic specialist immediately. Acquiring this information when timing is critical is less than ideal. This is one of the many reasons why I recommend having some form of all fundamental SOC services, even if they are on demand, meaning you just know who to call when needed.

Standing up a SOC service is not a simple task. It requires the right people, process, and technology. Before looking at developing new SOC services, you need to understand some of the expected challenges to overcome as you develop your SOC services. The last thing you want to do is invest in a new service and not be able to go live due to an unplanned bottleneck in the go-live process.

SOC Challenges

Before looking at common services found within a SOC, you must first understand some challenges you will face as you develop and eventually go live with a new SOC service. I find SOCs launching new services will experience one or more of the following challenges.

Challenge 1: People

The first challenge will be *finding the right people*. According to a survey of the cybersecurity workforce conducted by (ISC)^2 titled "(ISC)2 Finds the Cybersecurity Workforce Needs to Grow 145% to Close Skills Gap and Better Defend Organizations Worldwide," the estimated number of additional trained staff needed to close the skills gap came in at 4.07 million professionals worldwide. This translates to the cybersecurity workforce needing to increase 62% in the U.S. market. There are myriad reasons for this skills gap, which is a topic of Chapter 4, "People and Process," but know that SOC service–specific skillsets such as digital forensics, malware analysis, and threat hunting will be extremely challenging to find. The (ISC)^2 research relates to professionals with general cybersecurity skills, meaning it does not consider the subset of security professionals with SOC service–specific skillsets and experience. I believe there are even fewer of these types of professionals available in today's market, which makes recruiting the right people for your SOC extremely challenging. Chapter 4 also covers how to recruit and maintain the right talent.

Challenge 2: Low Maturity

Another challenge with launching a SOC service is doing so at a very *low maturity* level, leading to poor initial performance. As described in Chapter 1, "Introducing Security Operations and the SOC," low maturity means having ad hoc capabilities or having a tool to deliver a service but no defined processes to ensure the proper steps are taken each time the SOC service is needed. If a new SOC service doesn't show value, it will eventually be shut down and replaced by another option, such as outsourcing, or passed to another team within the organization outside of the SOC. My recommendation is to hold back a SOC service or run it as a test (often called a beta) to determine whether the service is running at full capacity. Another option is to outsource the service as it's being developed by the SOC and slowly transition responsibilities from the outside party to the SOC. Using an outside service allows the SOC to learn from a more mature service as well as take on responsibilities as the SOC is ready to handle them.

Challenge 3: Limited Tools

A third challenge can be *limited visibility and tools to properly deliver the service*. An example of limited visibility is not being able to track success due to a lack of management tools or processes. Another example is that a service such as incident response doesn't have network-wide visibility or visibility into specific systems, causing blind spots for the team delivering the service. Both of these examples are due to not having the proper tools.

Examples of not having the proper tools include not having network access control (NAC) enabled on certain devices because they don't support the technology, not having a proper sandbox for malware analysis, and not having a firewall that supports application visibility and controls, resulting in limited visibility to Layer 7 traffic. The problem of limited tools is a tricky one because organizations could always benefit from the latest and greatest tools but typically have to get by with what they currently have. Every new SOC service will be missing some tools, so you must decide whether you have the foundational technology to launch the new desired service or should wait to avoid prematurely launching the service, leading to immediate negative feedback.

Challenge 4: Lack of Experience

The fourth and final challenge I will highlight is *knowing how to build a new SOC service*. There are industry guidelines to help design a SOC service, but guidelines are not specific to your business. Also, guidelines that focus on SOC services typically do not take cost into consideration or compare commercial and open-source options, causing confusion regarding what your SOC should do as your team builds the service. For example, the concept of digital forensics has many guidelines recommending very expensive labs and extensive training. It will be hard for some SOC analysts developing a new forensics service to decide whether to invest in an expensive enterprise tool or leverage an open-source option. Some SOC services will require decisions about capability or capabilities that are hard to make until the service goes live and sees real data. My recommendation is that if you don't know how to start a service, either consult an outside specialist to help or outsource and learn from how another, more-qualified party delivers the service.

In this chapter, I will address these challenges with processes and methodologies. I align my recommendations with the FIRST Computer

Security Incident Response Team (CSIRT) Services Framework and the NIST Cybersecurity Framework (CSF). I will also reference research and consulting work performed both for small businesses and Fortune 500 organizations.

> **Note**
>
> Part of the learning process can be to consult an outside party with SOC service experience before you go live. I recommend speaking with other SOCs about how they stood up and currently run their services, if that option is possible. Some conferences to consider are RSA, Blackhat, DEFCON, BSides, ShmooCon, SecureWorld, and SANS.

Now that you understand some of the challenges you will face as you develop new SOC services, it's time to look at the process of developing SOC services. I'll start off by reviewing a methodology for standing up any type of new SOC service, which requires foundational SOC services that support the SOC before the SOC can provide services. I call the methodology regarding what the SOC needs before it can provide any value *The Three Pillars of Foundational SOC Support Services*.

The Three Pillars of Foundational SOC Support Services

Prior to planning the development of a service offered by a SOC, your organization must first ensure that all SOC support services are in place. I break these required support services into three pillars, as shown in Figure 3-1. A breakdown in any of these three pillars will cause a failure in all SOC services since these factors are required to keep the SOC operational.

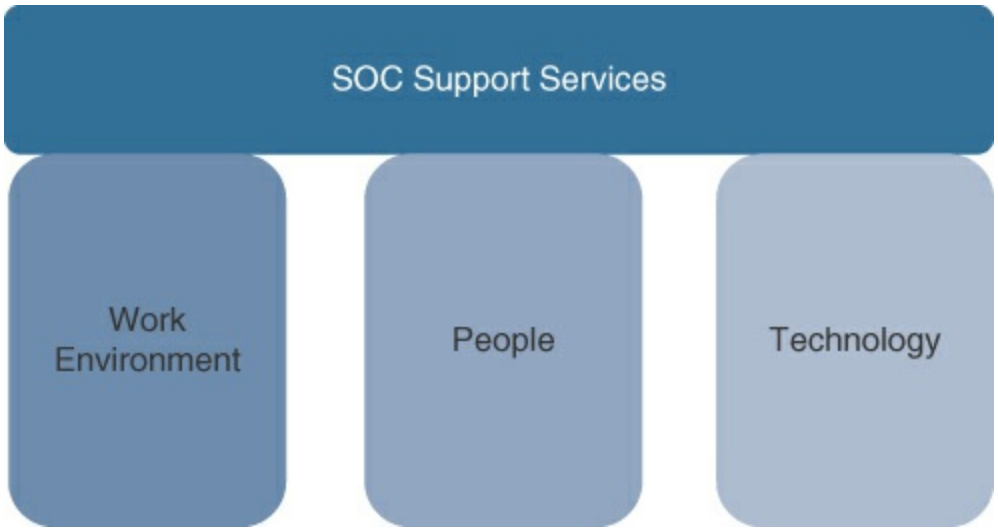

FIGURE 3-1 Foundational SOC Support Services

Pillar 1: Work Environment

Part of Chapter 2 focused on how to build a SOC work environment, including physical and logical design considerations as well as security considerations. Before a SOC service can be delivered, the SOC work environment must be established so that it can support the people and technology responsible for the SOC services. Obviously, employees need a place to work, so the real question is *where* should the members of the new SOC service work? Some service teams will be able to work effectively remotely while others will need to be located together in a specific location.

The work environment for the new SOC service can include one or more designated physical spaces at a physical location, space to accommodate required technology, and accommodations for how remote work is handled (if permitted). For example, if a SOC will be standing up a new vulnerability management service, the vulnerability management team may need a work area within the SOC that has room for team members to sit near each other for better collaboration purposes, space within the organization's or SOC's datacenter to host the vulnerability management equipment, and VPN agents installed on computers of any workers who will support the service remotely. More space may be needed for other vulnerability management requirements, such as standing up a lab for testing vulnerabilities or space within different

locations to host a vulnerability management server. Chapter 9, "Vulnerability Management," provides examples of a distributed vulnerability assessment architecture, which includes having either local vulnerability scanners at each location, using clients that connect over the cloud or using a cloud delivered scanning solution.

Planning a Work Environment

Know that there are many alternative options that can be considered as you plan the work environment. In the vulnerability management service example, the physical space for the vulnerability servers can be replaced with hosting the servers in the cloud using either an Infrastructure as a Service (IaaS) or a complete Software as a Service (SaaS). The SOC team members could work from a space in a different location than the organization's campus or could work completely virtual from home. Make sure to consider all options for the work environment, including a hybrid approach, rather than focusing on one way for the SOC to operate. The COVID-19 epidemic convinced many organizations that thought remote work would never work that not only can it work but it can work better than requiring employees to go into the office!

Some SOCs permit some or all of the people responsible for a service to work remotely. Doing so reduces the need for dedicating facility space to SOC personal. Some SOCs will manage the required technology to deliver the service locally, while others will use cloud services, once again reducing the need for physical rack space in the datacenter as well as maintenance needs. If physical space is currently being used or is being planned for in the future, your SOC needs to state those requirements up front and plan for growth. A key point from the Chapter 2 discussion of facility sizing is that oversizing initially and growing into a space is less expensive than undersizing and running out of space. The previous cloud use cases can also be options for accommodating growth, but it's best to plan for growth using both cloud and physical facility alternatives so that you aren't forced into cloud services due to lack of facility space.

Here is a general checklist to consider as you plan your work environment (Pillar 1):

- Where will the team members work from that deliver the service?
- What facility requirements are needed?
- Will technology need to be powered and maintained at a facility?
- Will remote work be permitted?
- Will the team need to meet and where will they meet?
- Will the team need access to a shared SOC facility such as the war room?
- Will additional physical security be needed at any of the facilities for this service?
- Are there any remote work limitations, such as data sovereignty requirements not allowing data to leave one's country?

Pillar 2: People

The second pillar for foundational SOC services focuses on people. Organizations or the SOC have to decide what types of skills and experienced people will run the SOC service that doesn't exist yet. This is a daunting task for many SOCs. SOC managers have difficulty trying to create documents for services that they will provide in the future before they have identified who will manage the services. This means they don't know who will manage daily tasks, what those daily tasks will involve, and how those tasks can change over time. This concept is by far the biggest challenge for developing a new service. It is hard to understand how a new SOC service will function daily without staff working the role and reporting back on the skill and experience expectations within the particular SOC and organization it supports. Not understanding the service leads to not understanding the specific required skills needed during the recruitment process. This leads to searching for the wrong people and asking the wrong questions during the interview of potential candidates.

Find the Right People

My recommendation is for SOCs to invest in people who understand general

cybersecurity concepts, understand how to use the basic tools and technology the SOC will be using, have strong problem-solving skills, and understand the organization's culture. You want people to understand what your organization does and how it accomplishes its goals before they analyze the services the SOC will provide and required skills needed to deliver them. That is why many SOCs start off with a simple log monitoring service. Once that basic service is operational, the SOC can build upon that basic service to handle more complex services with the goal of eventually offering some form of the eight fundamental SOC services listed at the beginning of this chapter. This approach enables team members to learn the fundamentals of the organization before taking on more complex tasks.

When I speak with SOC managers about recruiting challenges for new SOC services, I hear statements such as "I don't know who to hire because I don't have the service yet. I must first have the service before I know the type of person to run it." The good news is that there are resources available that can provide a general idea of skillsets associated with job roles. Chapter 4 explains how to develop a recruitment profile based on matching public profiles for similar jobs and provides other tips that can help overcome the challenges of hiring for services you are not yet delivering.

Here is a general checklist to consider as you plan Pillar 2:

- What job roles are required for the service?
- What is the estimated pay range for each job role?
- What skillsets are associated with each job role?
- What type of employee or contractor can and will be used (full-time employee, part-time employee, full-time contractor, managed service employee, etc.)?
- Who will manage the team?
- Where can people be recruited from?
- Is remote work permitted?
- Do certain roles require proximity to certain locations? For example, do team members need to be within a reasonable driving distance so they can attend in-person meetings? Does a datacenter administrator need to be within a reasonable driving distance of where critical services are

running?

- What are travel requirements for each role?
- What is the career path for each role?

Pillar 3: Technology

Technology is the third pillar of foundational SOC support services. Many security frameworks completely avoid mentioning specific technology because organizations tend to have radically different technology profiles, preferences, budgets, and needs. An example is a guideline for providing network access control. Every modern industry guideline highlights the need for NAC. The details for how NAC technology is delivered could include features within switches, external NAC technology that controls the switch, or host products that control access, meaning there is software on the host that is the NAC technology. In all three examples, different technology and approaches are used to accomplish NAC, yet they all end up providing the same high-level goals. This is why industry guidelines that reference NAC don't list the specific NAC technology, but instead speak of the desired outcome from the technology. You'll find the same approach to these recommendations in this section and throughout the book.

Regarding foundational SOC support services, all technology will warrant its own policies, staff, and procedures separate from those used for the technology supporting the services provided by the SOC. One good example of a common challenge that new SOCs run into regarding SOC support services is assigning responsibility to manage the SOC's datacenter. Some SOCs require their datacenter to be secured by SOC members, while other organizations hand off ownership to the organization's datacenter custodians. For this example, factors that influence the SOC's decision for managing the SOC's datacenter include budget, SOC service requirements, available resources, staff, and corporate policies, hence why every organization's SOC handles this function differently.

Securing SOC Technology

Regardless how the technology is hosted or managed, it is important to keep all SOC technology's data and management access to SOC tools segmented physically or logically from the organization. Some SOC services will require even further segmentation from the rest of the SOC data, such as digital forensics, which requires isolation of data involved with active investigations. Encryption can also be used to protect data that leaves the SOC's network.

Specific technology concepts for protecting the SOC are the same as those used to protect the organization. Chapter 1 described defense-in-depth concepts, including how to perform a SOC security capabilities assessment. The same defense-in-depth approach needs to occur regarding how the SOC itself is protected. My recommendation is to exclude the organization's defenses in the SOC security capabilities assessment if possible so if the organization is compromised, that threat won't impact the SOC itself.

Here is a general checklist to consider as you plan Pillar 3:

- What security tools are needed to protect the SOC and the new service being added?
- What additional technology is needed to add the new service (routing, storage, etc.)?
- What is the total cost (acquiring and installing the technology and training) for technology needed for the new service?
- What is the strategy to segment and/or isolate traffic for the technology?
- Are additional security or access restrictions needed for the new SOC service beyond what is included for internal SOC-related security practices?
- How long will the technology take to become operational?
- What redundancy requirements exist?
- What compliance and policies must be considered?

Evaluating the Three Pillars of Foundational SOC Support Services

As previously described, the required foundational SOC support services support the SOC itself. If the basic SOC foundation isn't ready, an alternative option is to use a service provider until the internal SOC is ready to take on the responsibilities associated with a service. When considering how to approach the three pillars of foundational SOC support services, the SOC needs to think about what services it can provide today, as defined in the mission statement and scope statement, as well as what it plans to eventually provide. When a new SOC service is identified, the requirements for the three pillars need to be planned out, including whether those requirements strain the resources of the SOC support services that are supporting the existing SOC services. You do not want to attempt to add a new service that leads to resource constraint on all existing SOC services!

The following are some general questions you need to answer regarding the impact of a new SOC service on the three pillars that support the SOC. Remember that if one of these three pillars is not supported correctly or is negatively impacted after you deploy a new SOC service, you will experience poor performance from the overall SOC itself.

- What additional people, process, and technology is needed?
- How will adding the new service impact the existing SOC services?
- What are the workspace considerations today and expected workspace needs in the future?
- Can cloud capabilities and/or remote work options help reduce the physical requirements for workspace needs?
- Are outsourced services needed at any point for the new service?
- What are the required skillsets and number of people needed?
- Is there a career path for SOC employees based on the new service?
- What are the policy, compliance, and data segmentation requirements for the new service?

SOC Service Areas

After a SOC decides to add a new service to its scope of offerings, the SOC

needs to define service areas. Service areas are a way of organizing the SOC service into more detailed and meaningful data by defining specifics on what the service will cover. For example, if the SOC will provide an incident response service, part of that service will include monitoring for threats. Otherwise, there would be no way to identify when a potential attack is occurring. Threat monitoring includes ingesting logs so that event analysis can occur when a security log violation is seen. This means log consumption and analysis is a subset of a SOC's incident management service and hence part of the service area supporting incident management.

SOC service areas and the technology that is used for each can spread across different SOC services. Log ingestion, event analysis, Windows host log monitoring, web filtering log ingestion, IPS log ingestion, and other specific items can be part of a SOC's risk management, vulnerability management, and incident response programs, hence these technologies and SOC service areas can be referenced more than once as a SOC defines SOC service areas for its core service offerings. SOC mission statements provide a high-level overview of SOC fundamental service offerings; however, the SOC service areas documentation must be much more granular to avoid confusion about what specifically will be offered and its relation to the rest of the SOC. This is why SOC service areas are not mentioned in the SOC mission statement or scope statement. Referring again to the example of having log ingestion within the SOC, the SOC service areas document for each SOC service will need to define if log ingestion is included, the purpose for the service's use of log ingestion, and which teams responsible for the SOC service are dependent on this technology. These details would be included in the procedures associated with a SOC fundamental service.

FIRST's CSIRT

Figure 3-2 comes from the FIRST's Computer Security Incident Response Team (CSIRT) Services Framework depicting service areas and services that a SOC can provide. Note that some of these services are labeled differently than how I label the eight fundamental services that all SOCs should offer. For instance, situational and security awareness is called Knowledge Transfer in Figure 3-2. Also, services such as research and development that I

recommend are not included in this example model. Nonetheless, Figure 3-2 provides an example of how service areas can be developed to represent what supports core SOC services.

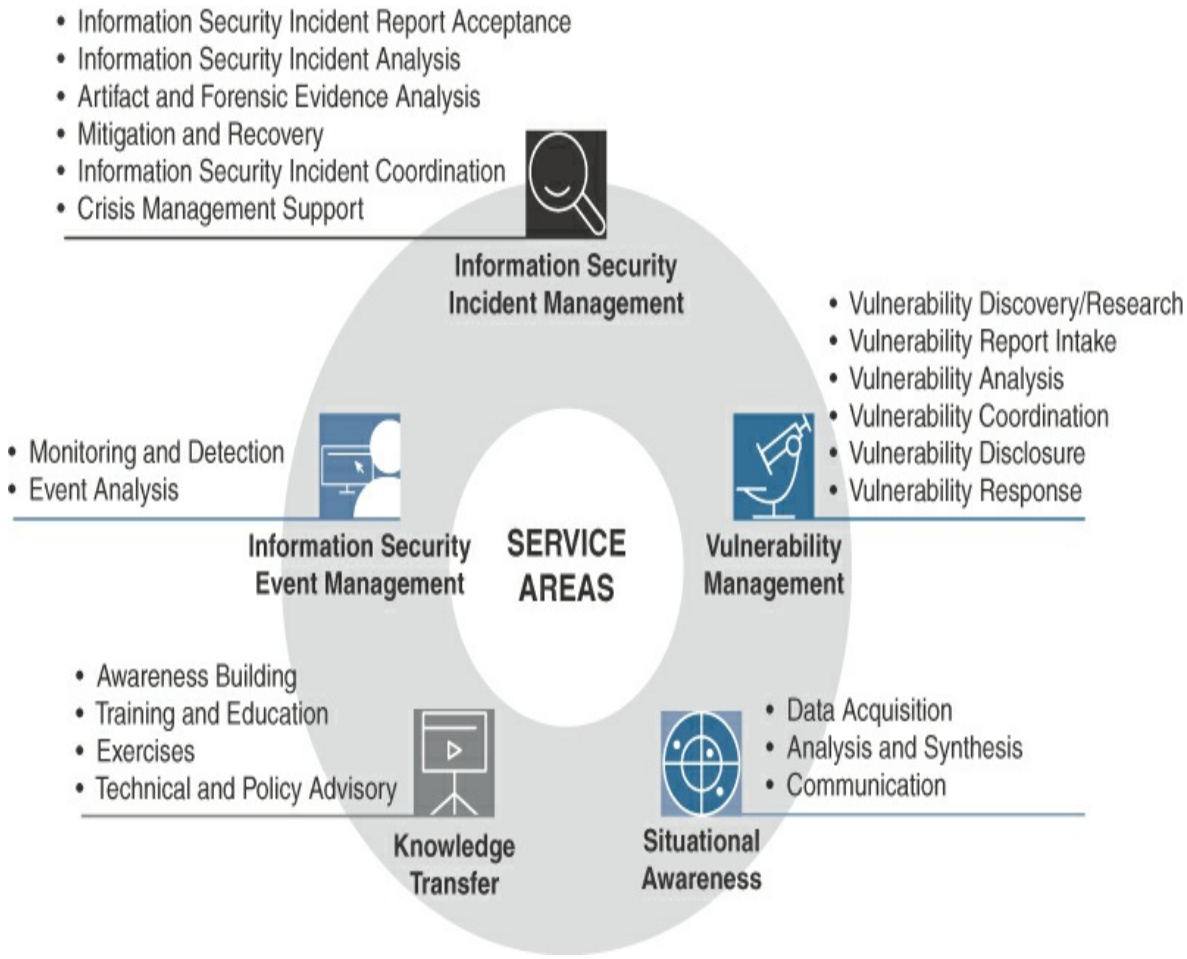

FIGURE 3-2 FIRST CSIRT Services and Service Areas

Note

As a reminder, most industry guidelines use different terminology for their SOC service models. Regardless, I find the concepts overlap, which is why I will continue to point out how my eight fundamental SOC services relate to other guidelines for SOC services. Doing so provides you multiple viewpoints of SOC service best practices.

Developing SOC Service Areas

Figure 3-3 shows how the concepts of SOC support services, the eight recommended SOC fundamental services, and SOC service areas all relate. The work environment, people, and technology are required for the SOC to function. These are the ingredients of the SOC support services. With those ingredients, the SOC is capable of delivering one or more of the eight fundamental SOC services. Finally, there are multiple SOC service areas that make up various services, technologies, and practices that can be part of one or more SOC fundamental services. Previously, I gave the example of a log digest service that collects logs for an analyst to review. That service area would be useful for multiple SOC fundamental services, including incident response, analysis, and risk management.

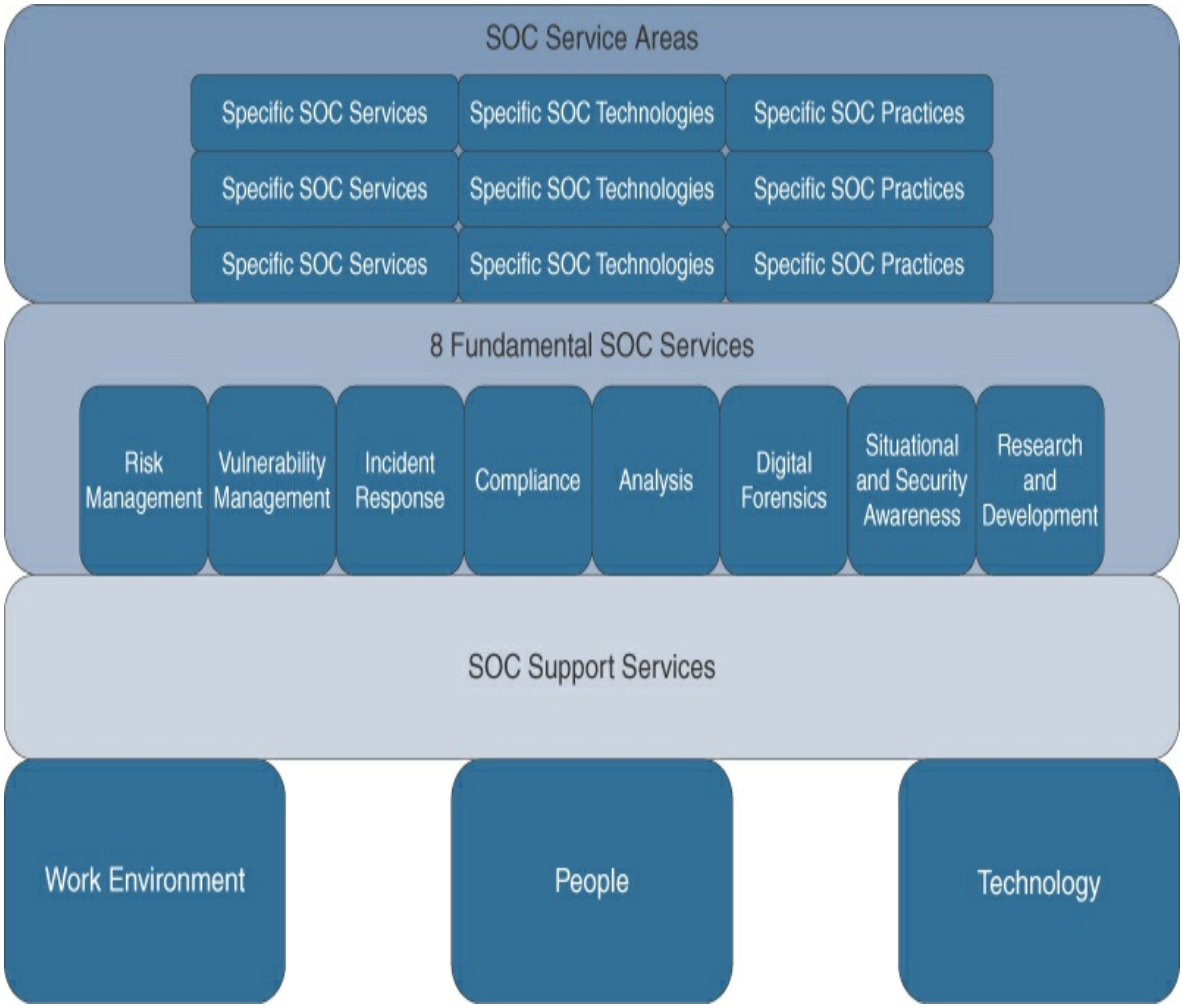

FIGURE 3-3 Building Fundamental SOC Services on Foundational Support Services and Adding a Pool of SOC Service Areas

Notice that Figure 3-3 also shows how the eight fundamental SOC services can't exist without the foundational SOC support services. Also notice that the SOC service areas should not be directly referenced by a specific SOC service but instead referenced by multiple SOC services. Regarding your SOC scope and mission statements, I recommend you refer only to the eight fundamental services. You can include details about SOC service area services in other, more detailed documents such as procedures.

> **Note**
>
> The documentation of SOC service areas is critical for standing up or documenting existing SOC services for planning purposes. When you are preparing high-level documentation that is public facing, such as mission statements, scope statements, and policies, I recommend that you refer only to the eight fundamental SOC services (not the service areas) to keep the language high level and outcome focused.

The service list that follows is an example of planning a basic SOC service, which I'm labeling as a SOC service area. The list for this example is not complete; however, this is the normal development procedure in the initial phases. As people are recruited for the service area and more requirements are understood, the SOC can finalize its plan to implement the service area, including making required adjustments to the service list. Also note that this example focuses on a basic service area that would become a subcategory for a fundamental SOC service such as the incident response service or analysis service. For this reason, this example service area doesn't provide any outcomes outside of data collection, which other SOC services will depend on to function properly. I sometimes see this type of service area developed within new SOCs as a steppingstone for more mature future services.

SOC Service Area: Ingest log data from security devices

- **Goal:** Ingest continuous real-time log streams from identified devices.

- **Outcome:** Logs can be queried, analyzed, archived, or inspected by SOC analysts or other security tools.
- **Dependencies:** Everything needed for the service to exist
  - Log collection tool
  - Storage
  - Facility to store and power technology
  - Support staff
- **Success criteria:** Collect, aggregate and make available a summary of up to 90 days of event data in a secure manner.
  - Archive older data for up to 2 years
  - Quality data as validated by the SOC analyst
  - Analyst can pull needed data within 10 minutes or less
- **Requirements:**
  - Identify skills and SOC personnel to support log collection
  - Determine amount of effort required to support log collection
  - Identify security appliances that will generate logs to be collected
  - Develop requirements for log format(s) accepted by SOC program
  - Determine whether log relay is needed on a central server to send logs to other devices
  - Normalize data and time for logs and ensure time zones, NTP, and other time sync protocol policies can be met
  - Determine what tools will collect logs
  - Determine any network, firewall, or routing configurations that will be needed
  - Understand log retention requirements with internal policies, compliance, and other requirements
  - Ensure log transfers work as expected

In-House Services vs. External Services

One important question a SOC will need to decide as it develops and evaluates services is whether the SOC will provide the service in-house or outsource it to a third party. The answer to this question comes from performing the service evaluation process covered in the previous section. Both in-house and outsourced approaches have respective advantages and disadvantages.

I covered comparing the value and disadvantages of in-house services against external services in Chapter 1, in the section "In-House vs. Outsourcing." Table 3-1 is a summary of the advantages of using both approaches.

> **Note**
>
> A partial internal, partial external worker can be seen as a contractor that works within a SOC.

TABLE 3-1 Benefits of Internal Services and External Services

Internal Service Benefits	External Service Benefits
Knowledge of business	OPEX costs that can be spread out
Data stored internally	No conflict of interest
Cross-department correlation	Scalability and flexibility
Tailored requirements	Leverage other customer trends

Table 3-2 is a summary of the disadvantages.

TABLE 3-2 Disadvantages of Internal Services and External Services

Internal Service Disadvantages	External Service Disadvantages
Cost	Limited business knowledge

People (hire/maintain)	External tools and data flow
Potential conflict of interest	Lack of communication
ROI concerns	Usually not detected people
	Limited customization
	Services are limited based on cost (e.g., tiered Gold, Silver, and Bronze services)

Many organizations developing new SOC services will first use external services while they build their own in-house services. The idea is to shadow the service provider and slowly outsource less until external services are no longer required. An example is to hire a few in-house experts to act as the program leaders and outsource all first-level tasks to an external party while the internal team is being built. This strategy attempts to balance the benefits of both in-house and external SOC approaches. If you believe your SOC is not ready to properly deliver a fundamental SOC service, outsource it until you are capable.

Contracted vs. Employee Job Roles

It is critical that all SOC services, regardless of in-house or outsourced status, have specific positions outlined and filled. For example, if your SOC is going to be responsible for providing forensics and incident response services, you need to ensure that you have positions dedicated to servicing roles within these offerings. If your SOC outsources those roles, you will still require an internal role with the skillsets to interact with the third-party service provider. Defining a role is important not only to enable the service provider to understand the role's skill requirements but also to give internal employees an opportunity to develop those skills to eventually take over the role. Recommended practice dictates that you develop an organizational chart showcasing all roles within the SOC, including any that are outsourced.

One other option that is pretty common within SOCs today is a hybrid model.

This approach blends outsourced and internal resources to run a SOC service. For example, the team lead can be an internal employee, with a team of analysts composed of both outsourced contractors and employees of the SOC team. This works well because, for example, there is a blend of analysts who are impacted by internal politics as employees of the organization and analysts from the outside that bring unique unbiased experience. My recommendation, however, is to ensure a SOC employee has responsibility over a contracted job role.

I recall an older job role in which I was the team lead for a government agency while working as a contracted resource. I honestly had too much power with few checks and balances, including creating job descriptions for other contracted roles within the government agency and being responsible for choosing what technology the government agency would buy. What was to stop me from creating special roles for people that only my real employer, the contracting company, could fill? What was to stop me from choosing only the tools my real employer, the contracting company, was selling? My ethical standards stopped me from taking such actions, of course, but contractors' ethical standards vary widely. This is why I stress the need for including SOC-employed staff even when contracted resources are being used.

SOC Service Job Goals

After the SOC identifies a service and roles, the next step is to define the goal for each role. Goals do not include specific tasks, which are explained within associated procedure documentation. Instead, goals are a high-level summary of what is in scope to be considered success for the service. For example, a vulnerability analyst in the SOC will be responsible for ensuring vulnerability management software is installed, up-to-date, and running on endpoints, evaluating scan results, tuning the vulnerability management solution, and working with the correct IT teams to patch vulnerabilities or create another mitigation plan. The analyst in that situation is going to need tools, software licenses, training, access to different teams, and the authority to enforce corporate policy for vulnerability management. Chapter 4 reviews common responsibilities found within roles for each SOC service covered in this chapter. Each job role that is listed in your organizational chart should

include a summary of what tasks are in scope for success, such as shown in Figure 3-4, which depicts typical roles and responsibilities found within a SOC. Chapter 4 dives deeper into SOC roles, but Figure 3-4 is a good reference point for gaining a general idea of the types of roles associated with different SOC services.

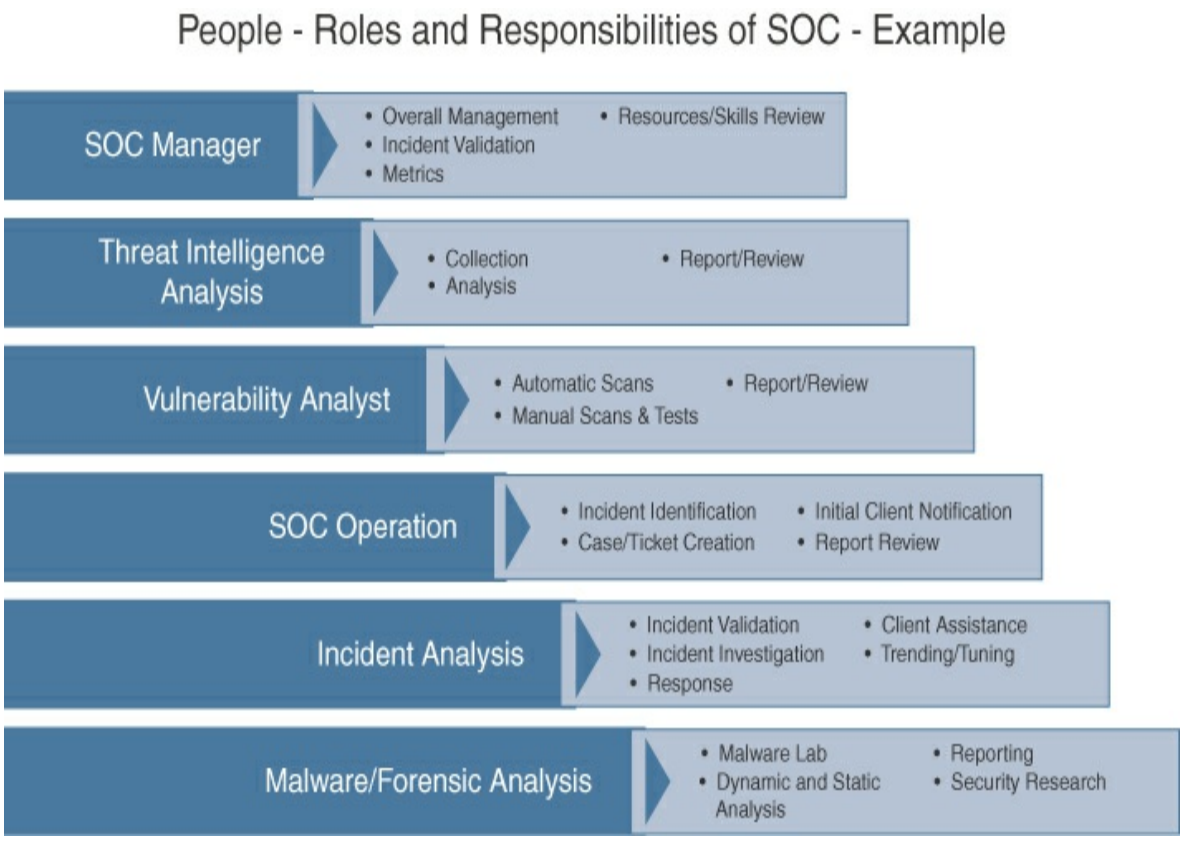

FIGURE 3-4 Sample SOC Roles and Responsibilities

Resource Planning

Another job role task is resource planning, meaning determining how many people should fill a job role. Resource planning for a new SOC service can be very challenging or very simple, depending on the size of the SOC and how complex the service is that is being planned. Estimating resources involves ensuring that the SOC has the right number of people not only to handle the volume of threats but also for job role redundancy to accommodate vacation,

loss of key personnel, or other job rotation processes. A SOC must avoid the scenario where one person runs a SOC service and decides to leave the organization, leaving the SOC without that service until that person is replaced. You also don't want to oversaturate a SOC service with resources, due to the high costs associated with staffing roles and keeping employees happy.

Many SOCs rely on technology to help solve the limited-resource issue, using automation to reduce the need for manual human tasks. For example, employees don't have unlimited time to weed through pages of alerts and consistently take appropriate action, so SOCs implement technology that can help prioritize the incidents that require manual review, making the process more efficient and minimizing the number of people needed to support the SOC. Chapter 10, "Data Orchestration," provides a closer look at automation and orchestration concepts. Know that technology will never replace people. Instead, technology will reduce the number of mundane and repetitive tasks so the SOC analysts can use their time on more complex work.

Unfortunately, there is no simple formula for calculating the number of people that should fill a job role based on all the factors that impact how your SOC handles the demand for a service. Recommended practice dictates planning for an estimated head count plus one for redundancy, running the service, and then evaluating success during a mandatory lessons-learned follow-up meeting on how the service is provided. The lessons-learned meeting evaluates the success of the service, focuses on which identified issues can be analyzed for potential mitigation adjustments (including recruiting more people), and considers ways of adjusting the process or acquiring technology that includes automation capabilities.

Note

I find that when a SOC overstaffs a SOC service, it always has other SOC services that can use more help. In contrast, I constantly hear how SOC teams feel understaffed and overwhelmed with work. Based on these observations, if you have the ability to recruit more people for a SOC service, I advise you to take those people while you can. It's a better option than later needing more people and not having them available!

Service Maturity: If You Build It, They Will Come

Maturity is a critical component of a SOC service. A SOC doesn't just obtain a service, they build it using internal or outsourced resources. For example, if I advise an organization to develop a vulnerability management service, doing so will take more than purchasing a vulnerability scanner. Figure 3-5 is a maturity model example taken from ISACA's COBIT 2019 guideline. Not having a vulnerability management service would mean a grade zero maturity ranking for the vulnerability management SOC service. Acquiring a vulnerability scanner would allow for ad hoc scanning, raising the maturity of the vulnerability management service to grade one. As a repeatable program is developed, the maturity will also increase.

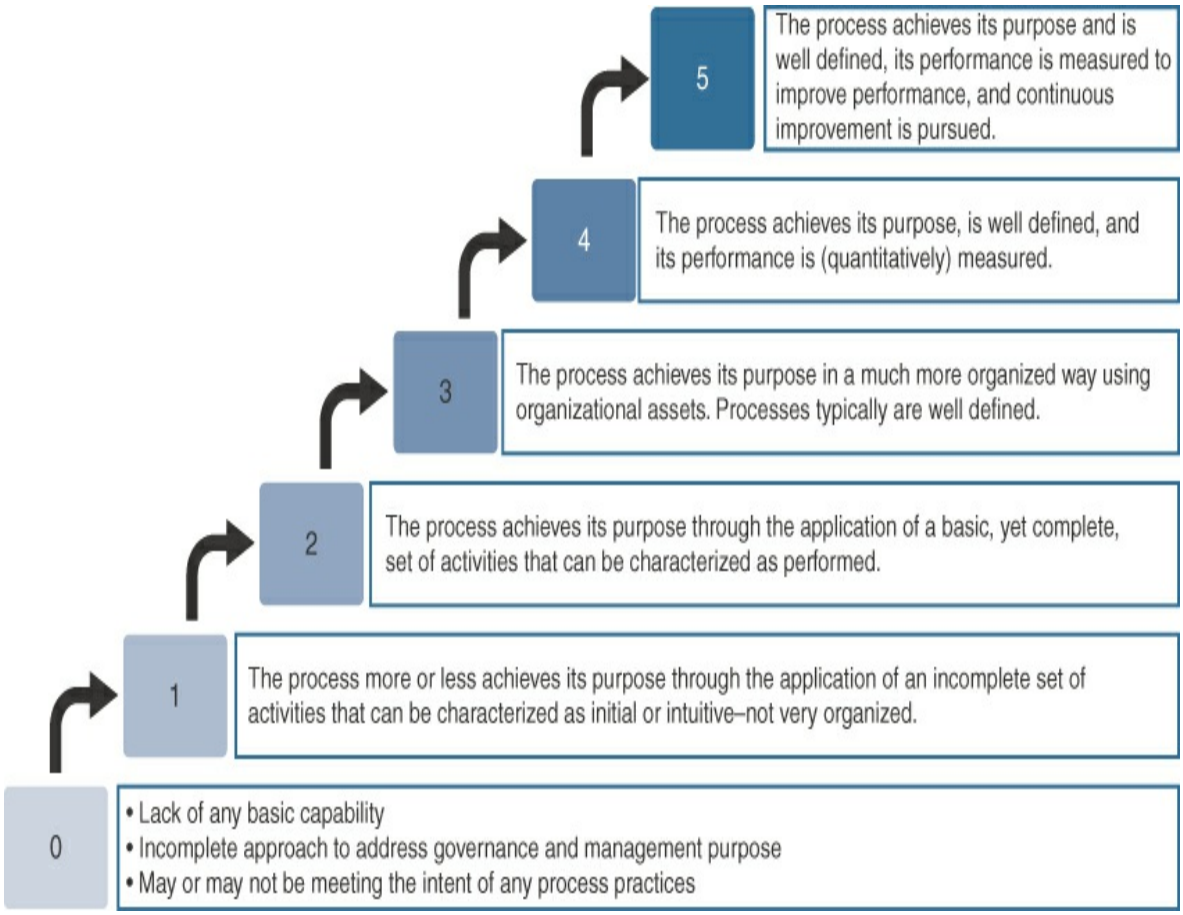

FIGURE 3-5 Maturity Grading Scale Example

Using a maturity measuring system is critical to the success of any SOC service for a few reasons:

- You have a way to explain to nontechnical decision-makers why continuous investments are required to make a service successful. This will help overcome statements such as "we already bought you the tools you need."
- You can create improvement goals, which can include reward systems when goals are achieved. Reward systems lead to career growth and employee dedication to the SOC's mission.
- Maturity models can provide a potential roadmap to improve the quality of a service.
- Maturity models highlight the important impact of people, process, and technology on a service, because all are required to meet higher-level maturity categories.

As you develop SOC services, you will need to evaluate at what maturity level they are operating. You can validate the results with the SOC sponsor, set goals for improving services, and plan for requirements in people, process, and technology to accomplish such improvements. Using this approach will allow for all technical and nontechnical parties to understand the current and potential future status of a SOC service. Chapter 1 covered how to evaluate SOC maturity.

Let's move into reviewing the eight fundamental services of mature SOC practices.

SOC Service 1: Risk Management

Found in most SOCs around the world, risk management is a fundamental service because it deals with all forms of risk. This means other services are essentially a subset of risk management. For example, vulnerability management deals only with technical vulnerabilities, while risk management covers all forms of risk including nontechnical vulnerabilities such as the risk

of social engineering attacks and the risk of physical security breaches.

Four Responses to Risk

Every organization has various forms of risk. The responsibility of the risk management team is to provide a response to encountered risk. An organization can take essentially four options as it responds to risk:

Note

There are other similar risk models that include the five options I have provided for responding to risk. Chapter 9 highlights how ISO 27005-2018 suggests four options for responding to risk using similar language: risk retention, risk avoidance, risk modification, and risk sharing. As with all guidelines, the language might be different, but the concepts are the same.

- **Reduce:** Apply a service, tools, or other countermeasures to reduce the risk
- **Accept:** Identify that the risk is acceptable and do nothing
- **Reject:** Identify that the risk is not acceptable but still do nothing
- **Transfer:** Share or outsource the ris

Note

Accepting risk and rejecting risk are similar but there is a key difference. *Rejecting* doesn't agree or accept the risk, however, nothing is done, while *accepting* the risk means the risk is within acceptable boundaries, so nothing needs to be done.

Reducing risk is likely what you expect as a response to risk. Reducing risk

can be any type of people, process, or technology change that lowers the level of risk. Ideally, you want to remove the risk through measures such as patching a vulnerable server. If a patch isn't available, you could place a security appliance in front of the vulnerable system, essentially preventing any external threat from exploiting the vulnerable server. Another option could be assigning a team to monitor traffic to the vulnerable server, allowing for a lower risk of the system being exploited without being noticed. Chapter 9 reviews how to choose which is the best option for addressing vulnerabilities.

Accepting risk means the risk is at a tolerable level. An example of this is not buying earthquake insurance for a building located where there are not any earthquakes such as the state of Florida. Yes, an earthquake could occur in Florida, but it hasn't happened in many years and likely will not happen anytime soon.

Rejecting the risk means the event is not acceptable but mitigation isn't worth the effort. Looking back at the earthquake example, imagine your office is located where there are frequent earthquakes, such as California. Your SOC will have to decide if it's worth mitigating the risk of earthquake damage, which can become extremely costly considering all that is required to secure a building. If it cost three times the effort to secure a building from earthquake damage, the decision could be made to reject the risk, meaning you do not accept it as within tolerance but it is not worth mitigating and, therefore, your SOC will deal with the potential damages if the threat occurs.

Keeping with our earthquake example, what many organizations do when faced with reducing the risk of natural disasters is *transfer* the risk through acquiring insurance. Rather than securing the building, the organization can acquire coverage that will reimburse losses if the event occurs, essentially transferring any concerns of the threat to another party, which in this case is the insurance company.

Note

I am often asked for my thoughts on cyber insurance. Personally, I believe that most organizations that don't have cyber insurance eventually will. The main reason for this belief is the fact that no C-

level executives for an organization would want to report they didn't think they needed insurance following a major cyber incident. C-level execs would rather say "Yes, we have a partner that provides that service," because when it comes to the public's view of who is responsible for a cyber breach, all eyes are on the C-level team. Does every organization have cyber insurance today? Certainly not; however, the trend I am seeing is a hypergrowth in the issuance of cyber insurance policies as organizations learn about the option to use cyber insurance. Chapter 9 looks closer at best practices for evaluating cyber insurance.

Reducing Risk

Reducing risk isn't always a simple fix, because new risk can be introduced based on what is being changed. Also, there will be a cost associated with each risk reduction option. If the cost outweighs the fix, it will not be worth mitigating the risk, leading to the choice of risk rejection. Risk isn't an all-or-nothing measurement, either. There are different models that can be used to calculate the amount of risk. One example is the expected monetary value (EMV) analysis. EVM analysis is a statistical technique that calculates the average outcome when the future includes scenarios that might or might not happen. EMV analysis takes the value of each possible outcome and multiplies it by its probability of occurrence. The probability-weighted values of the possible outcomes are later added together. This approach can be extremely valuable because you not only are able to see the total cost if all risk were to be exploited, but you can also generate a more realistic dollar value of the overall risk when considering multiple situations, which is more realistic.

EMV Approach

Let's work through an example of using the EMV approach to risk analysis. For this example, let's say there are eight different risks that your SOC must decide to deal with. Your SOC sponsor wants to understand the dollar value

associated with all of the risk. You can create an EMV chart that features all eight risk items. You can estimate what the probability of each risk occurring would be as well as estimate the cost if the risk was to be exploited. What is more important is understanding the cost compared against the risk actually occurring. Let's say for this example that accounting for all eight risks results in a required risk budget of $118,000. This large dollar value can be a huge ask for many organizations regarding a risk contingency budget; however, computing that the risk reality, known as *risk contingency*, is $33,500 may be more reasonable considering the likelihood that all events will not occur. It would be more likely that a few events occur, such as events C and E occur, which would mean $28,000 would be needed. Using the EVM budget, the $28,000 would be pulled from the $33,500 planned budget. Using the EVM in this manner allows for justification of a reasonable risk contingency budget that nontechnical leaders will need to understand. Table 3-3 shows the numbers from this example.

TABLE 3-3 EMV Calculation Example

Risk	Risk Probability	Cost Impact	Risk Contingency
A	80%	$10,000	$8,000
B	30%	$30,000	$9,000
C	50%	$8,000	$4,000
D	10%	$40,000	$4,000
E	30%	$20,000	$6,000
F	25%	$10,000	$2,500
Total		**$118,000**	**$33,500**

Risk Documentation

Another risk management technique is to properly document each risk so that estimates and decisions can be made regarding the associated impact. Table 3-4 is an example format you can use to ensure the risk management team captures each risk that the organization needs to address. The asset owner and security engineer will need to work together and agree on what goes into this

form to ensure that the results are calculated accurately. If the input is not accurate, using risk calculations such as the EMV technique would not provide useful results. It is also important to point out that the expected input for the empty Table 3-4 template could be pages of details.

TABLE 3-4 Risk Documentation Template

Name of Risk	Description	Source (with Explanation)	Likelihood of Occurrence (with Justification)	Severity of Impact (with Justification)	Controllability (with Justification)
1.					
2.					
3.					
4.					
5.					
6.					
7.					
8.					

Each risk must be tracked during its lifecycle within the risk management response process. This can be done using a *risk register* system, which documents essential information for each risk that is identified. There are different factors to consider including in the risk register you choose to create. Use the following template as a bare minimum of content that you will need to capture and successfully track risk events:

- Tracking Number/Date
- Assigned Engineer(s)
- Description
- Cause
- Impact

- Severity and Likelihood
- Risk Response
- Root Causes

Addressing Risk

The next risk management concept is addressing risk. For vulnerability management, the mitigation process will be technical, which can include patches or incorporating security tools to protect the vulnerable system. Risk management mitigation can be much broader since risk can be anything from environmental to physical in nature. To accommodate this wide range of potential mitigation options, the risk management team must be able to define a business contingency plan and a corresponding implementation plan.

There are many factors that can impact both plans depending on the type of risk, what it impacts directly and indirectly, compliance considerations, health factors, and other factors. The following two lists are examples of templates that can be used by a risk management team member to collect information that will be needed for the risk management team to make a proper response. Notice that these documents are much more detailed than what was collected in the risk documentation template. There are many variations of both the risk documentation template and the business contingency documentation that can be used and should be adjusted to the risk being addressed.

Business Contingency Plan

- Strategic Pre-Incident Changes
  - Sensitive Data
  - Normal Data Protection
  - Disruption Data Protection
  - Ethical Use of Data
  - Customer Records
  - Normal Security Measures

- Disruption Security Measures
- Ethical Use Protections
- Communication Plan
- Stakeholders
- Stakeholder Communications
- Restoration of Operations

Implementation of Contingency Plan

- Strategic Incident Changes
  - Communication of the CP
  - Monitoring and Testing of the CP
  - Adjustment of the CP
  - Communication of Changes

Risk Heat Mapping

One popular approach to quantifying risk is to use heat mapping, which provides a simplified view of the potential impact of an event. One method is to use the risk heatmap template shown in Figure 3-6. This template compares the probability of an incident against the severity, the two critical factors to understand the true threat of a risk. For example, a catastrophic risk that is extremely rare is a concern but can be seen as moderate because it is unlikely to occur. More elaborate heatmaps can include associated dollar costs for impact and mitigation using the EMV approach to calculate estimated risk contiguously planning.

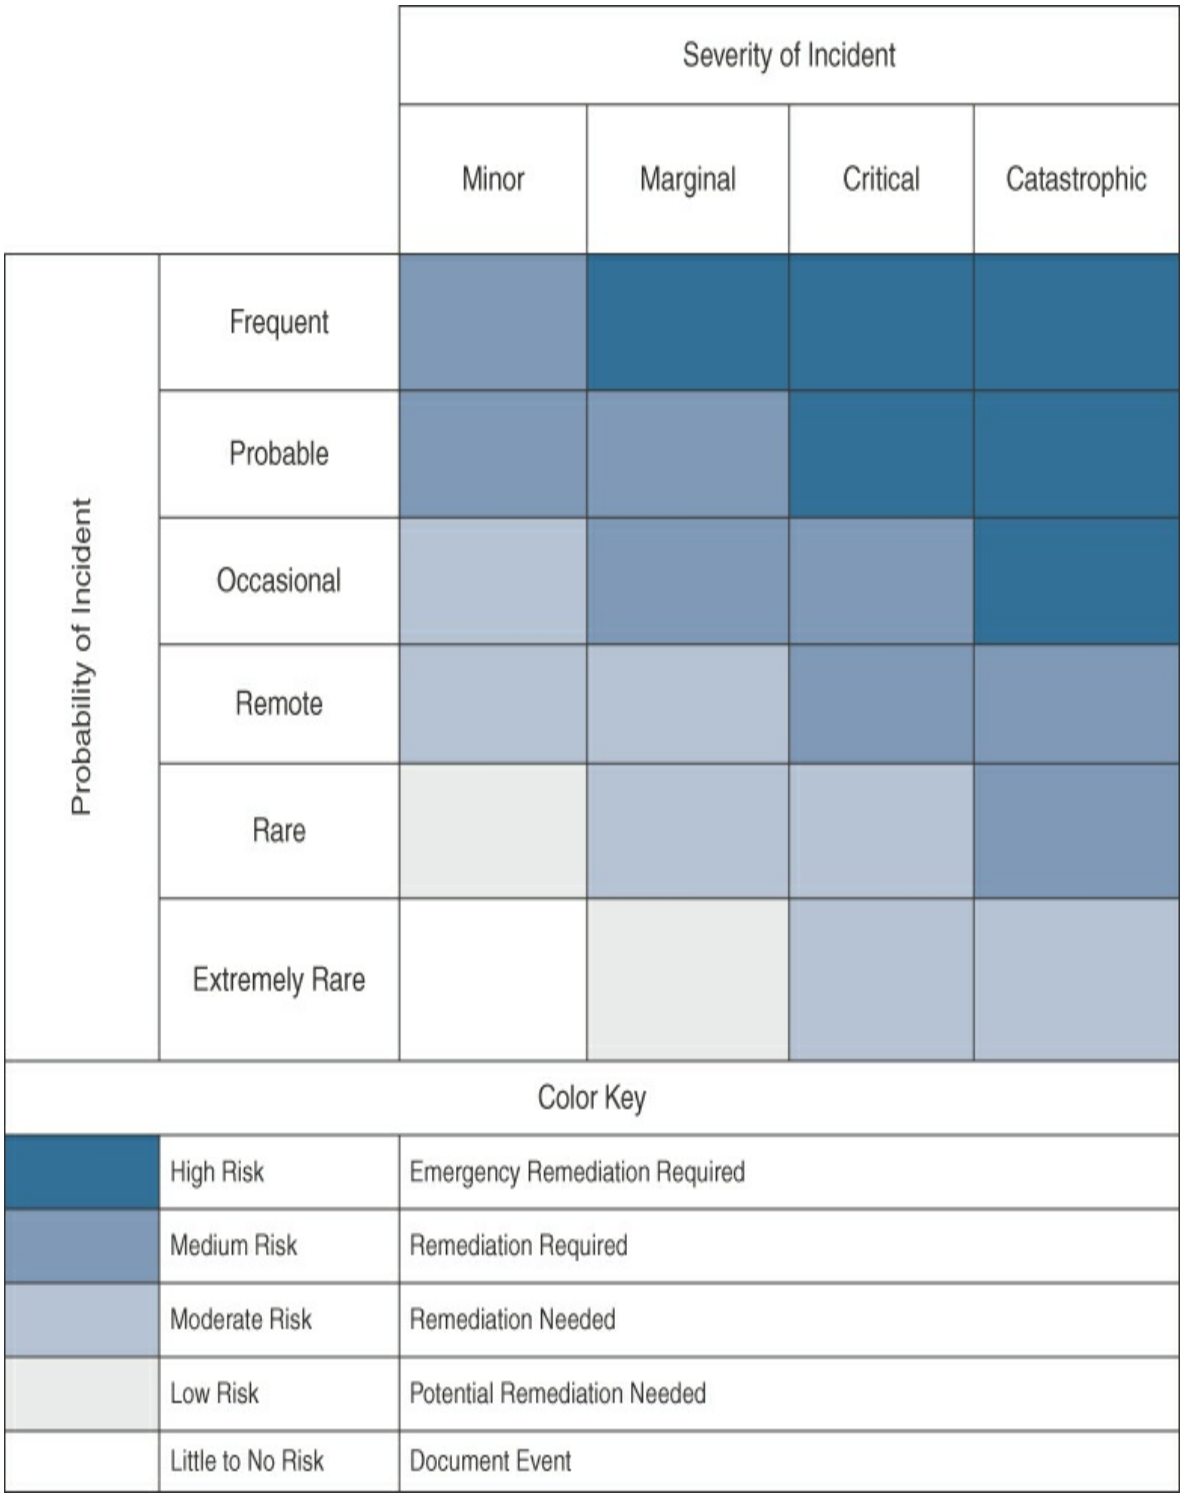

FIGURE 3-6 Risk Heatmap Example

There are more topics to cover that fall under risk management, including expectations for addressing and recovering from risk. Chapter 6, "Reducing

Risk and Exceeding Compliance," is dedicated to dealing with risk and risk recovery.

SOC Service 2: Vulnerability Management

Attackers use vulnerabilities to find cracks in security tools as well as the systems those tools are built to defend. If vulnerabilities don't exist, attackers have nothing to exploit. Unfortunately, every organization typically has some form of vulnerability, whether it's from missing software updates, flaws in how IT systems are designed, weaknesses in how systems work together, or any one of many other factors. It is common for organizations to focus on large, visible systems while overlooking secondary support systems that have access to similar important data. It is critical to subject all systems within the organization to vulnerability evaluation based on their value and potential risk. Additionally, maintaining compliance with the likes of Sarbanes-Oxley, PCI DSS, and others may require your SOC to evaluate all systems regularly for vulnerabilities.

Vulnerability Management Best Practice

A SOC needs both prevention and detection capabilities to prevent vulnerabilities from being introduced into the organization and to detect vulnerabilities that already exist. A vulnerability management program focuses on technical vulnerabilities, meaning weaknesses within IT systems. This is essentially a subset of risk management because only IT-related vulnerabilities are within scope. To develop a repeatable vulnerability management service, the SOC must take the following specific steps, depicted in Figure 3-7:

1. **Asset inventory:** Identify all assets within scope of being evaluated by the vulnerability management service.
2. **Information management:** Collect information about the asset, including who owns it, its value to the organization, and where it is located on the network.

3. **Risk assessment:** Prioritize and assess devices with the highest value to the organization before less valuable systems. An example is evaluating a datacenter before evaluating a user desktop. Also consider whether a system is in contact with higher priority systems, which if breached could lead to a much bigger breach.
4. **Vulnerability assessment:** Assess devices for vulnerabilities.
5. **Report and remediate:** Perform mitigation when applicable.
6. **Respond:** Follow up with asset owner, update and close out the ticket.

As indicated in Figure 3-7, these steps are cyclical. There is a set time that mitigation is performed before the entire process is repeated and vulnerabilities that were not addressed are added back to the "to be fixed" list.

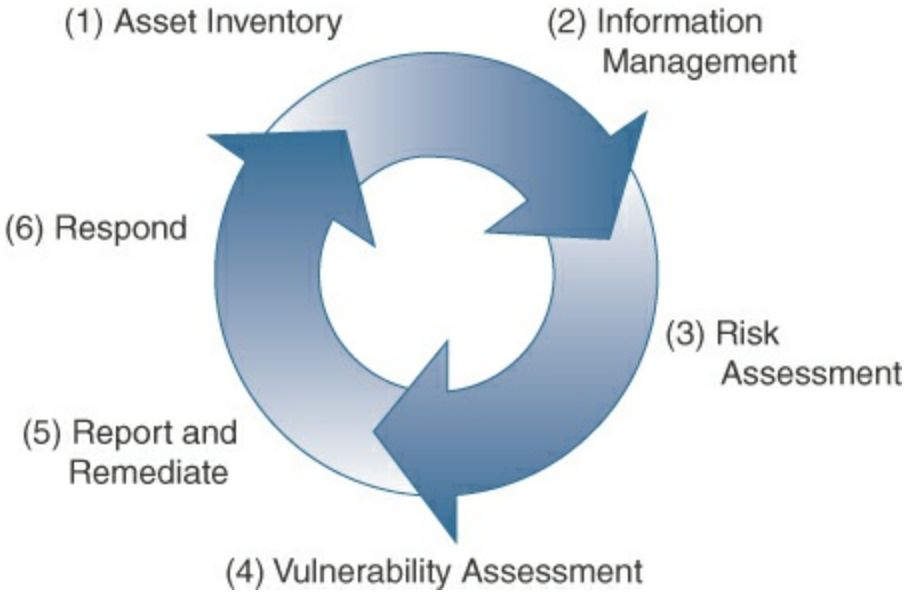

FIGURE 3-7 Best Practice for Vulnerability Management

There are different ways to identify vulnerabilities on IT systems, as shown as step 4 in Figure 3-7. The most common way is through vulnerability scanning, which relies heavily on vulnerability scanning tools. Vulnerability scanning is only a best guest regarding whether a system is truly vulnerable to attack. A more in-depth approach to evaluating systems for vulnerabilities is performing a penetration test. The benefits of a penetration test is obtaining

more accurate vulnerability data; however, the cost in time and risk to systems being tested, along with the needed skills to properly perform a penetration test, can be taxing to the SOC. It is common for vulnerability scanning to be performed as often as daily or weekly while penetration testing typically occurs less frequently, such as on a quarterly or semi-annual basis.

The next sections dive deeper into what is involved with vulnerability scanning and penetration testing. I'll start with vulnerability scanning, because it's the more commonly used service to identify vulnerabilities on IT systems. To better understand what is involved with vulnerability scanning, I first need to cover vulnerability scanning tools.

Vulnerability Scanning Tools

SOCs often run applications that perform vulnerability scans and audits in their organizations to understand where weaknesses occur in technical systems. This approach gives the SOC an idea of where the most risk lives in their organization based on a combination of a potential vulnerability and the value the vulnerable asset has to the organization. If SOCs cannot work with application and system owners to adhere to policies, they can still mitigate risks by segmenting vulnerable systems into untrusted network zones or include security tools to protect the vulnerable systems from being attacked. It is ideal to remove the vulnerability whenever possible rather than securing around it, which commonly is done by applying a patch or other fix to the system with the vulnerability. There might be situations in which a patch isn't available and *residual risk* must also be considered, which is potential risk associated with the unpatched system. Best practice is to first test any patch before applying it to a live system. If testing shows unreliable results or a patch is not available, applying security around the system, such as putting a security tool like an IPS in front of the system, will be the next best option.

Multiple tools and applications are available that can help SOCs keep track of vulnerabilities. Tools range from network access control technology, which can identify and evaluate devices upon connection, network scanners that perform network- or host-based scanning, and passive scanning tools, which compare existing device data against a database containing known

vulnerabilities. Recommending a commercial tool is a little bit like a religious debate, meaning there is no absolute winner or best tool. Whatever approach you use, my recommendation is to select a vulnerability management solution that has the capabilities to do the following:

- Automate and ensure any device connected to the network is accounted for
- Scan devices upon connection for vulnerabilities
- Periodically scan devices from a network viewpoint
- Periodically scan devices from a host or client viewpoint
- Track all vulnerable systems throughout the vulnerability lifecycle

Chapter 9 covers all of these recommendations in more detail. Most modern vulnerability management solutions include the capability to accomplish all of these recommendations. Tenable's tenable.sc, Rapid7's Nexpose, and Qualys VMDR are all examples of vulnerability management packages that offer network- and host-based vulnerability scanning as well as ticket tracking systems. Figure 3-8 shows a customized dashboard within tenable.sc that keeps track of vulnerabilities and changes.

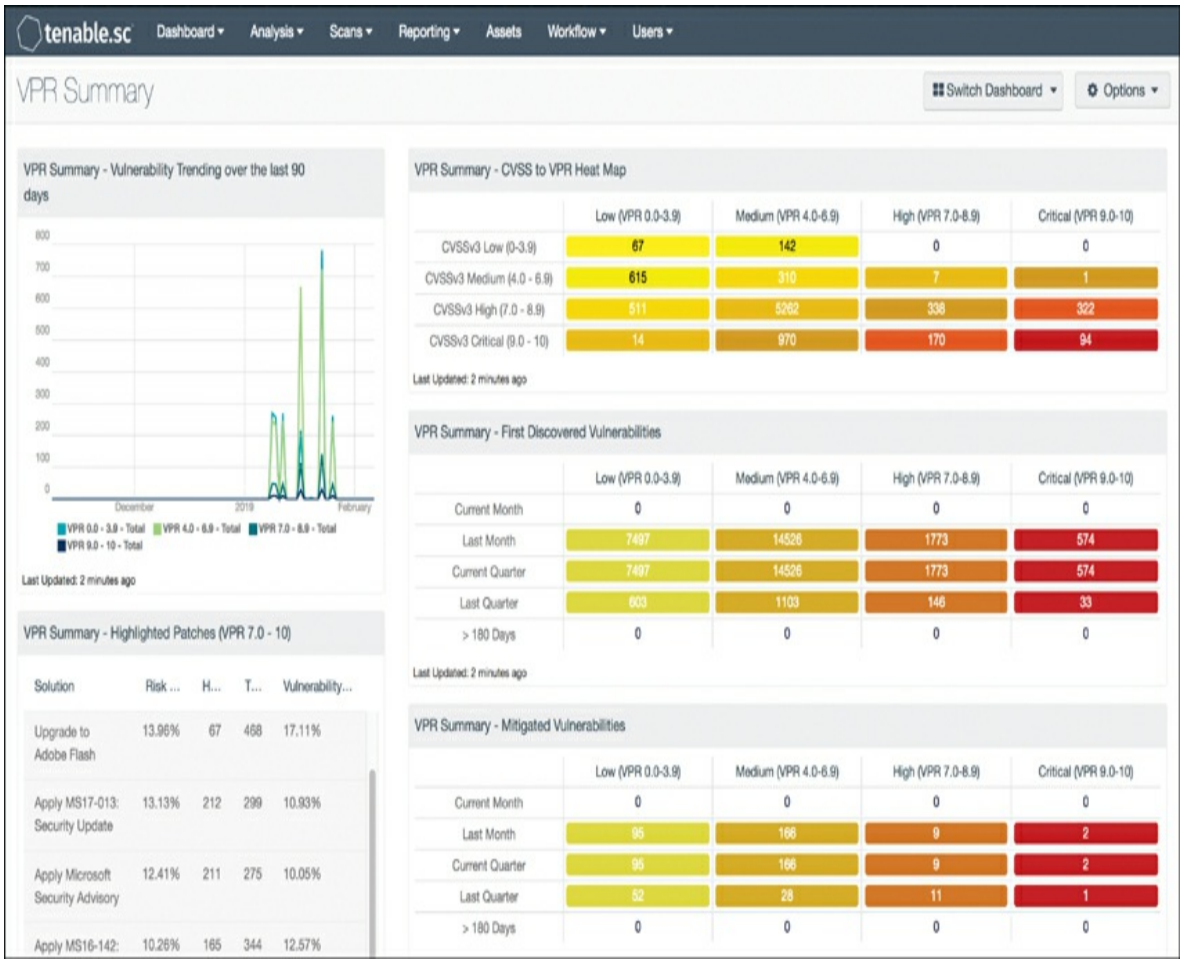

FIGURE 3-8 Tenable.sc Vulnerability Tracking

Open Source Option: OpenVAS

If you are looking for an open-source vulnerability scanner, OpenVAS is a popular choice. OpenVAS is technically a software framework that is composed of several software packages and tools. It has an open plugin architecture that allows its functionality to be expanded. The plugin language supports the Nessus Attack Scripting Language (NASL), which is extremely popular among security professionals for detecting attacks. OpenVAS has many of the same basic features for scanning and vulnerability detection as commercial vendors, but it lacks some of the configuration management features common with commercial products that keep track of changes. However, the open-source nature of the software makes it extremely customizable. Figure 3-9 shows output of OpenVAS using the Greenbone

GUI management interface.

FIGURE 3-9 OpenVAS GUI Example

Vulnerability Tracking

I continue to highlight the importance of including tracking within a

vulnerability management program as it is essential to maintain documentation of the configuration of network devices, security devices, and major systems pre- and post-vulnerability mitigation. Furthermore, it is important to maintain documentation of changes, when they occur, and why they occur. Archives in secure backup locations must be kept of all prior working configurations that can be used in an emergency prior to deploying any type of mitigation process. Well-established organizations have policies governing the rate of the changes that occur and who can perform changes.

Regardless of the tool that the vulnerability management service uses to track changes, the following list covers the process for requesting, evaluating, and implementing changes. I find that mature vulnerability management services include these steps as part of the change process associated with mitigating a vulnerability.

1. Request the change, with all required documentation. The request may originate from end users, helpdesk staff, or IT management.
2. Record the change request into the tracking system. This action formally enters the request into the change control process.
3. Evaluate the proposed change for its effect on the impacted information system, organizational security, and any related systems. This may involve testing in a development environment, if possible.
4. Approve or deny the change. This step is usually performed by the vulnerability engineer or equivalent authority. Regardless of the decision, it should be recorded along with the rationale for the decision.
5. Implement the change, if approved. Test and verify that the change had the desired effect and had no unintended effects on other systems or on its own security, known as residual risk.
6. Update the ticketing system and configuration backup solutions to reflect changes made. This is the most critical step, as a ticketing system that contains outdated information is arguably worse than not having one at all.

Chapter 9 covers all of these topics, ranging from vulnerability management to recovering from mitigation of vulnerabilities. Hopefully, this introduction to this topic demonstrates that a SOC vulnerability management program is much more elaborate than just acquiring a vulnerability scanner!

Penetration Testing

Penetration testing is a more focused, in-depth, and specialized subset of a vulnerability management program. Vulnerability scanning gives you a list of potential vulnerabilities but does not actually attempt to exploit those vulnerabilities. A penetration test performs the same exploit an attacker could use against a vulnerability to validate that the vulnerability is real. Because exploitation is being used, penetration testing has a much higher risk of damaging systems than does vulnerability scanning. It is essential to be authorized by the asset owner before performing any form of penetration test, known in the industry as a "get out of jail free card" because if something goes wrong, you are not liable based on what was agreed upon prior to performing the penetration test.

NIST 800-115

According to NIST 800-115, *Technical Guide to Information Security Testing and Assessment*, a penetration test is made up of four phases, as shown in Figure 3-10.

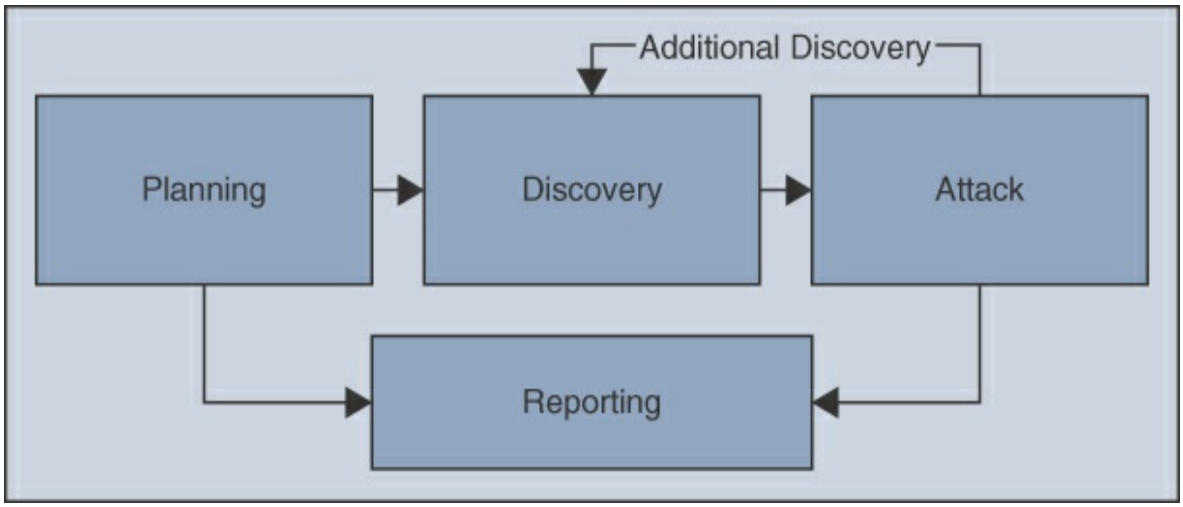

FIGURE 3-10 NIST 800-115 Penetration Testing Phases

The first phase of the NIST 800-115 penetration testing playbook is to plan how the penetration test will be delivered. This phase can include

determining what is considered within or not within scope of the test, identifying what is considered the success criteria, obtaining authorization, and conducting other activities prior to the actual penetration testing steps. If you are using an outside party, these planning steps would be part of your statement of work, which defines how the outside party will be instructed to perform the penetration test. If you are using internal resources, this would be your penetration test plan.

The following is a list of planning or statement of work penetration questions that you need to answer:

- **Target systems:** Which systems are within scope?
- **Timeframe:** How long will the test last?
- **Evaluation methods:** What tactics are within scope (rules of engagement)?
- **Tools and software:** What tools will be used?
- **Notified parties:** Will parties be notified about the testing or kept in the dark?
- **Initial access level:** Where does the penetration start? (You can be outside the network or assume the attacker has gained access to one part of the network.)
- **Authorization:** What authorization rights are given to the attacker?
- **Risks/critical operation areas:** What are the risks of testing, and which areas of operation must be out of scope to avoid negatively impacting the business?
- **Target space:** Which network space is part of the scope? Is it a specific subnet?
- **Define flag:** When does the attack stop, meaning when can the penetration tester consider they have won? Essentially, what is considered success criteria?
- **Deliverable:** What output do you want from the penetration test?
- **Expected remediation:** Do you want any remediation to occur if something breaks or if vulnerabilities are found?

- **Assumptions:** What can be assumed prior to the test? The answer to this means how much information is provided about the target prior to testing, also known as either white, black or gray box testing.

Discovery is the second phase of the NIST penetration testing playbook. This phase focuses on reconnaissance, which includes learning about the targeted environment and associated targets. The discovery phase tends to require the most time and effort, as the more data you know about a target, the more likely you will identify the best point of attack. I recommend to not have your penetration testers start with zero knowledge about the target, commonly referred to as a *black-box test*. The reason I recommend this is that certain aspects of your network and hosts are already publicly available. You can either pay a large amount of money to have somebody tell you what you know (assuming you are paying for a penetration testing service) or just give that information to the penetration tester and reduce the time spent on external reconnaissance research. This approach to penetration testing is commonly called *gray-box testing*, meaning you provide some information about the environment to the penetration tester, but only what you assume the attacker would easily find through open sources. This reduces the time spent on reconnaissance.

Phase three of the NIST penetration testing playbook is to attack, to perform the actual penetration test. Better-quality data that is obtained from the discovery phase will lead to much more effective attack results. If a critical vulnerability is identified during the discovery phase, the attack phase could occur within seconds, leading to a quick capture of the flag. It is common, however, for an effective attack to lead to a breach of the perimeter defenses, which means the penetration tester must step back to the discovery phase of the internal network before another attack can be launched. This is represented in Figure 3-10 as additional discovery, which can occur multiple times as the penetration tester navigates the target environment.

Reporting is the final phase (a repeating phase) of the NIST penetration testing playbook. The whole point of a penetration test is to learn about real vulnerabilities within the target environment. Leadership will want to know where the vulnerabilities are, how they were exploited, which systems are impacted, if any other systems were impacted during the test, and the likelihood that an exploit can occur or has happened already. A poor-quality penetration report will lead to an overall poor penetration experience

regardless of what was found, which is why reporting is the most important of all the phases.

According to NIST 800-115, the reporting phase occurs simultaneously with the other three phases of the penetration test. This allows for documentation to include not only what was found but also how it was found. It is extremely important to include such details as it represents evidence backing up the outcome from the penetration test. Many things can change between when a penetration occurs and when a report is delivered, including improvements in security. Documenting the process can reveal vulnerabilities that existed prior to patches or changes that were put in place post penetration test. Also, system owners that were found to have their systems compromised may attempt to argue against the results, which is when the penetration tester will want to prove his or her case by showing documentation of the penetration testing process that was used.

Note

I recommend as part of the planning phase to make sure to identify all parties that will be reading the penetration test report. If it includes nontechnical parties, make sure the core of the report, often referred to as the executive summary, does not use technical language. You can include technical details, but I recommend providing technical details after an executive summary so that parties who are interested in the technical details can obtain those while everybody, regardless of their technical level, will be able to benefit from the executive summary of the report.

Chapter 6 and Chapter 9 provide deeper looks at these penetration testing concepts. Next, I'll provide a short overview of a few penetration testing tools to help showcase how a penetration test can be executed.

Penetration Testing Tools

One of my favorite methods for SOCs to operationalize attacking

vulnerabilities as a penetration tester is using the MITRE ATT&CK framework, previously covered in Chapter 1. According to MITRE's website (https://attack.mitre.org/), "MITRE ATT&CK is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations." MITRE ATT&CK is a universal language a threat hunter can use to discuss attack techniques and tactics used by real-world adversaries. MITRE ATT&CK is composed of several matrices. Figure 3-11 shows a sample of the macOS Matrix of the ATT&CK Matrix for Enterprise (Windows and Linux matrices also are available). One of the reasons SOCs implement MITRE ATT&CK is that they can test their internal processes to determine if they can detect common attacks. ATT&CK is a great guideline you can use regardless of whether you are a newbie or an experienced penetration tester.

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Exfiltration	Command And Control
10 items	33 items	58 items	28 items	63 items	19 items	20 items	17 items	13 items	9 items	21 items
Drive-by Compromise	AppleScript	.bash_profile and .bashrc	Access Token Manipulation	Access Token Manipulation	Account Manipulation	Account Discovery	AppleScript	Audio Capture	Automated Exfiltration	Commonly Used Port
Exploit Public-Facing Application	CMSTP	Accessibility Features	Accessibility Features	Binary Padding	Bash History	Application Window Discovery	Application Deployment Software	Automated Collection	Data Compressed	Communication Through Removable Media
Hardware Additions	Command-Line Interface	Account Manipulation	AppCert DLLs	BITS Jobs	Brute Force	Browser Bookmark Discovery	Distributed Component Object Model	Clipboard Data	Data Encrypted	Connection Proxy
Replication Through Removable Media	Compiled HTML File	AppCert DLLs	AppInit DLLs	Bypass User Account Control	Credential Dumping	File and Directory Discovery	Exploitation of Remote Services	Data from Information Repositories	Data Transfer Size Limits	Custom Command and Control Protocol
Spearphishing Attachment	Control Panel Items	AppInit DLLs	Application Shimming	Clear Command History	Credentials in Files	Network Service Scanning	Logon Scripts	Data from Local System	Exfiltration Over Alternative Protocol	Custom Cryptographic Protocol
Spearphishing Link	Dynamic Data Exchange	Application Shimming	Bypass User Account Control	CMSTP	Credentials in Registry	Network Share Discovery	Pass the Hash	Data from Network Shared Drive	Exfiltration Over Command and Control Channel	Data Encoding
Spearphishing via Service	Execution through API	Authentication Package	DLL Search Order Hijacking	Code Signing	Exploitation for Credential Access	Network Sniffing	Pass the Ticket	Data from Removable Media	Exfiltration Over Other Network Medium	Data Obfuscation
Supply Chain Compromise	Execution through Module Load	BITS Jobs	Dylib Hijacking	Compiled HTML File	Forced Authentication	Password Policy Discovery	Remote Desktop Protocol	Data Staged	Exfiltration Over Physical Medium	Domain Fronting
Trusted Relationship	Exploitation for Client Execution	Bootkit	Exploitation for Privilege Escalation	Component Firmware	Hooking	Peripheral Device Discovery	Remote File Copy	Email Collection	Scheduled Transfer	Fallback Channels
Valid Accounts	Graphical User Interface	Browser Extensions	Extra Window Memory Injection	Component Object Model Hijacking	Input Capture	Permission Groups Discovery	Remote Services	Input Capture		Multi-hop Proxy
	InstallUtil	Change Default File Association	File System Permissions Weakness	Control Panel Items	Input Prompt	Process Discovery	Replication Through Removable Media	Man in the Browser		Multi-Stage Channels
	Launchctl	Component Firmware	Hooking	DCShadow	Kerberoasting	Query Registry	Shared Webroot	Screen Capture		Multiband Communication
	Local Job Scheduling	Component Object Model Hijacking	Image File Execution Options Injection	Deobfuscate/Decode Files or Information	Keychain	Remote System Discovery	SSH Hijacking	Video Capture		Multilayer Encryption
	LSASS Driver	Create Account	Launch Daemon	Disabling Security Tools	LLMNR/NBT-NS Poisoning	Security Software Discovery	Taint Shared Content			Port Knocking
	Mshta	DLL Search Order Hijacking	New Service	DLL Search Order Hijacking	Network Sniffing	System Information Discovery	Third-party Software			Remote Access Tools
	PowerShell	Dylib Hijacking	Path Interception	DLL Side-Loading	Password Filter DLL	System Network Configuration Discovery	Windows Admin Shares			Remote File Copy
	Regsvcs/Regasm	External Remote Services	Plist Modification	Exploitation for Defense Evasion	Private Keys		Windows Remote Management			Standard Application Layer Protocol
	Regsvr32	File System Permissions Weakness	Port Monitors	Extra Window Memory Injection	Securityd Memory					Standard Cryptographic Protocol
	Rundll32	Hidden Files and Directories	Process Injection	File Deletion	Two-Factor Authentication Interception					Standard Non-Application Layer Protocol
	Scheduled Task	Hooking		File Permissions Modification						
	Scripting			File System Logical Offsets						
	Service Execution									

FIGURE 3-11 MITRE ATT&CK Framework

A standard method used to test how effective a SOC is against attacks, and also to test the SOC's capabilities of operationalizing MITRE ATT&CK, is to use the Atomic Red Team tests. Figure 3-12 shows the Use Cases page of the Atomic Red Team website (https://atomicredteam.io/).

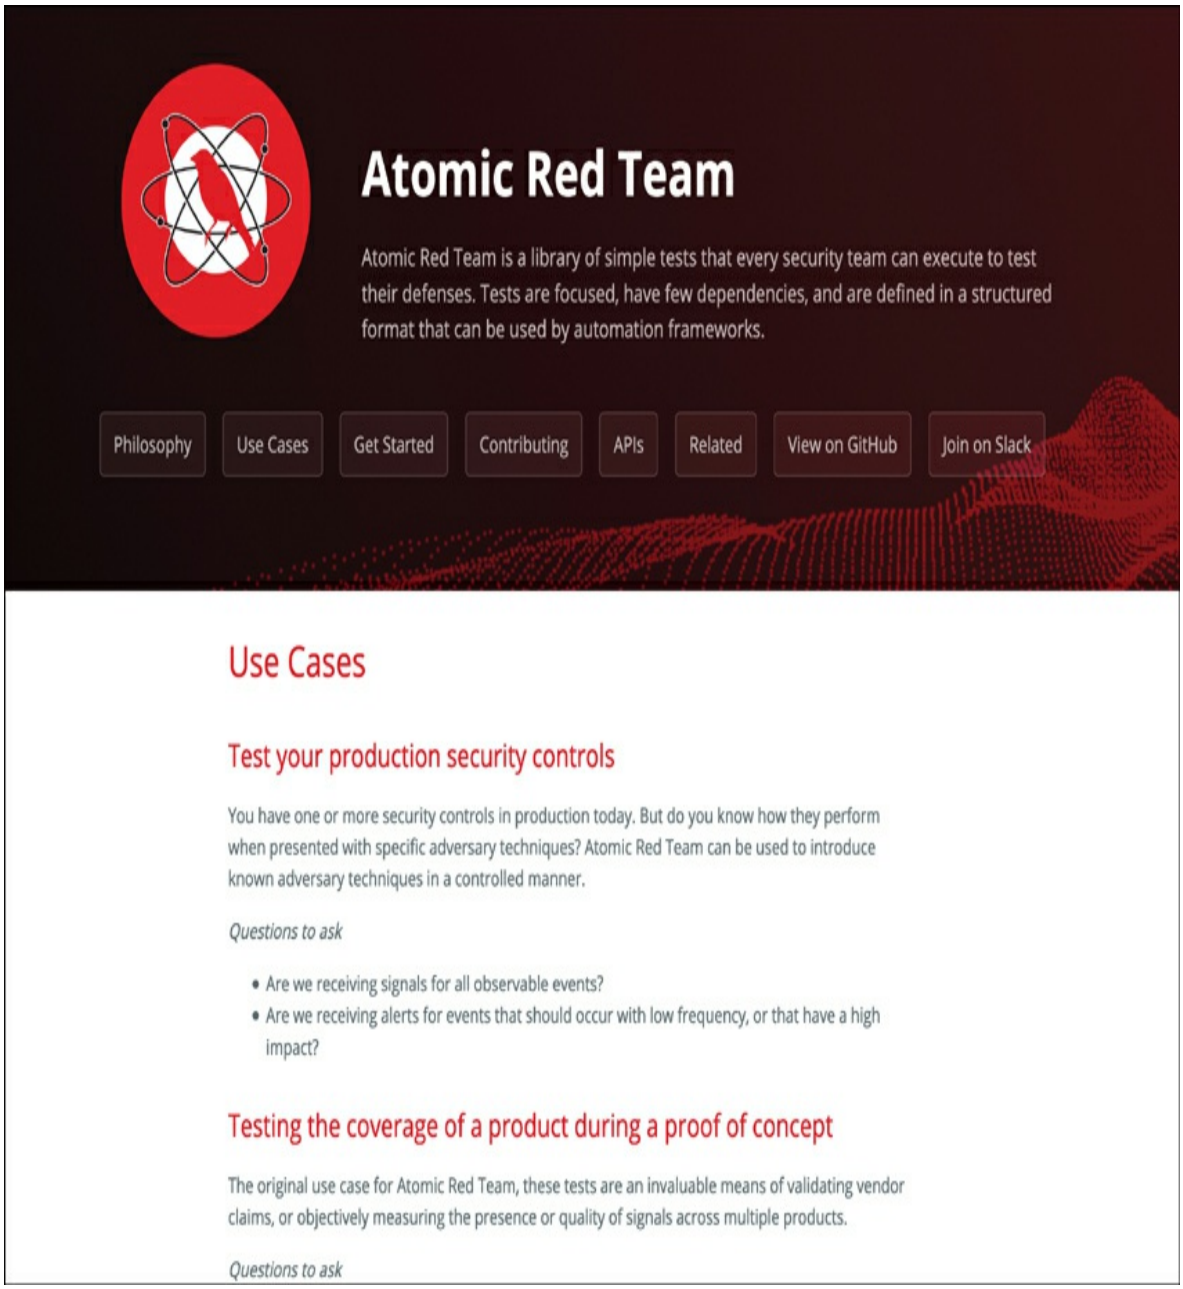

FIGURE 3-12 Atomic Red Ream Website

The Atomic Red Team is a collection of open-source techniques used by

attackers that individually test against attack phases from MITRE ATT&CK. Tests are divided by operating systems and support Windows, macOS, and Linux operating systems. The tests are available on Red Canary's Atomic Red Team GitHub page located at https://github.com/redcanaryco/atomic-red-team. SOC teams need to be able to test their technical controls and determine not only if they can detect an attack but also if they can prevent it from occurring.

When an attack does occur, SOC teams need to ensure they can collect forensic evidence when needed and restore systems to a working state. The Atomic Red Team tests provide a methodology to accomplish these goals in an open-source, publicly consumable format. Figure 3-13 displays a section of the open-source tests that are available for Windows. The tests and techniques are updated through community-driven efforts. Many SOCs that use the Atomic Red Team tests update them with new techniques or make corrections to the tests when they find flaws after using Atomic Red Team tactics within their own environment.

Windows Atomic Tests by ATT&CK Tactic & Technique

initial-access	execution	persistence	privilege-escalation	defense-evasion	credential-access
Drive-by Compromise CONTRIBUTE A TEST	CMSTP	Accessibility Features	Access Token Manipulation	Access Token Manipulation	Account Manipulation
Exploit Public-Facing Application CONTRIBUTE A TEST	Command-Line Interface	Account Manipulation	Accessibility Features	BITS Jobs	Brute Force
External Remote Services CONTRIBUTE A TEST	Compiled HTML File	AppCert DLLs CONTRIBUTE A TEST	AppCert DLLs CONTRIBUTE A TEST	Binary Padding	Credential Dumping
Hardware Additions CONTRIBUTE A TEST	Component Object Model and Distributed COM CONTRIBUTE A TEST	AppInit DLLs	AppInit DLLs	Bypass User Account Control	Credentials from Web Browsers CONTRIBUTE A TEST
Replication Through Removable Media CONTRIBUTE A TEST	Control Panel Items	Application Shimming	Application Shimming	CMSTP	Credentials in Files
Spearphishing Attachment	Dynamic Data Exchange	Authentication Package CONTRIBUTE A TEST	Bypass User Account Control	Code Signing CONTRIBUTE A TEST	Credentials in Registry

FIGURE 3-13 Atomic Red Team Example for Windows

There are prevention and detection tools that are incorporating software products based on MITRE ATT&CK and Atomic Red Team tests. One of my favorite commercial tools, Managed Detection and Response (MDR), comes from Red Canary, creator of the Atomic Red Team. Its products are designed to integrate with endpoint detection and response (EDR) solutions, such as

ones from Carbon Black. MDR maps attack, log, and alert data it receives from EDR solutions to MITRE ATT&CK. Additionally, MDR allows SOC operators to audit endpoints against specific and individual Atomic Red Team tests to determine how effective their solution is at stopping an attack.

> **Note**
>
> Red Canary has an open-source version of its software called Surveyor, available at https://redcanary.com/surveyor/.

Penetration Testing with Kali Linux and Metasploit

For those favoring open-source tools, one of the most popular open-source penetration platforms is Kali Linux, a Debian-based distribution aimed at advanced penetration testing and security auditing. Kali Linux has hundreds of tools covering various topics including exploitation, reconnaissance, forensics, vulnerability evaluation, password cracking, fuzzing, and many more. Figure 3-14 shows the main Kali Linux dashboard with all the tool categories displayed and the Exploitation Tool category expanded.

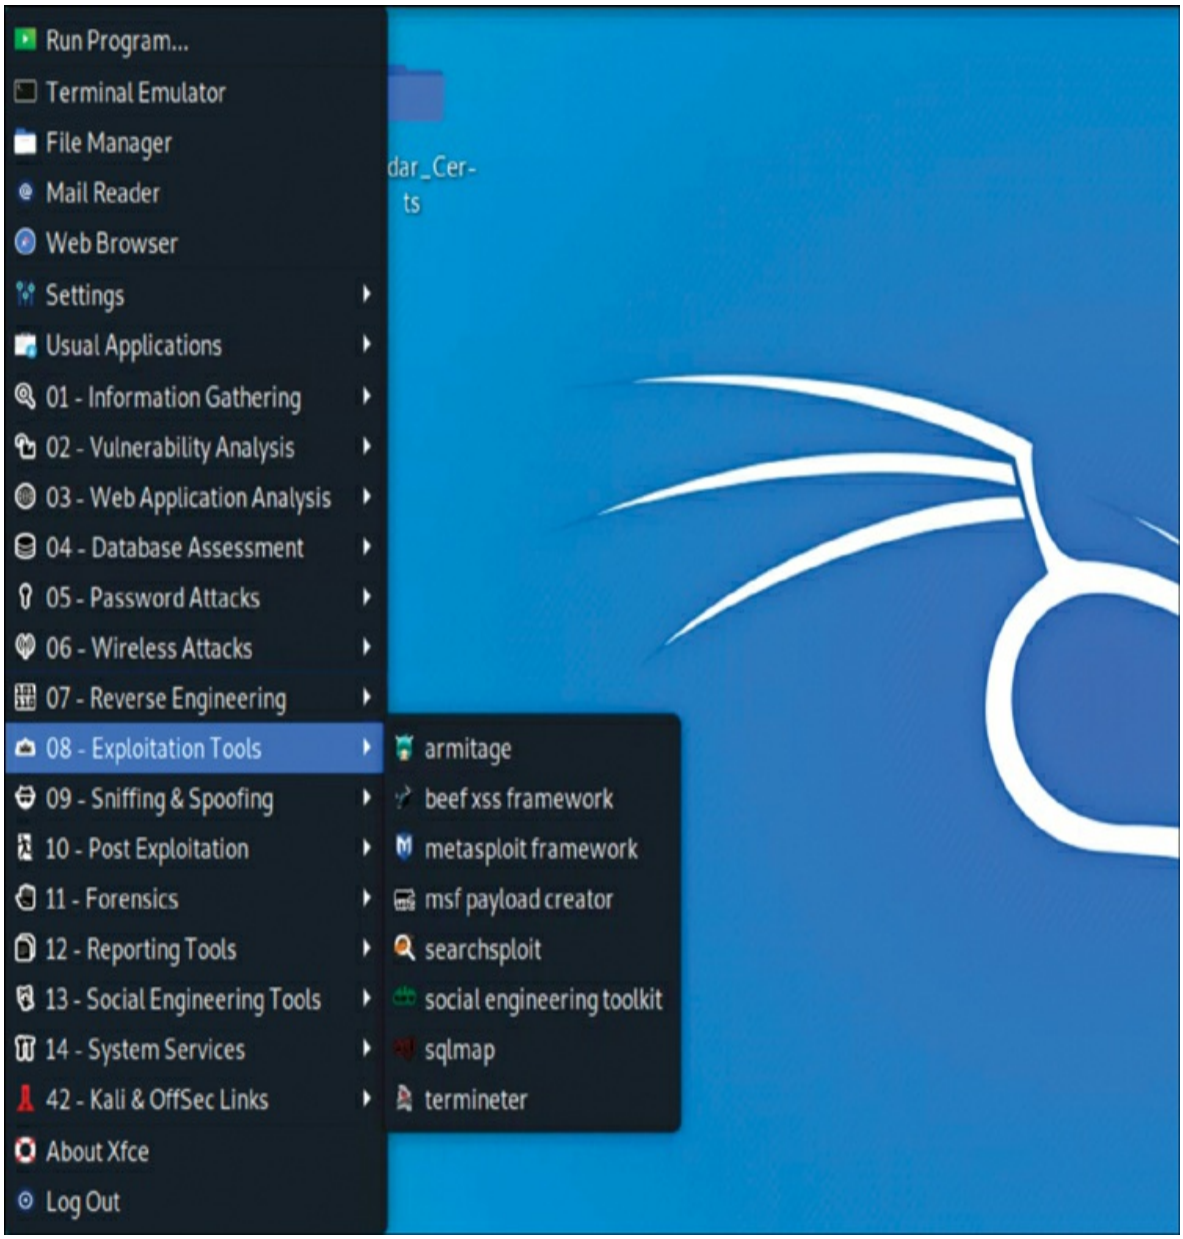

FIGURE 3-14 Kali Linux Tool Categories

One key penetration testing tool available both in Kali Linux and as an independent installation is the Metasploit Framework. Metasploit is a framework used by both cybercriminals and ethical hackers to probe systematic vulnerabilities on networks and servers. A penetration tester can use a vulnerability scanner to identify potential weaknesses within a system and later match the identified weakness to an exploit using Metasploit. For example, if you found a vulnerability within a version of Adobe, you could search Metasploit to see if there is a matching weaponized exploit against the

vulnerability of interest. Figure 3-15 shows an example of searching Metasploit for Adobe-related vulnerabilities.

FIGURE 3-15 Searching Metasploit for Adobe Vulnerabilities

Penetration testing can be a topic for its own book and there are many certifications you can pursue and obtain to gain these skills. Chapter 6, Chapter 9, and other parts of the book will touch on penetration concepts; however, my focus for vulnerability management will be on best practices for identifying and responding to vulnerabilities rather than exploitation techniques.

SOC Service 3: Compliance

Compliance is a service found in most organizations and is considered a

fundamental SOC service. Compliance at its core means to meet some set goals. Those goals can come from an organization's leadership, such as a corporate-mandated policy, a legal obligation, or just a general recommendation converted into a requirement. Not being compliant with a government-required policy leads to fines and potential consequences such as loss of investor confidence. Not being compliant with industry recommendations leads to gaps in security practices. Gaps become failures in security. Compliance is why many organizations include industry-recommended guidelines and government-required policies within their mandated corporate policy. Enforcing good security practices leads to a reduction in risk.

With all the reasons for compliance being covered, let me be clear: compliance does not equal security. Compliance is either a legal standard, a regulatory requirement, or a corporate-mandated policy. Although compliance is a fundamental service of a SOC, maintaining compliance with mandatory requirements should be considered the very minimum level of securing an organization, because compliance requirements are generic to an industry or business section and not specific to an individual organization. In other words, security compliance ensures that an organization is implementing the minimal set of capabilities that achieve defined security requirements; it doesn't ensure that you've looked at all security threats to your entire organization and considered all possible areas of risk. Compliance can help with improving security; however, compliance by itself should not be your only method of validating the quality of your security capabilities and response.

Including compliance within the SOC allows adding monitoring and reporting of compliance requirements within the existing SOC practice. For example, a compliance requirement for securing systems with PCI DSS data can be evaluated with the same tools that the SOC uses for its vulnerability management service. Sharing goals between different SOC services allows for combining budgets and reducing effort needed for each service.

Meeting Compliance with Audits

Conducting audits is a service that a SOC can provide for the organization to

validate that the organization meets compliance requirements. Organizations often ask me whether the internal SOC should be responsible for providing audits or an outside party should be hired to conduct audits. Either approach is fine as long as the auditor does not have vested interest in the outcome of the audit or a potential bias in the findings. Some compliance audits are required to be performed by an outside party; however, I recommend including internal efforts to ensure the organization meets requirements for a mandatory audit to reduce the risk of falling out of compliance prior to or during the audit. There are also products that can be used to assist with auditing various types of compliance.

Audits are not security services but can be built into security assessments as a report tied to specific goals such as meeting compliance standards. To be clear, properly delivered audit services require more effort than taking inventory of devices and reviewing configurations against suggested settings. Audits also include understanding business processes and procedures as well as the value provided by the systems. The goal of understanding this level of detail is ensuring the organization is getting the best return on investment (ROI) on the technology it has invested in. Auditing value from a system is not security related, but it is common to include checking for required configuration settings for both the purpose of meeting security requirements and the purpose of validating that an optimal version of the system is being used. Guidelines such as ISO/IEC 27000:2018, which is the standard for information security management systems (ISMS), can be used to help understand what an optimal version of a common IT system is. I recommended using the Plan-Do-Check-Act (PDCA) cycle, which is a lifecycle of implementing security that is recommended in ISO/IEC 27000:2018 and other guidelines.

ISO/IEC 27000:2018 describes the PDCA cycle as follows:

- **Plan (establishing the ISMS):** Establish the policy, the ISMS objectives, processes, and procedures related to risk management and the improvement of information security to provide results in line with the global policies and objectives of the organization.
- **Do (implementing and workings of the ISMS):** Implement and exploit the ISMS policy, controls, processes, and procedures.
- **Check (monitoring and review of the ISMS):** Assess and, if applicable,

measure the performances of the processes against the policy, objectives, and practical experience and report results to management for review.

- **Act (update and improvement of the ISMS):** Undertake corrective and preventive actions, on the basis of the results of the ISMS internal audit and management review, or other relevant information to continually improve the said system.

The PDCA lifecycle advises organizations to evaluate how they will implement technology and what should be the desired outcome (Do phase). The Check phase is put in place to ensure the technology of interest is meeting the goals of the service scope. This approach does not measure if the technology of interest is successful until it is put into operation and its value can be measured. Finally, the Act phase advises organizations to make adjustments if the goal and desired outcome for the technology of interest are not being met. Organizations should reconfigure or tune their technology products in this phase. Only if this phase fails and organizations cannot tune their products for the desired outcome, or if a new outcome is needed, should organizations consider new or different products. My recommendation is to ensure audits include a performance review similar to the PDCA approach to measuring success, along with security goals such as meeting required compliance. Chapter 6 covers compliance topics in much more detail.

SOC Service 4: Incident Management

SOCs need to assume that prevention capabilities will fail at some point and the organization will be compromised by a malicious party or experience an unwanted event. When systems are compromised, the SOC will need to be able to quickly identify what has been compromised and what type of modifications have occurred. Different industry breach reports show that the average time an organization takes to detect a compromise can range anywhere from weeks to months, which provides the attacker with tons of time to learn the environment and perform malicious actions.

Figure 3-16 is a summary of the findings from the Verizon 2020 Data Breach Investigations Report. Notice that the majority of breaches are external actors associated with organized crime and that they used hacking to accomplish the

breach.

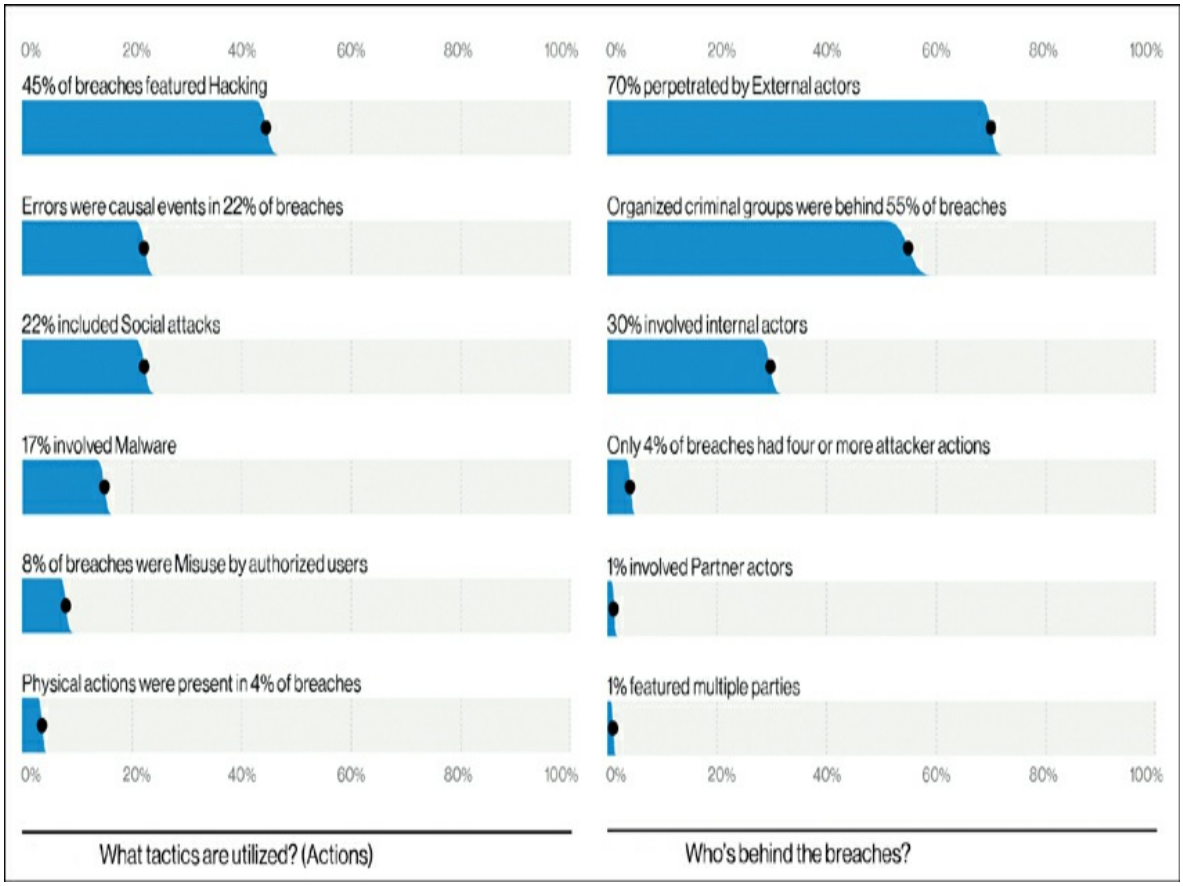

FIGURE 3-16 Verizon 2020 Data Breach Investigations Report Summary of Findings

There are many industry guidelines that explain the different steps involved in delivering a mature incident response service. In general, steps will include some form of preparation, meaning setting up the incident response program, steps to determine when an incident occurs, steps involved with responding to the incident, and post incident actions. One of the most popular guideline options for incident response is NIST Special Publication 800-61 Revision 2.

NIST Special Publication 800-61 Revision 2

When a SOC discovers that a breach has occurred, it immediately kicks off

the incident response process or calls upon the team responsible for the incident response service. An incident response service covers how a team prepares for, detects, responds to, and recovers from an incident. In each one of these areas, the SOC must have a well-documented understanding of the steps that need to occur to be successful. NIST Special Publication 800-61 Revision 2, *Computer Security Incident Handling Guide*, describes incident response plans as having a lifecycle with the following phases (see Figure 3-17):

- Preparation
- Detection and analysis
- Containment, eradication, and recovery
- Post-incident activity

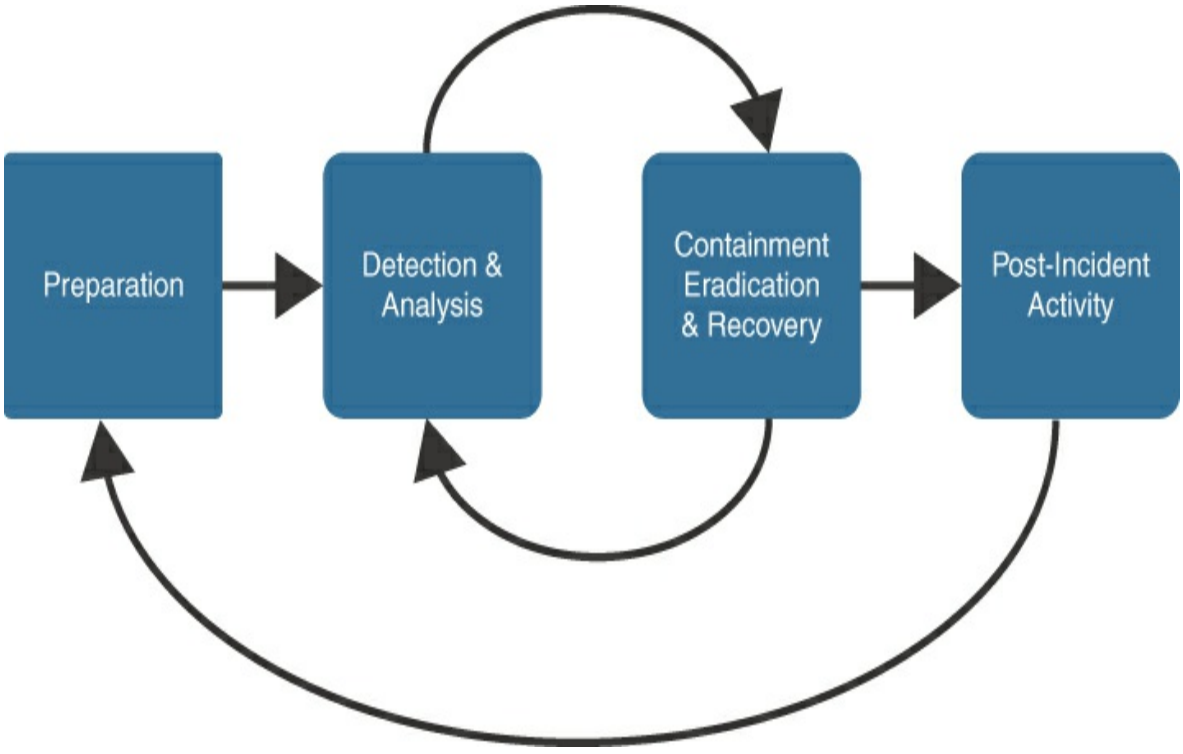

FIGURE 3-17 NIST SP 800-61 Rev. 2 Incident Response Process

Step 1: Preparation

Preparation is considered the most important phase of the plan. Not only does preparation cover how an organization will respond when an incident occurs, it also includes how organizations can prevent incidents from occurring by securing networks and applications. When organizations are preparing to handle events, they need basic policies that identify what actions will occur and resources that will be required such as contact and call lists; defined tracking systems; messaging protocols; access to forensic software; and physical access to war rooms and buildings. Incident response teams need access to multiple IT teams as well as procedures on how to engage legal, upper management, and public relation teams.

Note

A key part of preparation is preventing attacks from occurring. This is accomplished using risk assessments, user awareness training, vulnerability testing, host security, and other methods.

Step 2: Detection & Analysis

The *detection and analysis* phase prepares an organization to be able to detect a variety of different cyberattacks from different threat surfaces. Detection occurs through data collection from network devices, host machines, security tools, and people monitoring such tools. Failing at detection means threats can operate unnoticed within the network, systems, or services.

Analysis serves a few purposes. First, the SOC must use analysis to validate whether any events that are detected are a real threat or something else. The SOC can make four different determinations regarding the analysis of an event:

- **False positive:** When the detect tools incorrectly identify a security event the SOC needs to be concerned with
- **True positive:** When the detect tools correctly identify a security event the SOC needs to be concerned with
- **False negative:** When the security tools incorrectly don’t alert the SOC

about an event the SOC needs to be concerned with

- **True negative:** When the security tools correctly don't alert to an event the SOC does not need to be concerned with

The mature incident response program will have a strong analysis capability allowing for the reduction of noise made up of false positives. A mature incident response program will also have a strong detection capability allowing for a reduced amount of false negatives.

Step 3: Containment Eradication & Recovery

Containment focuses on stopping an attack from spreading. NIST SP 800-61 Rev. 2 often references the containment phase as "decision making." An example is how during a datacenter breach, the SOC will need to decide whether to shut down a server to prevent an attacker from gaining access to data or to leave the server running to avoid tipping off the attacker that the SOC is aware of the attacker's presence, enabling the SOC to study the attacker's methods and collect evidence. This raises a question regarding the need to preserve evidence of the attack versus minimizing the risk of attackers potentially damaging systems or stealing data. In other words, if an organization lets an attacker continue attacking a system, the organization may learn more about the attacker, their methods, and motivation. However, they do this at the risk of the attacker damaging systems, gaining a stronger foothold into the organization that may go undetected, and stealing or damaging data. The answer to this type of question will need to be addressed on a case by case situation. I will help you understand how to make the best decision regarding when to allow attacks to occur for research purposes in Chapter 8, "Threat Hunting and Incident Response."

Eradication focuses on responding to all impacted parties. Eradication can work only if proper containment has been performed. If a threat isn't contained, it can't be eradicated because it's possible it will always be spreading to new systems during the eradication of the systems identified as impacted by the threat. Eradication can be performed in many ways and will depend on the type of threat being handled. In certain situations, an advanced threat will be extremely difficult to remove, in which case the best action is a complete system wipe and rebuild. Other times, eradication can just remove

the threat without negatively impacting the infected system. Eradication can also include containment steps, such as shutting down services while the eradication process occurs. This is done to ensure containment of threats that can change behavior when the threat detects it is being eradicated.

Recovery can't occur until eradication has been performed, because recovery is based on the concept of returning things back to operation as they were prior to the incident. Often, containment and eradication include limiting services to prevent the spread of the threat. During the recovery phase, security measures implemented during the previous steps are removed and systems are monitored to ensure the threat has been completely eradicated. If any sign of the threat is identified, the entire process, starting with detection and analysis, is repeated until the recovery step shows no sign of the identified threat.

Step 4: Post-Incident Activity

Post-incident activity includes the lessons-learned session. What happened during the incident? What controls worked, and why? What controls failed and why? What changes need to be made to enable the entire organization to be more secure and more responsive? The answers to these questions are collected for the first phase of the incident response program to improve the preparation of the organization against future threats. Post-incident activity allows the organization to learn from experience and incorporate that knowledge in better protecting the organization.

Figure 3-18 shows a flowchart representing how a SOC can navigate through an incident. Figure 3-18 is a modified version of the NIST 800-61 Rev. 2 standard that includes guidance on how to proceed within each phase and what processes will come into play. You should expect that the approach your organization takes to implementing an incident response program will likely be slightly different from published standards.

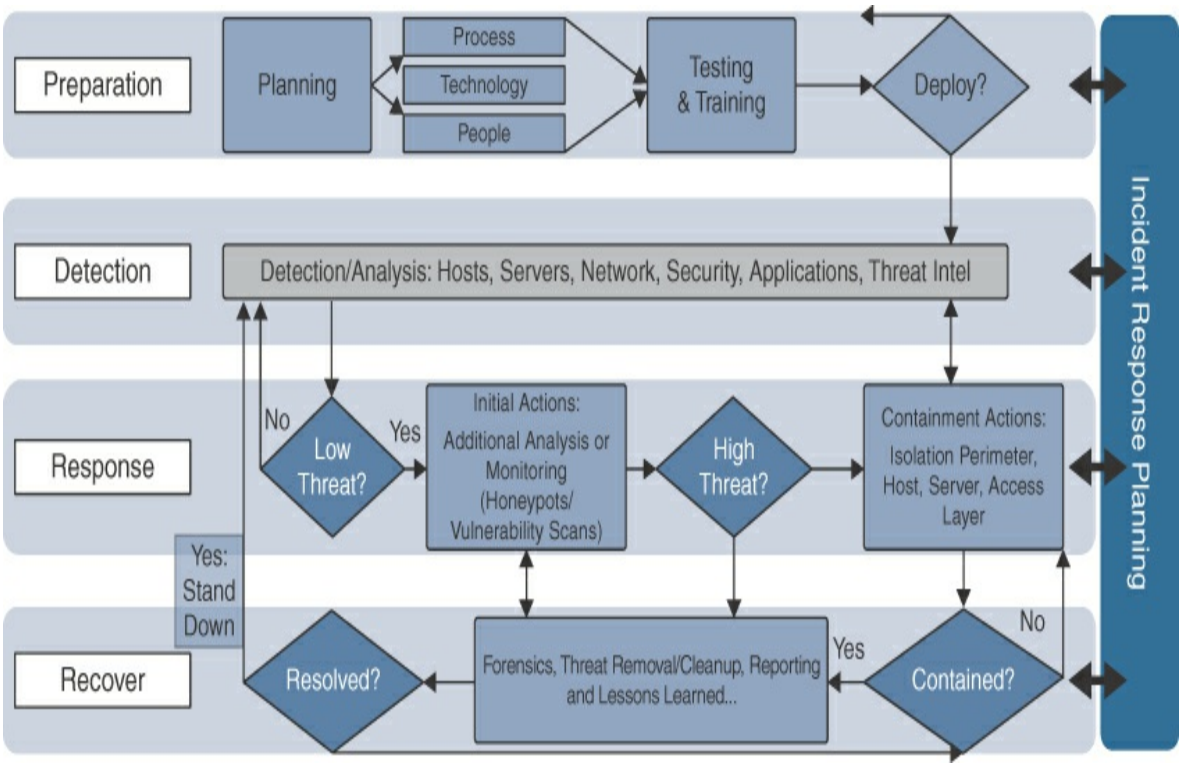

FIGURE 3-18 Incident Response Process

Incident Response Planning

A successful incident response plan needs to be documented, tested, and be repeatable by the security engineers in the SOC. It should allow for documentation, resolution, and reporting of all responses and findings for an incident. In other words, it should have a properly established communication chain of command. Many times, when an incident or a data breach occurs, organizations have technical, business, and political concerns that must be addressed. The SOC, as the incident response command center, will need established guidelines on when and how it communicates. According to NIST SP 800-61 Rev. 2, “It is also vital to build relationships and establish suitable means of communication with other internal groups (e.g., human resources, legal) and with external groups (e.g., other incident response teams, law enforcement).” Where communication can be challenging is when the incident is related to internal employees, for example. Confronting an employee, shutting down their access, or taking action against them might not be an option depending on their role and the situation at hand. In some rare

cases, the organization might actually be opening itself to liability from the employee if the organization takes actions that violate legal rights. This is just one example of why preestablished communication needs to be developed and part of an incident response program.

Another critical component of an incident response program is developing a method to report an incident to the incident response team. Anyone in the organization, or even outside the organization, must have the ability to report an incident. A SOC needs to set up a way to allow reporting that ensures communication occurs securely, because the context of the conversations will include sensitive information such as critical systems being exploited. An example of a secure method of communicating is using Pretty Good Privacy (PGP) keys for external sources as well as established procedures for internal employees on communicating. Avoid using any easily accessible public-facing message boards or websites outside of showcasing case-tracking data.

Incident Impact

One of the top questions that leadership will ask about an incident is its impact to the organization. Impact can be measured by a threat's severity, which is critical to understand so the proper response can be deployed against it. Establishing the severity of a specific incident enables the SOC to assign and dedicate the right resources it needs to respond to the incident appropriately. A low-severity incident might require reviewing documentation and operational procedures. A high-severity incident requires the SOC to investigate systems and applications the threat *could have interacted with*.

Note

I emphasize *could have interacted with* because many regulations require an organization to report the number of records that a threat could have interacted with, which may be a much higher number than the number of records the threat *actually* impacted. An example of an inflated number would be a threat that had two minutes of internal access to a datacenter holding 4 million records. It is very unlikely the

threat interacted with all 4 million records; however, because there is the "potential of access," the organization would have to report 4 million records impacted!

It is not uncommon for a SOC to be forced to reprioritize its resources and pull staff from existing projects when high-severity incidents occur. Determining severity is ultimately about determining damage to an organization. COBIT guidelines from ISACA have great recommendations on how to recognize the likelihood of a threat by looking at the severity of the vulnerability and likelihood of the threat agent to engage and exploit that vulnerability. Figure 3-19 shows an example of a COBIT severity model. This model takes into account not only how easy the vulnerability might be to discover and exploit but also the motivation, skill level, and other factors.

Threat Agent Factors				Vulnerability Factors			
Skill Level	Motive	Opportunity	Size	Ease of Discovery	Ease of Exploit	Awareness	Intrusion Detection
1	7	5	2	6	3	2	9
Likelihood of Threat = 4.375 MEDIUM							

FIGURE 3-19 COBIT Severity Model

Playbooks

Security operations centers use security frameworks not only to test and audit systems against industry best practice, but also to develop procedures to respond to specific threats. The Incident Response Consortium is an excellent resource SOCs can use to create playbooks on how to respond to specific threats against their organization. SOCs can have playbooks covering how to discover malware outbreaks, how to respond to phishing attacks, how to address elevation of privileges, and how to deal with many other threats.

Figure 3-20 shows a gallery of predefined playbooks that are freely available on the Incident Response Consortium website (https://www.incidentresponse.com/playbooks/). Figure 3-21 shows the download page for the Malware Outbreak playbook.

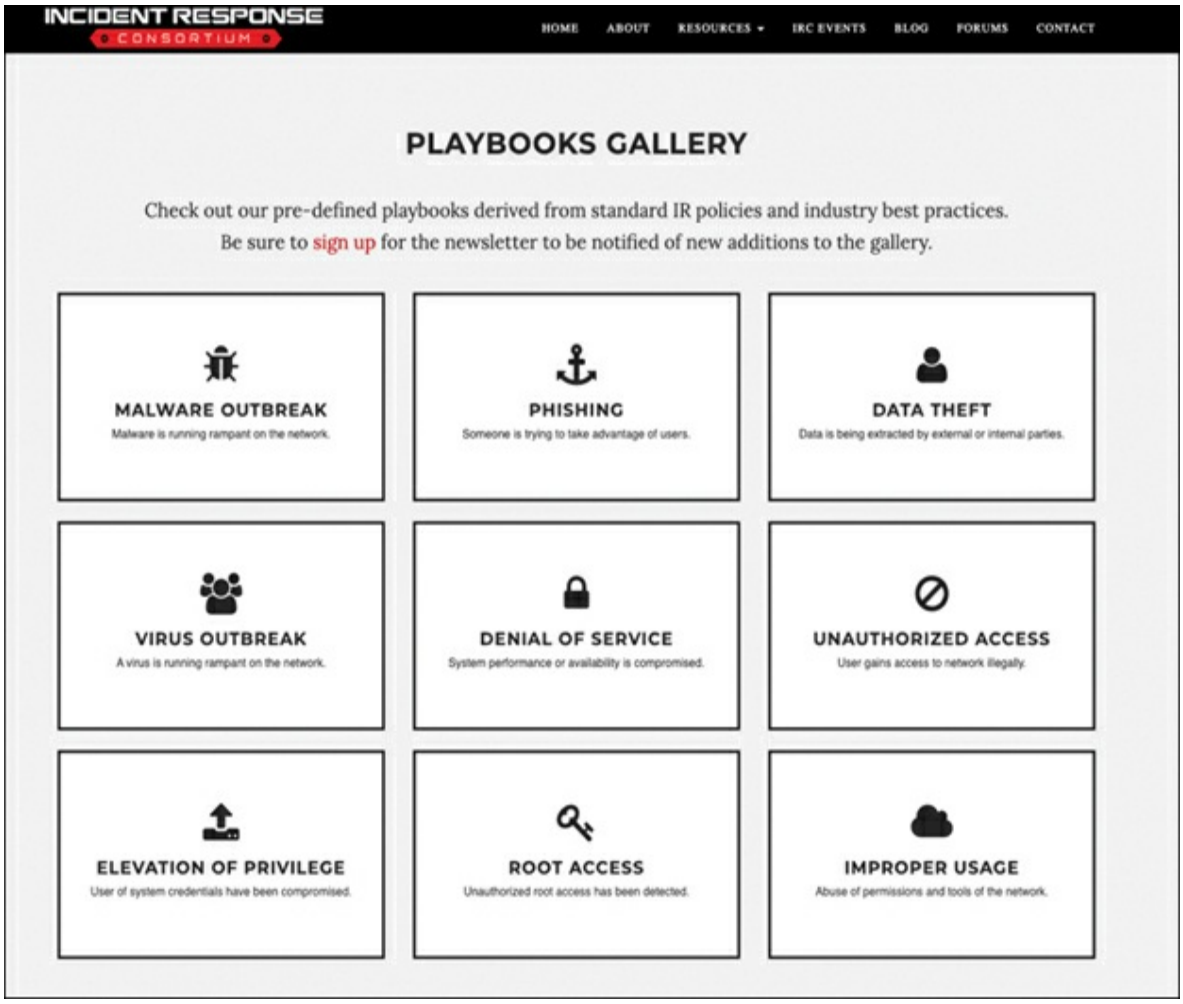

FIGURE 3-20 Incident Response Consortium Playbooks

FIGURE 3-21 Malware Outbreak Playbook

Playbooks need to cover technical and business aspects on how to prepare for, detect, analyze, contain, and recover from an incident. Playbooks also need to include which job roles are responsible for which tasks and define what tasks are performed. Playbooks can contain granular details about the tools used, all of which details are documented in the SOC incident response procedures. Essentially, playbooks are an automated form of a procedure document. Huge growth is occurring in the security orchestration, automation, and response (SOAR) market as SOCs convert their existing procedures into playbooks using tools offered in the SOAR market. Examples of SOAR solutions include RSA NetWitness Orchestrator, Splunk Phantom, and IBM Resilient. These industry trends are also leading to the need for a specific type of programming known as DevOps, which focuses on how to make tools work with other tools.

Chapter 8 covers more incident response concepts, including

recommendations for strategy, tools, and industry guidelines. Chapter 10 dives deep into the world of automation and orchestration, including how to convert SOC procedures into playbooks and later automate steps within the playbook.

SOC Service 5: Analysis

It is common for SOCs to have responsibilities for analyzing data and files. One situation is the need to understand what a file is and determine if it's a risk. For example, a user could see a file called evil.doc, which looks like a Microsoft office document. When the SOC analyzes the file with the right tool, such as TrIDNET, it finds that the file has a hidden extension revealing the actual file is evil.doc.exe. Malware authors use tactics such as this to trick a user into running an executable that leads to compromise of the system. The icon associated with the file can be changed to represent a word file and the .exe extension can be hidden revealing only the .doc extension. To the untrained eye, this file would appear to be a word document when really it is an executable! Figure 3-22 shows how to view hidden file extensions in Windows, how a .doc file can really be a doc.exe, and the main interface of TrIDNET.

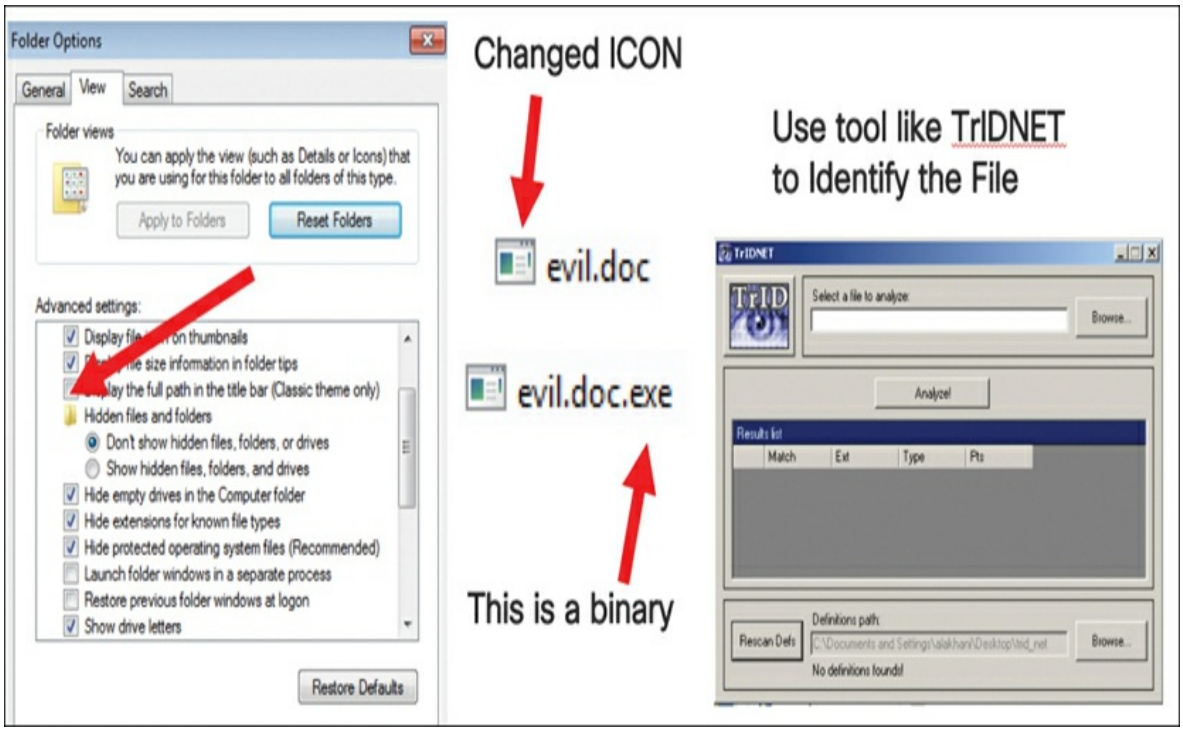

FIGURE 3-22 Diagram About Hidden Extensions

Note

You can find TrIDNET at https://mark0.net/soft-tridnet-e.html.

Static Analysis

There are two different types of file analysis that can be performed to learn if a file is malicious. The first type of file analysis is *static analysis*, analyzing the file without running it. Static analysis can be as simple as scanning a file with an antivirus program to see if it matches a known malicious file within the malware database. This tactic only works if the antivirus has the right signature to match the threat, meaning if the signature isn't matched, then the threat will not be identified. Malware authors are very aware of antivirus technology and use various techniques to change how malware looks so that it goes undetected when evaluated by security tools. One of the most common techniques to hide malware is to use a packer, which encrypts the

guts of the program. When malware is packed, only certain details can be seen within the header of the file. A decompression stub is needed to uncompress the file and see its contents. Figure 3-23 shows a high-level diagram of the file-packing concept.

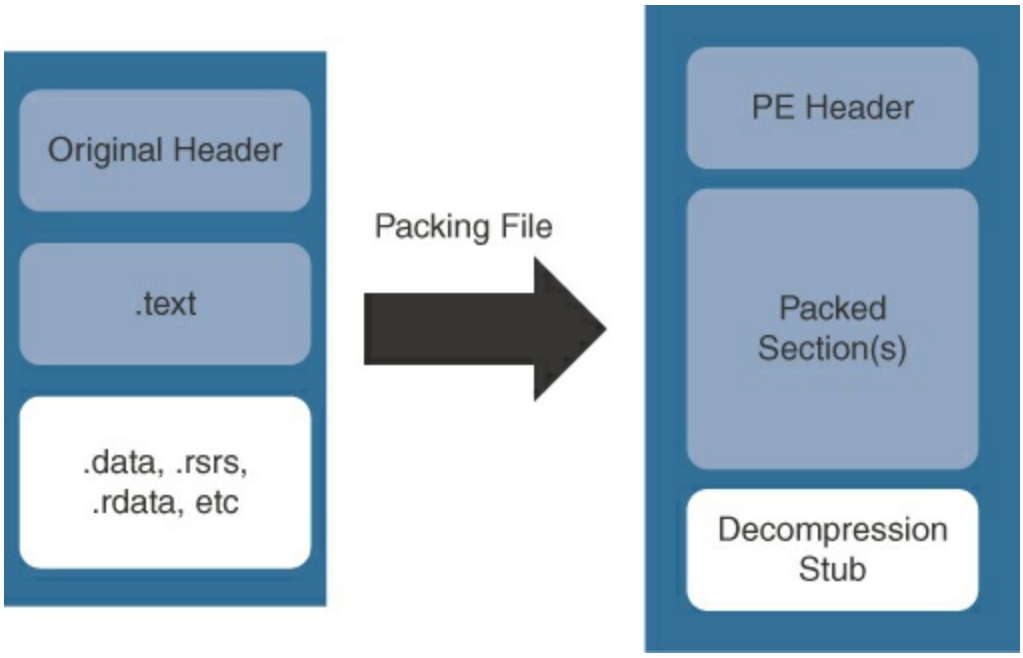

FIGURE 3-23 Packing Files

Tools are available that enable you to read a packed file with the goal of understanding if it's a risk regardless if you are unable to unpack the file and see its core contents. Figure 3-24 shows an example of running an open-source tool called Peframe (https://github.com/guelfoweb/peframe), which allows the analyst to collect details about a packed file to determine whether it is encrypted using packing and other attributes of interest. Indicators of a malicious file that Peframe could uncover include function calls to LoadLIbrary and Get ProcAddress (common within malware), abnormal entry points, very few import functions, memory with WRE permissions, and many more, covered in greater detail in Chapter 8.

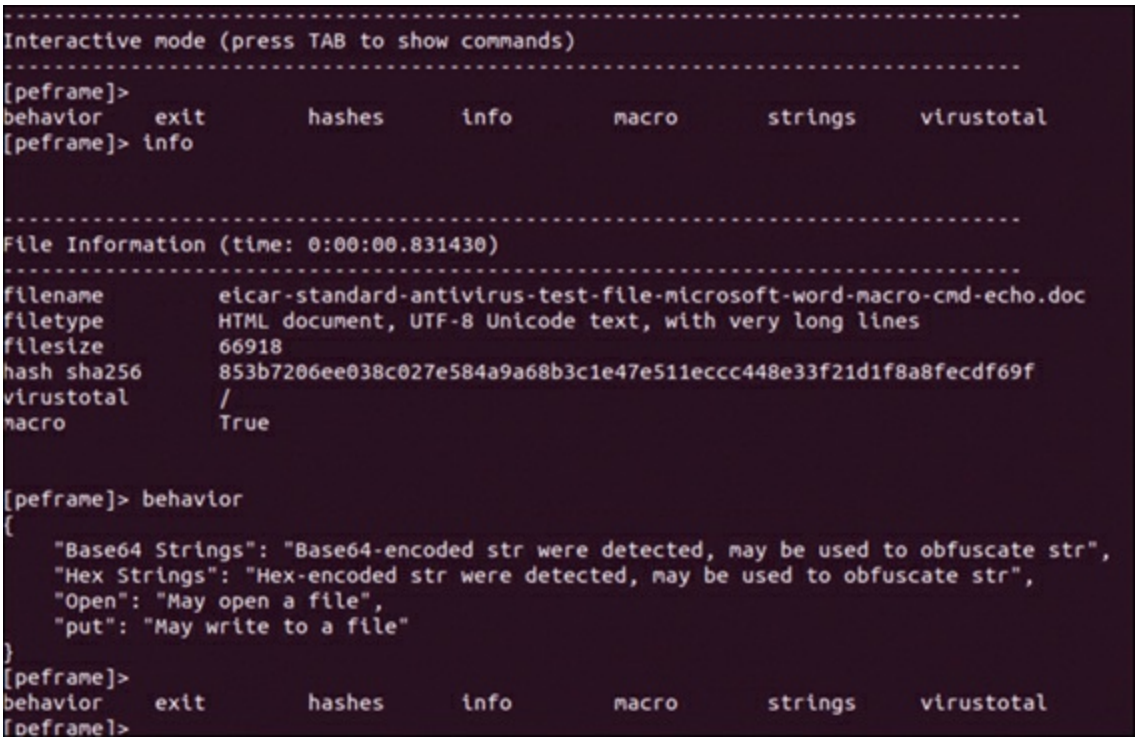

FIGURE 3-24 Peframe Analyzing a Packed File

Note

All packed files are not evil. Anybody that wants to hide the contents of a file can use packing. One common legitimate use for packing is a software developer hiding his or her work so that it can't be copied by unauthorized people.

An even deeper form of static analysis is reverse engineering a file. This concept works by first understanding that a malware author will develop malware using a programming language. You, as the analyst, will not have access to the source code since a program needs to be compiled in order to be run on a computer. By compiling the source code, the program is converted into machine code, which is only readable by computers. To convert the compiled code into something you can analyze, you use a *disassembler* that converts the machine code/compiled code into a lower-level language that can be used to reverse engineer the logic of the program. Figure 3-25 shows an example of how this process works. In this example, the programmer has

created the classic “hello world program” using C, which prints Hello World to the screen when compiled and executed on a computer. That code is compiled into machine code and sent to others to use. At some point, an analyst captures the compiled code and uses a disassembler to convert the compiled code into low-level code.

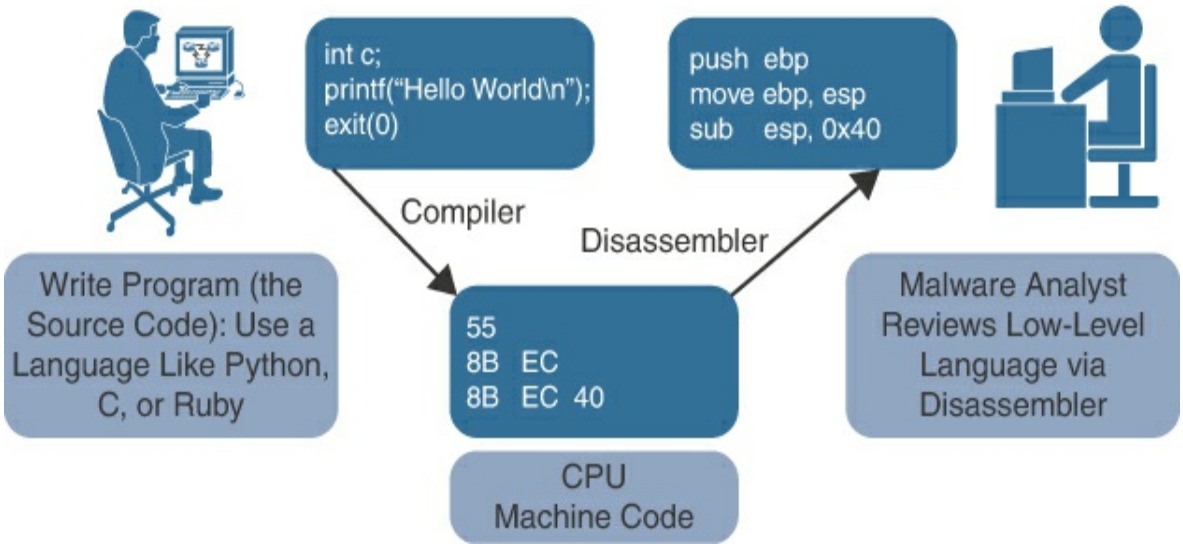

FIGURE 3-25 Overview of Programs and Disassemblers

There are many challenges malware authors invoke to not allow an analyst to disassemble code. If compiled code is packed, the file can’t be disassembled until it is unpacked. Other anti-disassemble technics include adding junk code and other statements that are designed to fool the disassembler into not correctly disassembling the file. Malware creators will study popular disassemblers and develop tactics that will cause bugs and other chaos when their malware is disassembled. Chapter 8 covers many of these topics in more detail.

Dynamic Analysis

Dynamic analysis consists of running the file and learning from its behavior. Dynamic analysis can be risky because you are allowing the malware to run, and many forms of malware have anti-detection capabilities. The most

common way to securely run malware is to use a sandbox. One very simple and free option is Joe Sandbox (https://www.joesandbox.com). An analyst can simply upload any questionable file and allow Joe Sandbox to run and report what is found. If you don't have any analysis capabilities, create a free account with Joe Sandbox so you at least have an ad hoc analysis capability.

Another option to perform dynamic analysis is to build your own sandbox using an enterprise or open-source option. Properly configured sandboxes can provide a ton of value, including capabilities to detect malicious behavior, simulate human activities to trick the malware into believing it has infected a real system, and simulate network services with the intent of learning what the malware wants to communicate with. The downsides of sandboxes include limited host environments and missing certain capabilities that you would get by creating your own virtual system. If you are okay with using a standard operating system built for your testing, a sandbox will be the best option. However, if you want the testing environment to mirror one of your customized systems, you will likely have to build your own sandbox-like environment. Chapter 8 covers how to build and use sandboxes.

Other forms of analysis include reading security logs to understand events that are occurring within your environment, which can be seen as another form of dynamic analysis since you are monitoring the behavior of other systems. Logs can come from various types of devices including firewalls, IPS/IDS systems, host systems, network tools, and even external sources such as threat intelligence. Collecting logs has little value if an analyst can't understand and respond based on what is seen, leading to the need for a centralized place to see logs from all systems. The tool used for this purpose is known as a security information and event management (SIEM) solution. Chapter 5, "Centralizing Data," looks at the different types of data an analyst will likely encounter and how to operationalize what is collected so services such as incident response can be more impactful.

SOC Service 6: Digital Forensics

Another more focused form of analysis is digital forensics. Digital forensics occurs after an incident and serves the purpose of understanding what happened as well as collecting evidence to prove a hypothesis of what

occurred. Proof allows evidence to be used for critical decisions and legal action. Key deliverables that can be expected from a digital forensic service include the ability to recover data, determine what a file is and who launched it, reverse engineer files, understand threat data, and convert it all into a storyline.

Note

The SOC must decide to perform either an incident response or a digital forensics investigation. Both options take a completely different approach to how an incident is handled. An incident response will modify impacted systems, leading to the loss of forensic evidence. A digital forensics response will focus on preserving evidence rather than returning systems back to normal option. Keep this in mind as you are developing your SOC services.

Digital forensics must be performed correctly during the entire lifecycle of the incident or evidence will be contaminated. Contaminated evidence will produce unreliable data, which will not be useful for legal actions and risky to rely on for any critical decisions. The following steps list, at a high level, what must occur during a forensic investigation based on my personal experience. Guidelines such as NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, are extremely useful for obtaining more details on the steps provided as my steps are based on industry best practice found in such guidelines.

1. **Identify a crime or event:** Identifying the need for forensics starts the investigation process.
2. **Collect preliminary evidence:** Preliminary evidence includes documenting the scene of the crime and collecting any information that justifies that a possible crime has been committed.
3. **Obtain court warrant for seizure if required:** Before action can be taken, a forensic investigator must ensure they are authorized to interact with any evidence. Each country will have different laws and rules regarding what is required before an investigation can proceed. For

crime within the United States, a warrant may be required, which authorizes the investigation.

4. **Perform first responder procedures:** Once the forensic engineer is authorized to proceed with collecting evidence, he or she will collect various forms of evidence without making any modifications. Evidence that is powered on will remain powered on unless there is no way to collect the evidence while powered. Before any evidence is collected, the crime scene is documented to ensure an understanding of what the crime scene looked like prior to the start of the investigation.
5. **Seize evidence at the crime scene:** Evidence is collected and secured to avoid contamination. Each piece of evidence is labeled and documented regarding where and how it was collected.
6. **Transport evidence to forensic laboratory:** Evidence is transported in a secure manner. Any transportation step is documented so the entire transportation process can be accounted for. All evidence is secured when it is not in transit. Proper security means using a safe rather than locking it in a desk drawer that could easily be opened by an unauthorized party.
7. **Create copies of evidence:** Bit-level copies of the original evidence are created and only the copies are investigated. The original evidence is secured and never altered.
8. **Generate images and exam for evidence:** Copies of evidence are investigated until a hypothesis is developed.
9. **Action taken:** Actions including legal, disciplinary, and mitigation are taken based on hypotheses concluded from the investigation.
10. **Return of evidence:** Original evidence is returned to owner.

Note

The NIST SP 800-86 forensic process offers a shorter explanation of the digital forensic process listed as data collection, examination, analysis, and reporting.

There is overlap in skills used to investigate malware during a forensic

investigation and what an analysis team would do to determine if an artifact is malicious. In both cases, there are challenges regarding coming up with a hypothesis about the artifact. The first challenge is that digital data is very volatile, making it hard to keep in its original form. An example is analyzing a flash drive, which has trim algorithms that continuously format the drive to maximize available space. Forensic investigators use hashing as a means to validate a file hasn't been modified, which is critical when proving that evidence hasn't been contaminated. If a flash drive permits trim, the data will be modified by the drive, causing a mismatch with any copy of data, leading to the assumption of corruption! Chapter 8 covers how to deal with investigating various types of evidence, how hashing works, and other best practices for performing digital forensics.

Note

I will provide limited digital forensic legal concepts in this book since the law you are responsible to follow will depend on where the crime was committed or the event occurred. I highly recommend leveraging legal counsel that is knowledgeable of local law before considering performing any steps that could violate a law such as one pertaining to privacy or human rights.

SOC Service 7: Situational and Security Awareness

As described in the article "Situational Awareness for Cybersecurity: An Introduction" by Angela Horneman,

> Situational awareness (SA) helps decision makers throughout an organization have the information and understanding available to make good decisions in the course of their work. It can be focused specifically on helping people and organizations protect their assets in the cyber realm or it can be more far reaching. SA makes it possible to get relevant information from across an organization, to integrate that information, and to disseminate it to help people make better decisions.

Essentially, Horneman's explanation means informing people helps them make better decisions through training.

People will be your weakest link in your security defense strategy. You can't apply a patch to people like you can with technology to reduce risk. You can't install firewalls or antivirus in people. You really have only two options regarding reducing the risk that people pose. Option one is providing training and awareness to help people avoid risky behavior. Option two is to automate defenses, such as automatically blocking risky websites when users attempt to perform a high-risk action. The challenge with the second option is that people will continue to work around the security if they don't believe the security is protecting them. A buddy that works for a large SOC told me he once received a ticket from a customer within his organization complaining they couldn't access a website that was flagged as being malicious. In the ticket, the customer not only complained about being denied access, but also explained all the steps taken to "try to get it to work," meaning attempting to bypass security. It didn't dawn on the customer that the website was being blocked for his or her protection because the customer wasn't educated on why it was being blocked. I have heard hundreds of similar stories.

User Training

The SOC is responsible for the security of the organization and therefore must also be involved with user training. Training programs need to mirror real events that have occurred recently or could impact the organization. Some very common threats that you must educate users about as part of your training program are social engineering, malicious files, and data confidentiality, which in my research and experience are the top-trending attacks targeting people. The purpose for phishing and social engineering attacks is to trick or coerce users into taking an action they wouldn't otherwise take. Your training program needs to educate how to identify and question any potential threat by assuming all communication could be a phishing or social engineering attack. Vendors such as Confense offer phishing education programs, including assisting with launching phishing attacks in a controlled manner to help educate users. You can also use open-source tools such as TrustedSec's Social-Engineer Toolkit (SET) to clone a

trusted website and send an email asking people to log into that website. Figure 3-26 shows using SET to clone accounts.google.com, which is the login screen for Gmail.

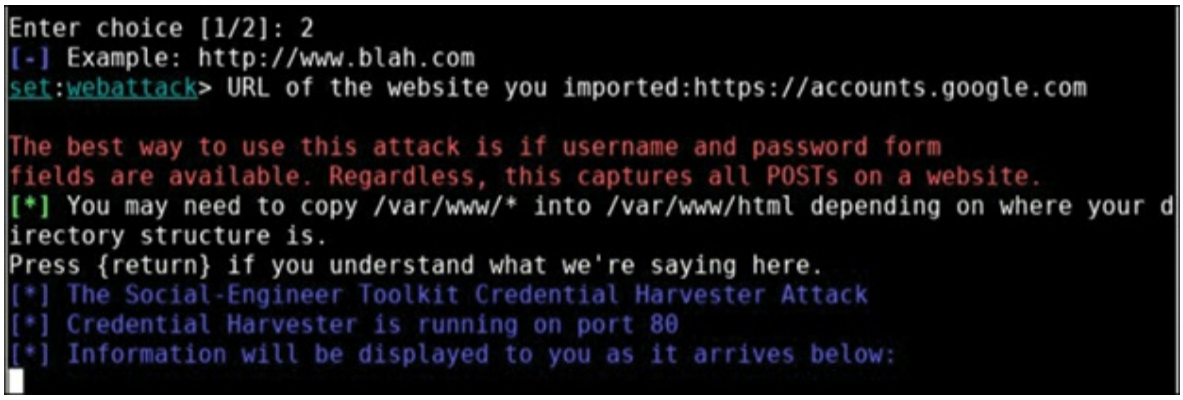

FIGURE 3-26 Using SET to Clone Gmail

Figure 3-27 shows the cloned Gmail login. Notice that 198.18.133.6 is the website of the attack server hosting the fake login for Gmail. In a real phishing attack, the attacker could purchase a website that looks similar to Gmail or use a website-hiding technique such as TinyURL (https://tinyurl.com/) to hide where the user would go if they click the hyperlink. A simple phishing email could state that somebody's password has been compromised and they need to log in to validate who logged in. A link to this fake website could be hidden in an icon planted in the phishing email that says "Click to Access Gmail Now," sending the victim to the fake Gmail website.

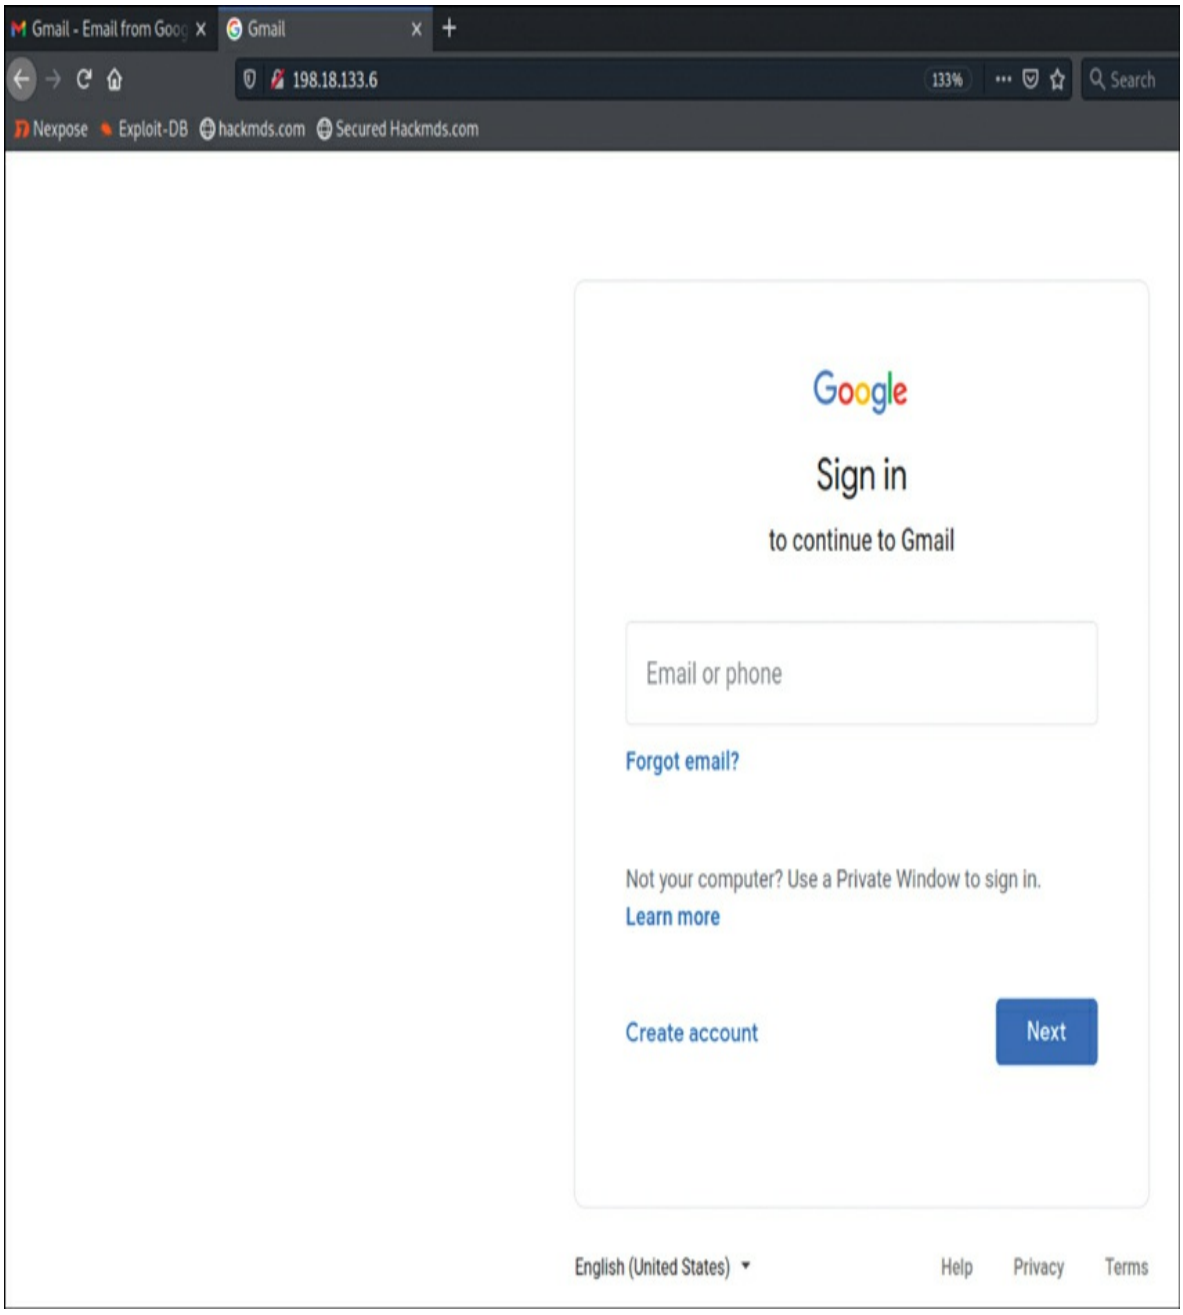

FIGURE 3-27 Cloned Gmail Website

I recommend performing simulated phishing attacks regularly so that users learn to question emails such as the Gmail phishing example provided. You will want to post the success of the campaign without shaming people that fell for the attack. Your goal is to continue regular phishing attacks and demonstrate improvement, meaning fewer people falling for the attack. I also recommend posting posters and periodically sending emails to your

organization about these threats to keep everybody aware of the threats that are targeting them. One example I found effective was an organization posting pictures of sad kittens in the breakroom with the slogan “Every time you click a phishing email, a cat dies somewhere.” As weird as this sounds, I found people in the break room noticing and making comments about the sign. I was also told that the customer heard some employees yell “Oh no, I killed a kitty!” when they were fooled by a phishing campaign being delivered by the SOC. The point of this example is to build your security message into something that fits your culture. Try to avoid using stale, artificial-sounding statements or all text warnings. Be creative regarding how to capture people’s attention or your message will be ignored. Chapter 4 delves into more training concepts.

Note

I recommend posting the success of a phishing campaign in a secure manner. Announcing that your users are suckers for phishing attacks won’t go over well and will have a negative impact on the organization.

SOC Service 8: Research and Development

The final fundamental SOC service to review is research and development (R&D). Research encompasses anything the SOC needs to understand. One research focus will be new technology and process. Let’s say the SOC needs to improve its SIEM technology. Research is required to understand which vendors are available and good, what the potential costs are regarding the lifecycle of the solution (install, training, maintenance, etc.), and who should be trained on the new technology. Vendors can be asked to provide a lot of these details, but the final decision will be the organization’s, which will rely on research for the best choices. Technology is constantly changing, so it is critical to have a team within the SOC dedicated to keeping up with the latest red and blue team tactics through continuous research of the latest attack and defend trends.

Another focus for research is collecting threat intelligence. Internal threat intelligence is useful but is limited to threats that have impacted your organization. There is a huge advantage to learn what other organizations are experiencing as well as monitor external resources to learn about potential upcoming threat actors. Chapter 7, "Threat Intelligence," looks closer at four different types of threat intelligence, summarized here:

- **Nontechnical intelligence:** This form of threat intelligence is designed to help leaders of organizations make decisions. It is also called *strategic threat intelligence*.
- **Tactical intelligence:** This form of threat intelligence focuses on the tactics, techniques, and procedures of threat actors so the SOC can learn about threat campaigns.
- **Operational threat intelligence:** This form of threat intelligence looks at specific attack campaigns, pulling data from various sources including social media, open-source intelligence, and industry expert blogs.
- **Technical threat intelligence:** Most SOCs think of the fourth type of intelligence when asked about threat intelligence, which are IP addresses of malicious sources, hash files representing malware, and other specific data. I see this form of threat intelligence commonly used to enhance blacklists within security tools.

All four of these different types of threat intelligence must be used properly in order to receive benefits or it becomes noise. The research team within a SOC needs to manage which sources to use and how they are leveraged within the organization. Without centralized control, your organization will end up having random threat feeds without purpose, leading to additional noise rather than value.

Other types of research a SOC can be tasked with performing are obtaining operational costs of the SOC services, reviewing lessons-learned data from other services to identify areas of improvement, meeting with other SOCs or consultants to develop improvement strategies, conducting legal research, or researching anything else that the SOC needs to know more about. It is highly recommended to specify a branch of the SOC that has responsibility for research or research tasks will become ad hoc extra work for employees who have other responsibilities, leading to low-quality research results. Research is very time consuming and must be seen as a job within the SOC in

order for the work to have the right level of focus.

Development involves creating new tools, tactics, procedures, or other useful things such as a training lab that benefit the SOC. Development projects can be extremely time consuming and resource demanding, and the SOC will need to be prepared to allocate resources toward new projects if they're to be valuable. Development often goes hand in hand with research since the SOC will need to understand what is available and how to use new concepts before they can be developed. Chapter 10 covers things to consider when building a new tool or using an existing open source tool. In certain cases, these options can provide a lot of value to the SOC; however, there are also times it makes more sense to outsource the development of tools to vendors.

Summary

In this chapter, you learned about the fundamental services that modern-day mature SOCs provide: risk management, vulnerability management, incident management, analysis, compliance, digital forensics, situational and security awareness, and research and development. These SOC services are the critical services that a SOC performs on a day-to-day basis for operations. You also learned about foundational SOC support services, which are required for the SOC to operate, and SOC service areas, which are subsets of services that relate back to one or more fundamental SOC services. The chapter next covered how to develop a new service, including when to use external resources versus developing in-house SOC service. The latter sections of the chapter provided an overview of each of the key SOC services. SOC services are deeply connected with business and technical aspects of a company and include all aspects of a cybersecurity lifecycle.

You need to aim at developing some form of all of these services and continuously improving the maturity of each service in order to keep up with the continuously changing threat landscape. In the following chapters, I will look much closer at each of these SOC services and provide various ways for you to start, deliver, and improve each service. To continue the journey, I move to a focus on people and process in Chapter 4.

References

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012, August). Computer Security Incident Handling Guide. Computer Security Resource Center. https://csrc.nist.gov/publications/detail/sp/800-61/rev-2/final

Endsley, M. (2000, January). Theoretical Underpinnings of Situation Awareness: A Critical Review. In Endsley, M. R. and Garland D. J (Eds.), *Situation Awareness Analysis and Measurement* (pp. 1–24). Lawrence Erlbaum Associates. https://www.researchgate.net/publication/230745477_theoretical_underpinnings_of_situation_awareness_a_critical_review

Forum of Incident Response and Security Teams. (2020). Computer Security Incident Response Team (CSIRT) Services Framework: Version 2.1. FIRST. https://www.first.org/standards/frameworks/csirts/csirt_services_framework_v2.1

Horneman, A. (2019, September 9). Situational Awareness for Cybersecurity: An Introduction. Carnegie Mellon University Software Engineering Institute. https://insights.sei.cmu.edu/sei_blog/2019/09/situational-awareness-for-cybersecurity-an-introduction.html

International Organization for Standardization. (2018, February 7). ISO/IEC 27000:2018: Information technology – Security techniques – Information security management systems – Overview and vocabulary. ISO. https://www.iso.org/standard/73906.html

$(ISC)^2$ (2019, November 6). $(ISC)^2$ Finds the Cybersecurity Workforce Needs to Grow 145% to Close Skills Gap and Better Defend Organizations Worldwide. $(ISC)^2$. https://tinyurl.com/ISCCyberWorkforce

Lord, N. (2018, September 12). What Is the Principle of Least Privilege (POLP)? A Best Practice for Information Security and Compliance. Digital Guardian. https://digitalguardian.com/blog/what-principle-least-privilege-polp-best-practice-information-security-and-compliance

National Institute of Standards and Technology. (2020, April 21). Cybersecurity Framework. NIST. https://www.nist.gov/cyberframework

Security Ninja. (2018, February 7). CIA Triad. Infosec Resources.

https://resources.infosecinstitute.com/cia-triad/

Tenstepadmin. (2015, January 7). Use Expected Monetary Value (EMV) to Determine Risk Impact. TenStep. https://tenstep.com/use-expected-monetary-value-emv-to-determine-risk-impact/

Verizon. (2020). 2020 Data Breach Investigations Report (2020 DBIR): Summary of Findings. Verizon. https://enterprise.verizon.com/resources/reports/dbir/2020/summary-of-findings/

$(ISC)^2$. (2019). $(ISC)^2$ Finds the Cybersecurity Workforce Needs to Grow 145% to Close Skills Gap and Better Defend Organizations Worldwide. https://www.isc2.org/News-and-Events/Press-Room/Posts/2019/11/06/ISC2-Finds-the-Cybersecurity-Workforce-Needs-to-Grow--145

Chapter 4

People and Process

Never forget what you are. The rest of the world will not. Wear it like armor, and it can never be used to hurt you.

—"Tyrion Lannister," *Game of Thrones* (George R. R. Martin)

This chapter focuses on the human element of the SOC. These are the people that deliver the services covered in Chapter 3, "SOC Services," and will be the highest cost of running the SOC. According to a 2018 survey of 620 IT and cybersecurity professionals conducted by Enterprise Strategy Group (ESG), as summarized by Jon Oltsik, a senior principal analyst at ESG, "cybersecurity represents the biggest area where their [survey respondents] organizations have a problematic shortage of cybersecurity skills." This means not only are good people hard to find, they are even harder to keep because the technology industry has more jobs than people to run them. This chapter looks at what skills are recommended for different SOC job roles, how to recruit the right people, and strategies to keep those people excited to be part of your SOC. Without the proper people, process, and technology, your SOC will experience failures in services. Also, remember from Chapter 3 that people are one of the three pillars (along with work environment and technology) of the foundational SOC support services that must be in place before any SOC service can be launched. Let's now spend a chapter focusing on your people.

Career vs. Job

My mother used to explain that the difference between a job and a career is the perspective of the person doing the work—that is, how serious the person

considers the work to be. For example, many teenagers look for a job simply to save enough money to purchase things they want. They don't care about advancements in their job because they are working just for the paycheck and typically don't even know or care about the mission of the organization they work for. By contrast, people who are career-driven are not showing up just for a paycheck. They also want career advancement, training to improve their skills, and the satisfaction of spending time working on something they enjoy doing.

The goal of this chapter is to help you not only plan to recruit career-driven people, but also develop and retain talent, because people are going to be your SOC's most important assets.

Developing Job Roles

Many different job roles fall under the categories "cybersecurity" and "information technology." Within those generic categories are roles that are responsible for presales, delivery of services, daily operations, and everything in between. Your SOC will have roles that fall under the cybersecurity and information technology categories; however, your SOC roles will require specific skills, knowledge, and experience based on the services your SOC offers. Sometimes skills, knowledge, and/or experience can be acquired on the job, while other times they are prerequisites for an employee to take on the associated responsibilities of a job role. Successful organizations clearly define job roles, compensation ranges, responsibilities, and paths for career growth because these elements are what attract and retain quality people.

In Chapter 1, I introduced the eight core services I find within mature SOCs. Each service has different types of job roles, which some can apply to multiple services while others are very specific to a single service. You will need to recruit and retain the right talent for the services you offer, which continues to be an extremely challenging task in today's competitive cybersecurity job market. Not only is it hard to find the right talent, but experienced talent will be expensive. You will have to decide when you can groom an internal employee for a role or seek external talent to fill a position.

One major factor that impacts these decisions is available budget for

recruiting talent. Leadership will need a general number of what the cost will be to fill a SOC position. The best way to determine a ballpark cost to fill a SOC position will be using publicly available pay scales. The general schedule pay scale is an example of such a resource.

General Schedule Pay Scale

The U.S. federal government uses a scale based on series and grade to categorize and define jobs. The series is a numbered system for grouping similar occupations. For example, a computer engineer is part of the 0854 series, while a nurse is part of the 0610 series. The grade refers to the General Schedule (GS) pay scale representing the pay level for the job. A job role with a higher GS grade will have a higher pay range. Employees with a high school degree and little experience fall under the GS-5 and lower range, while people with work experience can expect to be at least at a GS-7 level. Employees with a master's degree and special experience will expect a GS-9 or higher job role. People looking to work for the U.S. federal government can use this system to quickly understand the pay range for any available U.S. federal job request. Candidates can also refer to the standardized language of the GS pay scale jobs to ask about how the existing role can advance to higher GS grades as the candidate gains experience in the role.

Note

The U.S. pay scale is just one example of a grade scale format. I believe the concept of grade scales is useful for better understanding a pay range and what is involved with a job. I believe grade scales more accurately represent the responsibilities of a job role than do job titles. In my experience, I've worked with people who have fabricated fancy job titles when their official title as documented within the organization is different. For example, a salesperson who is responsible for northeast sales might use the title Director of North East Sales even though he or she is not performing what the industry would consider director job duties. I've found that people who haven't had an increase in responsibility tend to eventually create their own

made-up job titles. The most common example is using terms such as "Senior" to represent time served rather than an increase in responsibility. Time served does not automatically increase an employee's grade scale.

Formalizing Payscales

The GS pay scale is just one example of a pay scale you can use to standardize how compensation is distributed to each job role in your organization. You want to apply a formal pay scale to your SOC roles to set expectations for the pay range associated with your positions. You also need to be specific regarding what skills and other requirements are involved with each role to ensure potential candidates know what is required to qualify for the role. This also applies to advancements in a role. For example, as a SOC analyst gains experience, her title should change. A SOC analyst could start out as a grade 1 analyst. Once that analyst meets certain time, skill, and experience requirements, the analyst can request to be promoted to a grade 2, which will have a higher pay range. While skills are being obtained, salary increases should be provided that fall within the specific pay range. At some point, the candidate will hit the top of the pay range and must move to another pay range before any further increases in salary can be provided.

Formalizing pay scales enables employees to understand how their compensation will change as they increase in grade scale or switch roles, which will have their own assigned grade scale. Some job series will max out faster than others, encouraging an employee to switch roles if they desire a higher pay scale. An analyst series might max out at the role "analyst grade 5" while the pay for a analyst grade 5 is similar to a "architect grade 2" role. In this example, an analyst would not be able to make the same income as an architect grade 3 or higher, motivating the analyst to switch roles if he wants to be part of a higher pay scale than what an analyst pay scale could offer. Having certain job series max out at lower pay scales than other job series isn't a bad thing. Developing a job role structure with certain job pay scales maxing out lower than others encourages career development that is driven toward senior job titles. Companies that don't encourage career growth and just provide standard raises on an annual basis will not encourage employees

to invest time into developing their skills or career. As a result, employees will remain unmotivated and a flight risk.

> **Note**
>
> Learn more about the GS pay scale at https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/2019/general-schedule/.

IT Industry Job Roles

Job roles need to be clearly defined to identify a baseline of responsibilities as well as skills and experience expectations. The next section reviews the various types of jobs and their expected associated skills. It is up to your organization to customize and explain how the general skills associated with a job title relate to the specific job role and what additional skills and experience are desired for a potential candidate to be considered.

According to the employment website Indeed (https://www.indeed.com), the following items need to be included in a basic job title. Make sure to elicit responses to each of these categories with any job posting that you publish.

- **Job role:** Use targeted language rather than generic titles. Avoid lingo that is internally unique to your organization.
- **Job summary:** Sell your job with an attention-grabbing summary. Include the exact job location, including whether remote work is an option.
- **Responsibilities and duties:** Outline the core responsibilities. Highlight the day-to-day activities. Specify how the position fits within the organization and SOC.
- **Qualification and skills:** Provide a list of hard and soft skills. Keep the list concise.

To better understand job roles, let’s review common job titles and their associated skills.

Common IT Job Roles

Reviewing the common job roles that exist in the IT market space is a good place to start before focusing on the SOC-specific roles you will want in your organization. You can use the following list of IT industry job roles to better understand what type of skills are associated with a common IT title and determine if that role could apply to a SOC role you are looking to fill. Some roles will be tied to generating revenue, known as presales roles, while others will be supporting the organization in various fashions. Some job roles, such as a PCI DSS compliance officer, are tied to specific tasks, while others, such as a network engineer, are more generalized. The range and depth of skills will also vary between roles. A presales engineer might or might not have much hands-on experience with a technology depending on how the candidate utilized the technology in his or her previous role. It is best to qualify any skill during the interview process and validate experience through references.

Note

As you review these job roles, you might wonder how they relate to the SOC. As I've mentioned a few times in this book, the security industry is lacking sufficient qualified talent to fill all the jobs that are available. This concept holds especially true for people with SOC experience. I find that many organizations have to either grow SOC skills from within or expand their search to more generic IT skills in order to find available people. I will cover SOC job roles shortly, but it is extremely valuable to also know industry job titles as well. You might need to pull candidates from another IT field to find somebody for your SOC service.

- **Account manager (AM):** An account manager works in the sales and marketing department of a business and is responsible for managing client accounts. This job role requires very little technical knowledge, but it does require mature soft skills and a drive to execute on meeting or exceeding sales goals.

- **Sales engineer (SE):** A sales engineer combines technical knowledge with sales skills (a combination of hard and soft skills). Because many account managers lack technology knowledge, they require an engineer to handle technical-related tasks. Those tasks include understanding the customer's technical needs, explaining the technology or services those needs represent, providing demonstrations of technology, or possibly even installing technology to prove it can accomplish the desired goals so that a sale can be achieved. Sales engineers must be able to translate technical concepts into terms that nontechnical people understand.
- **Marketing engineer:** Organizations that sell products or services have teams dedicated to developing how those offerings are marketed to customers. Some marketing teams require creative people with a technical background to explain the value of the solutions being offered as well as validate if the marketing efforts meet their targeted customers' expectations. The level of technical and soft skills required for the marketing engineer position will depend on the type of products and services being offered as well as how the marketing engineer will be utilized.
- **Installation/post-sales engineer:** This role supports presales teams by delivering the products and services that were sold to the customer. Services could be short-term or long-term contracts and have various travel requirements. For example, an installation engineer could travel often to new customer locations for short projects or be part of a long-term deployment that spans across multiple locations.
- **Compliance officer:** Many organizations have compliance requirements that they must meet to offer certain types of services as well as to avoid the negative impact (such as fines) from not meeting mandatory compliance. Compliance officers are responsible for monitoring the current state of an organization's compliance status, obtaining proof that compliance is met, monitoring for changes in compliance, and performing other compliance-related tasks.
- **Manager:** Managers are responsible for addressing employees' needs. Fulfilling those needs can include operational requirements, such as providing tools and support to perform their jobs, or emotional support to encourage a positive working environment. Great managers help people

achieve goals as well as mentor employees so they can grow their skills and feel accomplished. When employees experience challenges, managers are responsible for representing their needs. Managers are expected to have strong soft skills and experience managing people.

- **Desktop support:** The desktop support group focuses on managing host-related services. This can include support needed for desktops, laptops, mobile devices, and sometimes servers. Desktop support can be responsible for issuing equipment, enforcing security within equipment, and supporting the equipment with updates or software requested by employees. Desktop support can also develop policies for endpoints and support the SOC's mission of enforcing security policies. Skills can range between operating system types and tools, depending on experience level.
- **Helpdesk:** The helpdesk team is responsible for anything related to supporting employees and their equipment. This role is typically the first layer of support for an organization's internal services. Examples of common helpdesk job duties include resetting passwords, provisioning hardware and software, and responding to security incidents, such as a user reporting that her computer might be infected with a virus. The desktop support role and helpdesk role can be the same role or have responsibilities divided between different teams. A SOC can include a helpdesk service to assist with responding to security incidents and to support SOC team members' technical requirements.
- **Database/cloud engineer:** Organizations create data and need a place to store it. Data can be stored locally on servers or on a cloud storage service provider's servers. A database or cloud engineer acts as a data custodian ensuring that data is protected and policies created by the data owner are enforced. Technical skills include setting up relational databases, designing queries and reports to access information in the databases, and administering backup and recovery procedures.
- **Network engineer:** Network engineers deploy and manage the organization's networks. Every organization has some form of network services such as LAN, VPN, and wireless. Even organizations that lead with cloud services need a network to enable employees to access the cloud. Network engineer skills range from configuring to monitoring and troubleshooting various types of network equipment.

- **Software engineer:** Computer programs are computer code created by software engineers. As IoT and other technology grows in popularity, the need for programmability and applications increases the need for software engineers. Many SOCs leverage customized applications that are built by software engineers or leverage open-source tools that can leverage programmable tools that modify how the tool works or how the data is used by the tool. Software engineers develop information systems by designing, developing, and installing software solutions.

Note

A question I'm often asked is, "How do I start a career in cybersecurity?" My answer is that it depends on where you see yourself in three, five, and maybe even ten years from today. As the preceding list of the IT job roles indicates, many different types of work fall under the category "cybersecurity" or more broadly "information technology." Know that all the types of IT and cybersecurity job roles are not a good fit for you. There are many types of jobs within the world of IT that require different skills and personalities. I recommend identifying the type of work you want to do and speaking with people in that job role. As you consider a future career, factor in requirements for expected travel, work hours, compensation, required skills, and associated culture, even if some of these factors will be based on the specific employer offering the job. Once you find a desired job role, work toward education and experience specific to that role.

Some of the preceding job roles could apply to SOC work, while others do not but could perform SOC work with some level of training. I also didn't cover every job role you will find if you search popular job recruiting resources using terms like "cybersecurity" and "information technology" since the list could take up the entire chapter. Many of these jobs are also feeder roles into security-related work, meaning jobs people do before they start working in a SOC or undertaking similar security-related work. Sometimes people find a job in security later in their career because the

candidate didn’t initially pursue a career in security after completing their education, found an opportunity in non-security-related work prior to performing SOC work, wasn’t qualified for security-related work, or other reasons.

SOC Job Roles

The expected career path for any job role in a SOC will depend on how the organization assigns responsibilities and pay scale to a job role. Roles in networking, software development, system engineering, and security intelligence can lead to entry-level SOC-related work. Entry-level SOC job roles such as junior analyst, consultant, or tester can lead to job titles such as senior architect or security administrator as responsibilities and pay scales increase. Know that there isn’t a set standard for job roles or how roles feed into other roles, meaning the role of analyst at one organization could require the same experience as the role of architect at another organization. One organization might require specific certifications, degrees, or experience to meet the requirements of a job role, while the same job role at another organization will have different requirements. Consider industry and SOC job role, pay scales, and expected experience as you develop your strategy for recruiting for any job role in your SOC.

The job roles covered in the sections that follow make up common SOC-related career paths. These roles range from entry-level to senior-level job titles. The specifics of the work will depend on the type of service offered by the SOC. I will attempt to group similar job roles and explain skills based on what I have encountered in SOCs around the world. Use the recommended skills and certifications listed as reference points for what training and certifications you could pursue if you work in one of these job roles.

Note

I will follow up this section with an industry guideline for job roles known as the NICE Framework, which is much more detailed than my general list covered next. I will not list everything found in the NICE

Framework but rather show you how to access and use that resource to research career path data.

Security Analyst

The security analyst role evaluates various types of data and plans and implements security measures to protect computer systems, networks, and data. Reviewing data can mean evaluating live network traffic or a copy of evidence such as event logs generated by security and network tools. In regard to a security operations center, a SOC analyst can be responsible for reviewing security logs and responding to events based on the services offered by the SOC. The skills associated with a security analyst can include reading logs and event data from various types of tools, implementing changes to security tools, such as configuring firewall rules, responding to incidents based on suspected events, and developing playbooks for the organization to standardize its responses to different events.

Table 4-1 outlines the responsibilities, skills, and certifications associated with the security analyst role. The security analyst role is ideal for the incident management SOC service but can also be part of the vulnerability management and research and development (R&D) services. Similar job titles include security engineer, security administrator, security specialist, and security consultant.

TABLE 4-1 Security Analyst Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Evaluate security measures and controls for vulnerabilities	Penetration and vulnerability testing, information security knowledge	CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security

		CPT: Certified Penetration Tester CEPT: Certified Expert Penetration Tester GPEN: GIAC Certified Penetration Teste CISM: Certified Information Security Manager
Establish plans and protocols to protect digital files and information systems against unauthorized access, modification, or destruction	Host security tools (antivirus, anti-malware, VPN), data loss prevention technologies, encryption concepts, identity management, access control	ECSA: EC-Council Certified Security Analyst Vendor NAC certification Vendor Data Loss certification

		Identity Management certification (e.g., Microsoft Active Directory)
Maintain data and monitor security access	TCP/IP, computer networking, routing and switching	GSEC: GIAC Security Essentials
		GCIH: GIAC Certified Incident Handler
		GCIA: GIAC Certified Intrusion Analyst
		CISM: Certified Information Security Manager
Perform security assessments and recommend security controls	Firewall and intrusion detection/prevention protocols	CISSP: Certified Information Systems Security Professional
		Vendor product

		certifications
Anticipate security alerts, incidents, and disasters and reduce their likelihood	Windows, UNIX, macOS, and Linux operating systems	Operating system certifications
Manage network and security systems	Network protocols and packet analysis tools. Windows, UNIX, macOS, and Linux operating systems	Vendor network certification (e.g., Cisco CCNA/CCNP/ CCIE) Operating system certifications
Analyze security breaches to determine their root cause and impacted parties	Digital forensics and threat hunting	EC Council Computer Hacking Forensic Investigator certification
Recommend and install tools and countermeasures	Understand industry frameworks, security tools, and security process	ISC2 CISSP CompTIA CySA+
Provide training to employees in security awareness and procedures	Developing training programs	SANS Security Awareness Professional (SSAP)

Penetration Tester

The penetration tester role is focused on identifying vulnerabilities and

testing those vulnerabilities in a similar manner to how an adversary would. Assessment officers and others that are responsible for identifying vulnerabilities tend to leverage automated tools and focus on identifying potential vulnerabilities but do not validate how realistic the vulnerability may or may not be. Penetration testers invest additional time validating that vulnerabilities exist using the same tools used by adversaries. Penetration testers attempt to exploit the vulnerability and then document the results. A penetration tester must be knowledgeable in how to identify vulnerabilities as well as common tactics used to exploit a vulnerability to achieve the same outcome a potential adversary could obtain. This skillset is commonly referred to as red team skills.

Table 4-2 outlines the responsibilities, skills, and certifications associated with the penetration tester role. A penetration tester is ideal for the vulnerability management SOC service but can also work in the compliance, risk management, and R&D services. Similar job titles include security analyst, security engineer, threat researcher, ethical hacker, red team member, and tester.

TABLE 4-2 Penetration Tester Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Perform penetration tests and assessments of web-based applications, networks, and computer systems	Exploitation, assessment, and audit skillsets; technical writing; legal and compliance understanding	CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT: Certified Expert

		Penetration Tester GPEN: GIAC Certified Penetration Tester
Conduct physical security assessments of servers, systems, and networks	Vulnerability and physical security assessment capabilities Lock picking	A+ and other hardware certifications
Design and create new tools and tests for penetration testing and assessments	Network servers, networking tools, security tools and products	OSCP and PEN-200 from offensive security CEPT: Certified Expert Penetration Tester
Probe targets and pinpoint methods that attackers could use to exploit weaknesses and logic flaws	Computer hardware and software systems; vulnerability management and exploitation tactics	GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security

		CPT: Certified Penetration Tester CEPT: Certified Expert Penetration Tester
Employ social engineering to uncover security holes	Web-based applications and behavior science	OSCP: Offensive Security Certified Professional
Incorporate business goals into security strategies and policy development	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.)	CISSP: Certified Information Systems Security Professional CISM: Certified Information Security Manager
Research, document, and review security findings with management and IT teams	Vulnerability analysis and reverse engineering	CCFE: Certified Computer Forensics Examiner
Improve security services, including the continuous enhancement of existing methodology material and supporting assets	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.)	CISSP: Certified Information Systems Security Professional

Provide feedback, support, and verification as an organization fixes security issues.	Communication and writing	College degree

Assessment Officer

An assessment officer is responsible for identifying potential vulnerabilities or gaps in corporate policy, compliance requirements, or general security best practices as defined in popular frameworks. Unlike a penetration tester, an assessment officer works within specific scopes as defined by policies, compliance, or frameworks, meaning he or she must be aware of the latest requirements and continuously validate the organization is meeting those requirements. Any vulnerabilities out of scope of such requirements will be overlooked by the assessment officer because the focus of an assessment officer is auditing rather than general security validation. An assessment officer's skills are focused on business and operations with a strong understanding of industry frameworks, compliance, and laws associated with cybersecurity as it relates to the organization.

Table 4-3 outlines the responsibilities, skills, and certifications associated with the assessment officer role. An assessment officer is ideal for the compliance and risk management services but can also work in the vulnerability management service or assist other services such as incident management and R&D. Similar job titles are compliance officer, policy officer, security officer, and infosec officer.

TABLE 4-3 Assessment Officer Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Incorporate business goals into security strategies and policy development	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.)	CISSP: Certified Information Systems Security Professional CISM: Certified Information

		Security Manager
Conduct physical security assessments of servers, systems, and networks	Vulnerability and physical security assessment capabilities; lock picking	GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT: Certified Expert Penetration Tester
Interview employees, obtain technical information, and assess audit results	Management and strong communication skills	College degree or special communication skills training CISM: Certified Information Security Manager
Understand industry data security regulations	Understand HIPAA, PCI DSS, etc.	Specific industry data security certification and experience
Develop and execute tests based on regulations being audited	Critical-thinking skills	College degree and/or programming certification
Research, document, and	Critical-thinking skills	College degree

review security findings with management and IT teams		and/or programming certification
Understand organization policies and procedures	Critical-thinking skills and experience with SOC policies and procedures	College degree
Provide feedback, support, and verification as an organization fixes security issues	Critical-thinking, project management, and communication skills	College degree

Incident Responder

An incident responder is a cyber first-responder or a higher-tier resource responsible for responding to a security incident. This role involves providing rapid initial response to IT security threats, incidents, and cyberattacks on the organization. The role can also include some penetration and vulnerability testing, network management, intrusion detection, security audits, network forensics, and maintenance of IT security systems. The primary responsibility may be monitoring traffic for any unusual activity or unauthorized access attempts and initiating the appropriate response when a potential event is identified. The response can include patching systems, initiating segmentation, isolating systems, alerting all associated parties, and assisting with returning impacted systems back to an operational state. The incident responder can work through the entire lifecycle of the incident or handle one part of the incident while higher-tier responders or other teams take over responsibilities, depending on the severity of the incident and how the SOC runs the incident management practice.

Table 4-4 outlines the responsibilities, skills, and certifications associated with the incident responder role. An incident responder is ideal for the incident management service but can also work in the situational and security awareness service or vulnerability management service. Similar job titles include incident response engineer, computer network defense, IT network defense, incident analyst, intrusion detection specialist, and network intrusion

analyst.

TABLE 4-4 Incident Responder Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Actively monitor systems and networks for intrusions	Windows, UNIX, macOS, and Linux operating systems	Operating system certifications CompTIA CySA+
Identify security flaws and vulnerabilities	Computer hardware and software systems; vulnerability management and exploitation tactics	GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT:

		Certified Expert Penetration Tester
Perform security audits, risk analysis, network forensics, and penetration testing	Exploitation, assessment and audit skillsets; technical writing; legal and compliance understanding; TCP/IP-based network communication	GCFE: GIAC Certified Forensic Examiner GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT: Certified Expert

		Penetration Tester
Perform desktop security assessments and update/patch potential vulnerabilities	Computer hardware and software systems; vulnerability assessments	GPEN: GIAC Certified Penetration Teste
		CEH: Certified Ethical Hacker
Develop a procedural set of responses to security problems	Operating system installation, patching, and configuration	CISSP: Certified Information Systems Security Professional
		CISM: Certified Information Security Manager
Establish protocols for communication within an organization and dealing with law enforcement during security incidents	Critical-thinking, project management, and communication skills	College degree
Create a program development plan that includes security gap assessments, policies, procedures, playbooks, training, and tabletop testing	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	CISSP: Certified Information Systems Security Professional

		College degree
Produce detailed incident reports and technical briefs for management, administrators, and end users	Critical-thinking, project management, and communication skills	College degree
Liaison with other cyberthreat analysis entities	Critical-thinking, project management, and communication skills	College degree
Handle case management duties of an incident and be involved with lessons-learned post-incident meetings	Case management experience and tools	CompTIA CySA+ CISM: Certified Information Security Manager College degree

Systems Analyst

A systems analyst is responsible for monitoring and interpreting different forms of data. Data can include logs from security tools, alerts from networking equipment, or other event data. A systems analyst might also be responsible for analyzing various types of artifacts, including files and programs, the goal being to determine whether there is any potential risk to the organization and discover the purpose of the artifact (meaning why it was created). For example, a word document might have a rootkit included, so the purpose of the document is to trick a user into running it and installing the rootkit.

Systems analysts that work in the incident management service spend time monitoring SIEM/SOAR/XDR systems, looking for potential threats within hundreds of thousands of event data points. A system analyst either addresses events directly or passes them to a member from the incident management service group. Systems analysts that work in the analysis service have isolated labs dedicated to containing potentially threating artifacts and learning what artifacts do. Common duties for analysts involved with the analysis service include performing static analysis, such as scanning or disassembling artifacts, and performing dynamic analysis, such as running artifacts in a sandbox to learn their behavior.

Table 4-5 outlines the responsibilities, skills, and certifications associated with the systems analyst role. A systems analyst is ideal for the analysis service or incident management service but can also work in the digital forensics and risk management services. Similar job titles include operations analyst, business systems analyst, business intelligence analyst, and data analyst.

TABLE 4-5 Systems Analyst Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Actively monitor systems and networks for intrusions	Windows, UNIX, macOS, and Linux operating systems	CCE: Certified Computer Examiner
Identify security flaws and vulnerabilities	Computer hardware and software systems; vulnerability management and exploitation tactics	GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified

		Penetration Tester CEPT: Certified Expert Penetration Tester
Perform security audits, risk analysis, network forensics, and penetration testing	Computer hardware and software systems; vulnerability management and exploitation tactics TCP/IP-based network communications	GPEN: GIAC Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT: Certified Expert Penetration Tester
Perform malware analysis and reverse engineering	Computer hardware and software systems	GCFA: GIAC Certified Forensic Analyst
Experience working with SIEM and SOAR orchestration and automation	DevOps and playbooks skills	Certification in DevOps
Reverse engineer/disassemble malware and other artifacts	Disassemblers, debuggers, and other static-analysis tools	GIAC Reverse Engineering Malware (GREM)
Develop sandboxes and analyze software behavior	Sandboxes and other dynamic analysis tools	GIAC Reverse Engineering

		Malware (GREM)
Analyze logs and other data sources	Security tool logs (firewall, IDS/IPS, etc.), SIEMs, and SOAR	CCNA Cyber Ops, CompTIA Cybersecurity Analyst (CySA+)
Liaison with other cyberthreat analysis entities	Forensic software applications (e.g. EnCase, FTK, Helix, Cellebrite, XRY, etc.)	CREA: Certified Reverse Engineering Analyst
Understand assembly language and how computer systems operate (RAM, ROM, storage, etc.)	IDA Pro, Ghidra, RAM/ROM dumps	GIAC Reverse Engineering Malware (GREM)

Security Administrator

A security administrator is responsible for managing IT-related security and safety issues within a company. Tasks can include developing policies and procedures as well as overseeing that policies are followed by employees. Security administrators also oversee the implementation of solutions that prevent cyberthreats and protect data's confidentiality, integrity, and availability. Tasks include administering security controls to reduce the risk associated with potential vulnerabilities.

Table 4-6 outlines the responsibilities, skills, and certifications associated with the security administrator role. Security administrators are ideal for compliance, risk management, and situational and security awareness services. Similar job titles include security manager, information security manager, network security administrator, systems security administrator, information systems security officer, and IT security administrator.

TABLE 4-6 Security Administrator Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications

Protect systems against unauthorized access, modification, and/or destruction	Windows, UNIX, and Linux operating systems; system security capabilities	CompTIA Security+ (popular base-level security certification)
Perform vulnerability and networking scanning	Computer hardware and software systems; vulnerability management and exploitation tactics	CCNA: Cisco Certified Network Associate
	TCP/IP-based network communications	CEH: Certified Ethical Hacker
Monitor network traffic for unusual or malicious activity	Strong understanding of firewall technologies	ECSA: EC-Council Certified Security Analyst
		CompTIA CySA+
Configure and support security tools such as firewalls, antivirus software, and patch management system	TCP/IP, computer networking, routing and switching	CISSP: Certified Information Systems Security Professional
Implement network security policies, application security, access control, and corporate data safeguards	Network protocols and packet analysis tools	CISM: Certified Information Security Manager
		CISSP:

		Certified Information Systems Security Professional
Train employees in security awareness and procedures	Critical-thinking, project management, and communication skills	College degree
Perform security audits and make policy recommendations	Intermediate to expert IDS/IPS knowledge; vulnerability evaluation; security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.).	CISSP: Certified Information Systems Security Professional
		College degree
Develop and update business continuity and disaster recovery protocols	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	College degree

Security Engineer

This role is similar to a security analyst, with responsibilities of performing security monitoring, security and data/log analysis, and forensic analysis. The goal of this role is to detect security incidents and launch a response. A security engineer can also have responsibilities for identifying which security technologies are used by an organization, maintenance of existing security technologies, development and maintenance of security policy, and developing methods to improve policies.

Table 4-7 outlines the responsibilities, skills, and certifications associated with the security engineer role. A security engineer can work in the incident

management, analysis, digital forensics, and R&D services, depending on the specific skills and experience the engineer has acquired. Similar job titles include security analyst, security administrator, security architect, security specialist, and security consultant.

TABLE 4-7 Security Engineer Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Configure and install firewalls and intrusion detection/prevention systems	IDS/IPS, penetration testing, and vulnerability testing	CISM: Certified Information Security Manager CISSP: Certified Information Systems Security Professional CEH: Certified Ethical Hacker
Perform vulnerability testing, risk analyses, and security assessments	Firewall and intrusion detection/prevention protocols	CCNP Security: Cisco Certified Network Professional Security CEH: Certified Ethical Hacker
Develop or work with automation scripts to handle and track	Secure coding practices, ethical hacking, and threat modeling	GSEC: Security Essentials

incidents		
		GCIH: GIAC Certified Incident Handler GCIA: GIAC Certified Intrusion Analyst
Investigate intrusion incidents, conduct forensic investigations, and launch incident responses	Windows, UNIX, macOS, and Linux operating systems	CISSP: Certified Information Systems Security Professional CompTIA CySA+ CCFE: Certified Computer Forensics Examiner
Collaborate with colleagues on authentication, authorization, and encryption solutions	Critical-thinking, project management, and communication skills; encryption technology concepts	Systems Security Professional College degree
Evaluate new technologies and processes that enhance security capabilities	Critical-thinking, project management, and communication skills	College degree

Deliver technical reports and formal papers on test findings	Communication and technical writing skills	College degree
Supervise changes in software, hardware, facilities, telecommunications, and user needs	Critical-thinking, project management, and communication skills	College degree
Define, implement, and maintain corporate security policies	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	CISSP: Certified Information College degree
Analyze and advise on new security technologies and program conformance	Critical-thinking, project management, and communication skills	College degree
Recommend modifications in legal, technical, and regulatory areas that affect IT security	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	CISSP: Certified Information CISM: Certified Information Security Manager Systems Security Professional College degree

Security Trainer

A security trainer is responsible for implementing standardized training programs based on the organization's policies and the current threat landscape. Security trainers develop and schedule training needs based on feedback from interviewing leadership and employees. Responsibilities include developing the training material, coordinating and monitoring enrollment, schedules, costs, and equipment, and delivering training metrics to leadership. Other duties include researching industry training concepts, training people to deliver training content, and updating content as needed.

Table 4-8 outlines the responsibilities, skills, and certifications associated with the security trainer role. A security trainer is ideal for the situational and security awareness service but can also work in the risk management and R&D service groups. Similar job titles include training instructor, information assurance analyst, training analyst, security service training manager, and security training and development manager.

TABLE 4-8 Security Trainer Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Develop a schedule to assess training needs	Experience with technologies and best practices for instructional manuals and teaching platforms	Certification from talent and training associations
Ensure strict adherence to company philosophy/mission statement/sales goals	Understanding policies, procedures, and industry guidelines, standards, and frameworks	CISSP: Certified Information Systems Security Professional
Deliver training to customers or other trainers	Excellent verbal and written communication skills	College degree
Manage security awareness program based on threat research	Strong project management skills with the ability to supervise multiple projects	College degree
Deliver technical reports	Identity and access	College degree

and formal papers on test findings	management principles	
Test and review created materials	Critical-thinking, project management, and communication skills	College degree
Maintain a database of all training materials	Basic database and program management skills	College degree

Security Architect

A security architect oversees the implementation of network and computer security for an organization. This role is typically a senior-level employee responsible for creating security structures, defenses, and responses to security incidents. Additional responsibilities may include providing technical guidance, assessing costs and risks, and establishing security policies and procedures for the organization.

Table 4-9 outlines the responsibilities, skills, and certifications associated with the security architect role. The security architect is ideal for the risk management service but can be part of other services such as compliance, situational, and security awareness, and research and development. Similar job titles include information security architect, IT security architect, and senior security analyst.

TABLE 4-9 Security Architect Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Plan, research, and design robust security architectures for any IT project	Risk assessment procedures, policy formation, role-based authorization methodologies, authentication technologies, and security attack concepts	CISSP: Certified Information Systems Security Professional
Perform vulnerability	Computer hardware and software	GPEN: GIAC

testing, risk analyses, and security assessments	systems; vulnerability management and exploitation tactics	Certified Penetration Tester CEH: Certified Ethical Hacker OSCP and PEN-200 from offensive security CPT: Certified Penetration Tester CEPT: Certified Expert Penetration Tester
Research security standards, security systems, and authentication protocols	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	CISM: Certified Information Security Manager CISSP: Certified Information Systems

		Security Professional
Develop requirements for LANs, WANs, VPNs, routers, firewalls, and related network devices	Security controls such as firewall, IDS/IPS, network access control, and network segmentation	CISM: Certified Information Security Manager
Design public key infrastructures (PKIs), including use of certification authorities (CAs) and digital signatures	Security and encryption technologies	CISM: Certified Information Security Manager EC-Council Certified Encryption Specialist (ECES)
Review and approve installation of firewall, VPN, routers, IDS/IPS scanning technologies, and servers	Security concepts related to DNS, routing, authentication, VPN, proxy services, and DDOS mitigation technologies	GSEC: GIAC Security Essentials GCIH: GIAC Certified Incident Handler GCIA: GIAC Certified Intrusion Analyst
Provide technical supervision for security team(s)	Critical-thinking and communication skills	College degree

Define, implement, and maintain corporate security policies and procedures	Network security architecture development and definition	CISSP: Certified Information Systems Security Professional College degree
Oversee security awareness programs and educational efforts	Critical-thinking and communication skills	College degree
Update and upgrade security systems as needed	Windows, UNIX, macOS, and Linux operating systems	A+ Security CISSP: Certified Information Systems Security Professional

Cryptographer/Cryptologist

A SOC that uses encryption to secure information or to build a system will assign these requirements to a cryptologist. A cryptologist researches and develops stronger encryption algorithms. A cryptologist may also be responsible for analyzing encrypted information from malicious software to determine the purpose and functions of the software.

Table 4-10 outlines the responsibilities, skills, and certifications associated with the cryptographer/cryptologist role. Cryptologists are ideal for digital forensics and analysis services but can work in other services based on the need for implementing, understanding, or identifying crypto.

TABLE 4-10 Cryptographer/Cryptologist Responsibilities, Skills, and

Certifications

Responsibilities	Skills	Certifications
Protect information from interception, copying, modification and/or deletion	Computer architecture, data structures, and algorithms	The cryptologist field is new and only has programs in universities and special learning programs. Certification programs include cryptology aspects, but dedicated certifications are not available at this point in time.
Evaluate, analyze, and target weaknesses in cryptographic security systems and algorithms	Linear/matrix algebra and/or discrete mathematics	EC-Council Certified Encryption Specialist (ECES)
Develop statistical and mathematical models to analyze data and solve security problems	Probability theory, information theory, complexity theory, and number theory	EC-Council Certified Encryption Specialist (ECES) College degree in math and cryptologist certification
Investigate, research, and test new cryptology theories and applications	Principles of symmetric cryptography and asymmetric cryptography	EC-Council Certified Encryption Specialist (ECES) College degree in math and cryptologist certification
Probe for weaknesses in communication lines	Principles of symmetric cryptography and asymmetric cryptography	EC-Council Certified Encryption Specialist (ECES) College degree in math and cryptologist certification
Ensure financial data is securely encrypted and	Network Access Control concepts	Operating system certifications

accessible only to authorized users	Data loss prevention technologies, encryption concepts, identity management, access control	Vendor security certifications Authentication vendor certifications
Ensure message transmission data is not illegally accessed or altered in transit	Principles of symmetric cryptography and asymmetric cryptography	EC-Council Certified Encryption Specialist (ECES) College degree in math and cryptologist certification
Decode cryptic messages and coding systems for military, political, and/or law enforcement agencies	Principles of symmetric cryptography and asymmetric cryptography	EC Council Computer Hacking Forensic Investigator Certification College degree in math and cryptologist certification
Advise colleagues and research staff on cryptical/mathematical methods and applications	Principles of symmetric cryptography and asymmetric cryptography	College degree in math and cryptologist certification

Forensic Engineer

Many organizations will experience a breach, and they will need to understand how the breach occurred. Digital forensics is the art of collecting evidence regarding a security incident. Evidence can be used for legal actions, to remediate the vulnerability used to cause the breach, or as part of a

lessons-learned exercise. Forensic engineers require specific skillsets focused on collecting data without creating changes to what they are collecting. These engineers may also have legal knowledge to assist with investigations that lead to legal actions.

Table 4-11 outlines the responsibilities, skills, and certifications associated with the forensics engineer role. This role is ideal for the digital forensics service but can also work in the analysis and incident management services. Similar job titles include forensic scientist, forensic consultant, and digital forensics engineer.

TABLE 4-11 Forensic Engineer Responsibilities, Skills, and Certifications

Responsibilities	Skills	Certifications
Conduct data breach and security incident investigations	Network skills, including TCP/IP-based network communications	CCE: Certified Computer Examiner
Recover and examine data from computers and electronic storage devices	Windows, UNIX, and Linux operating systems	CEH: Certified Ethical Hacker
Dismantle and rebuild damaged systems to retrieve lost data	Windows, UNIX, macOS, and Linux operating systems; digital forensics concepts	EnCE: EnCase Certified Examiner
Identify systems/networks compromised by cyberattacks	Computer hardware and software systems	GCFE: GIAC Certified Forensic Examiner
Compile evidence for legal cases	Operating system installation, patching, and configuration	GCFA: GIAC Certified Forensic Analyst
Draft technical reports, write declarations, and prepare evidence for trial	Backup and archiving technologies; technical writing	GCIH: GIAC Certified Incident Handler

Give expert counsel to attorneys about electronic evidence in a case	Cryptography principles; legal experience; digital forensics experience; strong communication skills	CCFE: Certified Computer Forensics Examiner
Advise law enforcement on the credibility of acquired data	eDiscovery tools; strong communication skills	CPT: Certified Penetration Tester
Provide expert testimony at court proceedings	Forensic software applications (e.g. EnCase, FTK, Helix, Cellebrite, XRY, etc.)	CREA: Certified Reverse Engineering Analyst
Stay proficient in forensic, response, and reverse engineering	Data processing skills in electronic disclosure environments	CCFE: Certified Computer Forensics Examiner
		College degree

Chief Information Security Officer

Also called a CISO, this role is part of high-level management and is positioned as the person responsible for the entire information security division of an organization. A CISO is responsible for all assurance activities related to the availability, integrity, and confidentiality of customer, business partner, employee, and business information in compliance with the organization's information security policies. A CISO works with executive management to determine acceptable levels of risk for the organization.

Table 4-12 outlines the responsibilities, skills, and certifications associated with the CISO role. It is common for the CISO to be responsible for the risk management service but can also oversee all other SOC services.

TABLE 4-12 Chief Information Security Officer Responsibilities, Skills, and

Certifications

Responsibilities	Skills	Certifications
Appoint and guide a team of IT security experts	Practices and methods of IT strategy, enterprise architecture, and security architecture	CISA: Certified Information Systems Auditor
Create strategic plan for the deployment of information security technologies and program enhancements	Security concepts; critical-thinking and communication skills	CISM: Certified Information Security Manager
Supervise development of corporate security policies, standards, and procedures	ISO 27002, ITIL, and COBIT frameworks	GSLC: GIAC Security Leadership College degree
Integrate IT systems development with security policies and information protection strategies	PCI DSS, HIPAA, NIST, GLBA, and SOX compliance assessments	CCISO: Certified Chief Information Security Officer
Collaborate with key stakeholders to establish an IT security risk management program	Network security architecture development and definition	CGEIT: Certified in the Governance of Enterprise IT
Anticipate new security threats and stay up to date with evolving infrastructures	Knowledge of third-party auditing and cloud risk assessment methodologies	CISSP: Certified Information Systems Security Professional
Develop strategies to handle security incidents and coordinate investigative activities	Critical-thinking and communication skills	CISSP-ISSMP: CISSP Information Systems Security

		Management Professional
Act as a focal point for IT security investigations	Critical-thinking and communication skills	CISSP: Certified Information Systems Security Professional
		College degree
Prioritize and allocate security resources correctly and efficiently	Critical-thinking and communication skills	College degree
Prepare financial forecasts for security operations and proper maintenance coverage for security assets	Critical-thinking and communication skills; contract experience	College degree
Work with senior management to ensure IT security protection policies are being implemented, reviewed, maintained, and governed effectively	Security frameworks (e.g., ISO 27001/27002, NIST cybersecurity framework, etc.); critical-thinking, project management, and communication skills	College degree

Every job role you recruit for will have an associated learning curve to onboard an employee into your SOC environment. Every SOC has its own unique networks, processes, and capabilities that can only be taught while in the job role. The next section looks at role tiers to better understand how job titles can change as employees gain experience and knowledge.

I opened this section with the caveat that a wide variety of different names are used for similar job roles. What you believe a security analyst does, for example, may be different from what others think that job role entails. To help standardize job role concepts, next I'll cover a U.S. government guide regarding responsibilities associated with cybersecurity industry jobs.

NICE Cybersecurity Workforce Framework

The previous section defined SOC roles found in SOCs around the world. Another approach (among many) to exploring these roles and alternative names for them is the U.S. government resource known as the National Initiative for Cybersecurity Education (NICE) Cybersecurity Workforce Framework (NICE Framework). I include this reference as an alternative to how I see job roles within the SOC, since different people will interpret job titles differently.

The NICE Framework is part of the Cybersecurity and Infrastructure Security Agency's National Initiative for Cybersecurity Careers and Studies (NICCS) and is described on the NICCS website as "a nationally focused resource that establishes a taxonomy and common lexicon to describe cybersecurity work, and workers, regardless of where, or for whom, the work is performed." You can use the NICE Framework to develop job requirements for recruiting, to prepare questions for interviewing potential candidates, and to get an idea of the skills associated with common cybersecurity job tiles. The rest of this section describes how to drill down to specific job roles on the NICE Framework web page at https://niccs.cisa.gov/workforce-development/cyber-security-workforce-framework.

Nice Framework Components

The NICE Framework is composed of the following components:

- Seven categories representing a high-level grouping of common cybersecurity functions
- Thirty-three Specialty Areas representing distinct areas of cybersecurity work
- Fifty-two Work Roles representing the most detailed groupings of cybersecurity work and composed of specific knowledge, skills, and abilities (KSAs) required to perform tasks in a Work Role

Figure 4-1 shows the seven categories of the NICE Framework as presented on the NICSS website. Notice that the description for each category focuses

on the type of work from a high level regarding the type of skillsets people have that work within the category's field of focus. The descriptions are developed this way to accommodate multiple specific skillsets that may fall under a more generic category. For example, suppose I need an analyst for my incident management SOC service and I want to identify specific job requirements for purposes of recruiting an analyst. I would start with the Protect and Defend category based on the description "Identifies, analyzes, and mitigates threats to internal information technology (IT) systems and/or network" that indicates people in this category have skills in evaluating and responding to events based on security logs or other event logs, which is what incident management is all about. Categories are outcome focused, meaning the field of work, so I would need to drill down deeper to identify associated job roles.

FIGURE 4-1 NICE Framework Seven Categories

To better understand the job skills in the Protect and Defend category, I can

click the category's Specialty Area button. Figure 4-2 shows the Protect and Defend category and its four Specialty Areas. Because I am looking for a description of the skills of an analyst for my incident management service, I can narrow down the Specialty Areas to two of the four based on their descriptions: Cyber Defense Analysis and Incident Response. I believe the Vulnerability Assessment and Management Specialty Area could also be useful but would be more relevant to the vulnerability management service than the incident management service for which I need to recruit an analyst. The Incident Response role would be the best choice, but the Cyber Defense Analysis could also do the job based on the number of similar skills as seen with an Incident Response job role. In order to see the specific skills associated with a job role, I will need to click into that role.

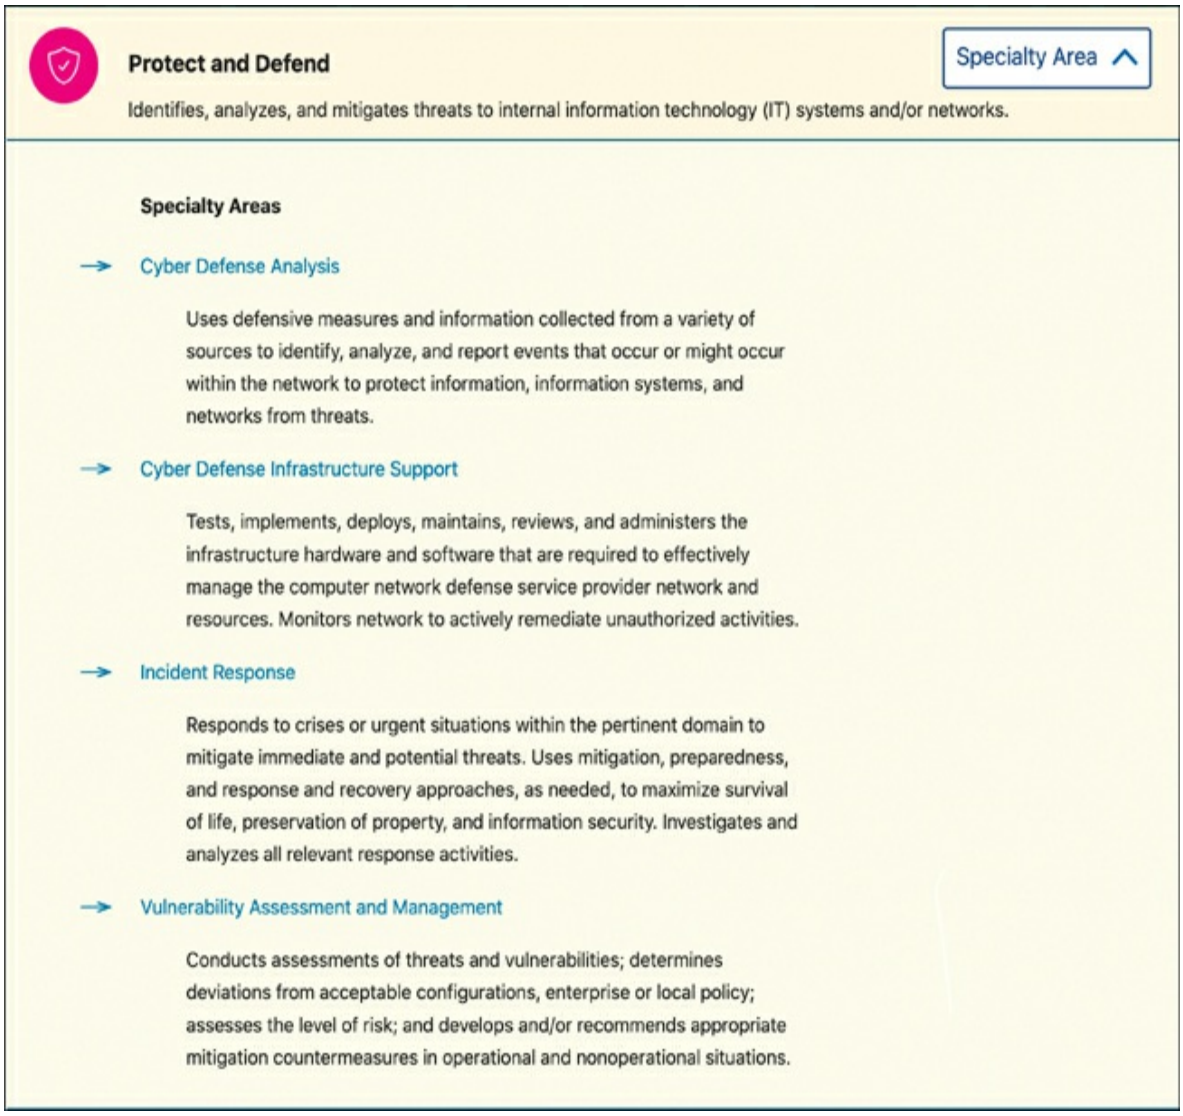

FIGURE 4-2 NICE Framework Protect and Defend Category with Four Specialty Areas

Next, I'll go with my first pick, which is Incident Response specialty area. To see the details of a specialty area, I click the specialty area to bring up the Work Role details. Figure 4-3 shows some of the details of the Cyber Defense Incident Responder Work Role, including a description of the role and the required abilities. As Figure 4-3 indicates, details regarding the knowledge, skills, and tasks of a Cyber Defense Incident Responder can be displayed by clicking the drop-down arrows. The language used by NICE to explain the job role is much more specific, allowing a better understanding of what tasks this type of employee would be expected to know how to do.

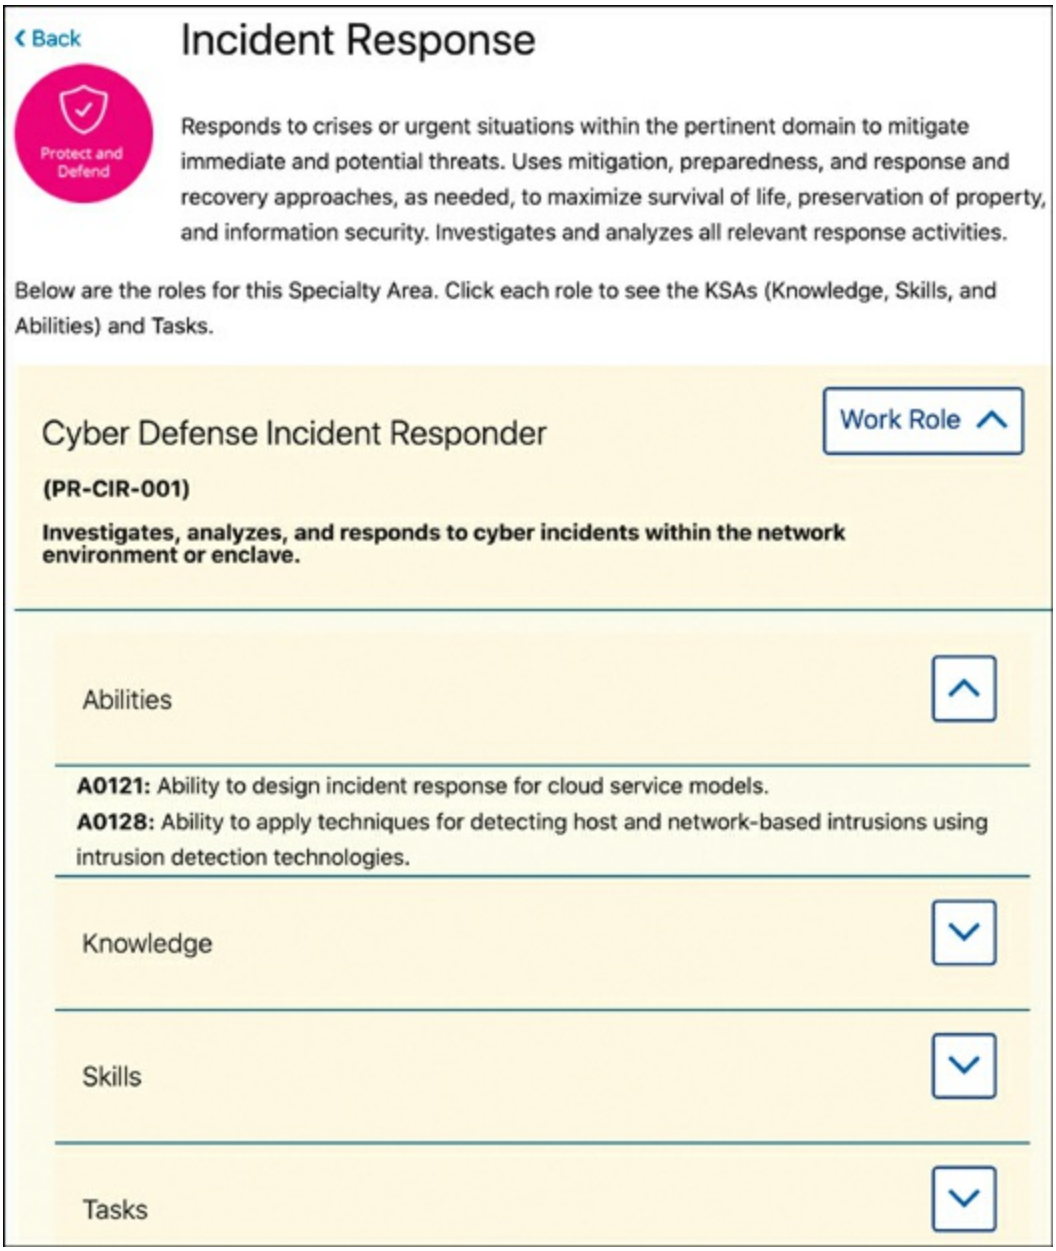

FIGURE 4-3 NICE Framework Cyber Defense Incident Responder Work Role Details

Clicking the Knowledge tab in the Incident Responder job role reveals tons of knowledge concepts, as shown in Figure 4-4. These concepts can be extremely useful when creating a job profile for the candidate you plan to recruit for. In Chapter 3, I pointed out that many SOC managers who are responsible for starting a new SOC service don't know what skills they will need until the service goes live, making it challenging to develop a job profile

for a service before it exists. Using the NICE Framework not only can help you develop requirements for job roles based on industry trends but also provides you with a validation point for the type of job titles you should seek out based on what the NICE Framework lists as expected skills associated with a job title.

K0001: Knowledge of computer networking concepts and protocols, and network security methodologies.
K0002: Knowledge of risk management processes (e.g., methods for assessing and mitigating risk).
K0003: Knowledge of laws, regulations, policies, and ethics as they relate to cybersecurity and privacy.
K0004: Knowledge of cybersecurity and privacy principles.
K0005: Knowledge of cyber threats and vulnerabilities.
K0006: Knowledge of specific operational impacts of cybersecurity lapses.
K0021: Knowledge of data backup and recovery.
K0026: Knowledge of business continuity and disaster recovery continuity of operations plans.
K0033: Knowledge of host/network access control mechanisms (e.g., access control list, capabilities lists).
K0034: Knowledge of network services and protocols interactions that provide network communications.
K0041: Knowledge of incident categories, incident responses, and timelines for responses.
K0042: Knowledge of incident response and handling methodologies.
K0046: Knowledge of intrusion detection methodologies and techniques for detecting host and network-based intrusions.
K0058: Knowledge of network traffic analysis methods.
K0062: Knowledge of packet-level analysis.
K0070: Knowledge of system and application security threats and vulnerabilities (e.g., buffer overflow, mobile code, cross-site scripting, Procedural Language/Structured Query Language [PL/SQL] and injections, race conditions, covert channel, replay, return-oriented attacks, malicious code).
K0106: Knowledge of what constitutes a network attack and a network attack's relationship to both threats and vulnerabilities.
K0157: Knowledge of cyber defense and information security policies, procedures, and regulations.
K0161: Knowledge of different classes of attacks (e.g., passive, active, insider, close-in, distribution attacks).
K0162: Knowledge of cyber attackers (e.g., script kiddies, insider threat, non-nation state sponsored, and nation sponsored).

FIGURE 4-4 NICE Framework Cyber Defense Incident Responder Knowledge Tab Details

I highly recommend using the NICE Framework if you don't know the type of skills a person needs to have to work for your SOC service. This same concept can apply as you develop interview questions for potential candidates.

Role Tiers

Roles within each SOC service can be broken down into different tiers or skill levels, which signify associated responsibilities. For example, a first-tier SOC analyst may be responsible for detecting, identifying, and troubleshooting security events that come into the SOC. Often this is the tier that communicates with the affected party. Responsibilities include detection, classification, and escalation of events. A second-tier analyst may have mitigation responsibilities over any event escalated by a first-tier SOC analyst. If the event requires even further support, a more-experienced third-tier analyst may be involved to remediate the situation. The third-tier analyst might also build tools and processes to improve capabilities within the SOC, including the processes followed by lower-tier analysts. Higher-tier roles have higher compensation but require deeper technical skills and experience. The same tiered approach can apply to other job roles with SOC services, such as a first-tier developer handling basic coding while a higher-tier developer would have responsibilities over the project's direction.

Each job role you create for your SOC should have a tier structure to promote career growth. A pay scale should also be assigned to each tier of a job role to inform employees what the expected compensation range is with an associated job role. The specific requirements advertised for the job role that includes the associated tier can reference lower tiers along with including the additional experience and skills needed to be considered for the higher-tier job role. Using this structure not only weeds out candidates that do not have the associated skills for the job tier being requested but opens the door for those same candidates to consider a lower tier of the same job role that might be more appropriate for their skill and experience level. For example, a SOC

might post open job roles for multiple analyst jobs at different tier levels. A candidate who interviews for the tier 3 analyst role might be informed that he is not qualified for that role but should consider applying for a tier 1 or 2 analyst role, with the goal of eventually gaining the experience to be promoted to a tier 3 analyst. Using this approach will provide direction for career growth, open your recruiting efforts to more candidates, and keep expectations for hiring and promotions clear to all employees.

It is important to validate industry pay ranges and experience expectations against any job role you create as well as the tier you associate the role with. With publicly available sources of pay ranges online, job candidates have expected pay ranges for specific job titles. The same expectation applies to associated tiers with a job title. For example, job and recruiting website Glassdoor estimates an average base pay for a tier 1 analyst at $77,665 per year USD, while an experienced analyst salary range increases to $99,898 per year USD. Aligning with industry trends for pay ranges will reduce the risk of not capturing quality candidates as a result of not advertising acceptable pay scales in your job posting. You should apply similar research to expectations for skills and experience. Online employment resources such as Glassdoor and Monster not only provide expected pay ranges for job roles, but also suggest years of experience in the role and a generic view of expected skills. Use these expectations as you list out what requirements for skills and experience are needed for your job roles, keeping in mind that your job roles will be based on the services provided by your SOC and will be different than a generic industry explanation of a job title and associated experience tier.

Note

Other factors, such as geographical considerations, can impact the salary for a job role. For example, a candidate for a job in San Francisco will expect a much higher salary than a tier 1 analyst candidate in Des Moines.

SOC Services and Associated Job Roles

The roles and job skill requirements for your SOC will depend on the different services the SOC is responsible to deliver to its customers. People are required for any SOC service regardless of the type of technology being used. Even the value from advanced security analytics such as artificial intelligence boils down to how the SOC uses the technology. Software does not provide the answers to what problems your SOC faces; it provides the tools and delivers the data needed to discover answers. Essentially, people are needed to run technology and to interpret the results of the tools used in the SOC.

The following sections review how people relate to the different services that can be offered by a SOC. Each service will be made up of one or more job roles previously described in the chapter.

Risk Management Service

The risk management service is responsible for managing all aspects of risk to the organization. This includes analyzing risk, calculating the potential impact of risk, and making decisions based on the organization's risk appetite. Employees responsible for risk management must have great communication skills, enabling them not only to ensure that everybody in the organization understands any significant risk but also to explain the organization's risk management strategy. Working for the risk management service also requires a solid understanding of business, because decisions of the service will impact various internal and external elements of the organization. Successful employees responsible for risk management are skilled at negotiation and diplomacy. They can work under pressure and are able to modify strategies as various factors change the current state of the organization's risk status.

Possible job titles include chief information manager, chief information security officer, security officer, risk management analyst, and analyst.

Vulnerability Management Service

Successful employees responsible for vulnerability management have experience in and understanding of network and computer security. They can analyze hardware, software, networks, and communication to discover and address vulnerabilities. SOC members involved with vulnerability management have solid communication skills so they can explain identified vulnerabilities as well as work with various parties to validate findings, including third-party vendors and other external experts. Employees responsible for vulnerability management are detail-oriented, have strong problem-solving skills, and can adapt methods used to manage vulnerabilities based on the ever-changing threat landscape.

Possible job titles include penetration tester, vulnerability engineer, ethical hacker, red team tester, security analyst, and security engineer.

Incident Management Service

SOC employees responsible for incident management actively monitor systems and networks for intrusions. The incident management team develops a procedural set of responses to security problems and oversees their execution. This team is also responsible for restoring services back to a normal state following an incident as quickly as possible while minimizing the impact to business operations. Communication and diplomacy skills are required to produce incident reports and provide technical briefings to various parties about incidents in a diplomatic fashion. Employees are required to be able to work under pressure while coordinating all activities required to perform, monitor, and report on the incident management process.

Possible job titles include incident responder, security analyst, computer network defense, IT network defense, incident analyst, intrusion detection specialist, and network intrusion analyst.

Analysis Service

A security analyst is responsible for detecting and preventing cyberthreats to

an organization. Members of the analysis team review security logs from various types of devices and work with the team responsible for incident management when a threat is confirmed. In addition to dealing with real-time threats, the analysis team analyzes and responds to undisclosed hardware and software vulnerabilities when a dedicated vulnerability management team isn't present. The analysis team can also take on responsibilities as a security advisor and develop security strategy based on data captured and analyzed. Members of the analysis team must be analytical and detail-oriented with specific skills in understanding how devices generate logs and how to work with network and security tools that generate logs. Analysis engineers can also be responsible for analyzing and reverse engineering various types of artifacts, requiring a different set of analytical and technical skills than an analyst that works with security logs. Analysis engineers are technical, detail-oriented, and specialized in the types of data they are responsible for analyzing.

Possible job titles include security analyst, security engineer, security administrator, security specialist, security consultant, network engineer, operations analyst, business intelligence analyst, and data analyst.

Compliance Service

The most fundamental skill for employees responsible for compliance is the ability to deal with risk and conflict management. A compliance officer uses specific factors for scoring risk, which will be based on the requirements for the type of compliance being enforced. A compliance officer will encounter situations requiring explaining and defending their point of view to internal employees as well as external agencies such as regulators. Communication and analytical thinking are critical for this role as well as a willingness to learn, as the world of compliance is continuously changing. Other skills associated with successful members of the compliance team are being detailed-oriented, being capable of interpreting data, and having strong problem-solving skills.

Possible job titles include compliance officer, assessment officer, policy officer, and infosec officer.

Digital Forensics Service

Roles in digital forensics are technology-focused, requiring a desire to learn, deep analytical skills, and the ability to work with various technologies ranging from desktop computers to mobile devices. Digital forensics requires acute attention to details and a comprehension of cybersecurity fundamentals. Communication skills and an understanding of law and criminal investigation are important because the results from a forensic investigation might be used in court, in which case the investigator will be required to defend his or her work. Digital forensics requires working with different groups, from legal to technical, as well as tolerance for disturbing material that might be discovered during an investigation. Successful digital forensic engineers have experience in both legal and technical matters related to cybersecurity.

Possible job titles include forensic engineer, forensic scientist, forensic consultant, and digital forensic engineer.

Situational and Security Awareness Service

The key purpose of this service is to address the human element of security. The goal of the work performed by the situational and security awareness team is to change the behavior of employees so that they operate with security in mind, reducing their risk to the organization. Duties include everything regarding security awareness and developing an education program. Roles responsible for situational and security awareness require strong written and verbal communication skills. Members in this role must be able to interpret all industry regulations, standards, and compliance requirements as well as ensure that everybody understands the organization’s risk management strategy. Successful situational and security awareness officers can accomplish these goals using a positive and engaging approach, which includes creating a metrics framework that can effectively measure results of the program.

Possible job titles include security trainer, training instructor, information assurance analyst, training analyst, security service training manager, and development manager.

Research and Development Service

SOC members of the research and development service are responsible for researching, planning, and implementing new programs and protocols for the organization. Duties include market research, tracking costs related to the creation of new programs and protocols, and making decisions on which projects are worth investing in. This group also validates if current programs, procedures, and technology being used are up to date with current and advanced industry standards. Members in this role have project management experience, are able to manage a budget, and are detail-oriented and creative.

Possible job titles include researcher, threat researcher, threat analyst, analyst, security analyst, programmer, software developer, and DevOps engineer.

Soft Skills

Another important element that must be included in your job roles is a description of the required soft skills, or nontechnical capabilities. Soft skills are different from having the technical ability to perform a job role and are just as important as technical skills when considering candidates for your SOC. Let’s look deeper into the concept of soft skills.

Soft skills are a combination of people skills, social skills, communication skills, character and personality traits, career attributes, emotional intelligence, and other human-based factors. Identifying the ideal candidate for any job role must include considerations for your position’s soft skills along with the expected technical skills (also known as “hard skills”) to ensure a successful match is made. For example, if an employee is shy and can’t communicate well, he or she would not be ideal for a role that requires that responsibility. I see many companies make the mistake of promoting a person into a manager or team lead role just because that person has many years invested in the company or is a top performer in his or her current position. The soft skills associated with a manager are unique and require leadership attributes, which some employees will not have based on their personalities and social skills. Not considering soft skills when recruiting new people or promoting employees will lead to underperformance in your SOC.

Certain job roles in a SOC require mature soft skills. Any role that involves communication with executives, public relations, or legal parties requires brevity and clarity of communication in both digital and in-person communication. Soft skills must also include adjusting what is being communicated based on the impact it could have on the target audience. SOC roles that interact with executives must also include soft skills that can provide respectful pushback and constructive feedback when necessary.

Certain roles within a SOC are responsible for developing escalation procedures for events and executing those procedures when an event occurs. These types of SOC roles require soft skills for communication to ensure the accuracy of data that is provided as the escalation process occurs. Soft skills also include deciding when to escalate an event, how often the event should be escalated, and how to identify the severity of an incident. Mistakes in communication can cause a breakdown of the escalation process ranging from overlooking severe incidents to wasting resources on non-severe incidents.

Evaluating Soft Skills

What soft skills should you look for as you recruit candidates for your SOC? According to the LinkedIn article “Hiring Without These Critical ‘Soft Skills’ Is a Recipe for Disaster” by Lou Adler, creator of the Performance-based Hiring methodology, several key hiring mistakes that are not related to technical or soft skills contribute to failure. The first mistake is a mismatch between a manager’s style and the new hire’s need for management and coaching. Some employees will want guidance and will feel isolated if left alone, while other employees interpret guidance as micromanagement and will not approve of being continuously monitored and managed. It is important for a hiring manager to explain their management style and identify if candidates would be comfortable working in that type of environment. A simple question you could ask candidates to identify their desired management style is, “Are you more comfortable with a hands-on manager or a hands-off manager?” Essentially, you are asking potential new hires if they prefer having periodic interaction or continuous interaction with their direct manager. Experienced managers will be able to adjust their management

styles to how their direct reports want to communicate with their manager, the expectations for which can be set upfront during the interview process.

Another soft skill conversation hiring managers should have with potential candidates is about the pace of the organization and expected motivation factors to complete tasks. Organizations work at different speeds, sometimes putting pressure on people to meet specific timelines or encouraging people to work late hours. For example, some organizations may claim to work 9 a.m. to 5 p.m. business hours but frown upon people who leave right at 5 p.m. if work requirements are behind schedule. A hiring manager should be upfront during the interview about how aggressively work schedules are enforced.

People are accustomed to different types of communication styles and expectations. Mismatching communication expectations between a manager and employees can lead to misunderstandings and team underperformance. For example, some people view text messaging as real-time conversation that requires immediate response, while others treat text messages similarly to email, responding to incoming messages when time permits. In this example, somebody with expectations for real-time responses to text messages may interpret not receiving a prompt response as being ignored or as the receiver not wanting to respond, whereas the reality might be that the receiver of the text message believes that text messages should be treated like any other form of communication and prioritized based on importance. Communication style should be confirmed between the hiring manager and candidate, including how often communication should occur and what type of details should be communicated. Examples of reports and data expected to be delivered by employees to their direct manager are great items to go over with a potential candidate to identify if the candidate meets the required soft skills to complete the tasks.

SOC Soft Skills

Specific roles in the SOC have corresponding soft skills expectations, many of which were identified earlier in this chapter as I described skills involving how people communicate and work. I pointed out that some roles have strong organizational and operational skill requirements. Some roles require critical

thinking and problem solving. Roles involving interacting with team members require the ability to collaborate with others. Technical writing skills are needed for roles that create reports or develop training. Many of these skills are not developed through technical training but rather are gained through work experience or general education or are just part of a person's personality or natural abilities.

Many SOC managers and directors I speak with are less concerned about a new SOC member's knowledge of specific tools. SOC leaders want a new SOC member to have an understanding of underlying functions, systems, networks, and processes and be able to fit into the SOC culture. Along with a strong work ethic (discussed in the next section), soft skills are a critical evaluation point for many SOC roles. Soft skills tend to be more important than technical skills for many roles.

The following is a list of soft skills that I find are common in members of a SOC regardless of which service they provide. I recommend including these soft skills in job profiles when recruiting.

- **Problem solving:** Industry and market knowledge
- **Analytical skills:** Troubleshooting complex issues
- **Communication:** Business understanding
- **Negotiation and diplomacy:** Work under pressure
- **Detail-oriented:** Organizational skills
- **Teamwork:** Documentation and presentation

Security Clearance Requirements

In addition to the previously discussed hard and soft skill requirements, another factor to consider as you develop a job description is that some roles in the SOC may require certain levels of security clearance in order to have access to specific content. Security clearance can be mandated by the organization and/or by law and is a license issued by an agency, the head of a department, or a branch of the federal government. Many U.S. federal employees and many employees in the private sector are required to obtain

security clearance. The amount of time required to obtain any level of security clearance depends on different factors, but according to one source, Security Degree Hub (https://www.securitydegreehub.com), obtaining a U.S.-based security clearance on average takes six months to a year. During a clearance evaluation, various aspects of a candidate are verified, including their identity, where they were born, where they live, who lives with them, any previous or current financial troubles, or anything else that could represent a risk of granting the candidate enough trust for the specific level of clearance they are applying for.

Security clearances have different levels, which grant specific levels of access to classified content. Regarding the U.S. federal government clearance stages, there are three levels, corresponding to the potential impact data loss at that level could have on the government and associated parties:

- **Top Secret:** Highest level of classification. Exposure would cause "exceptionally grave danger."
- **Secret:** Second highest level of classification. Exposure would cause "serious danger."
- **Confidential:** Lowest level of classification. Exposure would cause "damage."

It is important to point out that the U.S. federal government has additional language and classification levels used in classified communities. Some Top Secret clearances indicate the employee has passed a Single Scope Background Investigation (SSBI). This means the employee needs Top Secret clearance and access to sensitive compartmented information (SCI) in order to do their work. This clearance is not the same as an employee granted Top Secret SCI, which represents a SCI program run by a specific agency. SCI programs can ask for additional validation, including polygraph examinations, as part of the screening process, but it is inaccurate to assume that all Top Secret SCI employees have had a polygraph or additional validation beyond what is required for a Top Secret clearance. The requirements for a SCI program are specific to the agency it is assigned to, meaning even if you have Top Secret clearance, you would not be granted access to any material deemed Top Secret SCI unless you have been granted SCI access by the specific agency behind the SCI program. If one SCI program grants an employee Top Secret SCI clearance to its agency's SCI,

that does not grant the same employee Top Secret SCI clearance access to any other agency's SCI.

> **Note**
>
> Learn more about the United States Security Clearances program at https://www.state.gov/m/ds/clearances/c10978.htm.

Countries in the European Union (EU) use a similar classification system known as the European Union Classified Information (EUCI) system. The EU approach breaks classified information into four levels. Like the U.S. classification system, each level is based on the potential impact data loss could have on the government and other associated parties.

- **Très Secret UE/EU Top Secret:** The unauthorized disclosure of this information could cause exceptionally grave prejudice to the essential interests of the EU or one or more of the member states.
- **Secret UE/EU Secret:** The unauthorized disclosure of this information could seriously harm the essential interests of the EU or one or more of the member states.
- **Confidentiel UE/EU Confidential:** The unauthorized disclosure of this information could harm the essential interests of the EU or one or more of the member states.
- **Restreint UE/EU Restricted:** The unauthorized disclosure of this information could be disadvantageous to the interests of the EU or one or more of the member states.

Certain groups, such as the General Secretariat of the Council (GSC), provide approval lists for the types of cryptographic products that can be used on certain levels of EUCI classified data. The same policies apply to people, process, and technology associated with EU classified information. Learn more about the EU classification system at https://www.consilium.europa.eu/en/.

The type of clearance your SOC or the organization protected by your SOC will or will not require will be based on the laws governing your organization

and the data it is associated with. In some situations, access to content can be granted while a clearance is being processed, known as being in an “interim status” or temporary status. Other times, the clearance process must be completed before access to protected content can be granted. Most security programs require a periodic reinvestigation after a specific length of time, which time will be shorter as the level of clearance is increased. You will need to validate requirements for clearance with somebody that specializes in security clearances, such as a security clearance officer, before you consider providing specialized clearance to any of your employees.

Pre-Interviewing

At this point, I have covered how to create a job role, the different types of roles that exist in the industry, the job roles associated with SOC services, and how both soft skills and technical skills (and perhaps security clearance) should be considered for a job role. You can use all of these factors to develop job requests for the positions that you need to fill as you launch new SOC services or grow existing SOC services. Now it is time to look at how to fill job roles in your SOC with the right people by executing a successful interviewing process.

You will want to create a filtering system to avoid wasting time interviewing unqualified candidates for any job role you are looking to fill. According to a study by ISACA, 57% of respondents note the lack of qualification of half of the candidates they have hired. This feedback translates to half of the candidates seen by ISACA’s survey were found to not be able to perform the skills advertised on their resume during the interview process! Qualifying skills is a critical step of the interview process and must be done for any skill required to perform the job you are looking to fill. Candidates will list anything on their resume, from how long they worked in a position to the type of work that they performed; however, it is up to you to validate whether the provided information is true. Make sure to do this early using a prescreening process that includes one or more knockout questions to filter out unqualified candidates.

Candidates can provide proof of their skills through certifications and degrees, which might or might not be current, valid, or completed. Verifying

industry-recognized certifications and degrees from accredited universities will be easy and can be done by visiting the provider's website or using a validation service as long as you have the candidate's full name, certification number, and date of graduation if applicable. For example, you can consult the National Student Clearinghouse (https://www.studentclearinghouse.org/) to verify a degree from an accredited school was obtained by a candidate. Verifying certifications and degrees can also be used as part of the knockout process.

Note

A certification does not mean a skillset exists! Certifications show the required skills were validated at a specific point in time. Skillsets must be practiced or they are lost. Many recertification programs do not use the same rigor as the original certification, meaning a recertification date would not reflect the same skills existed as when the original certification was achieved. Some certifications can be cheated through the use of brain dumps, which publish the answers to the exams required to achieve a certification. For all of these reasons, it is important to use your own validation system to verify skills exist rather than depending on an external certification program.

Verifying work experience can be more challenging based on what is provided as a reference point. Things will change over time, including the status of people who worked with the candidate and the status of the organization the candidate worked at, sometimes causing a reference to no longer be available. Some candidates will also ask that you not contact their current employer until an offer is provided, prohibiting any validation of their current skillsets. If you can't speak with the direct manager of a current candidate, ask the candidate if you can contact a coworker or other party that can validate the skills you are looking for in your potential candidate.

Note

It is important to be mindful that some of your strongest candidates

will have skills that are represented on the resume in their job experience and not through education or certification programs. I have worked with very capable engineers who do not have a high school degree or certifications. I have also encountered unqualified candidates who have listed dozens of certifications and years of experience on their resume. Make sure to use your own skill assessment process when validating skills.

Avoid using language during any job postings or during a live interview that includes preference for a particular gender, race, age, religion, or other such status. For example, posting “we are looking for a *young* and *energetic* team member” would suggest age discrimination based on the use of “young” and “energetic.” You can highlight your organization’s view of providing an unbiased recruiting process externally by stating you are an “equal opportunity employer” or stating “nothing in the job posting or description should be construed as an offer or guarantee of employment” in your job posting and during a live interview. Keeping your hiring process unbiased will not only attract a diverse pool of candidates but also help avoid unwanted legal matters in regard to violating people’s rights.

Interviewing

Once you have created and posted requirements for your SOC role to be filled, you will need to evaluate the potential candidates. The initial conversation can be a phone call, video conference, or web chat. The focus of the first interview is to exchange information about what is being offered by both the recruiter and the potential candidate to see if a potential match exists. According to Monster, a common mistake made by hiring managers is spending too much time describing or “selling” the position. It is important to also spend time listening to candidates so that you can assess their qualifications, skills, and personal characteristics. Not doing this leads to wasting time with follow-up interviews with candidates that are interested in the opportunity but not qualified or not a good match for the role. It is ideal to include a member of the team that has the job role opening to assist with the

interview process, not only to help validate that the candidate's skills are a fit but also to look for potential team chemistry. Candidates' answers regarding specific qualifications or skills should be assessed by experts in those areas to ensure candidates are properly evaluated. Lastly, ensure that any special constraints related to the role are covered upfront, such as required travel or potential overtime.

Interview Prompter

One tool that can be used to standardize the questions delivered during the interview is an interview prompter template specific to the job role. Questions within the prompter can be developed and validated by internal team members and experts in the associated technology prior to the prompter's usage. Experts can also be used later to review the responses that are provided by candidates during the live interviews.

An interview prompter template can include questions about the following:

- General skills
- Specific technical skills
- Educational background
- Years of experience and what the experience involves
- Details about past projects and job roles
- Work the candidate enjoys and doesn't enjoy being involved in
- Career and personal goals
- Limitations and constraints, including salary and overtime availability
- If employed, reason for leaving their current role and considering this role
- Availability to start
- Descriptions about the role
- Overview of the position
- Describe the team

- Company business and culture
- Company benefits
- Compensation system
- Associated projects and expectations

> **Note**
>
> The last seven questions are focused on selling the SOC position while the first questions are designed to learn about the candidate.

The interview prompter template is very helpful for organizing questions, but asking questions in the specific order listed in the prompter isn't required. It is common for an interview to start with the interviewer providing an overview of the opportunity and then letting the conversation flow naturally from topic to topic as questions are asked by either the interviewer or interviewee. The interviewer can check off items on the interview prompter to ensure that all topics are covered within the interview time slot regardless of the order in which the answers are obtained. The prompter also helps ensure that the interviewer covers required topics within the allocated time for the interview using the task checkoff process.

Post Interview

After first-round interviews are conducted, qualified candidates might be asked for a face-to-face follow-up interview. Among the purposes of the second interview are to enable the candidate to meet with the team members or direct manager, to permit the candidate to assess the environment they would be working in if hired, and to have the candidate perform additional skill tests. Skill tests can include hands-on work with tools or applications, logical exams, or other methods to validate the expected knowledge and skills meet what is required to perform the job role. If both parties remain interested after the second interview, the hiring manager should provide a target date for a formal decision regarding whether the candidate will be

offered the position. The offer can also occur at the end of the second interview and be verbal if time is required to develop a formal draft of the offer. There may be other circumstances that would postpone a formal offer, such as ensuring the candidate meets substance testing requirements before being formally offered a position.

When developing a formal offer letter, make sure all details are clearly defined. This includes the position, expected tasks, total compensation package, and start date. The offer letter should include the name of the new hire's immediate manager and any additional document(s) that must be brought in on the first day. It is standard practice for the human resources department to develop and provide the offer letter to the new employee rather than the recruiter or hiring manager.

After providing an offer letter, the next stage of the hiring process is onboarding the new employee.

Onboarding Employees

Once a job role is filled, the hiring manager will need to prepare to bring the new employee into the job role. This process is also called *onboarding* the new hire. It is critical to properly prepare for a new employee, both to ensure that the new employee's time isn't wasted waiting beyond the designated day to start work and to ensure that the new employee has a positive first impression of the new job. A new employee will be frustrated if he or she arrives the first day ready to start working but doesn't have a workspace and computer allocated—basic essentials which should be prepared before the new employee arrives. The following list are requirements a hiring manager needs to prepare prior to the new employee's arrival. Provisioning of these items can be done by other team members such as desktop support and human resources, but it is the overall responsibility of the hiring manager to ensure these items are available prior to the arrival of the new employee.

- Allocated physical space within facilities, such as a desk and chair if applicable to the role
- Expected office supplies

- Computing equipment
- Employee identification and credentials such as telephone numbers, user IDs, and passwords
- Special software or tools
- Scheduling of education or overview of job role, if necessary
- Printed or electronic documents on processes, policies, methodologies, and other items relevant to the job role

The hiring manager also needs to prepare other internal team members for the arrival of the new employee. Information such as the background of the new hire should be shared with the direct team. Additional information such as personal interests can be shared to promote a positive chemistry, if disclosure of those details is authorized by the new hire prior to his or her arrival. Skills and duties associated with the job role should be shared and validated with the direct team so expectations for the new hire are clear to everybody.

Onboarding Requirements

Certain SOC roles will have specific onboarding requirements. Those requirements can include obtaining authorization to access sensitive resources, learning existing processes, attending training for new hires, and signing off on required compliance documentation. Some training might involve shadowing employees, with the goal of switching from a monitor to an interactive role as the new hire learns skills and processes. For example, a new hire might be assigned to monitor the incident handling procedure the first month on the job or work on a fake incident before being responsible for interacting with a real incident. Senior team members can review how a fake incident is handled by a new hire and provide coaching and reference to procedures as the new hire transitions into an operational role.

SOCs that follow industry guidelines should have new hires study the guidelines relevant to the job role. An example is having a new hire who is part of the incident response program first review the NIST 800-61 (Rev. 2) *Computer Security Incident Handling Guide* or the FIRST PSIRT Services Framework (introduced in Chapters 1 and 3, respectively). Required reading

can also be provided before a new hire starts, which the hiring manager could reference and even quiz the new hire about to ensure learning objectives were achieved the first week they started the new role. Expectations for this material can be shared as part of the expected onboarding process as a new hire's first few weeks schedule is developed by the hiring manager.

It is critical to ensure a smooth transition into a position for any new hire. The first few weeks will determine if the candidate is a fit for the role and will be capable of handling the associated responsibilities. Part of creating a welcoming environment for employees is properly setting career expectations.

Managing People

Failure to properly manage people will lead to a SOC whose employees are a flight risk, ready to leave for another organization if a better offer comes along. The current IT market is strong, and it will take effort to retain top talent. Great managers know what drives the people who report to them and act as an enabler for those goals. Career-driven people are not focused only on how much money they make. Table 4-13 lists the top five things that make employees happy at work and the top five reasons why employees are not happy and eventually leave a position. This data comes from Monster and BioSpace.

TABLE 4-13 What Makes a Happy or Disgruntled Employee

Happy	Disgruntled
Feel accomplished	Are disengaged
Receive positive reinforcement	Are stressed out
Like their co-workers	Have a negative mindset
Have some level of autonomy	Have poor relationships with managers and colleagues
Are proud of what they are	Not fully using their intellect or strengths

As a SOC manager, you will want to identify what motivates each of your SOC employees as well as help guide which future position and goals would be most ideal for them to target. This includes identifying that an employee is working a stepping-stone position with the goal of taking on a more senior role once they acquire the appropriate skills and experience. Goals should be documented in an employee development plan and must benefit both the organization and the employee. Before setting goals for an employee, consider the business goals and how that employee aligns to short-term and longer business objectives. Make sure to consider whether certain roles need to be filled in the future and, if so, whether this employee could be groomed for that needed role. Having a business goal aligned with an employee goal helps justify investment in training and experience so that both parties benefit from the promotion.

Once business goal alignments are identified with potential employee goals, speak with the employee and confirm career aspirations. When a career goal is confirmed by the employee that aligns with the business goal, assess the potential and readiness for the employee to take the role by asking the employee the following questions:

- Would you be able to gain the skills required for the role?
- What skills and experience do you currently have or lack that are required for the role?

Look at the gaps in the readiness assessment and develop a potential timeline to achieve those missing skills and experience. The results of this exercise will provide a development plan that leads to achieving a goal that is good for both the employee and the organization.

Common factors that can act as motivators or discourage an employee from working within a specific job role are as follows:

- Income
- Geographical location
- Travel
- Work/life balance

- Type of work (technical, social, sales driven, etc.)
- Benefits
- Training and experience opportunities

It is important to identify how each of these factors impacts every employee as you create development plans to ensure their personal goals are met along with professional goals. For example, one position might have a higher pay but require more travel, posing an unwanted work-life balance for a particular employee. Another job might pay less but provide the opportunity to live where the employee desires and offer teleworking opportunities, which might be more important than higher pay to a particular employee. Not covering personal goals can lead to moving employees into roles that negatively impact their personal lives, causing the employee to leave regardless of the benefits of the promotion. Consider these personal factors when creating a development plan for your employees.

Job Retention

Often, the term *competitive workplace* refers to competition between existing employees. This view translates to an employee struggling with separating themselves from the other career-driven employees all vying for attention from management to gain a promotion. In the field of cybersecurity, the dynamic has switched in favor of the employee. Now the competition exists at the organizational and corporate levels, leading to organizations shopping for talent within other organizations based on the huge demand for talent. Employees with the right experience and skills will be bombarded with job offers daily, making job retention extremely difficult to maintain across all SOC positions. You need to focus on job retention or all of your time invested in finding the right employees will end up benefiting somebody else that poaches them!

Salary.com suggests a few benefits you can offer to your employees to improve job retention. Offering some of these items might not be possible or cost effective for every employee, but providing them when possible is ideal. The first offering is good health coverage for all employees, including part-time workers. Good healthcare includes wellness benefits such as gym

memberships, healthy snacks, and other ways to keep employees healthy and strong along with traditional healthcare services. Healthcare services include preventative benefits, which cover all aspects of physical health, dental care, and vision. Healthcare also includes self-care benefits such as legal services to help with personal matters, discount programs, and mental healthcare. All of this will help reduce personal distractions and keep employees happy.

Salary.com also suggests offering telecommuting opportunities and flexible hours when a job role allows either. These benefits can be offered in small doses, such as once a week or more frequently, depending on whether employees in the role can provide the same value as they provide working in the office on a fixed schedule. Offering telecommuting and flexible hours also increases the geographical reach for recruiting people as well as retaining employees that have to move but still want to be part of the SOC, since they can continue work from their new living location. Factors such as the hours and availability of the SOC as well as location requirements will impact these offerings.

Another benefit suggested by Salary.com is encouraging employee training, workshops, and other forms of education. Offering training opportunities not only develops talent within the organization, but keeps employees motivated to stay and improve their capabilities. Completing training can be used as milestones for raises and other rewards, giving employees a clear direction on how they can advance their career. Make sure to consider all types of training and development, ranging from formal classroom training to on-demand online courses. Other development options outside of training include shadowing senior members for over-the-shoulder training, one-on-one coaching and mentoring, local networking groups, and adding members to special projects outside of their normal job duties. Also consider group rates if a specific skill or certification can be applied to multiple employees to save on training costs.

Some organizations use their investment in training to retain employees by offering to pay for training if the employees commit to remaining in their role or employed by the organization for a specific period of time. The benefit of this approach is that it discourages employees from leaving the organization since they would forfeit having the organization pay for their training. The downside of this approach is possibly discouraging some employees from pursuing training due to their unwillingness to sign a retainer agreement. An

alternative to training retainers is to provide compensation awards in the form of stock or pay that pays out over a specific period of time based on the employee remaining within a role or at the organization. This approach also encourages employees who don't want to commit to a retainer to obtain training.

Training

Training is the action of teaching a particular skill or type of behavior. SOC employees need to be trained to be able to perform their jobs and keep up with the changing threat and IT landscape. When an incident occurs, a common corrective action is more training. I already covered how training is used by companies to retain top talent. Considering all of these reasons for investing in training, the costs for training can quickly become a fortune and training results can be hard to measure if specific objectives are not defined. The following are recommended steps and considerations when developing a training program for any SOC employee:

Step 1. Create the business case: How does this training impact the SOC and employees that will be attending it? Does the training target a specific SOC service need or is it for career development? A cost-benefit analysis might be needed to justify the requested training.

Step 2. Define objectives and learning outcome: Describe what knowledge should be obtained via the training and how to judge if learning objectives were met. This could be achieved in several ways, such as having the employee obtain a certification or demonstrate the new skill.

Step 3. Select a training method: There are many methods to deliver training. The traditional in-person class may be more effective than delivering training online, but a live class will cost more both in time and money. Using recordings will reduce the cost of delivering training, but students will not be able to have discussions with the trainer, potentially reducing the quality of the training. Consider all options, including over-the-shoulder training, video, video conferencing, and live classes.

Step 4. Identify resources: Who will provide the training? Will it be in-house or an external resource? Are there any qualifications required for somebody to deliver the training properly? Some certification programs require a certified proctor to deliver content, limiting available resources to provide the training.

Step 5. Develop training material: Make sure the content that is developed is in line with the training objective identified for the business case of the training. This includes meeting all learning objectives so that candidates who complete the training can be properly qualified as successfully trained.

Step 6. Deliver training and evaluate effectiveness: Deliver the training and include a way to obtain feedback. Feedback should come from both the trainer and trainees to best understand both parties' experience of the course.

Step 7. Improve the training: The final step is to grade how well trainees accomplished the learning objectives as well as review the feedback from both trainers and trainees. Use these results to adjust the class so that it becomes more effective.

An example of going through this process is considering training for using a specific tool. The business case can be based on the impact the tool will have to a SOC service once the users are properly trained. The outcome of the training could be a certification from the tool vendor as well as the trainee's ability to showcase how they use the tool. The method of training could be a live boot camp delivered by the vendor's training resources or some other method that accomplishes the desired training outcome. The resource and material could be provided by the vendor, but a SOC sponsor can also be involved to help with running the class and obtaining feedback. The cost for this entire process can be computed and weighed against the value of the outcome to properly justify the training before any investments are made.

Training Methods

There are many variations of training, the quality of the results for which will be impacted by the method used. Many cybersecurity concepts require hands-

on experience with potentially illegal tools. Certain divisions of the U.S. military such as the U.S. Cyber Command (USCYBERCOM) request contracted training to include working within real-world scenarios that replicate the actual challenges organizations are likely to encounter. Expectations are that the USCYBERCOM candidates will have experience dealing with real malware and defending against genuine exploitation tactics. USCYBERCOM not only requests real-world scenarios but also sets expectations for persistence as part of their success criteria. Persistence means training must be regularly scheduled as well as sometimes unannounced to continually hone skills.

Training might not be project specific. Your SOC employees might want to take on different roles that have certain training requirements to perform properly. Encouraging career growth is key to developing a relationship with employees, leading to employee retention and savings on in-house promotion versus the costs to replace lost employees. Enabling career growth can be accomplished not only through formal training but also using informal over-the-shoulder shadowing of other employees. This approach not only saves in training costs, but also develops redundancy for skillsets and critical personnel. Formal training can also be offered, which can be tied to agreements for a trainee to remain within their role at the organization for a period of time in exchange for the training being paid for by the organization. A violation of the agreement could require the trainee to pay for the training, reducing the likelihood of the employee leaving their role during the agreement period. Promotions and other awards can also be tied to training milestones, which milestones can align with expected skills of more senior job roles defined within your organization. Aligning training and career paths will improve employee retention since employees will have a reason beyond a paycheck to remain in the organization.

Another training consideration is to develop a *cyber range*, the purpose of which is to simulate a real environment and the types of threats that an analyst could encounter. A cyber range might not be tied to a specific learning objective, but can be viewed as a practice ground to help members test out various types of scenarios that will come up as the SOC operates as well as customized scenarios based on specific learning objectives. A cyber range should have a student utilize tools to solve challenges in real time using a similar environment to the SOC’s real network. The cyber range should be

isolated from the real network, providing a safe, legal environment to gain hands-on skills with tools used by the SOC and expected situations the SOC will encounter. Many guidelines, including the National Initiative for Cybersecurity Education (NICE), define recommendations for a cyber range. One military-based saying that highlights the importance of using a cyber range to gain experience with cyberthreats is “the battlefield is the last place you want to meet your enemy for the first time.” It is best to fail in a range rather than in the SOC.

I recommend considerations for training based on real-world scenarios and including criteria for persistence to ensure that employees not only learn skills but retain them. It is not unusual to find that a candidate certified in specific skills isn’t able to perform those skills after a prolonged period of time of not using them. This brings us to an important concept, which is understanding the relationship between certifications and training.

Certifications

An IT certification validates that the certified professional has competency in a specific aspect of technology. Each certification program has its own method to validate a candidate’s skills, which range from combinations of test takers answering multiple-choice questions to performing hands-on exercises. After a candidate’s skills are validated through a program’s assessment process, the program issues a certificate signifying the person met the program’s requirements and the specific date on which the certificate was issued. Many programs require a recertification assessment within a certain period of time after initial certification. Recertification requirements vary from program to program and can involve either performing the same skills required for the initial certification, using a condensed version of the testing system, or just paying a fee, typically used to fund a membership program. Do not assume that a certification validates a person’s *current* skill level; take into consideration when the individual was certified, what was involved to get certified, and how often recertification occurs. The best approach to validate any skill is to have the person perform that skill in your own real-world scenario.

One common challenge I hear from SOC managers is determining which

certification is the best option for their employees. My advice is to consider aligning the purpose for a certification program with the SOC position looking to get certified. Certain certifications and training are designed for specific job roles. For example, the CompTIA CySA+ is designed for a cybersecurity analyst, while the EC-Council Certified Penetration Tester is obviously targeting the penetration testing market. I included suggested certifications for each job role related to SOC work earlier in this chapter.

Evaluating Training Providers

Different training providers will offer their own version of a certification program. For example, EC-Council, SANS, and Offensive Security all offer a penetration testing certification. Some of the content will be similar, while other parts of the program will be unique based on how the provider develops its material. It is recommended to consider the following when evaluating a program:

1. What steps/efforts are required to learn and achieve a certification?
2. What are the upfront and annual costs following completing the certification? Some programs require recurring fees.
3. What are the recertification requirements?
4. How respected is the certification/program based on industry feedback?
5. Do the learning objectives align with your own learning objectives?
6. Who will be developing the content and teaching the content? Some programs push live classes with generic teachers that provide little value for the high cost of the course.
7. When is the training offered and does it meet your training timeline?
8. Are there better competitive training options that accomplish similar learning objectives?
9. Does the training and testing format mesh with your learning style?

Training should not be limited to individual learning or technical knowledge. The SOC should also train as a unit to improve its services. One popular approach to accomplish SOC training is performing tabletop exercises.

Company Culture

One key factor that is outside of the power of an employee's manager that will encourage or discourage an employee to stay within a role is the company culture. Company culture is the personality of a company. Company culture is a mix of various ingredients including the work environment, company mission, ethics, and values. Some organizations operate in a very casual manner, while others enforce strict rules and regulations. The Balance Careers (https://www.thebalancecareers.com), a service covering career advice, points out that employees enjoy work when their needs and values are consistent with those in the workplace. This leads to employees developing better relationships with coworkers and being even more productive. The Balance Careers also points out that if you don't fit in with company culture, you are likely to take far less pleasure from your work. Forcing an employee that prefers to work independently to work in a team environment will not yield a happy employee.

I have seen organizations attempt to create, and sometimes force employee participation in, what leadership believes would be considered "fun and desired" exercises, which sometimes works very well but often has the opposite effect. For example, an organization might invest in team-building events rather than training, or offer free lunch rather than more paid time off. Some organizations might attempt to push the concept of work culture by changing the language used about the work being done. An example is labeling a call center a "customer satisfaction center." Some organizations might develop sales or service contests such as having the sales team perform customer sales pitches to team members for a chance to win the best sales pitch award. All of these processes are designed to impact people with hopes of improving the organization's culture. I highly recommend any of these actions as long as they align with a business goal that can be measured. If running a team-building exercise or sales contest, make sure to also establish

a goal that can be measured following the event. If free lunch is going to be provided, what is the return on this investment and, more importantly, does this have the impact intended and, if so, is it the best option to obtain that impact? Make sure to use a combination of the business strategy alignment techniques covered earlier in this book along with employee surveys to develop processes and other activities that will lead to a great culture. Don't force events for the sake of culture or you will upset some employees as well as waste time and money on efforts that do not positively impact the organization.

Summary

This chapter opened by highlighting the importance of the people within your SOC. You learned about industry job roles to give you an idea of expected skills based on common job titles. Next, you learned about the different SOC services and focused on the expected skills of the people that provide those services. Another important topic covered was the concept of soft skills and how they should be considered as you recruit employees for your SOC. You also learned about security clearance requirements. All of this data is designed to develop job requirements to fill your SOC with the right people as you launch or mature different SOC services.

The second part of this chapter provided recommendations for developing an interview plan, including using an interview prompter to ensure that all questions are covered during a formal interview. It also covered many topics that need to occur after interviewing, including recommendations for bringing new hires into the organization and management best practices to ensure you retain top talent. The chapter closed with a look at training, certificates, and company culture, which all impact retaining top talent.

Next up is Chapter 5, which reviews all the types of data that will be generated by a SOC and how to centrally manage and benefit from those results.

References

Adler, L. (2019, April 29). Hiring Without These Critical “Soft Skills” Is a Recipe for Disaster. LinkedIn. https://www.linkedin.com/pulse/hiring-without-critical-soft-skills-recipe-disaster-lou-adler/?trk=eml-email_feed_ecosystem_digest_01-recommended_articles-5-Unknown&midToken=AQGIjIs7uSiggQ&fromEmail=fromEmail&ut=2DzrO73wb7TUI1

BDC. (n.d.). How to Hire the Right People for Your Business. BDC. https://www.bdc.ca/en/articles-tools/employees/recruit/pages/7-steps-recruiting-right-people.aspx

Berkowitz, M. (n.d.). Think Before You Hire: Maintain a Legal Hiring Process. Monster. https://hiring.monster.com/hr/hr-best-practices/recruiting-hiring-advice/acquiring-job-candidates/legal-hiring-process.aspx

Cotter, T. (n.d.). Evaluate Candidates with a Pre-employment Assessment Test. Workable. https://resources.workable.com/blog/skills-assessment

GetEducated. (n.d.). 13 Highest Paying Technology Careers. GetEducated. https://www.geteducated.com/careers/highest-paying-technology-careers

Indeed. (n.d.). How to Write a Job Description. Indeed. https://www.indeed.com/hire/how-to-write-a-job-description

Jones, G. (2014, June 16). 8 Key Elements of an Effective training Program: Design, Measurement, Continuous Improvement Important. Canadian Occupational Safety. https://www.thesafetymag.com/ca/news/opinion/8-key-elements-of-an-effective-training-program/187061

Mayhew, R. (n.d.). Legal Requirements of Job Descriptions. Chron. https://work.chron.com/legal-requirements-job-descriptions-20506.html

Monster. (n.d.). How to Write a Job Description. Monster. https://hiring.monster.com/employer-resources/recruiting-strategies/talent-acquisition/writing-job-descriptions/

Monster. (n.d.). Keep the Interview Legal. Monster. https://hiring.monster.com/employer-resources/recruiting-strategies/interviewing-candidates/legal-job-interview-questions/

Oltsik, J. (2018, January 11). Research Suggests Cybersecurity Skills Shortage Is Getting Worse. CSO. https://www.csoonline.com/article/3247708/security/research-suggests-

cybersecurity-skills-shortage-is-getting-worse.html

Pingboard. (n.d.). 10 Ways to Encourage a Healthy Work-Life Balance for Employees. Pingboard. https://pingboard.com/work-life-balance/

Security Degree Hub. (n.d.). What Is a Security Clearance? Security Degree Hub. https://www.securitydegreehub.com/what-is-a-security-clearance/

ISACA. (2020, February 24). ISACA's Cybersecurity Study Reveals Struggles with Hiring and Retention Persist, More Diversity Progress Needed. https://www.isaca.org/why-isaca/about-us/newsroom/press-releases/2020/isacas-cybersecurity-study-reveals-struggles-with-hiring-and-retention-persist-more

Chapter 5

Centralizing Data

It is a capital mistake to theorize before one has data.

—Sherlock Holmes (Arthur Conan Doyle)

Data is the foundation of everything we do in IT. According to his 2018 *Forbes* article, “The Mind-Blowing Stats Everyone Should Read,” Bernard Marr states there are 2.5 quintillion bytes of data created every day, a pace that is accelerating with IoT growth. Marr provides lots of statistics to support that massive number, a few examples of which are that (as of May 2018) Google processes 3.5 billion searches every day, and that every minute people send 16 million text messages and 156 million emails, Snapchat users share almost 600,000 photos, and Tinder users swipe almost 1 million times. You are probably thinking, “People generate lots of data,” and you are absolutely correct! Also, data production globally is increasing each year. The industry calls dealing with large amounts of data that is too big to store within a database “working with big data.” Big data has become a common denominator for many fields, including researching cyberthreats. A 2018 Cisco Talos whitepaper titled “Talos Intelligence” states that Cisco Talos has insight into more than 17 billion web requests each day—and web requests are just one of many data sources Cisco Talos lists as part of its research in this paper.

Big data has raised many concerns regarding how data is used and if privacy is being violated in the process of collecting so much data. Organizations such as Google and Facebook are facing government scrutiny about how they are using everyone’s data; in fact, in 2019, the U.S. Federal Trade Commission (FTC) fined Facebook $5 billion over privacy violations. One could argue that your entire digital footprint is stored within databases owned by search engine and social media service providers. Based on the data they

have collected about you, they know what you like, who you know, what you have done, and how likely you are to think a certain way. Knowledge is power, and search engine and social media service providers possess and profit from that data, which we all give them in return for the "no cost" services they provide.

This leads us to the conversation about something you can control, which is how you use your data (I won't talk further about how others are using your data since that is a touchy topic and out of scope for this book). The main point of this chapter is that the value of data is not *what* you store but *how* you use it. When you consider the vast amounts of available data from external sources, such as threat intelligence mixed with the various data-producing resources in your environment, you quickly eliminate any possibility of being able to manually view and understand everything you control that produces data. We are at a point in information technology where you must have a strategy regarding controlling what you collect, properly managing it, and how it will be used in order to successfully use the data that is available to your SOC.

This chapter looks at the different types of data that you can collect, various methods that you can use to tune how that data is sent and received, and ways to filter what you need to avoid working with useless data. You will learn how to manage data after you have collected it and how you can apply it to various services within your SOC. You will also learn about some best practices for centralized data management, which will provide you with a foundation of useful data that you can leverage for concepts that are explored in subsequent chapters, including Chapter 10, "Data Orchestration," which covers how to automate actions based on data. Keep in mind that if you fail to collect data properly, any services that use that data will also fail. I've said this in a much simpler way to my customers: "Garbage in, Garbage out!"

Note

Many security orchestration, automation, and response (SOAR) tools require a security information and event management (SIEM) solution as their data source. You need to understand this chapter before you read Chapter 10.

Data in the SOC

Data tells you a story and it enables you to make decisions. The success or failure of using data is based on whether or not a SOC can convert data into actionable intelligence in a reasonable amount of time. This means that it is critical to be intentional about how data is collected and provided to your analysts based on the mission of the service that those analysts support. Being intentional leads to categorizing different data types and aligning what is collected with desired context. One analyst might want to know about malware, while another analyst might be more interested in risk to operations. The results of properly leveraging data should be very specific, filtered-down data that a certain user can digest in a reasonable amount of time. Failure to properly filter data will cause your analyst to be overwhelmed with information, produce tons of false positives, and lead to many missed critical events that are buried in the data noise. For example, looking back at the massive breach Target experienced in 2013, although its security tools recorded that the breach had occurred, Target's SOC took no action precisely because it failed to properly leverage its security-related data.

Chapter 1, "Introducing Security Operations and the SOC," raised the topic of threat models. Threat models are used to understand how adversaries can attack and what the SOC can do to reduce the risk of such attacks. Threat models like MITRE ATT&CK reference different aspects of an attack, including understanding the adversary launching the attack, which represents bigger-picture thinking, and identifying specific techniques used to accomplish certain objectives, which represents very specific actions. Collecting data for both of these goals will require different types of data based on how the SOC intends to use the data.

Strategic and Tactical Data

One data type that aligns with bigger-picture thinking is *strategic*—the data is designed to understand the whole battlefield. Strategic data involves problem solving and creative input used to inform and develop conclusions. Having strategic guidance parallels how organizations can use abstract data, such as unstructured data, for research and discovery of questions both in threat

hunting and in developing guidance for the best placement for future responsive capabilities. I talk about data structures shortly, but know that abstract data such as social media can be useful for learning about upcoming adversaries if the analyst's goals are strategic.

Another data type aligned with accomplishing objectives is *tactical*, representing very specific data used from short-term actions. Tactical data is never abstract. An analyst or tool must be able to have clear direction for what is being seen within the data, such as "block *x*" or "look for *z*." Tactical data is grounded in short-term use, which limits its use for any bigger-picture concepts, meaning any strategy can be weighted down if it has too much of a tactical focus. During the response to a security breach, the incident response team will need tactical details of what to look for as they contain, eradicate, and recover from the breach. Tactical details such as a list of hash values matching malware used by the incident response team, however, do not help the risk management team determine what other threats could be associated with the breach. The risk management team will need a different set of data, including where the threat came from, how it spreads, and trends with threat actors that use similar tactics to breach organizations. Chapter 7, "Threat Intelligence," takes a closer look at threat intelligence, including methods to identify and use both strategic and tactical data.

Note

In Chapter 7, I introduce four types of threat intelligence. Each type has both strategic and tactical datapoints as these are the fundamental elements of data found within threat intelligence. For example, a Operational Threat Intelligence feed can provide both static details about the attacker and tactical details about what they are using as their method of attack.

The main takeaway about data is that there are many data types and different situations that call for specific data. You don't just collect anything security-related and dump it on a SOC analyst to review, or send everything to a security tool such as a SIEM tool and expect the tool to provide the right data for the many use cases the SOC will need it for. I have seen this mistake

made often, leading to failures in SOC services. I referenced the Target breach earlier in this section. There could be many reasons why the SOC did not see the breach alarms, but I would guess that it had something to do with the analysts responsible for monitoring security events not being provided the right data. This happens when an analyst is overloaded with data, typically caused by security tools not filtering data for specific needs.

Data Structure

All data has some form of structure. What determines whether you see it as having structure is whether the data has a predefined data model organized in a predefined way that your tools are capable of leveraging. Any data that doesn't have a structure that your tools support is essentially seen by your SOC as unstructured data. This means the first way to categorize data is to group it into the following types:

- **Structured data:** Data that resides in a fixed field within a record or file. The predictability of structured data means you can develop processes and procedures for how to store, process, and access the data because the tools you leverage know where to find what they need within the record or log file. An example is a file containing data that a programming language such as SQL can read and write to based on a variable category, such as Day, followed by a response to that category, such as Monday.
- **Unstructured data:** Data that you can't simply place in a category for your tools to understand. Examples include photos, videos, and PowerPoint presentations, which all contain data, but not structured in a way that something such as a SQL program can easily digest and modify.
- **Semi-structured data:** Parts of the data can be recognized by tools while other parts cannot. An example is using SQL to pull metadata from an unstructured data source, such as the author of a PowerPoint presentation, but the rest of the data can't be abstracted, making it unstructured data. This can be useful if you are investigating a few hundred PowerPoint files and want to filter out any that don't relate to the case. By pulling the metadata with a tool, you can quickly categorize details about the author of the documents and reduce your investigation list to only those that matter. From there, you would have to manually investigate each file

because the PowerPoint data would be unstructured, meaning you can't use a tool to identify what you are looking for within images, videos, and other unstructured data types.

> **Note**
>
> Unstructured data *can* be useful; however, it must be treated differently than structured data.

Data Types

Your SOC will encounter various forms of structured and unstructured data. Within these categories are basic types of data, called *primitive data types*, that serve specific purposes. Many programs and security tools that digest data include methods to identify primitive data types and convert between them when necessary so that the programs and tools use the data types properly for their intended purposes. As you leverage new security tools and data sources, be aware that some tools might not convert all primitive data types properly, causing misalignment in expected results. A common example is leveraging SIEM software to digest data from multiple tools. If the SIEM's parsing engine misunderstands a data type, it will produce the wrong results for the analyst. Imagine that a SIEM tool is expecting one of the seven days for the variable Day but aligns the results of Day to people's names. Another common example is using scripts or other programs to read in data. If the data type is incorrect, bad things can happen. Hackers (in terms of people looking to cause systems to do things they were not intended to do) will introduce the wrong data types to applications on purpose, which is why input validation is a critical step to securing input into such applications. Later in this chapter, you'll see a few examples of security tools misunderstanding data and you'll learn about some recommendations for adjusting data to fix the issue.

Data Type Examples

The following are examples of common primitive data types:

- **byte:** The byte data type is an 8-bit signed two's complement integer. A byte has a minimum value of –128 and a maximum value of 127 (inclusive). Byte data can be useful for saving memory in large arrays, where the memory savings matters. Byte data can also be used in place of int data where their limits help to clarify code.
- **short:** The short data type is a 16-bit signed two's complement integer. It has a minimum value of –32,768 and a maximum value of 32,767 (inclusive). You can use a short to save memory in large arrays, similar to a byte, in situations where the memory savings matters.
- **int:** By default, the int data type is a 32-bit signed two's complement integer, which has a minimum value of –231 and a maximum value of 231 – 1. Use the Integer class to use the int data type as an unsigned integer.
- **long:** The long data type is a 64-bit two's complement integer. The signed long has a minimum value of –263 and a maximum value of 263 – 1. In Java SE 8 and later, you can use the long data type to represent an unsigned 64-bit long, which has a minimum value of 0 and a maximum value of 264 – 1. Use this data type when you need a range of values wider than those provided by int.
- **float:** The float data type is a single-precision 32-bit IEEE 754 floating point. As with the recommendations for byte and short, use a float if you need to save memory in large arrays of floating-point numbers. This data type should never be used for precise values, such as currency.
- **double:** The double data type is a double-precision 64-bit IEEE 754 floating point. For decimal values, this data type is generally the default choice. This data type should never be used for precise values, such as currency.
- **boolean:** The boolean data type has only two possible values, true and false. Use this data type for simple flags that track true/false conditions. This data type represents one bit of information, but its "size" isn't something that's precisely defined.
- **char:** The char data type is a single 16-bit Unicode character. It has a minimum value of '\u0000' (or 0) and a maximum value of '\uffff' (or 65,535 inclusive).

Data Context

Primitive data by itself doesn't provide much value. It's all about its context. Having a value of "true" doesn't mean anything until you know what is true. This concept reminds me of the storyline from *The Hitchhiker's Guide to the Galaxy* where a supercomputer finds that the "Answer to the Ultimate Question of Life, the Universe, and Everything" is the number 42, but nobody knows the question, hence millions of years of research goes to waste. Adding useful context to data in regard to security operations is known as *data enrichment*. Data enrichment turns raw data into meaningful insights. Among the many data enrichment possibilities, data can be enriched with contextual information from the following sources: user directories, to answer who is associated with the data; asset inventory, to reveal which devices are linked to the data; geolocation tools, to answer where the data is being generated; and third-party threat intelligence databases, to provide insight into how data relates to other organizations. Many SOCs use data enrichment techniques through integration with and consumption of data feeds to tune their security tools for better threat detection, threat hunting, and incident response services. The following are some examples of data context and enrichment enhancements commonly used in SOC security tools:

- **Identity context:** Linking data to users is useful for linking an asset owner to the system that is generating data. Common sources for adding identity context to data include identity and access management (IAM) systems, directories, enterprise resource planning (ERP) systems, and Microsoft Active Directory (AD).
- **Asset information:** Adding asset information can provide additional details about a device such as who is the owner, where it is located, and what version of software it is running. Asset details can be obtained from tools such as a configuration management database (CMDB), which stores details stored about known assets, or a network access control (NAC) tool, which collects asset information upon connecting to or while using the network.
- **Access privileges:** Tools such as NAC technology enforce access control based on user and device privileges. Access privileges associated with user and system data are commonly held in identity databases, such as Active Directory group memberships. Some recent NAC concepts such as

security group tagging (SGT) link access privileges with each data packet sent by a device so any network or security device on the network is capable of allowing or denying the packet based on predefined conditions that evaluate the access privileges tagged to the data. SGT essentially can enforce access privileges down to the packet level.

- **Nontechnical feeds:** There are many nontechnical data enrichment resources, such as background checks, time, or badge data, that can change the context of data. For example, it may be okay for an employee to access a building, but if the badge data shows access at 3 a.m., that could mean something completely different. Nontechnical data can either be part of a tool or be converted into technical data using converting technology. For example, Chapter 7 covers how to scrape social media data and convert it into a threat intelligence resource.
- **Vulnerability context:** Vulnerability management and risk management are two common services provided by a SOC. Flagging data associated with known vulnerabilities can help deliver these services more effectively. Resources such as vulnerability scanner reports and vulnerability databases are common resources to pull vulnerability context about systems.
- **Social and online context:** Most people, regardless of age, that use the Internet have social media footprints, which can provide a lot of interesting context about who they are and what they do. Many users see social media as their only resource for current news and interact with some of their friends only through online chat. Chapter 7 covers how to use this extremely valuable form of data with your SOC.
- **Network maps and geolocation:** Where data is produced, where it is sent, and where it has been are all examples of context points that can enrich the value of data. For example, it is common for organizations to block or flag as a high risk for cyberattack any data that comes from parts of the world they do not do business with. Internal network classification or cross-border analytics can provide a lot of context value to data.
- **Process and operational context:** The “how” things are done can be a critical value point for implying potential breach conditions that are not commonly seen. This data context category can also help tie together some of the other data context details. Examples include if the same login

occurs from two globally diverse locations or if an IoT device starts interacting with users. The context might not be explicit regarding it being a threat, but how the data source is used, sourced, and directed across the organization can speak to behavior traits leading to a new conclusion.

> **Note**
>
> In Chapter 7, I speak further about how data enrichment is used to convert threat data into threat intelligence.

Summarizing the concept of data context, the value of data is based on how it is used rather than what it represents. To leverage data properly, you need to understand the context to which it applies. Based on your understanding of that context, you are able to make decisions. Keep this in mind as you leverage tools that consume data. If you are not receiving the value you intended from that resource, it is possible that either the wrong data is being sent, the tool is unable to understand what is being sent, or the tool is not set up to deliver what you want. The goal for this chapter is to address all three of these areas.

When I am asked to consult for a customer and evaluate how data is being used, I can't jump in and know if there are problems with data being sent, problems with data being received, or problems with how tools are using data until I can establish a baseline understanding of the environment. One tactic that is helpful for collecting the information that is needed to improve how data is being used is to perform a data-focused assessment.

Data-Focused Assessment

I will continue to state in this chapter the simple fact that the value you receive from data is *how* it is used, which means in what context you apply the data. Keep this in mind as you evaluate tools for your SOC. As simple as this concept might sound, I find that many people skip the *why* when evaluating tools and go right to the *what*, meaning people focus on what a

tool can do rather than why they need the tool to do something. Referring back to Chapter 1, you learned how to perform a capabilities assessment to evaluate what capabilities you have, don't have, and could replace. Figure 5-1 is an example of what could be used to perform a network and endpoint best practice mapping of security capabilities. Chapter 1 focused on the capabilities of tools, but now let's look at the same approach to assessing capabilities but from a different viewpoint. Because each tool or capability generates structured data based on how it is designed to assess an information flow, you can evaluate each of your tools in terms of what type of information it examines and the structured information and/or specific actions (and resulting data) that will be derived by the specific tool. Essentially, this section demonstrates how to perform an assessment of the types of data you can obtain from your existing or desired tools.

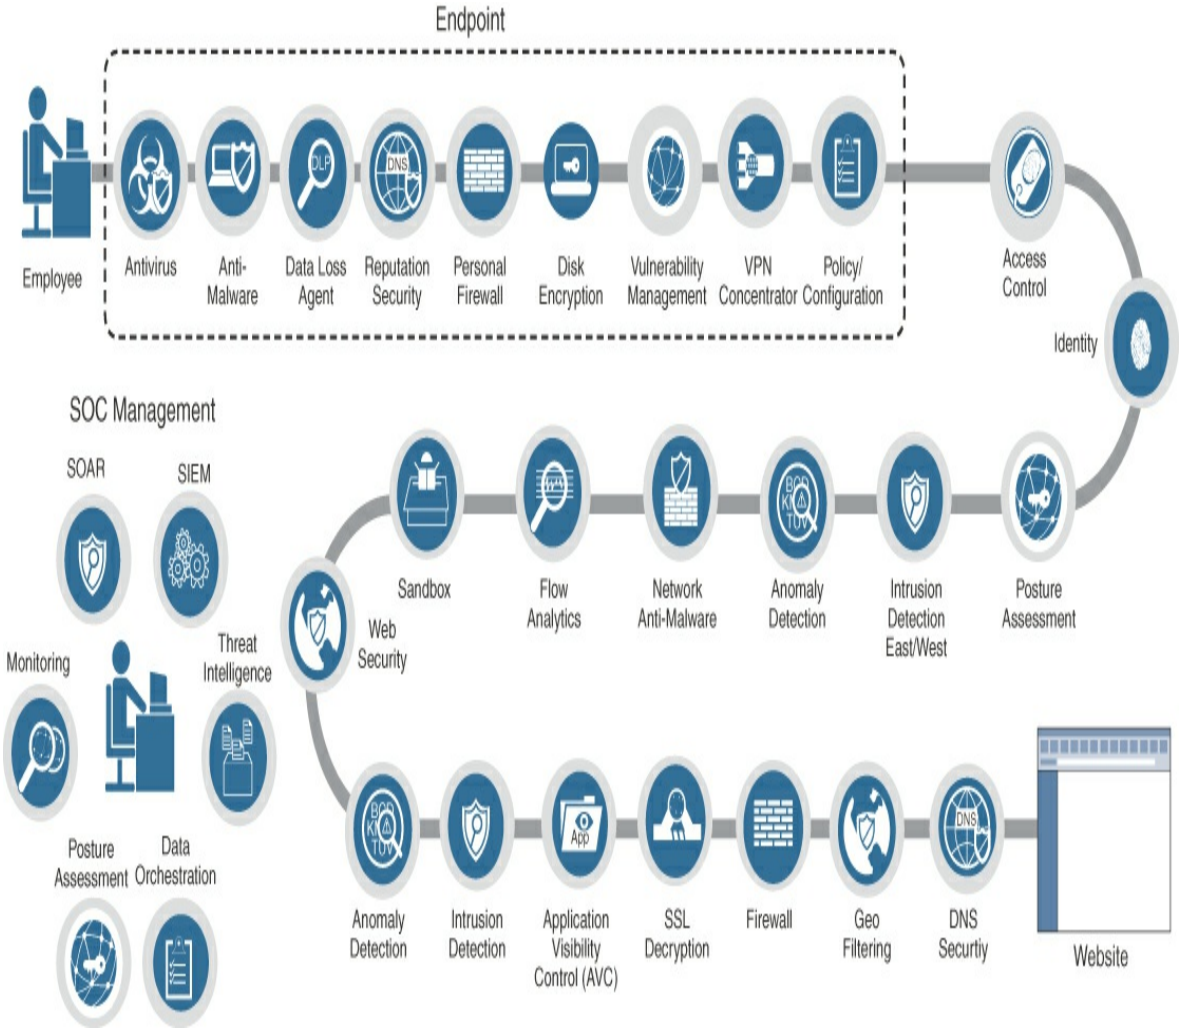

FIGURE 5-1 Example Security Capabilities Assessment Diagram

When evaluating which tools to use within your SOC services, you need to consider the data each tool can produce and how the context from that data would complement what the SOC management tools will already know. The value of doing so is that it allows you to filter out unnecessary data sources, consolidate where data overlap is being seen, identify gaps in desired data, and plan for future investments. As a consultant, I find many SOCs lack an understanding of how all of their data sources are being utilized and simply don't know how to tune their technology because they don't know which data sources are and are not needed. Rather than guessing or turning something on/off and seeing what happens, conducting a data-focused assessment can give you much more clarity up front before you make any changes.

Data Assessment Example: Antivirus

Let's go through an example of evaluating a typical antivirus tool. This example leaves out any vendor details and assumes that the tool provides only antivirus capability rather than a bundle of host-based security capabilities. In this example, the organization uses a host-based antivirus tool that is managed by a cloud management system, also called a Software as a Service (SaaS) antivirus offering. The organization is interested in leveraging the data produced by the antivirus product within its SIEM solution but has not researched if that is possible or, if so, what type of data could be pulled from the antivirus tool.

After researching the antivirus tool, a data-focused assessment for the antivirus tool is as follows:

- Antivirus produces alerts when host-based threats are identified.
  - Host events are created by the host antivirus agent and sent to the cloud management tool.
  - Automatic removal of malware is enabled on endpoints.
- Host events are managed by the cloud manager.
  - Host alerts are based on severity levels, with different actions taken

based on severity data tags.

- Limited details can be queried, such as what version of antivirus is running and the hostname.
- Exporting of data is done using APIs. No industry-standard log exporting such as syslog is available at this time.

- The organization's SIEM solution can receive standard log formats.
  - With the antivirus data, the SIEM solution can produce a widget within the SOC dashboard that monitors antivirus events. This allows analysts to monitor antivirus events without having to log into the manager of the antivirus solution.
  - Host signature threat data provides visibility of host posture and recent threat activity.
  - Host antivirus agents can be remotely updated and tasked to validate if a signature is enabled on a host system.

The key data point to focus on is that the antivirus tool can provide visibility into the security posture of host systems. Details are limited to security capabilities, versus a general assessment of the endpoint. This is an important concept to understand. If an antivirus tool blocks a link to a malicious advertising (*malvertizing*) ad on the Internet, that does not mean the system was or was not compromised. You are only being alerted of the view of the antivirus tool, meaning what it was able to identify as evil. It is possible the malvertizing website used multiple exploits, and the antivirus was able to block one; however, another exploit could have used a different tactic to successfully compromise the system. This is why taking in multiple data sources is critical to gain a full view of a situation. Not doing so can cause a false sense of security based on the limited data you are receiving about a situation.

In short, the following data points could be pulled from the antivirus tool:

- Signature-based host events
- Limited host details (antivirus installed/running, version of antivirus, hostname of system)
- If actions have been taken against a piece of software such as quarantine

- Search hosts for signatures that are and are not enabled

Threat Mapping Data

Once you perform a data assessment of a security capability, the next question is to answer *why* its data matters to the SOC. Looking back again at threat models from Chapter 1, the value of reviewing threat models is to evaluate the tactics, techniques, and procedures (TTPs) used by adversaries against your capabilities. This includes taking into consideration the combination of TTPs used during a campaign versus just considering each specific attack technique that could be used by an adversary. Many of the less mature security teams I have worked with fall prey to myopic event analysis—they look at each condition as a distinct and mutually exclusive condition. Adversaries use chained exploitation techniques to accomplish their goals, which means you need to view events holistically rather than individually.

I built a simulated ransomware lab based on the TeslaCrypt ransomware that uses the EternalBlue exploit to compromise a system and then uses the Mimikatz exploit to capture domain passwords from that system. Once domain passwords are captured, the ransomware can spread to other systems, regardless of whether the other systems are vulnerable to either EternalBlue or Mimikatz. Some antivirus tools might be able to detect EternalBlue; however, Mimikatz runs in memory, a characteristic that can evade many antivirus products that evaluate only files. Antivirus would not see the lateral movement as the ransomware spreads to other systems, and there will be a communication to an external source as the ransomware part of the malware encrypts files. To properly identify this attack campaign, the SOC's data should include exploitation behavior against hosts, anomalies detected within hosts, unusual lateral communication between hosts, unusual communication from hosts to external resources, and unusual usage of domain accounts. If all of these data points are combined, an attack campaign can be identified; however, individually viewing each of these data points might fall under the radar for a group of analysts buried with security alerts to investigate.

A threat model of my simulated ransomware would suggest a SOC obtain the data points that I covered as part of its security capabilities assessment. Antivirus would make up only part of the suggested data points, which would

lead to the need for additional capabilities to meet the requirements for modeling my threat. Other threat models such as MITRE ATT&CK can be used to help apply the security capabilities and data visibility they produce against what would be ideal to defend against modern threats.

Applying Data Assessments to SOC Services

Another use for data assessments and threat mapping is to identify how the data from capabilities can benefit fundamental SOC services. Looking again at my threat model, I would suggest at least two SOC services should use antivirus data as part of their daily monitoring tasks:

- **Vulnerability management:** Identify host system signature defenses, which include protecting known vulnerabilities from exploitation. For example, if new malware is identified by the SOC, any associated recommended signatures could be deployed to systems that are not protected. A search can also validate if a signature already exists to protect hosts from the malware.
- **Incident response/analysis:** Antivirus can automatically remove identified malware. The incident response team can push new signatures to hosts to defend against current malware.

A general antivirus solution isn't perfect by any means, but does offer value to the SOC. The value of traditional antivirus has been decreasing as attacks become more sophisticated, leading host security vendors to offer more capabilities, which also means more possibilities for the type of data they produce. This is why it is important to be intentional regarding who within the SOC receives antivirus data; some teams that do not need it would see antivirus alerts as a distraction.

Looking back at my simulated ransomware example, there are a handful of missing data points and capabilities that the organization should consider as it evaluates the antivirus tool. These missing capabilities can be acquired by using another tool, such as an anti-malware tool, to complement the antivirus or by adding capabilities from the antivirus vendor, assuming the vendor offers more capabilities for an additional cost. Comparing these desired data points and capabilities against the existing tool can justify change for future

investments that are aligned to SOC business goals:

- **Vulnerability management:** Include host-based vulnerability assessment capability. This will provide host-level vulnerability assessment data that allows the SOC's vulnerability management program to proactively identify and address vulnerabilities. Doing so will reduce the need for leveraging signatures for vulnerabilities that have been remediated.
- **Incident response, analysis, digital forensics, research and development:** Include ability to query host systems for what software is installed and what processes are running to gain more visibility into what is running on host systems. This will allow response teams to search hosts for hash values of artifacts of interest (such as a known malicious file), identify what defenses are in place, and better understand the state of host systems.
- **Compliance:** Include the capability to collect additional details on assets. Also include possible data loss prevention (DLP) features to identify whether sensitive data exists on the system.

Using a data-focused assessment will clearly identify why the SOC needs specific tools as well as what additional data elements would be desired to improve the maturity of each SOC service. A data-focused assessment also allows the SOC to identify which teams should receive specific data from a tool, allowing for better up-front filtering of alerts and other logs that a security tool produces. Combining a data assessment and a threat model allows for a better understanding of both how security capabilities provide a defense-in-depth approach to reduce the risk of exploitation and how the SOC should expect to view the data from the combination of tools. This is extremely important, as having the tools is only part of the solution. Your SOC must take action, or you might end up like the Target exploit where the tools had the data but nobody saw or responded to the event!

These assessment and modeling concepts work at a high level for planning security capabilities; however, as you evaluate individual tools, you will find limitations in how it can deliver data to your SOC's data management tools. Looking back at the antivirus example, the cloud management system for the antivirus didn't specify how events it sees could be exported outside of offering APIs. Will this system offer a universally accepted method to export logs in the future, or are you expected to leverage an API or specialized tool

to obtain its data moving forward? I will cover all three of these approaches to obtaining data from common security tools in the following topics.

First, let's focus on one of the most universally accepted exporting structure for exporting data, which the industry calls a *log*.

Logs

Almost every computing system generates logs. Logs in the computing world represent a file that records events that occur, ranging from what the operating system is doing to how the system is used. Logs are structured data designed to be machine readable (and often human readable) and easily parsed. If an event occurs on a host system, a host system log can include information about the host, the time of the event, what type of event has occurred, the severity of the event, and details on the event.

Abstracting context from a log is critical to being able to successfully consume large amounts of log data. Tools that digest logs can parse details into tables and take some action depending on specific data points within a parsed field. If an organization receives 30,000 logs every few hours, it would be extremely challenging to review each log. Instead, a field such as Severity can be used to sort all logs so that the most severe events are brought forward to the SOC analysts, enabling them to maximize their time by reviewing only the more critical logs. Organizations will never review every log generated by their tools, hence the importance of tuning what sends and receives logs as well as what ends up in front of an analyst through different deduplication and filtering techniques.

Log Types

Various types of systems within an organization can generate logs. The data assessment approach I covered earlier will allow you to categorize the expected data from different systems, which results can be used to decide which data makes sense to send to your data collection tools. The following are general categories of logs you will find within your organization. Using a

category grouping approach to segment log types will help you compare and contrast all of your possible sources for a specific data category. Looking back at my host-based antivirus example, you may find there are four other endpoint log sources, which other sources may have more relevant data to your SOC services than what could be pulled from the antivirus cloud management tool. An example could be a host management tool or anti-malware tool that includes many of the desired data points not found within the antivirus tool. The results of this type of evaluation might lead to the decision to not pull data from the antivirus management system or to limit what is pulled to only specific antivirus events that are a specific severity rather than all events seen by the antivirus management platform.

Here is a list of general log data categories:

- **Endpoint logs:** Endpoints generate multiple types of logs from different levels of their software stack, ranging from the operating system to the applications installed. Endpoint logs are useful for understanding the status of the device and activity it is involved with. Endpoints can range from mobile devices to user desktops, allowing for a wide range of data formats to exist. One of the most common sources for endpoint logs would be the Windows event logs.
- **Network device logs:** The backbone of networks is composed of the routing, switching, and load balancing components. These devices produce logs that provide critical data about traffic flow, including the source and destination IP addresses of devices using the network. Details within these logs include users, protocols used, traffic volumes, and other useful details about what is going on within the environment.
- **Application event logs:** Applications run on servers and generate their own logs that differ from the logs generated by endpoints that use the applications. For example, the Windows operating system provides a centralized event log that collects startup, shutdown, heartbeat, and run-time error events from running applications. Linux posts similar data within the /var/log folder. These types of logs can contain information about application performance and the services they offer, including email, web, or database server services.
- **IoT logs:** IoT logs are a growing source of log data. Some IoT device logs are similar to endpoint or application logs based on the type of OS

installed and how the IoT device is used. Other IoT devices either do not offer logging or require events to be recorded from a cloud management system. In this cloud example, if a disconnect occurs between the cloud management system and the IoT device, log events would be missed. The expected format for IoT logs can vary based on the wide range of use cases IoT devices address.

- **Security tool logs:** Security products such as firewalls, IDS/IPS, sandboxes, and honeypots all generate logs. These logs are critical to maintain security awareness across the organization. Although security tools have their own management interface, it is ideal to export logs to a centralized tool so that the SOC analyst has one place to view security events from all tools.
- **Directory service logs:** User privilege management is commonly performed using a centralized identity management tool. This tool generates logs for authentication and authorization as part of the accounting aspect of the AAA framework.
- **DNS server logs:** Domain name servers link IP addresses to websites, which is extremely useful to the SOC for many reasons. Security teams want to know when malicious sources are being referenced, and investigation teams want to see what resources were accessed during an event. DNS server logs can provide these details.
- **Replication logs:** It is common practice to back up systems, data, and services to ensure business resilience. Systems involved with the backup process will generate logs, which need to be monitored to ensure resilience is maintained.

You might be wondering at this point of the chapter, “Why not just send every type of log to my log collection tools? Don’t the tools provide sorting and filtering for me? The topics covered seem to require a lot of manual effort!” Good points; however, controlling what data you use has huge benefits over relying on a log management tool to do this work for you. I will cover these benefits shortly, but a quick answer as to why you should be smart about what data you leverage is that it will allow for cost savings, reduce the time spent on tuning, and increase your control and understanding of your security environment. If your security tool charges you for the number of logs it receives, being tactical about what you send the tool could

mean the difference between spending or saving thousands of dollars. Less noise means less wasted time tuning. People are and will continue to be a critical factor in security. You can't just automate the solution to all of your problems!

Another key point is that sending data without specific intent will lead to event fatigue. Too much data, even structured data, takes on an unstructured characteristic in large volume. An analyst needs a way to quickly abstract useful elements for a specific search, the data structure for which allows tools to simplify the filtering process. Without structure, many security tools will fail to show value. I'll provide a specific example of this shortly by demonstrating a SIEM solution (Splunk) attempting to process data in the wrong format. Before I provide this example, I need to cover the importance of log formats.

Log Formats

Computing systems generate logs; however, the format in which they generate logs can vary. There are industry-standard formats that vendors follow so that different tools can work together. Security tools designed to collect logs support common log formats such as comma-separated values (CSV), syslog, and JSON. Using a different format means the data fields will be in different places, but the same data will be present. One format might have a category such as Time followed by numerical time value. The category and the value can be separated by a comma, dash, slash, blank space, or some other method. It all depends on which log format is being used.

When a tool attempts to read a log, it uses some method to verify which log format is being used, such as looking at a header, and attempts to parse the data based on how it expects the data to be formatted. If a comma is expected to separate a value and result for a log file, the tool parses the data in that manner. This is an important concept to know for troubleshooting when a tool does not correctly parse log data. Figure 5-2 shows an example of a Cisco Advanced Malware Protection (AMP) log that isn't properly parsed by Splunk, the SIEM solution. This data sample provides a ton of useful artifacts; however, Splunk isn't categorizing the data properly. As a result,

the data context can’t be abstracted, leading to little value obtained from the data source. Buried within all of this data are details such as IP addresses, states of systems, ports, and protocols; however, there isn’t a predictable and referenceable field for most of this data, so, outside of searching for specific key terms, you won’t be able to find what you are looking for. Also, many fields such as user_agent equal zero, meaning either that nothing is included or that incorrect data parsing has led to variables pointing at the wrong associated data. The solution to this problem may be to adjust how the log is being collected or how the tool generating the log is formatting the log before it exports it to Splunk. We will cover troubleshooting log parsing within SIEM solutions later in this chapter in the section, “Troubleshooting SIEM Logging.”

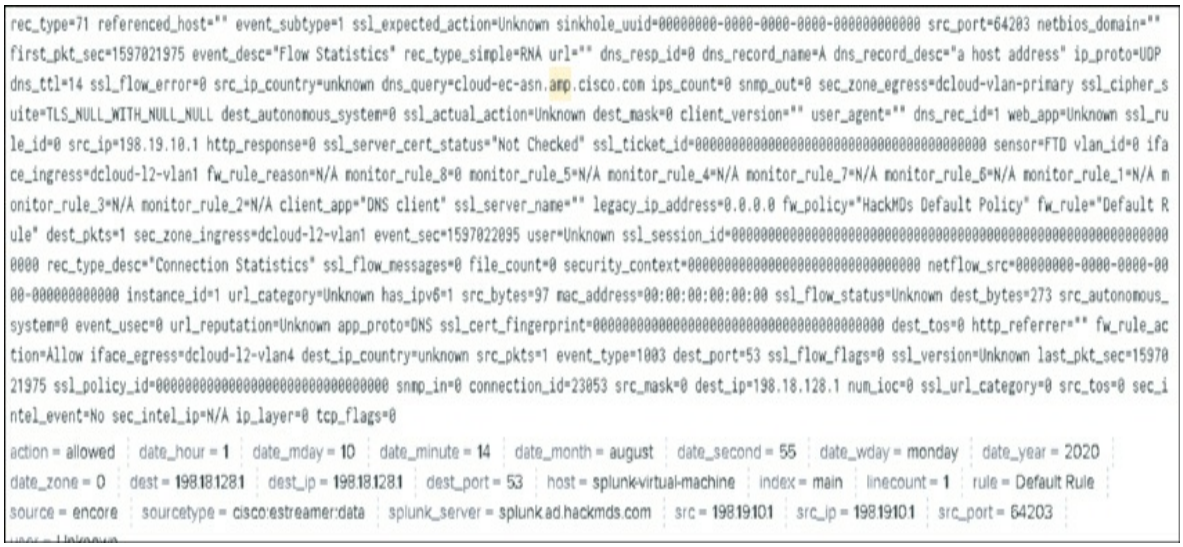

FIGURE 5-2 Poorly Parsed Log Within Splunk

Before delving into tools, you first need to understand the universal log types that are used by the majority of security tools designed to digest external log data. There are too many log types to cover every possible format, so I will introduce several of the most popular ones: syslog, JSON, Windows Event logs, Common Event Format, Common Log Format, and Extended Log Format.

Syslog

One of the most common and universally accepted protocols used in the IT world for decades to generate data is the syslog protocol. Syslog is a way to transport messages from devices to a logging server, known as a syslog server. A syslog server provides a way to consolidate logs from multiple sources. Syslog servers have a listener component designed to receive logs as well as a database for storing and retrieving what was collected.

Syslog is supported on major operating systems such as macOS, Windows, Linux, and UNIX, which your SOC tool's OS are based on. Syslog messages are made up of the following three elements:

- **Header:** Includes the version, timestamp, hostname, priority, application, process, and message ID.
- **Structured data:** The data is structured in an 8-bit Unicode transformation format.
- **Message:** Can be pretty much anything, which is why syslog is popular.

Syslog has limitations regarding what can be gathered. You can't perform the same actions on a network device using syslog as you could using other protocols such as Simple Network Management Protocol (SNMP). Syslog only supports the ability to send messages to a defined location when certain events happen. Also, the flexibility of syslog means it doesn't have a defined standard format for message content, leading to endless ways messages could be formatted. Another concern about syslog is that it doesn't include a way to guarantee that a message is delivered, allowing for the potential of losing some log messages. One final concern about syslog is that it lacks security features, opening the door for messages being sent from unauthorized sources. Regardless of these limitations, syslog continues to be an extremely popular logging option.

Example 5-1 demonstrates the default syslogd service.

EXAMPLE 5-1 Default syslogd Service

Click here to view code image

```
# Options to syslogd
# -m 0 disables 'MARK' messages.
# -r enables logging from remote machines
# -x disables DNS lookups on messages received with -r
```

```
# See syslogd(8) for more details
SYSLOGD_OPTIONS="-m 0"
# Options to klogd
# -2 prints all kernel oops messages twice; once for klogd to
decode, and
#    once for processing with 'ksymoops'
# -x disables all klogd processing of oops messages entirely
# See klogd(8) for more details
KLOGD_OPTIONS="-x"
#
SYSLOG_UMASK=077
# set this to a umask value to use for all log files as in
umask(1).
# By default, all permissions are removed for "group" and
"other".
```

JSON

The JavaScript Object Notation (JSON) format is a highly readable data-interchange format that is another standard for structured logging. JSON is easy for both humans and machines to read based on a compact and lightweight format. Pretty much every programming language and security tool is capable of parsing JSON-formatted logs, including Windows-, macOS-, and Linux-based tools. JSON logs are richer than most other log formats and can easily be further enriched with extra context and metadata. An example use case is having a tool search for the term “ERROR” to find any errors associated with a specific system. Similar to syslog, a downside of JSON is that because of its flexibility, its contents can contain just about anything and the size of the logs can quickly grow as data is included with large amounts of log generation.

Example 5-2 demonstrates a JSON log.

EXAMPLE 5-2 JSON Log

Click here to view code image

```
{ "timestamp": "2020-05-24 22:19:41", "id": 0, "class": "audit",
"event":
"startup", ...
```

```
{ "timestamp": "2020-05-24 22:19:45", "id": 0, "class":
"connection", "event":
"connect", ...
{ "timestamp": "2020-05-24 22:19:45", "id": 1, "class":
"connection", "event":
"disconnect", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 0, "class":
"connection", "event":
"connect", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 1, "class":
"general", "event":
"status", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 2, "class":
"connection", "event":
"disconnect", ...
{ "timestamp": "2020-05-24 22:19:53", "id": 0, "class": "audit",
"event":
"shutdown", ...
```

Windows Event Logs

Windows owns a large market share for endpoints and servers, making it one of the important systems to collect logs from. Windows event logs provide detailed records of the operating system, application, and security, and event notifications. Logs are captured and stored by Windows, and a system administrator can view those logs to diagnose potential issues or prevent future problems. Windows log files can track events such as application installations, system setup operations, errors, and security issues, including details on the core OS as well as running applications.

A Windows event log includes the following data points (this list includes system events, setup events, security events, application events, and forward events):

- Date the event occurred
- Time the event occurred
- Username of the user logged onto the machine when the event occurred
- Name of the computer

- Event ID, which is a Windows identification number that specifies the event type
- Source, which is the program or component that caused the event
- Type of event, including information, warning, error, security success audit, or security failure audit

Figure 5-3 demonstrates a Windows Event log.

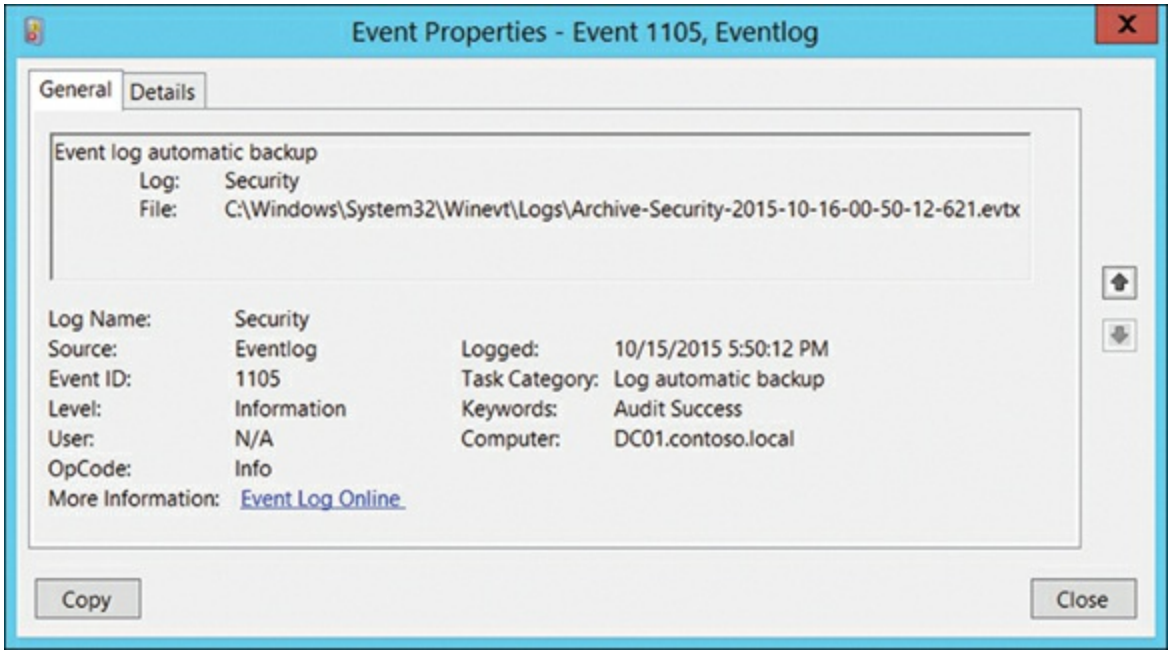

FIGURE 5-3 Windows Event Log

CEF Format

The Common Event Format (CEF) is an open logging and auditing format created by ArcSight. CEF is a text-based format that contains event information. CEF is popular due to its human- and machine-readable format. The original purpose of the CEF format was to enable various vendor security tools to send data to ArcSight, but many tools leverage CEF beyond ArcSight or for working with ArcSight. CEF is yet another transport mechanism for delivering information between systems.

CEF uses the UTF-8 Unicode encoding method, meaning the entire message

must be UTF-8 encoded. The CEF format includes a standard header and a variable extension constituted by several fields logged as key-value pairs. All CEF messages have a header that includes the date and hostname, such as "Dec 21 11:06:34 host message." The data fields are formatted using a common prefix composed of fields separated by characters, such as the following example:

> CEF:Version|Device Vendor|Device Product|Device Version|Signature ID|Name|Severity|Extension

The extension part of the CEF message acts as a placeholder for additional fields, allowing for huge flexibility regarding what can be included.

Common Log Format

The NCSA Common Log Format is a fixed log format used by web servers when they generate server log files. This format is named after NCSA_HTTPd, which is a discontinued web server software, but it is still used by many administrators today. Every line in Common Log Format is stored using the following standardized syntax:

> *host ident authuser date request status bytes*

The following is an example using the common log format.

Click here to view code image

```
192.168.2.1 user-identifier joey [15/Oct/2019:21:54:12 -0700] "GET
/apache_pb.gif
HTTP/1.0" 200 4782
```

Extended Log Format

Extended Log Format (ELF) is similar to Common Log Format, but ELF files are more flexible and can contain more information. Example 5-3 demonstrates an ELF file.

EXAMPLE 5-3 ELF File

Click here to view code image

```
#Version: 1.0
#Date: 20-Jan-2020 00:00:00
#Fields: time cs-method cs-uri
00:34:23 GET /foo/bar.html
11:21:16 GET /foo/bar.html
11:45:52 GET /foo/bar.html
11:58:34 GET /foo/bar.htm
```

Other Log Formats

There are other log formats you will encounter as you work with computing systems. Other formats may be also considered universally accepted, allowing for easy parsing of data. Some vendors use proprietary formats for their logging, which is becoming a taboo characteristic because it can cause vendor lock-in. In the past, customers would tolerate vendors using proprietary logging (designed as a way to control how a vendor's tools are used). Today, many vendors offer universal logging as a selling point against any system that can't function in a vendor-sharing format. Any time you evaluate a security tool that you plan to import data from, *I highly recommend* that you include questions about how easy it is to produce logs with the tool, what data can be included in a log, and if universally accepted logging formats are available. This applies even if you don't have a centralized log management solution today. You might find a need for collecting logs in the future or supporting vendor integration, and you don't want to discover that you are stuck in a vendor lock-in situation.

I've referenced a centralized log collection tool many times in this chapter but haven't called out a specific category of technology. Now that I have covered the common log formats along with the data they contain, it's time to look at the most universally used tool for centralizing security logs, which is a SIEM solution. The term SIEM has already been defined in this book based on the value it provides, but I have not peeled back how a SIEM solution is configured to provide that desired value. In the next section, I will cover how a SIEM solution collects data and converts it into actionable intelligence. It is useful to understand individual log impact within the SIEM; however, a more

important use of logging is applying concepts such as normalization and correlation across multiple logs, which is one of the values obtained from using a centralized security tool such as a SIEM solution.

Security Information and Event Management

Security information and event management (SIEM) is a category of tools that collect log and event data generated by IT systems, security devices, and applications throughout an organization's infrastructure. SIEM tools centrally collect data, normalize it, categorize it, and analyze it with the goal of providing actional intelligence to an analyst. The focus on security information management (SIM) means offering the ability to quickly search through large amounts of information, also known as *data mining*. Imagine needing to find details only on systems doing a specific thing within thousands of system log files containing tons of data. The focus on security event management (SEM) is more related to dealing with events associated with an incident response service. I find that all SIEM solutions are not created equally—some are better suited for SIM functions, while others are better at SEM. For example, Splunk and QRadar are quite different in their approach to offering SIEM solution capabilities. Both offer SIEM solution, but Splunk leans more toward SIM while QRadar leans more toward SEM, in my opinion.

SIEM Data Processing

Different SIEM solutions offerings approach data mining and event management differently; however, they all have a similar process for collecting and digesting data. Understanding this process is important not only when you are configuring a SIEM to digest data, but also if you have to troubleshoot a SIEM because it is returning results from the collected data that are not what the data analyst expects to see. Possible errors could occur at any point of the SIEM solution data processing lifecycle. I find it is easiest to remediate common errors early in the SIEM usage lifecycle, meaning adjusting how data is being sent or modifying how the data is parsed by the

SIEM solution versus trying to make good results with bad data. Remember what I stated at the beginning of this chapter: garbage in, garbage out. You always want to first ensure that the right data is being sent to the SIEM solution and then ensure that the system knows how to process it based on how the system can accept data.

Let's walk through a generic SIEM solution data-processing workflow:

Step 1. Data parsing: When the SIEM solution starts receiving data, it first attempts to identify the data format. Identifying the format enables the SIEM solution to parse data fields and link associated values, meaning a field for a numerical amount would have an associated number value. The format used by the data sender will be different based on the type of system and the available options for exporting data, which can cause problems for a SIEM solution if it can't determine the structure data is exported in.

A SIEM solution includes a log parser that converts imported data into a structured data format that the tool can understand. Many SIEM solutions provide parsing for common data sources; however, sometimes parsing has to be customized in order for the data to be converted to the proper structure that the SIEM solution can work with.

Step 2. Data normalization and categorization: The normalization process merges events containing different data elements into a consistent format that contains only common and useful event attributes. Event attributes include time, network address, operation performed, and so on. The categorization process involves adding context to the events. This can include system events, authentication data, local operations, and so on. For example, a SIEM solution can group data about hosts, data about network devices, and details from security tools into different buckets based on how the data will be analyzed. This sorting process is critical for converting data into actionable intelligence.

Step 3. Data enrichment: The SIEM solution applies additional information to make the data more useful. Data enrichment examples include associating data to a known malicious source, flagging unusual times of use for data, or adjusting new data being processed based on any

number of factors the SIEM solution has previously encountered.

Step 4. Data indexing: I opened this chapter by explaining how organizations generate tons of data. In order for a SIEM solution to offer data mining, it needs a way to develop an index of common attributes across all the tons of data collected. A data index enables the SIEM solution to search for key terms and quickly collect associated results instead of scanning all raw data. This capability is critical for a SIEM solution to be able to search for things within a reasonable timeframe. If a SIEM solution takes a long time to mine data, it is likely not indexing data correctly or doesn't offer indexing.

Step 5. Data storage: Once the data is indexed, it is stored so the SIEM solution can access it. Traditional storage meant storing all events on a centralized system for a specific period of time and archiving data that passes that specific point in time. Today, organizations are leveraging data lake technology such as Amazon Simple Storage Service (Amazon S3) to allow for cloud storage of data and more flexible data backup. Data retention requirements will dictate how long data is stored, the format in which it is stored, and where it is stored.

The previous section covered different possible log formats that can be used by IT systems. Most security tools designed to import logs (such as SIEMs) support universally accepted formats such as CSV, JSON, or Common Event Format (CEF). Whatever format is used, the SIEM solution needs to identify the fields for specific data items, such as the time of the log and the IP address associated with the log. Sometimes the format has the variable followed by the result, and other times these items are flipped around. For example, a CSV log can show the date and time an event occurred as 3/31/20 9:40:55 EST, represented as date, time, and time zone. If the SIEM solution expects these three items in a different order, it might show the wrong details, cause an error message, or have other negative impacts. Remember that many formats offer flexible messaging, which means you have to customize the data in the right format, or you will have bad results.

The example of not properly formatting time represents an issue that likely will lead to a really bad situation. Time has a major impact on how log data is leveraged by a SIEM solution. If an attacker is able to change the time

associated with logs, the attacker's digital footprint can go unnoticed as well as be completely erased. Imagine an attacker changing the time of logs to show that events that occurred in the last few hours occurred a few years ago. A SIEM solution would see these logs and respond differently, which would include archiving or flat out deleting the logs based on its belief that it's old data! Not parsing data correctly can cause a similar situation where the SIEM solution is configured to store only current logs; it might interpret the date in the time field within the log as being old and delete the logs. Following proper troubleshooting steps will ensure that you catch a situation like this as you set up the use of a new data source.

Data Correlation

SIEMs will perform additional actions once data is indexed to help an analyst obtain value from what was collected. A SIEM solution performs data deduplication and event consolidation post indexing. The goal for this step is to reduce the amount of data delivered to an analyst as they run reports or monitor events within real-time widgets. This can be especially helpful for repetitive data, such as a repeating alarm. Instead of sending the same alarm log to its dashboard, a SIEM solution typically converts multiple alerts into a single event and includes a counter of how many times the alert has been seen.

> **Note**
>
> There are some situations where you will *not* want a SIEM system to leverage deduplication, such as for forensic purposes.

SIEM solutions also apply *event correlation* to data, which involves comparing events against other event data to identify a larger event. Event correlation works by aggregating multiple events from a single source or from multiple sources based on specific criteria such as IP address, UserID, or a time window. SIEM solutions use event correlation rules to help the decision-making process based on a set of aggregated data or the sequence of

specific conditions in an event stream. Many correlation rules are created using if/then logic, meaning that if certain events occur, then flag that as a correlation event. For example, if port scans and data exfiltration alerts are associated with the same IP address, it's possible both events are part of a larger incident, which would be labeled with a higher severity of risk because, when combined, these two events represent a network breach. This example correlation rule can be coded as "if a port scan flag and data exfiltration flag are seen with the same data, produce a new alert named correlation event." Event correlation is a critical value obtained from centralizing data within a SIEM solution. I stated earlier that one huge mistake less mature security organizations make is having an isolated view of individual security events. Event correlation attempts to automate the process of looking across multiple events, allowing for better event data to be delivered to the analyst.

To give you a better idea of how correlation rules can be leveraged within a SIEM solution, here are several correlation rule examples specific to firewall event logs:

- **Rogue name server:** Best practice is to have organizations configure their users' endpoints to leverage internal corporate DNS servers. Correlation rules in a SIEM solution can be set up to monitor UDP/TPC port 53 or the DNS applications and alert the SOC when the destination is not the internal DNS server.
- **Rogue proxy servers:** Perimeter/gateway security tools should include controls to filter traffic attempting to connect to malicious locations or content that violates corporate policy. Web proxies can be used by end users to bypass gateway security, putting the organization at risk. Correlation rules can be set up to look for rogue proxy communication behavior by monitoring outbound traffic with known proxy resources.
- **Spam bots:** All corporate email must be delivered through corporate SMTP relays unless your organization is using a SaaS service. Correlation rules can be set up to look for SMTP traffic seen from a different SMTP server. Seeing this behavior could be an indication of an email compromise and should be investigated by the SOC.
- **Server compromise:** Client/server technology allows clients to connect to servers on TCP ports up to 1024 using a source port of 1024–65535. A

correlation rule can be set up to monitor for clients using 1024–65535 and connecting to TCP 80 or 443 with the purpose of identifying malicious parties leveraging internal resources over uncommon ports. Tuning will be needed to accommodate installation and updates, but there is no harm in monitoring for this activity.

Data Enrichment

Some SIEM solutions can apply even more data enrichment to data that is processed, on top of the context rules and data consolidation. Compliance checks are an example of a specialized correlation that certain SIEM products offer. For example, a SIEM solution can include the ability to validate if a set of numbers matches the algorithm used for a USA driver's license as it processes data from various systems. In this example, the SIEM solution is not just doing a pattern match; it is using a pattern match as the initial check, but once a potential set of numbers looks similar to a USA driver's license number, the SIEM solution validates if that number could exist as a driver's license within one of the states based on how each state has a special algorithm used to generate a driver's license number. It is important to include this level of granularity in compliance-checking capabilities to avoid overloading the analyst with false positives. Popular compliance validation rules found within tools such as data loss prevention or data at rest technology include checks for Health Insurance Portability and Accountability Act (HIPAA), General Data Protection Regulation (GDPR), and Payment Card Industry Data Security Standard (PCI DSS) compliance, along with more specific checks for credit card numbers, Social Security numbers, and custom data tags configured to be searched for.

Figure 5-4 represents a generic overview of how a SIEM solution processes data. Notice the number of logs that are initially collected compared to what the analyst is presented with after the processing is completed. There would be no way for an analyst to review the volume of data sent from security tools without the SIEM solution reducing the data to a digestible workload. Any issues with how a SIEM solution processes data may result in either too many events being presented to the analyst's desktop; events that matter being filtered out of the final result; correlation rules not triggering, causing events

to be missed; or overall poor reporting from the SIEM solution.

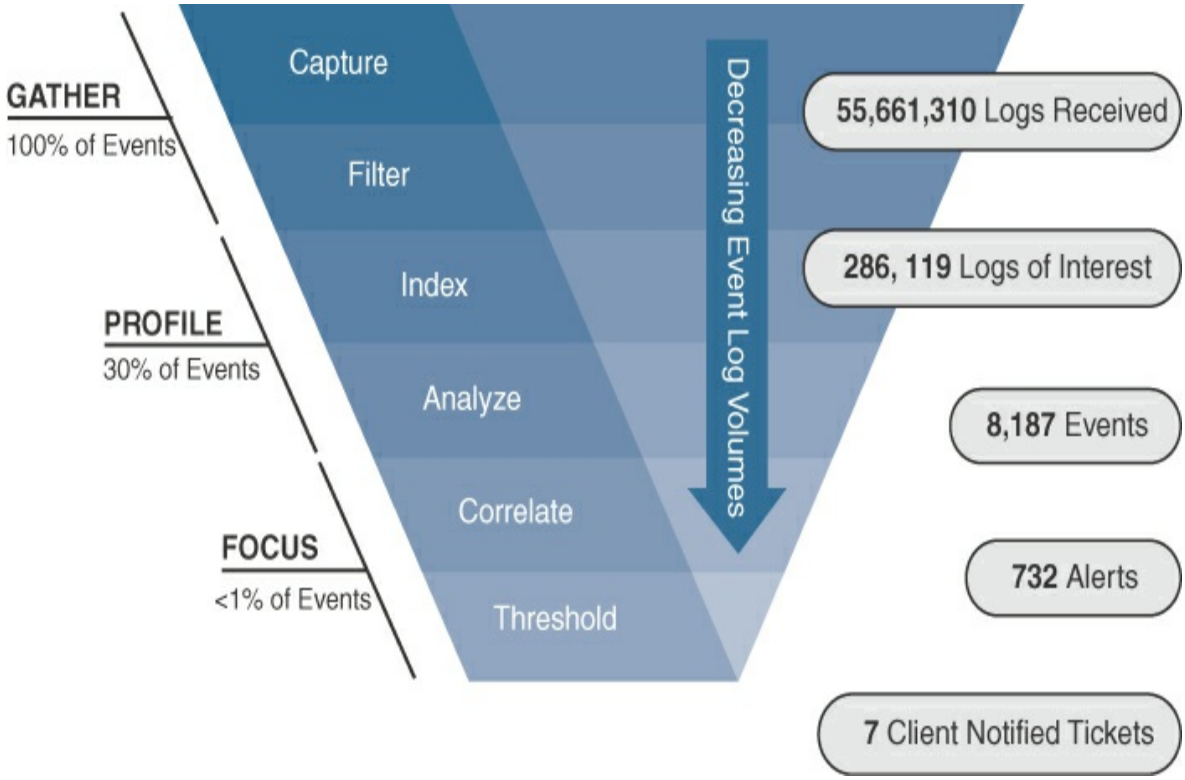

FIGURE 5-4 SIEM Data Digest Flow

SIEM Solution Planning

The key for successfully abstracting value from a SIEM tool is to have the best possible results end up at the bottom of the inverted pyramid shown in Figure 5-4. What are the best results? That depends on why you are using the SIEM solution. If you don't have a plan, you don't have a way to quantify the value provided by the SIEM solution. As simple as this sounds, I have encountered organizations that have only very generic goals for their SIEM solution, which leads to difficulty quantifying its value because they can't identify the purpose of using the SIEM, much less whether the tool is fulfilling that purpose. A better approach is to limit your focus for your SIEM project to a few specific goals that can be easily measured for value. There should be a team assigned the SIEM project, which will develop the goals and ensure they are accomplished.

Every SIEM project should follow these steps to ensure success for the project:

Step 1. Assign a team: This team owns the responsibility for achieving the following steps.

Step 2. Define clear goals, use cases, and requirements: As with any tool, the team needs to answer *why* the SIEM solution is needed. This includes plans for how to run, administer, and use the SIEM solution. The scope should be created prior to acquiring the SIEM solution.

Step 3. Target a few use cases: Do not try to accomplish too much in the early stages of using the SIEM solution. The team should pick a few specific use cases and scope the project around what the results should be for those use cases.

Step 4. Define data collection, retention, and reporting: The team should identify what will send data to the SIEM solution, where and how long the data will be stored once collected, and what type of reporting will be produced by the SIEM solution. The team should specify the number of devices sending data and how the SIEM solution should scale in terms of events per second, storage, and compute power. The team should also ensure that the data is structured properly and the data is directed at the specific target use cases for which it is intended.

Step 5. Develop an initial 6- to 12-month roadmap: The roadmap explains the goals of the SIEM solution for each phase of implementation, with a result of accomplishing five to seven use cases by the 12-month marker.

Step 6. Follow an output-driven model: The team should develop a plan for what the results will look like once data is processed by the SIEM solution and how it will impact the service that will be using it. Simply put, does the SOC receive the intended value from the results? If so, why and what data is needed to obtain that value? Sometimes this question can’t be answered until the SIEM solution is functioning and the results are seen by the analyst. Regardless, you need to have a general idea of what you are trying to produce

and let the analyst provide feedback as you tune the SIEM solution.

Note

One common question I often receive is how to properly size a SIEM solution. There is no universally correct answer because SIEMs are based on the number of events rather than on something easy to track such as the number of devices sending data. This means you can't say that an organization with 1000 devices will always need a SIEM solution that is larger than what is needed by an organization with 100 devices. If the organization with 100 devices generates more events than the organization with 1000 devices, the former will need a larger SIEM solution. A more accurate way to answer the events per second question is to deploy a tool that can measure an accurate events per second count. I recommend working with a SIEM solution provider and running a proof of concept or, if you are using an open source option, being prepared to size up or down depending on the events per second you capture during testing. Also make sure to accommodate future growth in events, which should be at least a 25–50% growth flexibility on top of whatever number you feel is correct for your organization today.

SIEM Tuning

Another important practice is to periodically evaluate how your SIEM solution is functioning, even if the SOC analysts report that the results they are producing are acceptable. There are many ways a SIEM solution can achieve the same results, and some approaches are more ideal for maximizing your SIEM solution investment. Tuning your SIEM solution could produce cleaner results and require less processing power, leading to better SIEM performance. Also, cleaning up redundant and poorly structured rules in your SIEM solution will allow for more flexibility when reports and dashboards need to be modified. I have seen situations where customers can't change how their SIEM's dashboards look without having to rebuild all of the

customized rules that have been pieced together over a long period of time.

The following are four foundational tuning points to consider regardless of the quality of the results produced by your SIEM solution. Keep these points in mind as you tune your SIEM solution by evaluating what tools are sending data to your SIEM solution as well as the type of data it is receiving.

- **Key tuning point 1:** Many SIEM solution providers bill based on events per second or the amount of data being stored. This concept translates to the more you send the SIEM solution, the more you will have to pay to use it. SIEM vendors can be a bit sneaky regarding what they offer by providing low-cost hardware and various capabilities at little to no cost with the intent of attracting customers to send all of their data to the SIEM solution offering. By doing this, the SIEM vendor is able to increase the SIEM bill or license cost based on the amount of data received. As companies become more reliant on the SIEM solution, they send more data to it, and the SIEM solution bill continues to increase.

  When a SIEM vendor charges in this manner, tuning what is being sent can save you a ton of money. Consider performing a data assessment and limiting what data is being sent to a SIEM product that charges based on number of events or data size. Another option is to use a different tool that can perform similar processing and data reduction as the SIEM solution, so that you send only useful data to the SIEM tool that bills you for data load. Looking back at Figure 5-4, if you have a tool that can perform the capture and filtering tasks at no cost, having the results sent to your SIEM solution would dramatically reduce the overhead on the SIEM solution as well as your cost to use it! For example, NetFlow-based security tools such as Plixer and Cisco Secure Network Analytics (formerly Stealthwatch) both offer data deduplication capabilities. Rather than sending raw NetFlow data to the SIEM solution and having the SIEM solution perform deduplication, which increases the events per second, you can use a dedicated NetFlow security tool to clean up the NetFlow data before it is sent to the SIEM solution, dramatically reducing the event per second load on the SIEM solution.
- **Key tuning point 2:** Manually tuning and adjusting a SIEM solution takes time and sometimes specialized skills. I've found that some SOCs use complicated customized scripts and widgets that filter data into actionable intelligence, which can produce good results; however, any

changes to data sources or modifying how the data is processed will be painful. Rather than overcomplicating things by sending all data to a SIEM solution and filtering out what you don't need by using complicated logic, it is much easier to tune how data is being filtered upon connection or prior to sending the data so that you reduce required filtering at the SIEM solution processing level. I've seen the nightmare many SIEM solution managers experience—something changes regarding how data is received, and all of the SIEM solution displays and reports become instantly useless! Following this key tuning point also addresses the concern of overloading the SIEM solution, which will cause additional problems outside of overbilling.

- **Key tuning point 3:** Using only the default correlation rules provided by the vendor does not maximize your SIEM solution investment. Vendors such as IBM QRadar offer over 500 default correlation rules, which are good rules, but they all apply to generic use cases. In contrast to key tuning point 2, you do want to take the time to develop the skills to customize correlation to meet your needs (but without overcomplicating things). The combination of key tuning points 2 and 3 means you are tuning what comes into your SIEM solution as well as adjusting how that data is being processed to maximize your investment. What you don't want to do is allow for tons of false positives and other noise to exist within your final result or be satisfied with subpar results based on what is produced by default correlation rules. If that is occurring, you need to customize the correlation rules so that you have actional intelligence relevant to your SOC's needs.
- **Key tuning point 4:** Assess and reassess your results. I've found that organizations that do a few hundred threat investigations over a month typically have a common set of data that is used within the SIEM solution. That common set tends to make up a certain percentage, and if it's in the 30–40% range, that means a large majority of data within the SIEM solution is never used. By continuously reassessing the results of the SIEM solution, the SOC can determine if certain data being collected is relevant to the services being provided. Also, new data types can be considered and introduced as part of the assessment process, slowly improving the SIEM's alignment to the SOC service goals over a period of reassessment cycles.

Remember these four key tuning points as you evaluate a newly installed SIEM solution or tune an existing SIEM deployment. A short way to summarize utilizing these four key tuning points is *run, watch, tune, and reassess*. This means you pick a use case, run the data through the SIEM solution, and watch to see if the results meet your business goals. After you receive the results, you make any necessary adjustments, and repeat the process. Tuning includes all of the concepts covered in this section, ranging from adjusting how data is being sent to the SIEM solution to how the SIEM handles the data. If all efforts lead to poor results, you will need to troubleshoot or call upon the vendor for help. This leads us to the topic of how to address when your SIEM solution is producing the wrong data.

Troubleshooting SIEM Logging

One of the most frustrating situations is investing time and money into building a SIEM solution only to find that it produces poor results. Poor results could mean having too many false positives, leading to the SOC ignoring alarms based on the belief that the SIEM solution isn't providing alerts about real problems. I have seen this happen in SOCs I have consulted for. Customers will point to alarms on their SIEM dashboard and say, "Oh, that happens all of the time . . . it's nothing." I am baffled about why a SOC would spend the time to display data if it means nothing to the SOC. What value does an alarm that you ignore provide? I have also seen a similar situation where too many general alarms are displayed in the SIEM solution dashboard, overloading the analyst with event data to the point where no actions can be taken. If a data overload problem is occurring, the SOC is better off not using the SIEM solution until tuning can be applied to reduce the results to a reasonable level. A SIEM solution is a distraction and serves little to no value whatsoever.

Yet an even worse problem I have encountered while meeting with organizations is a SIEM solution that produces inaccurate results. I have heard C levels of large organizations claim they have a good handle on their security events based on the low trends of event data reported by their SIEM solution. The truth behind such low trends typically tends to be one of several possibilities: the SIEM solution is not consuming data from all of the

organization’s security tools, the parsing within the SIEM solution is misconfigured, or certain parts of the organization do not have security tools available to generate data for the SIEM solution, representing blind spots in the SIEM’s view of the organization’s security awareness. One customer I evaluated had a misconfigured SIEM solution that was not parsing data correctly, causing important events to be missed and even dropped altogether. A situation like this benefits attackers because the organization they are targeting will have a false sense of security awareness. Any of these SIEM situations is bad for the organization and needs to be dealt with.

SIEM Troubleshooting Part 1: Data Input

Troubleshooting poor results from a SIEM solution starts with determining the type of problem you are encountering. Are you seeing the wrong results or no data whatsoever? Let’s start by addressing the situation in which you are not seeing any data within the SIEM solution. Anytime you encounter a lack of data, you need to start by troubleshooting the basic connectivity between the data sender and the data receiver. Every engineer I’ve worked with has a story about how they spent hours performing technical troubleshooting before realizing that a physical cable or power outage was the problem, which they could have identified within a few minutes by testing basic connectivity first. Start your troubleshooting by validating that you can ping each device and that no new security or networking changes have recently occurred that could impact the connection between the data sender and your SIEM solution. This includes configuration changes to tools such as firewalls and IPSs that might allow ping but block other protocols being used by tools to send data to the SIEM solution. One important test to ensure that communication is occurring between the data sender and the SIEM solution is to search the raw logs in the SIEM solution and filter on the IP address of the data sender to ensure that something is being received. If not, you likely have a network problem, a problem with the tool sending data, or a problem that is causing your SIEM solution to no longer accept data from the sender. If you see any data from the sender within the SIEM, you can rule out the network and start looking at what data is being sent.

After you have validated connectivity, the next step in troubleshooting poor

results within your SIEM solution is to evaluate what data is being sent to the SIEM solution. Most tools that can export logs include steps to do so in universally accepted formats. Even tools that use proprietary formats often include universally accepted options as well to accommodate third-party tools like SIEMs. Security tool vendors who market their tools as being capable of adding value to a SIEM typically not only explain how to do so but also include configuration examples of sending logs to market-leading SIEM providers such as Splunk, IBM QRadar, LogRhythm, and ArcSight. If a tool vendor doesn't provide that information, I find that between searching Google and YouTube, I am always able to find examples of how a configuration is performed along with how to troubleshoot common problems. This concept leads to how to start step two for troubleshooting poor results within a SIEM solution, which is to do some research to find what formats are available in the sending tool, how to configure the SIEM solution to accept that format, and any examples of the tool of choice being configured to send data to the SIEM solution you are working with. Research can also include contacting the vendor if the problem seems to be more complicated than what your research reveals about the situation. I have saved countless hours of troubleshooting SIEM configurations by simply finding somebody else's research and workarounds on the Internet.

Research will help answer questions such as what formats are available to send the data to your SIEM solution, how those formats should be adjusted, how to configure the SIEM solution to accept data, and how to solve any common problems. Looking back at the discussion of how data can be formatted, many universally accepted formats such as syslog allow for messages to be delivered however the sender desires. This flexibility can also lead to poor results if the SIEM solution doesn't understand how log data has been formatted. You will need to verify during your research if a data template is available to help accommodate expectations the SIEM solution has for receiving data. Do not assume the templates provided by the sending tool are the best option to use as you configure the tool to send data to your SIEM solution of choice. You must test and tune any template. I find that many security tool templates often do not work well within the SIEM solution they are sending data to without some customization and tuning. Later in this section, I will show a real-world example of this that I had to troubleshoot for a customer. Remember to *run, watch, tune, and reassess* as you work to convert data in a SIEM solution into actional intelligence.

SIEM Troubleshooting Part 2: Data Processing and Validation

Once you validate that your SIEM solution is properly collecting data, you need to verify how that data is being parsed by the SIEM solution and validate that the data is being delivered to the parties for which it is intended. One way to view how data is parsed is to search for all data from only the sender of interest and see how the SIEM solution has parsed and indexed the variables. A properly parsed and indexed dataset will have data points of interest properly separated and searchable, allowing for context rules, dashboard widgets, reports, and other uses to be more valuable to the SOC. If the results of how the data is being used are not correct, try evaluating how the rules, widgets, and reports are set up. You might also need to go back to the last step and verify that the format of the data being used is one that is accepted by the SIEM solution. It is possible that a modification to the SIEM configuration, update to the SIEM code, or change in data is causing the logic behind those configurations to not function correctly. A common situation I encounter is one in which existing rules, widgets, or reports that once worked start producing poor results. A common situation I encounter is one in which existing rules, widgets, or reports that once worked start producing poor results. This typically results from using customized rules that rely on very specific factors that can change. To troubleshoot when this occurs, you will need to open up the search that is used to create the desired results and validate whether it is still filtering in on what you expect it to see.

The second part of the data process troubleshooting is the validation of the results. I mentioned how very specific customized rules used for filtering can cause a problem if the dataset is changed; however, that does not mean customization and filtering is a bad thing. Remember that the value of a SIEM solution is based on how the data is used. This means you need to validate the right parties are getting their intended data. During a new SIEM deployment, you might find that the right data is being sent, but that it is being sent with lots of additional data, causing the analyst to have to manually filter out the noise. Part of the validation process needs to include adding customization and filtering to how the SIEM solution processes data to improve the quality of what is delivered to the analyst. Over time, the tuning will allow the SIEM solution to provide value to the analyst more

quickly by delivering only what they need for their mission.

One concern I hear when I speak to some organizations about filtering out data is that data that is not useful now might be useful in the future. There is nothing wrong with that mindset; however, your analysts shouldn't have to manually filter out data that is not being used as part of their daily tasks. An option is to place customized filters on data delivered to a specific analyst while using a catch-all dashboard for displaying general data. This allows the analyst to flip between data specific to their analysis and a more generic dataset. This is common practice and a way to accommodate different views of the same larger dataset quickly.

Figure 5-5 shows a dashboard I built as a lab project for the SOC of a fake hospital called hackMDs. This dashboard shows data from NetFlow; however, it filters out all IP addresses outside of what the SOC analysts are responsible for monitoring. Another dashboard can be created that takes into account all hospital IP addresses, or an analyst can click Search & Reporting and search across all networks if they need to view NetFlow outside of what is prefiltered in the example dashboard. This example allows an analyst to start any investigation with a dashboard that has data filtered specific for their mission; however, the analyst also has the flexibility to search outside of the dashboard if operational tasks require them to do so.

FIGURE 5-5 Splunk Customized Dashboard Example

The following steps summarize the process of troubleshooting poor results within a SIEM solution:

Step 1. Test basic connectivity between the data sender and the SIEM solution.

Step 2. Verify that your SIEM is receiving some form of data from the data sender.

Step 3. Research which data formats the SIEM solution supports.

Step 4. Verify that the data sender is sending data in a format the SIEM supports.

Step 5. Validate within the SIEM solution how the data is parsed and indexed.

Step 6. Focus on the rule, widget, or report that is not functioning properly and verify the search logic being used to produce results.

Step 7. Verify the results meet your intent for the data.

Step 8. Make any adjustments you can make to resolve any of the previous issues.

Step 9. If you still find you cannot produce your desired results, contact the vendor for additional support.

SIEM Troubleshooting Examples

The best way to see the preceding SIEM troubleshooting steps in action is to work through some examples. This section demonstrates how to add data to and troubleshoot data issues within two marketing-leading SIEM solutions, which are Splunk and IBM QRadar. I choose these two vendors because I believe they have differing approaches to the SIEM market. Splunk tends to lead with a SIM approach, including a heavy focus on community applications, while QRadar has more of a SEM focus, with a particularly heavy focus on IT operations, including really nice asset-management features. To be clear on these statements, this is my personal opinion; I am a fan of both offerings and want to remain vendor neutral when possible. As for the data sender, the following examples use Cisco Stealthwatch and are based on a configuration that I set up and troubleshot for a customer. For these examples, focus on the process rather than the details specific to the technology being used. Although your setup likely differs, the general troubleshooting process should be very similar.

SIEM Troubleshooting Example: Data Storage

Let’s start with how a SIEM solution receives and stores log data from different types of systems. Looking at Splunk first for this example, the process starts with identifying the data source of interest, which Splunk sees as a data input. Splunk indexes whatever it receives and saves it either locally, remotely, or to the cloud depending on your deployment. If you plan to store the SIEM’s data on a remote or cloud system, you need a Splunk forwarder to send the data back to Splunk. From a troubleshooting connectivity viewpoint, make sure to verify where the data is being stored and validate that storage location is functioning properly when evaluating any issues regarding lost data. Splunk SmartStore is a popular data management

option for Splunk users that allows for compute and storage elasticity, meaning storage is adjusted as needed. Data can be stored locally or using AWS S3 or S3 API compliant object stores. Figure 5-6 represents what the SmartStore architecture looks like at a high level.

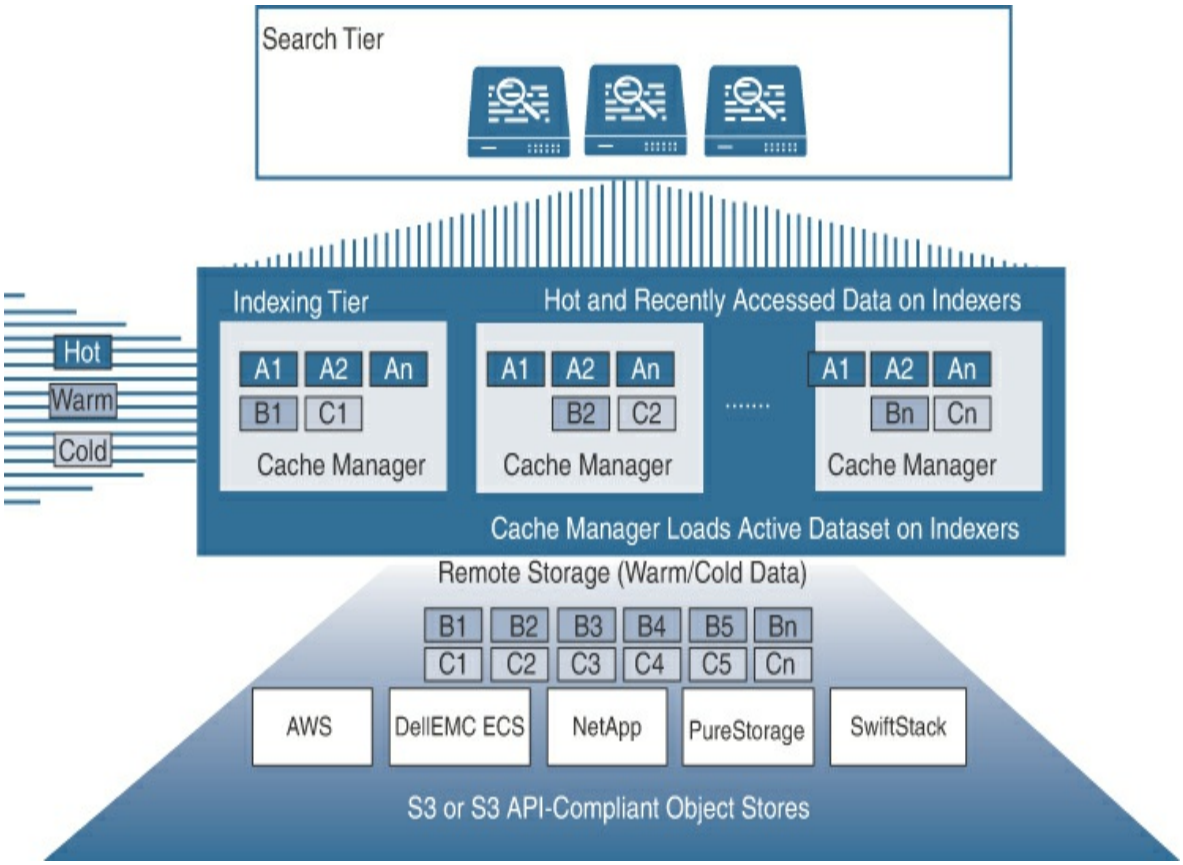

FIGURE 5-6 Splunk SmartStore Diagram

You might have other flavors of remote storage in your deployment, but the key point here is to ensure that data is making it to and from the remote storage option. For Splunk troubleshooting, examining event logs will provide insight into the communication that is occurring. If you are troubleshooting a SmartStore setup, the following list provides examples of specific logs that you can use to validate connectivity based on events being recorded:

- Investigate **splunkd.log**:
  - **S3Client:** Communication with S3

    - **StorageInterface:** External storage activity (at a higher level than S3Client)
    - **CacheManager:** Activity of the cache manger component
    - **CacheManagerHandler:** Cache manager REST endpoint activity (both server and client side)
- Investigate **search.log**:
    - **CacheManagerHandler:** Bucket operations with cache manger REST endpoint activity
    - **S2BucketCache:** Search-time bucket management (open, close, and so on)
    - **BatchSearch, CursoredSearch, IndexScopedSearch, ISearchOperator:** Search activity related to buckets
- Investigate **audit.log**: Contains information on bucket operations, such as upload, download, evict, and so on
- Investigate **metrics.log**: Contains metrics concerning operations on external storage.
- Investigate **splunkd_access.log**: Contains a trail of the search process activity against the cache manger REST endpoint

Note

See the docs.splunk.com community for more details on troubleshooting this specific issue.

A similar approach to troubleshooting remote storage can be applied with other SIEM vendors. IBM QRadar offers Fibre Channel, iSCSI, and NFS storage and cloud options for external storage. Troubleshooting involves validating that communication and data transfer is occurring between the /store, /store/arieal, and /store/backup file systems. See IBM QRadar for those details. If your SIEM isn't reviewing data, validate basic connectivity as well as how data is being stored if external storage is being used.

SIEM Troubleshooting Example: Data Input

After you have established your approach to storing data and connectivity and have validated that it is functioning properly, you need to set up how the SIEM accepts your desired data input. Looking at Splunk, Splunk categorizes data inputs as either files and directories, network events, Windows sources, or other sources. These data sources are configured as an input within Splunk. Some systems can require multiple inputs based on the type of data they will send. You must verify with the vendor of the tool that will be sending data which data ports the SIEM needs to have open to collect data from the tool. As stated earlier, you are likely able to find these details by searching for the SIEM solution and the tool, such as "Splunk Cisco Stealthwatch Configuration." Make sure to do your research before starting your configuration to avoid wasting time using the wrong configuration or data format! Figure 5-7 shows a screenshot of the possible inputs for my example Splunk deployment.

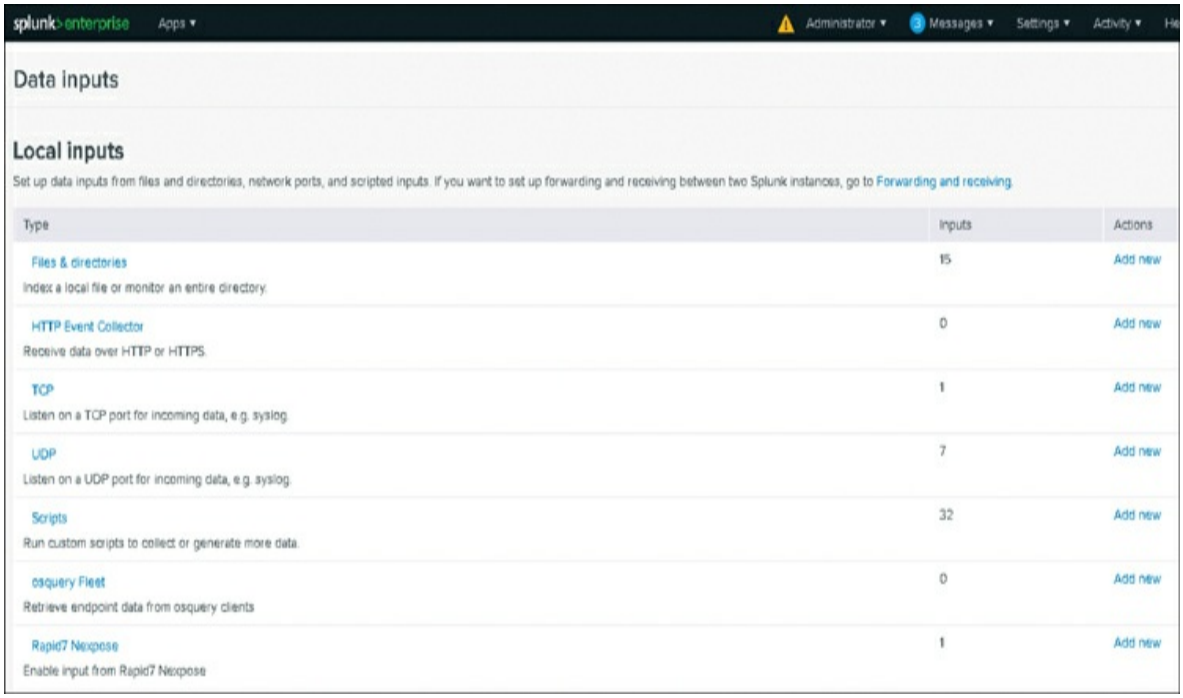

FIGURE 5-7 Splunk Data Input Options

Configuring and troubleshooting inputs starts with matching the data source with the type of input that is configured. For the next setup and

troubleshooting example, I'll review how to configure Cisco Stealthwatch to send syslog events to Splunk. Cisco Stealthwatch collects NetFlow from various tools in such a way that the network effectively becomes a security sensor. The Stealthwatch Management Console (SMC) is the centralized manager that can export logs to Splunk. There are other data export options; however, I will focus on this specific syslog use case.

There are two parts to performing and troubleshooting any configuration like this: validating the sender of the data and reviewing how the data is being received by the receiver. First, let's look at setting up the tool that will be sending data. I'll summarize these steps with a focus on what is common across many tools rather than diving into specific details regarding how to set up Cisco Secure Network Analytics (formally Stealthwatch) to export data.

Note

I highly recommend tuning your tools to send only the data that is needed, based on the results of a data assessment. This can save you many hours of tuning how the data is parsed, save SIEM process power, and potentially save SIEM costs. Do this up front rather than post deployment.

The list that follows provides a summary of the process to configure Cisco Stealthwatch to send syslog data to Splunk. The primary purpose of this example is to give you a basic overview of how many tools are configured to send syslog data to a SIEM solution.

Step 1. Log into the SMC Java applet.

Step 2. Navigate to **Configuration > Response Management**.

Step 3. Click **Syslog Formats**.

Step 4. Fill in the following required fields to set up the exporter:

Name: **Splunk**

Enabled: **Yes**

Facility: **16 - Local Use 0 (local0)**

Severity: **6 - Informational: Informational Messages**

MSG Part: **Use the message format needed for the add-on**

Step 5. Click **OK**.

Step 6. Click **Actions**.

Step 7. Fill in the following required fields to configure:

Name: **Splunk**

Enabled: **Yes**

IP Address: 198.19.10.15

Port: **514**

Format: **Splunk**

Step 8. Click **OK**.

Step 9. Click **Rules**.

Step 10. Change the action of your rules and add Splunk.

Figure 5-8 is a screenshot of building this example configuration in Cisco Stealthwatch. A key point to notice in this screenshot is the format I selected called Splunk_New, which is something that I had to configure prior to setting up this page or it would not have been available. Also, notice under the Format drop-down menu the Syslog Formats button, which you can click to create new or modify existing syslog formats. You need to pay special attention to this.

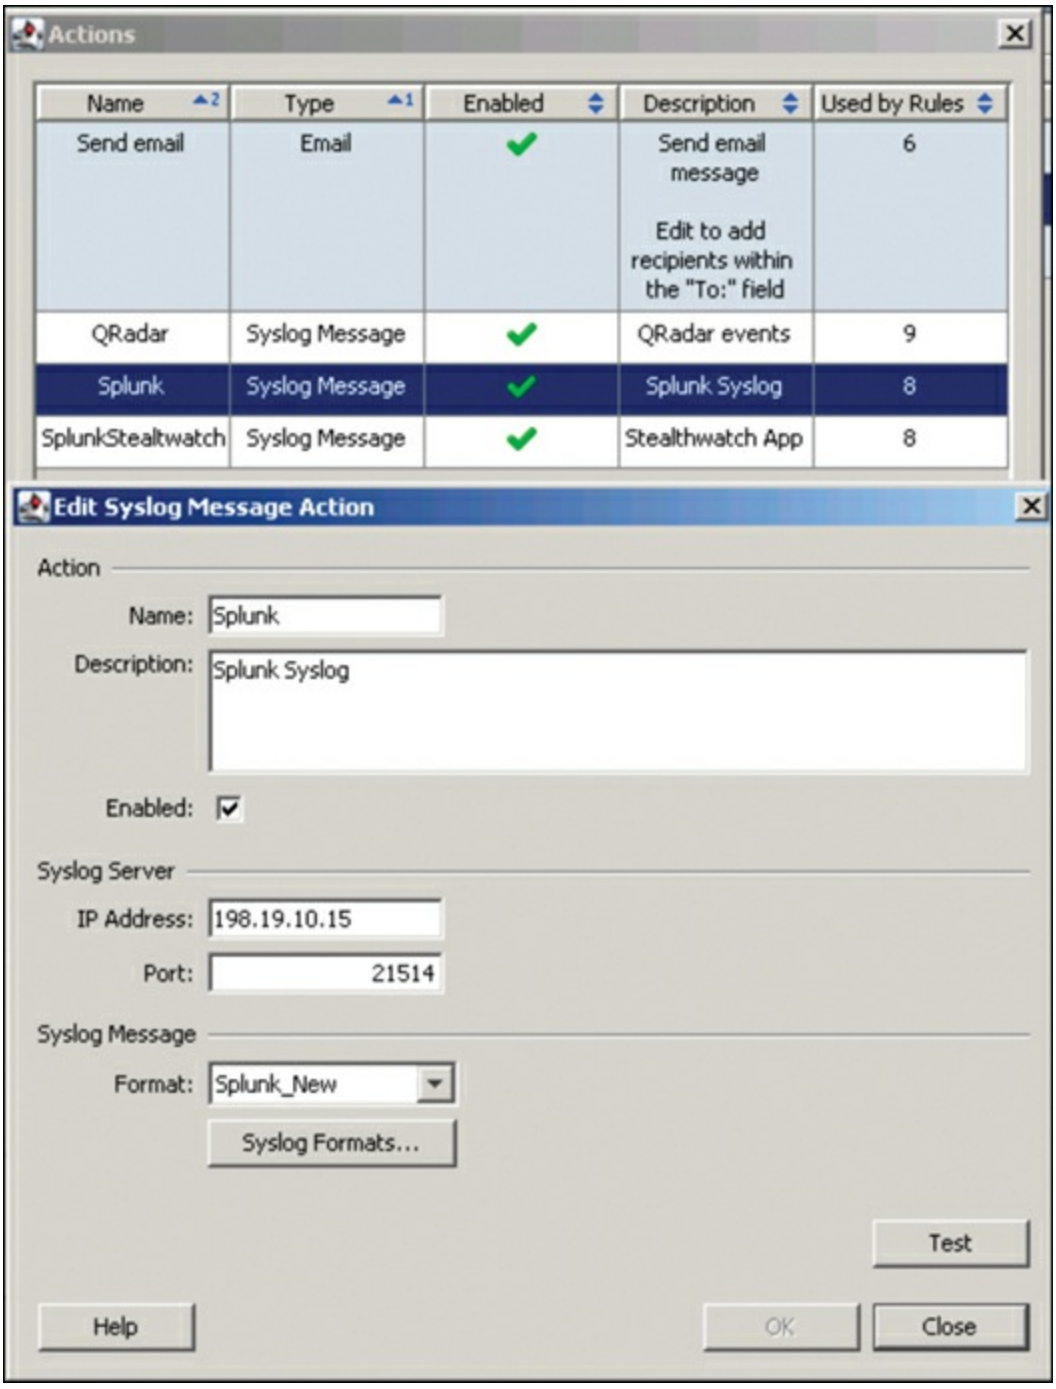

FIGURE 5-8 Cisco Stealthwatch Configured to Syslog Data to Splunk

Properly configuring the Format field is extremely important to the success of how the data is sent to a tool such as a SIEM solution. Remember that syslog is a universally accepted format that enables you to adjust content however you see fit. This is good for flexibility but bad for a tool such as a SIEM

solution that is attempting to accommodate the various ways syslog data can be sent to it. The IT market addresses the syslog format problem by offering templates. That is why you should always research how data should be sent from a tool rather than assuming default templates will work for whatever SIEM solution you are setting up to receive data. Regarding this configuration example for Cisco Stealthwatch, there are default template items you can use to build a syslog template. Looking at Figure 5-9, there are data options to parse, which you can select on the right side and copy to the left using the <- button. This approach is an attempt to control how the syslog will be formatted, but it doesn't work all of the time. When I used this approach to send syslog data to Splunk during a deployment I was contracted to do, I ended up with a bad syslog format, leading to poor SIEM output. Figure 5-9 is an example of doing it incorrectly using the default Cisco Stealthwatch parsing templates. If you highlight everything on the right side, click the <- button to add it over and use that format, a SIEM solution like Splunk will not know what to do with the data.

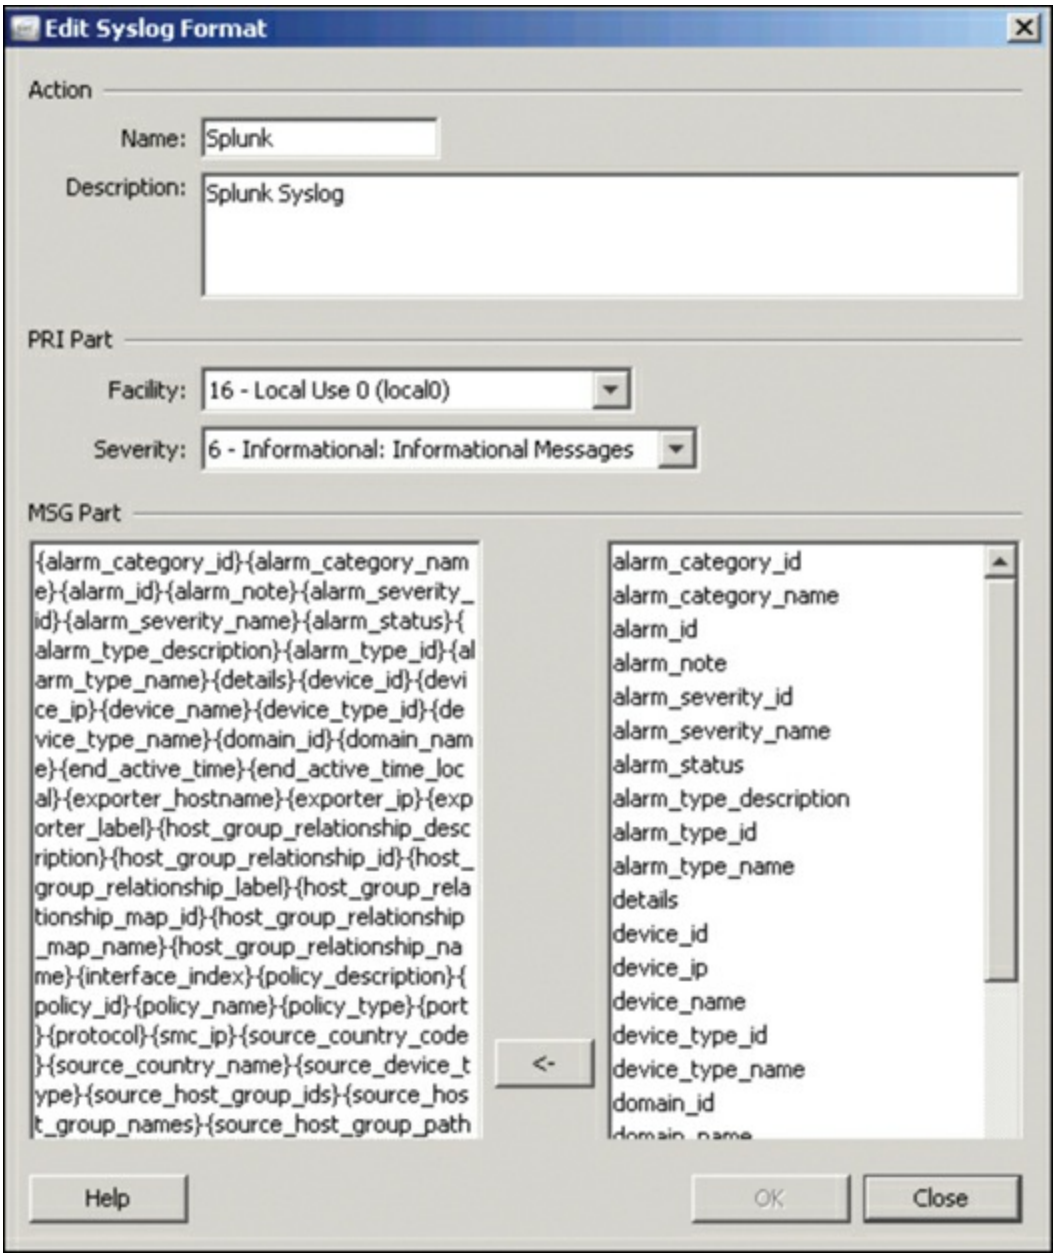

FIGURE 5-9 Using Default Parsing Template Within SMC

Figure 5-10 shows what happened when I used the SMC default syslog format. Although Splunk sees the data, it doesn't parse the data properly, so I can't use the data for any rules, widgets, or reports. The arrows in Figure 5-10 identify useful data within the raw data, but the SIEM solution doesn't see these important items as data points. Keep this key concept in mind as you troubleshoot a SIEM solution that has data indexing problems. Rather than creating complex custom parsers to accommodate a poorly formatted dataset,

first try to adjust the format of how the data is being sent so that the SIEM solution can accept it using its default data indexing and parsing. At this point, I can at least validate that connectivity is occurring and data is being received by the SIEM solution. However, I need to fix how the SIEM solution is digesting data from the SMC.

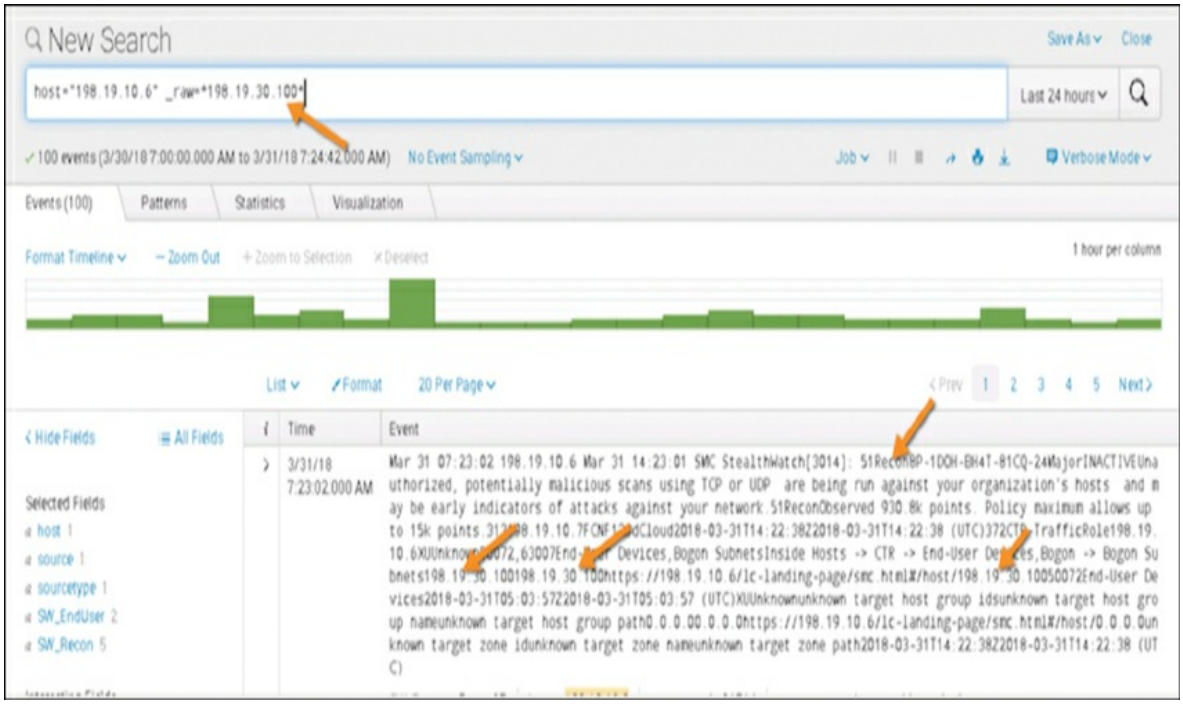

FIGURE 5-10 Results from Poorly Formatted Syslog

The best way to fix this problem is to identify how Splunk expects to accept syslog data. I found through researching Splunk and other SIEM vendor support pages regarding how syslog should be formatted that templates for syslog are available and recommended by various customers that have dealt with the same problem. Ironically, I found a very useful template on the IBM QRadar support page that works well within Splunk, which looks like Figure 5-11. I copied the template and pasted it into the left part of the SMC Syslog Format rather than using the SMC default template tool.

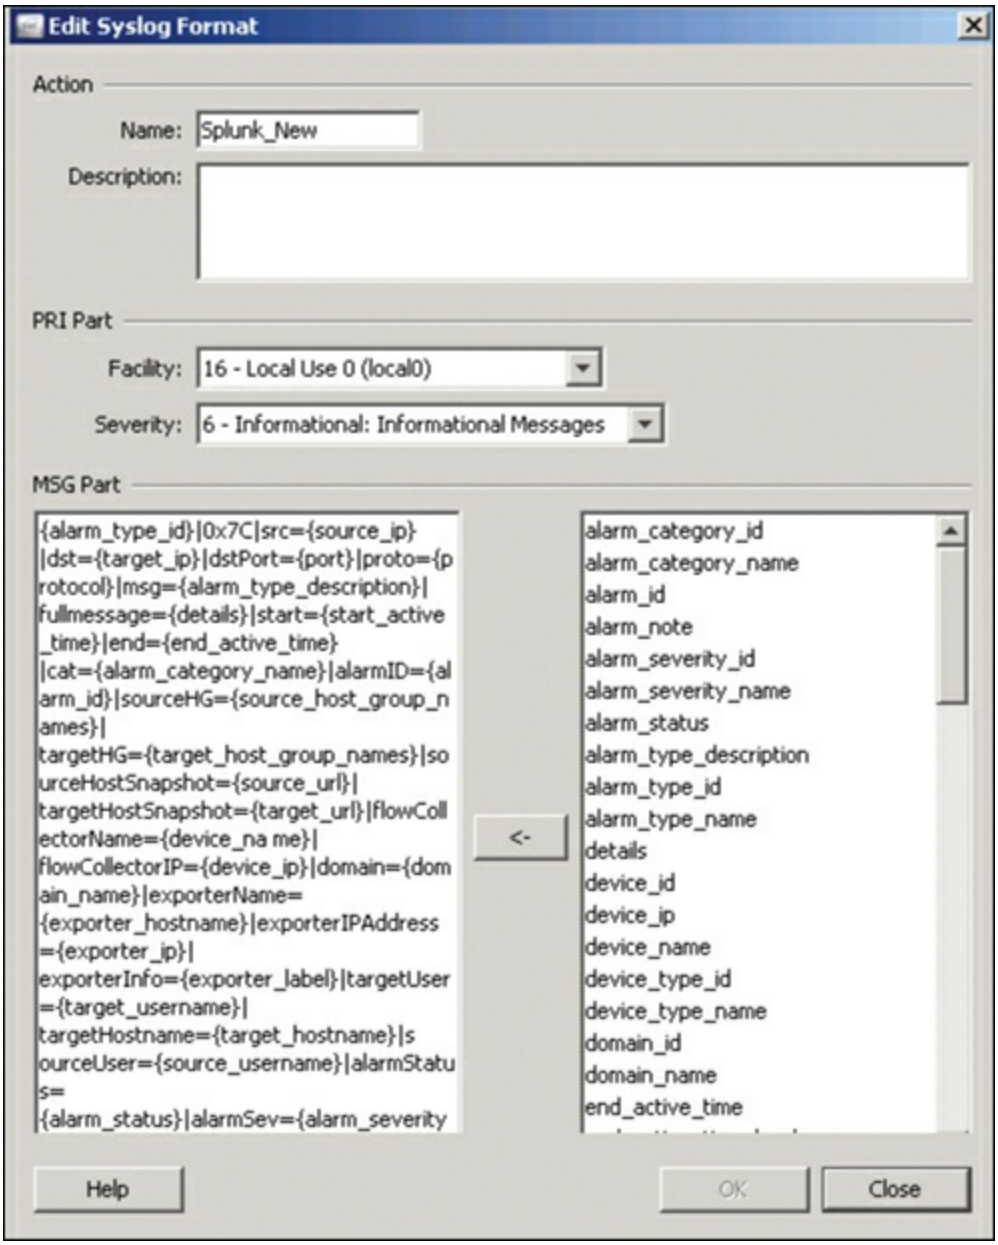

FIGURE 5-11 Using Custom Syslog Template in SMC

The difference between using the wrong and right syslog format resulted in the SIEM solution being capable of turning data into reports and widgets, such as what is shown in Figure 5-12, without any customization of how the SIEM solution accepts the data. After using the proper syslog format, Splunk was capable of identifying key terms, which I easily searched for and converted into widgets. Looking at Figure 5-12, I created these widgets by searching for terms such as Recon, Data hording, Exfiltration, and Concern

within the data being sent from SMC. It would be an absolute nightmare to create these widgets using the same data not indexed correctly due to the poor syslog format. Plus, any changes to how the SMC sends data could break the logic behind customized parsing and search rules, causing any of these widgets to stop showing useful data.

FIGURE 5-12 Converting Syslog Data into Reports and Widgets

When rules, widgets, and reports are not showing the correct data, always start your troubleshooting by validating how the data is being sent and received. If you find that data is being seen but not parsed or indexed correctly by the SIEM solution, try adjusting how the data is being sent to the SIEM solution before modifying or customizing the receiving end.

SIEM Troubleshooting Example: Validating Results

There are other problems that could occur outside of the situations covered in the previous section. For example, it is possible the wrong ports or protocols are selected by the sender of the data or the SIEM solution, resulting in unusual data being collected. Also, data being sent and received by the SIEM solution could be functioning properly but the data that is sent could be bad. Anytime everything regarding configuration seems to work but your results are not good, make sure to validate the raw logs and ensure that the data being sent is what is supposed to be sent. I once was called in to troubleshoot

a situation where a report was producing the wrong data. After some troubleshooting, I found that IP addresses of different tools had switched during a reboot, resulting in the SIEM solution being confused regarding expected data it was receiving. I also found that one tool stopped sending logs and instead was sending error data to the SIEM solution, causing associated reports to show random characters. The best way to validate the raw logs is to first look at a log being sent from the sender and compare it to how the SIEM solution sees the same data once it has been collected and indexed. Figure 5-13 shows the Log Activity page within IBM QRadar. I can search here for what data is being received by the SIEM solution by using the Quick Filter option or by performing more advanced searches.

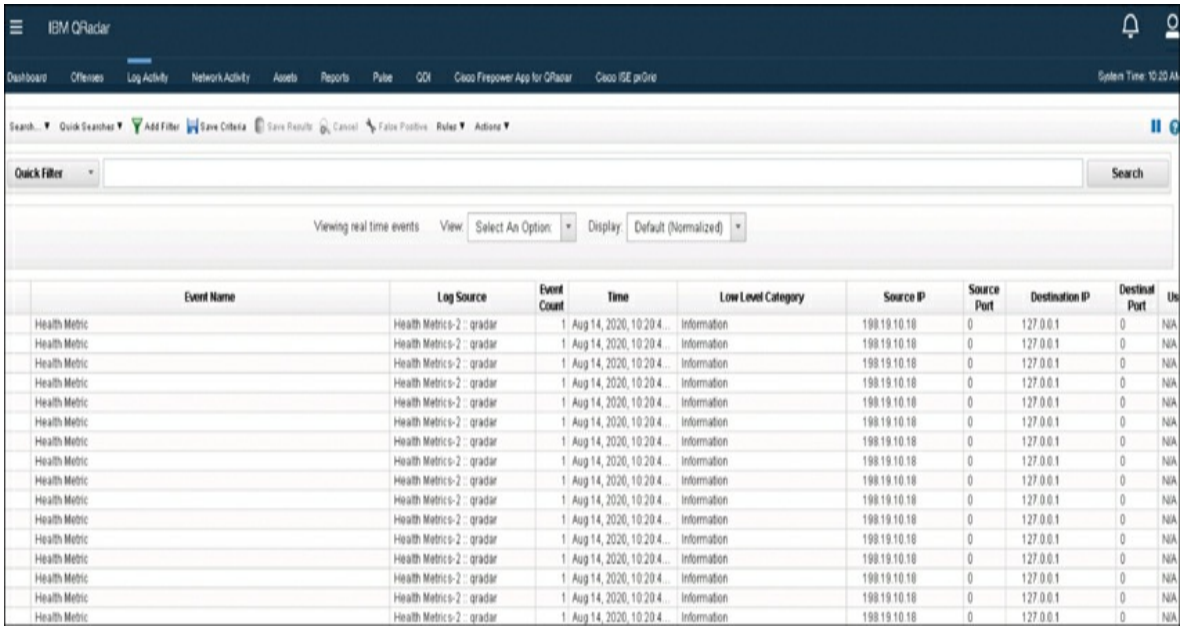

FIGURE 5-13 IBM QRadar Search Screen

If I want to see if data is making it to the SIEM from a device, I would start by searching for any data from a specific IP address. Next, I would limit the time to the last 10 minutes and narrow down on data items of interest that are needed in what I want for my success criteria. If my goal is to identify data exfiltration activity from a log coming from the SMC, then I would search for the term that would contain this value, which is likely "exfiltration." Figure 5-14 shows the results of running a search for the IP address of the SMC and searching for "exfiltration" as part of the search. If I didn't see any results

based on this search, I would revert to searching for all data from the SMC IP address and reviewing how the data is being indexed as well as what key terms are available for me to search for based on how IBM QRadar is categorizing data variables found within SMC data.

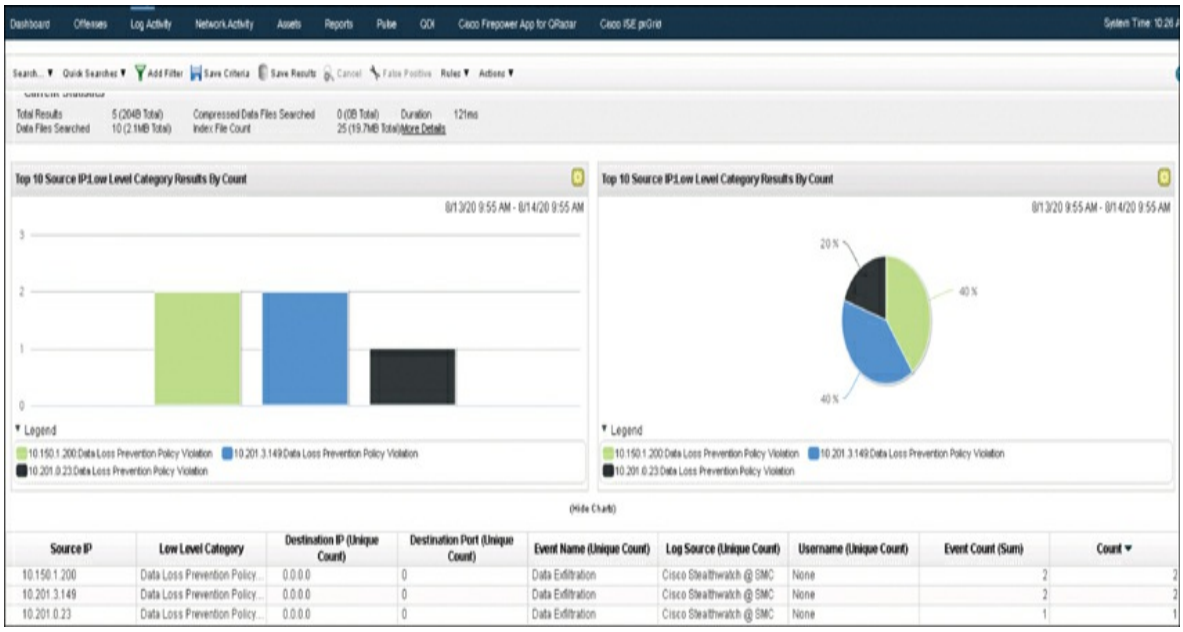

FIGURE 5-14 Searching for Data Exfiltration in IBM QRadar

SIEM Troubleshooting Example: Actionable Intelligence

The final phase of setting up and troubleshooting data being sent from a tool to your SIEM solution is converting your desired data into actionable intelligence. For many SIEM solutions, this means creating widgets and reports that filter based on what the SOC analyst needs from the security tools sending data to the SIEM solution. Dashboards can be created that contain multiple widgets. This data can be exported as a report, displayed on a monitor, and emailed to the analyst to give an update on the state of the environment. Figure 5-15 shows an example of creating some simple widgets within IBM QRadar that an analyst focusing on security incidents would benefit from. Each of these widgets represents the results of using a specific filter criterion. If any of these widgets stopped producing useful data, the analyst could click the Edit button and look at the results of the search to validate everything is still working properly. Maybe something has changed

or maybe there are not any results populating the search after filtering is applied.

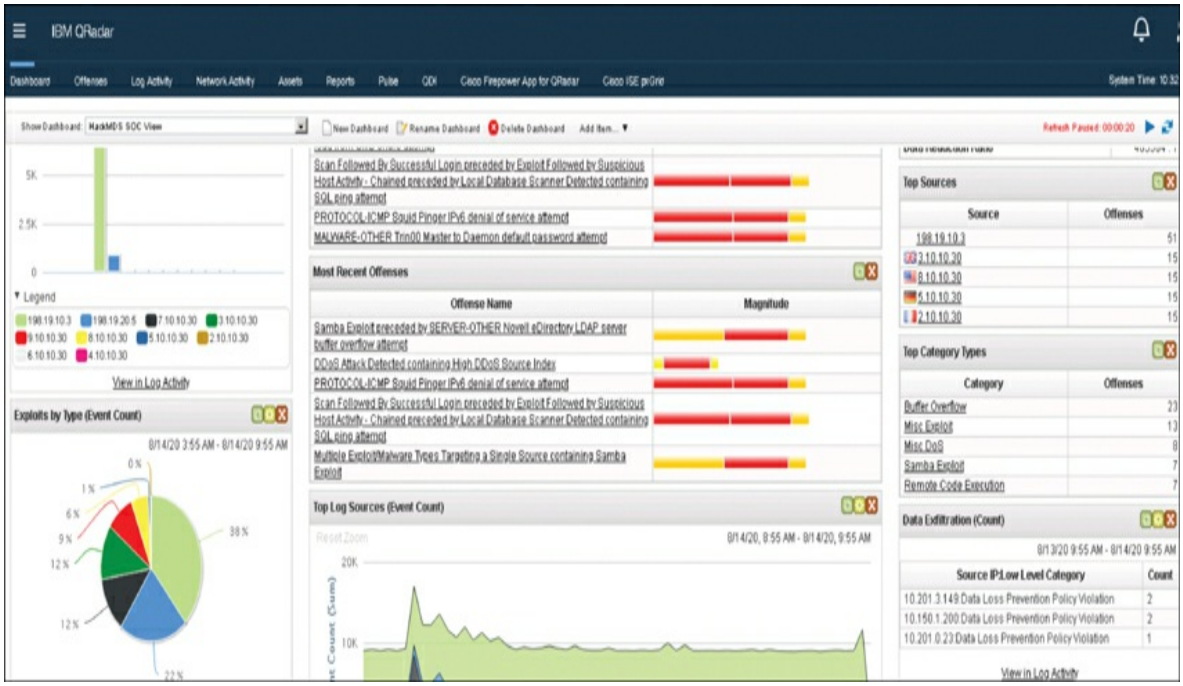

FIGURE 5-15 IBM QRadar Dashboard Example

Most modern SIEM solutions offer a similar approach to developing widgets, dashboards, and reporting. As an example of how this dashboard works, in Figure 5-16 I can click the widget labeled High Concern and select to open one of the 31 items listed to use as a dedicated search. Notice for one of the widgets labeled High Concern, how I can click into it and select to open what data is used within this widget as a dedicated search. The example shows how a dashboard is essentially a bunch of filtered searches displayed on one screen.

FIGURE 5-16 Splunk Dashboard Example

Additional SIEM Features

SIEM solutions come in many forms and offer many special features beyond basic data digesting and display. Splunk is an example of a SIEM solution that heavily focuses on applications. Figure 5-17 shows the Cisco Network Security Analytics (formerly Stealthwatch) application used within Splunk. This application enables the analyst to view in Splunk a dashboard similar to the SMC dashboard without having to log into SMC. Using applications such as this saves the analyst time not only by not having to jump between tools to obtain actionable intelligence but also by being able to customize within the app which data to receive from the tool sending data rather than customizing the tool itself. Troubleshooting applications involve reviewing the documentation associated with the application because the approach to install will vary based on vendor. If you are including SIEM applications as part of your deployment, you will need to include setup and troubleshooting the

applications as part of your research process.

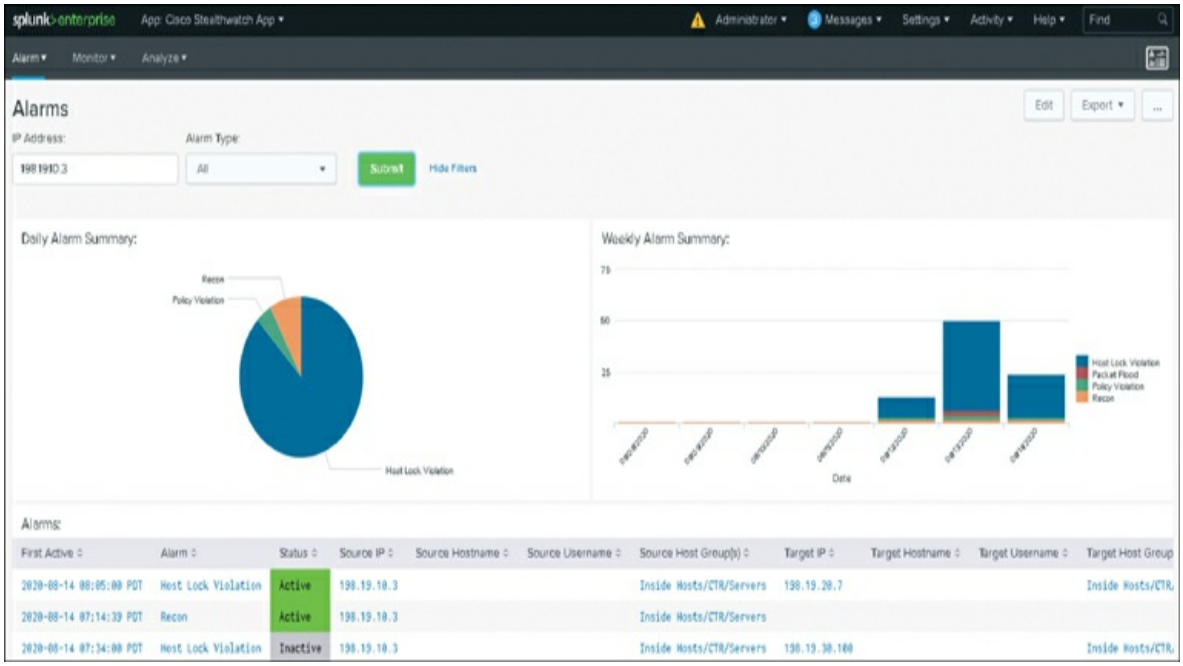

FIGURE 5-17 Stealthwatch Application within Splunk

Another example of a useful feature found in IBM QRadar is the capability to manage assets and vulnerabilities. QRadar can collect asset data and apply vulnerability scanning against it. Vulnerability data can be obtained either by using a vulnerability scanning capability within QRadar or by integrating with an external vulnerability scanner. Figure 5-18 shows IBM QRadar identifying different assets and associated vulnerabilities using an integration with the Rapid7 vulnerability scanner Nexpose. This is yet another example of specialized features that SIEM solutions with a focus on SEM offer.

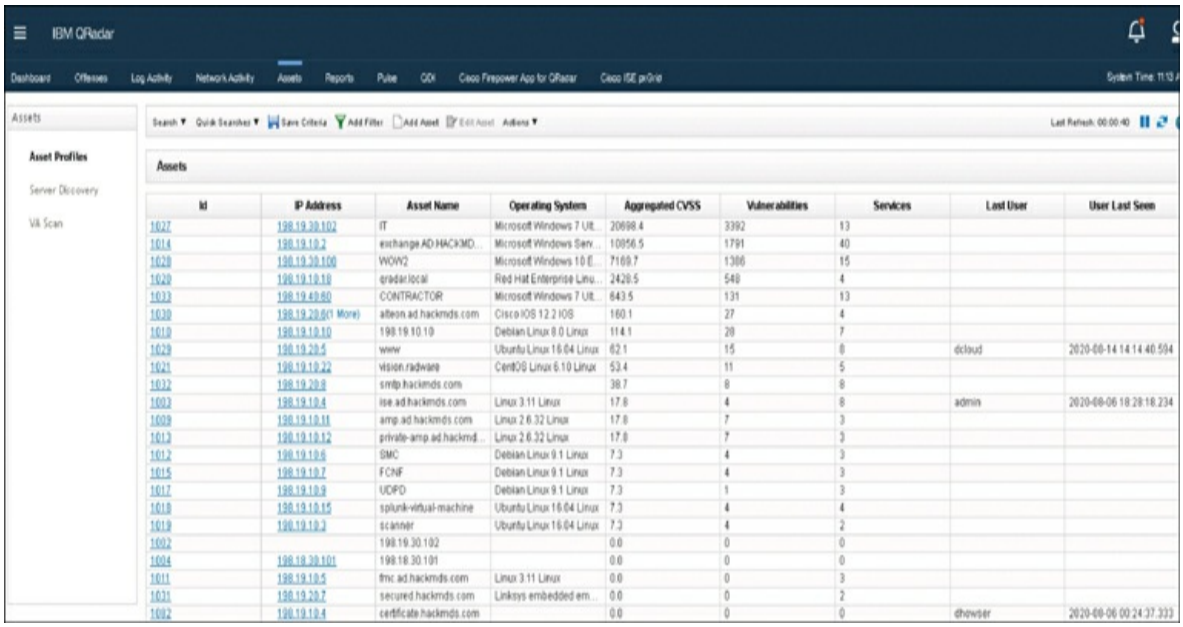

FIGURE 5-18 IBM QRadar Asset and Vulnerability Management Example

I showed the feature in Figure 5-18 to demonstrate how SIEMs provide value not only from pulling log data but also through integration with other tools. Integration means the SIEM solutions are capable of interacting with other tools, such as QRadar pulling vulnerability data from Nexpose as shown in Figure 5-18. Other SIEM vendors use different approaches to integrate third-party tools. One common approach is the use of APIs, described next.

APIs

An application programming interface (API) is a software intermediary that enables two applications to talk to each other. APIs define specific functions that are available and independent of their respective implementations. Good APIs allow simplification of programming by providing functional and useful building blocks designed to accomplish tasks that would otherwise require much more work. Developers can use those building blocks to develop complex processes that require very little new code. The biggest benefit of using APIs is that programmers do not have to reinvent the wheel, so to speak, meaning they do not have to develop code for existing functionality

but instead just call upon it.

An analogy to the API process is ordering food at a restaurant. Instead of cooking the food yourself, you send instructions through a waiter, who relays those instructions to the cook. The cook processes those instructions by cooking your meal, which is then brought to you by the waiter. In this example, the waiter is acting as the API. If the waiter understands and communicates your instructions correctly, you end up with the meal you desired. If the waiter misunderstands your instructions or does not communicate them correctly, you receive a different meal than what you ordered or your order is rejected by the cook.

A technical example of an API involving the IBM QRadar SIEM and Rapid7 Nexpose vulnerability scanner integration works by Rapid7 sending vulnerability data to IBM QRadar using API instructions. Those instructions specify when the data is sent, how often data is sent, the data format, and any other requirements for both parties to complete the data sharing correctly. If the instructions are not accepted, either IBM QRadar will produce an error message, not show any vulnerability data, or show the wrong data. Figure 5-19 shows a high-level diagram of how this example integration works.

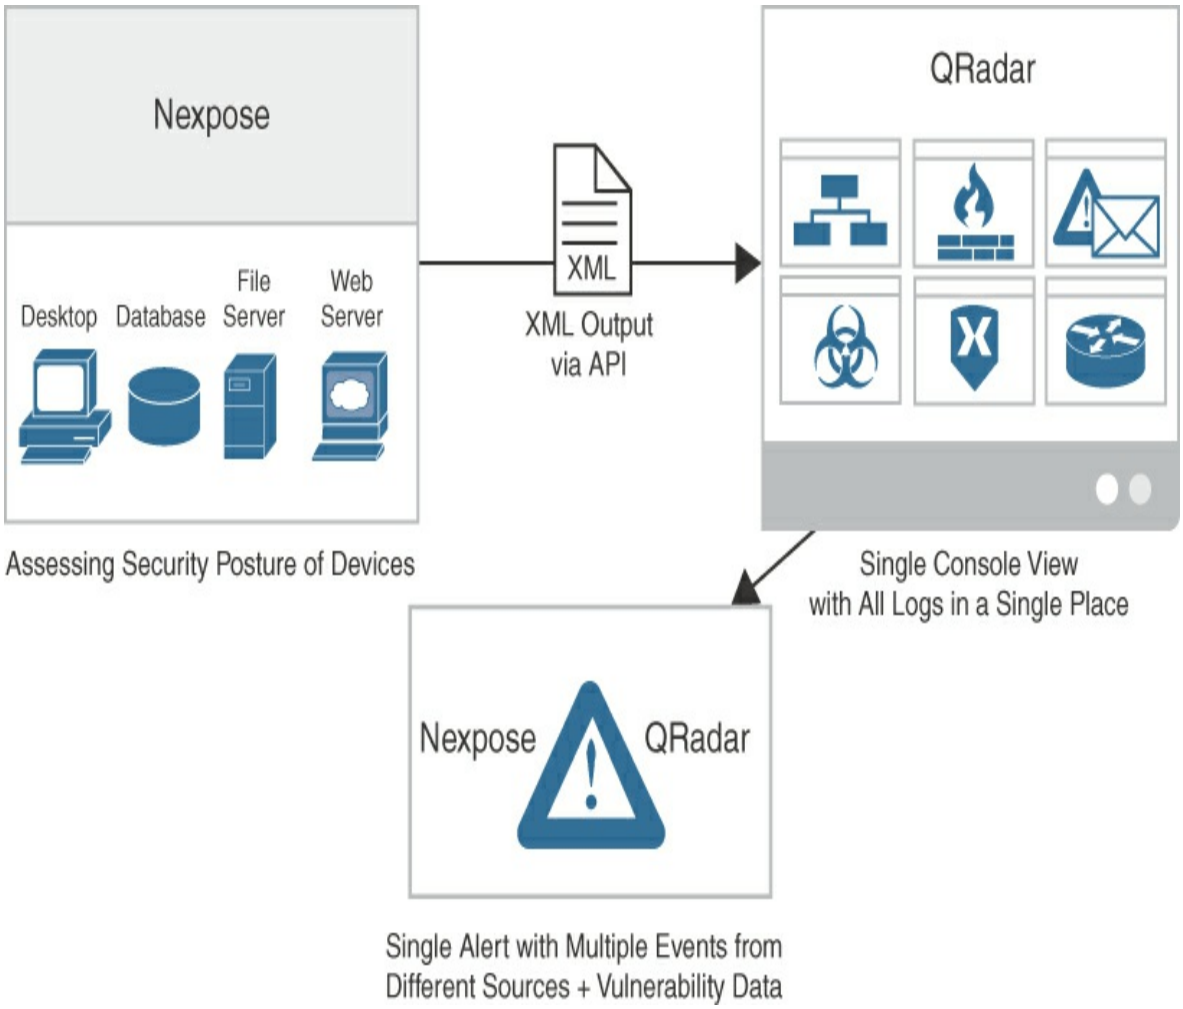

FIGURE 5-19 Rapid7 Nexpose and IBM QRadar Integration Diagram

Leveraging APIs

At a high level, leveraging APIs looks easy, but there are many important factors that you must evaluate before using APIs. First, consider how the API is being used. System-based APIs allow one system to access data from another system. In the security world, I see this type of API used often for pulling user information from an Active Directory system, sharing vulnerability data from one system database to another, and pulling security alerts from a system. Process-focused APIs interact with and shape data within a single system or across systems with the purpose of breaking data down into data silos. This approach is common for creating reports, which

can use different APIs to collect and modify data. Sometimes filtered APIs are applied, meaning the intent of the API is to reduce data to a specific format for a system to digest or view, such as a customer portal. The approach taken by an API will be defined by the vendor, including what tasks are expected to be purposed by using the API. Figure 5-19 has the API pulling XML output from Nexpose and sending it to QRadar.

Another factor to consider regarding APIs is if the API is functioning and accessible internally or externally. Some APIs are available only to a system developer, which means that once the system goes live postproduction, users of the system won't be able to leverage the API. Externally available APIs are APIs that can be called by any user and system while the system is actively being used postproduction. Figure 5-19 shows an external API, since active postproduction systems are using the API to communicate with each other. It is extremely important to be aware of the risk associated with both internal and external APIs so that they are not abused by attackers. A malicious party could modify a system in an undesired way if certain APIs are left exposed and unsecured.

API Architectures

There are common API architecture styles that are used in the industry to help simplify expectations of what tools will need to interoperate with other systems. The following are some of the most common API architectures used within the IT market:

- **Representational State Transfer (REST):** REST is an architecture approach that separates the API consumer from the API provider by relying on commands that are built into the underlying networking protocol. Clients use specific links and forms to perform actions such as a read, update, share, or approval. It is common for HTML to use REST APIs since it is flexible regarding supporting data formats such as JSON and XML, among others. Many HTTP-based graphical user interfaces (GUIs) have REST API options limited to specific functions that other third-party tools can call upon.
- **Remote Procedure Calls (RPC):** RPC enables developers to execute specific blocks of code on another system or network without having to

understand the network's details. RPC is used to call other processes on the remote systems, similar to how communication occurs between local systems. RPC uses a client/server model to accomplish this. In the RPC model, the requesting program is a client and the service-providing program is the server.

- **Event-driven/Streams:** These APIs don't wait for an API consumer to call upon them for delivering a response. Instead, a response is triggered by the occurrence of an event. Clients subscribe to receive updates when values of the service change. There are a few variations for this API style, including reactive, publish and subscribe, event notification, and Common Query Responsibility Segregation (CQRS).

API Examples

In the IBM QRadar and Rapid7 Nexpose integration example, a REST API is being used. The complicated work to make this communication occur is simplified by QRadar when you choose to add a third-party vulnerability scanner. By doing this, predefined API calls are used to allow communication between QRadar and Rapid7 Nexpose.

The following list summarizes how this configuration looks according to Rapid7 documentation:

1. Nexpose performs a security assessment.
2. An XML report is generated with vulnerability findings.
3. Nexpose is added as the VA scanner within QRadar.
4. A scheduled task is created to pull data on a periodic basis.
5. An XML report gets generated with the latest vulnerability data.
6. Data gets imported automatically and normalized within QRadar.

The results of this integration can look like Figure 5-20, showing QRadar representing vulnerability data found with Rapid7.

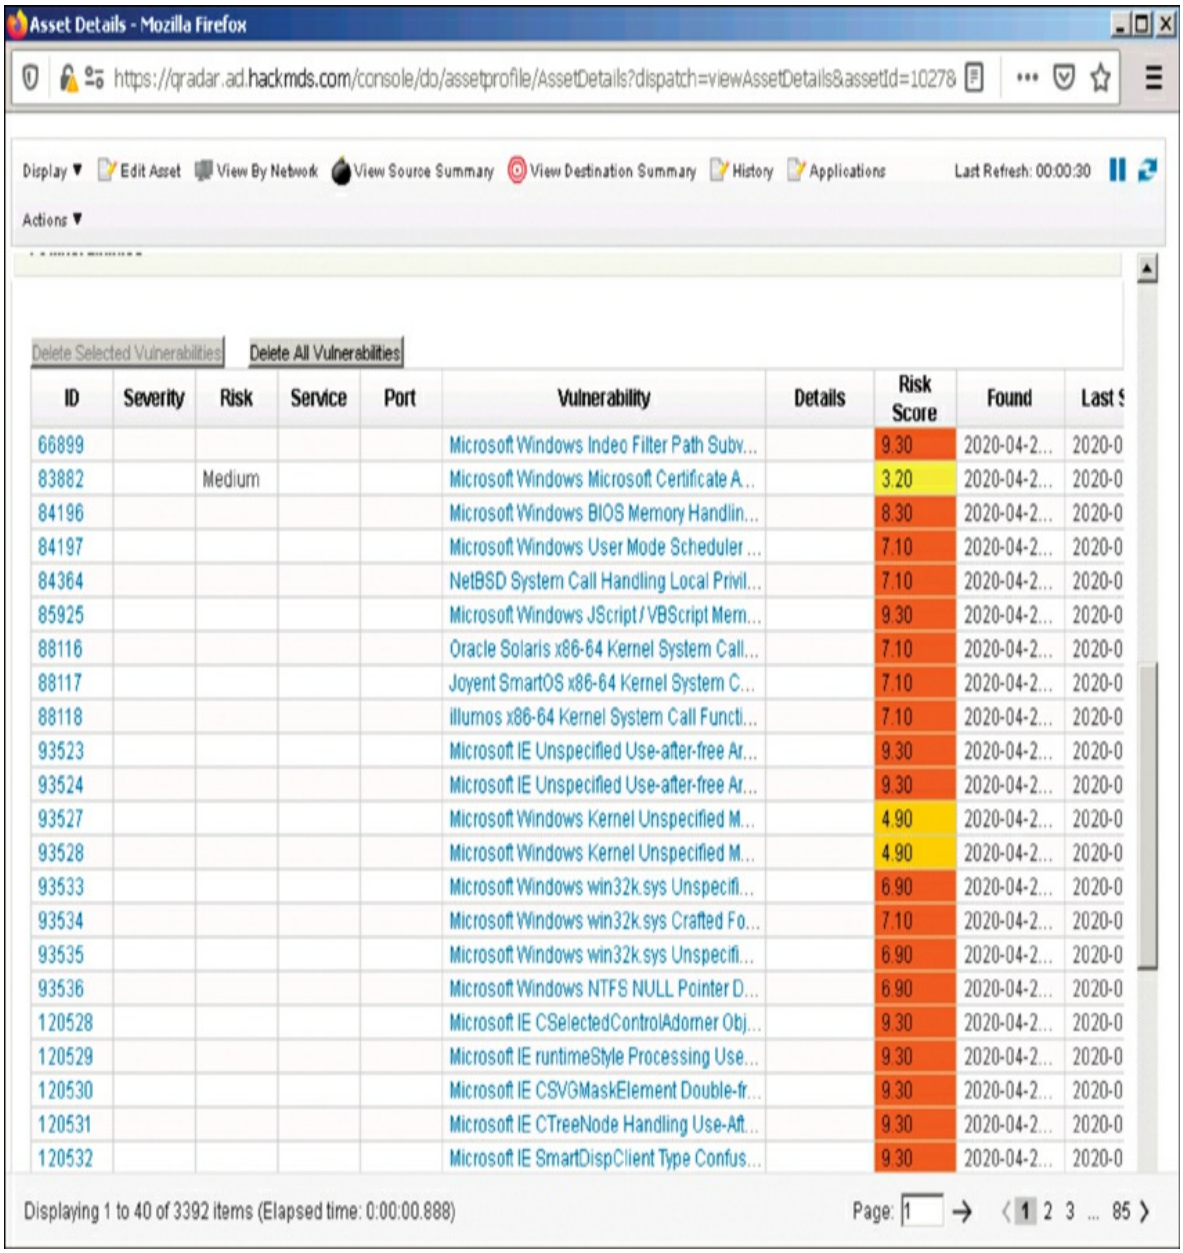

FIGURE 5-20 QRadar Showing Rapid7 Vulnerability Data

A similar example is having a next-generation firewall such as Cisco Firepower leverage APIs to pull the same vulnerability data that was used in the last example from QRadar into Cisco Firepower. In this example, Cisco Firepower offers an option to add a third-party vulnerability scanner. By doing this, prebuilt APIs are leveraged to allow the commutation to occur between Firepower and the Rapid7 vulnerability scanner. The steps to perform this configuration as specified by Rapid7 are as follow:

1. Rapid7 InsightVM performs a security assessment.
2. An XML report is generated with the latest vulnerability findings.
3. The InsightVM connector connects to Cisco FireSIGHT Management Center and pushes a CSV file with the latest vulnerabilities and asset details.
4. FireSIGHT Management Center adds the corresponding vulnerabilities to its Host Map database and pushes it out to each sensor.
5. Rules can be enabled to stop the corresponding attack.

The results of this integration allow Cisco Firepower to leverage vulnerability data captured by the Rapid7 vulnerability scanner as shown in Figure 5-21.

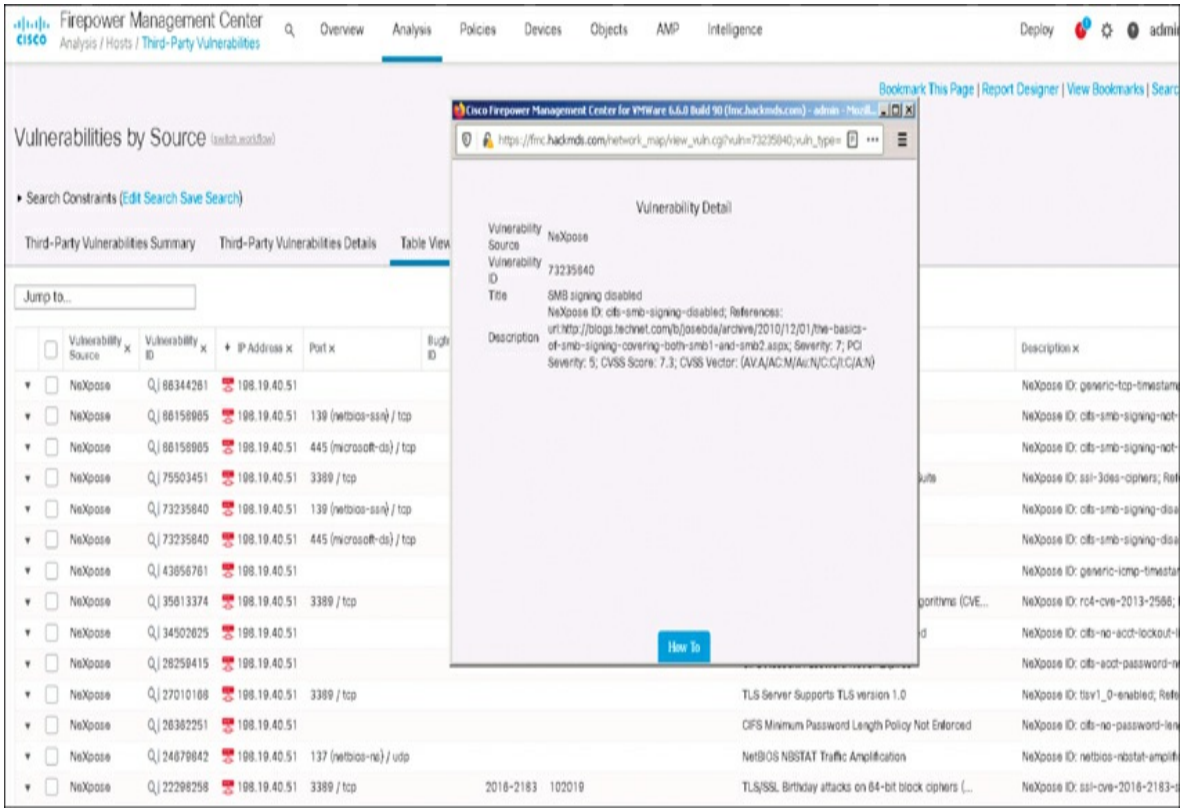

FIGURE 5-21 Firepower Leveraging Rapid7 Vulnerability Data

The benefits of using APIs open the door for many additional use cases that could not be accomplished without allowing for integration between different tools. Just like with setting up any other capability, you should first research

how to use a vendor's API and develop specific goals before attempting to set up an API integration. You also need to validate any associated risks and develop benchmarks for success to ensure the API usage is accomplishing a business goal.

Another useful integration is pulling in data feeds. One very common form of data feed is leveraging big data.

Big Data

Throughout this chapter I have described how to abstract value from various data sources. All topics assumed the data being collected has specific limitations regarding the types of systems producing the data and the amount of data being received by the SIEM solutions or other tool. In practice, there are millions of systems producing data, and the amount of data some systems produce can be trillions of times larger than the amounts assumed up to this point. I touched on what the industry has coined as "big data" in the opening of this chapter, but now I want to address this topic in regard to using big data in a SOC.

First, let's look back to the big data example introduced at the beginning of this chapter. Recall that the Cisco Talos whitepaper "Talos Intelligence" reports that Cisco Talos has insight into more than 17 billion web requests each day. The same whitepaper points out that other data sources Talos inspects include more than 300 billion emails each day. If you imagine the number of data sources Talos uses on top of these two examples, which I'm guessing include data from Cisco products, data from security research, data from tools such as honeynets, or taps, within dark and deep networks, and data from Cisco customers, and then add the total data from all of those sources, I am talking about a number larger than I could write down without searching for how to state something that is many zeros after a billion. You should get the idea of how big I'm speaking about by now. Big data is data you can't assign a number to because it's just too big to do so.

If you were to direct big data at an average security tool such as a SIEM solutions, you would kill the tool based on data volume. Also, a good portion of big data is unstructured, meaning your tools will not be able to understand

how to parse and index what is contained within the data. This brings up the question of how a SOC can work with big data. Big data is a field that focuses on methods to analyze, systematically extract information from, or otherwise deal with datasets that are too large or complex to be dealt with by a traditional data-processing solution. The concept makes sense; however, how can a security tool such as a SIEM solutions work with data that it can't consume directly? The answer to this question is to use a tool that is capable of processing big data. One of the most common tools used in the industry to consume big data is Hadoop.

Hadoop

Apache Hadoop (https://hadoop.apache.org/) is an open-source, Java-based framework used for storing and processing big data. Hadoop is a popular option based on its distributed file system that enables concurrent processing and faulter tolerance. The key values that Hadoop provides over traditional databases are *speed* and *capacity* based on performing tasks and storing data across multiple servers. For big data, Hadoop is critical because it collects, stores, and organizes data in a way that enables a security tool to abstract meaningful patterns from the data without having to process the entire load of data. Hadoop can accommodate the unstructured data diversity expected from big data at a low cost based on its open-source framework.

Hadoop is a framework made up of an ecosystem of components. First, the Hadoop Distributed File System (HDFS) component maintains the distributed file system. HDFS allows Hadoop to store and replicate data across multiple systems, making it the foundation to the framework. Yet Another Resource Negotiator (YARN) manages and schedules the resources and decides what should occur within each node. MapReduce is the programming that allows Hadoop to split data into smaller sets. MapReduce is critical for digesting large amounts of data in a reasonable fashion. Together, these are the core components of Hadoop. Figure 5-22 provides a high-level diagram showing Hadoop 2.0 and YARN.

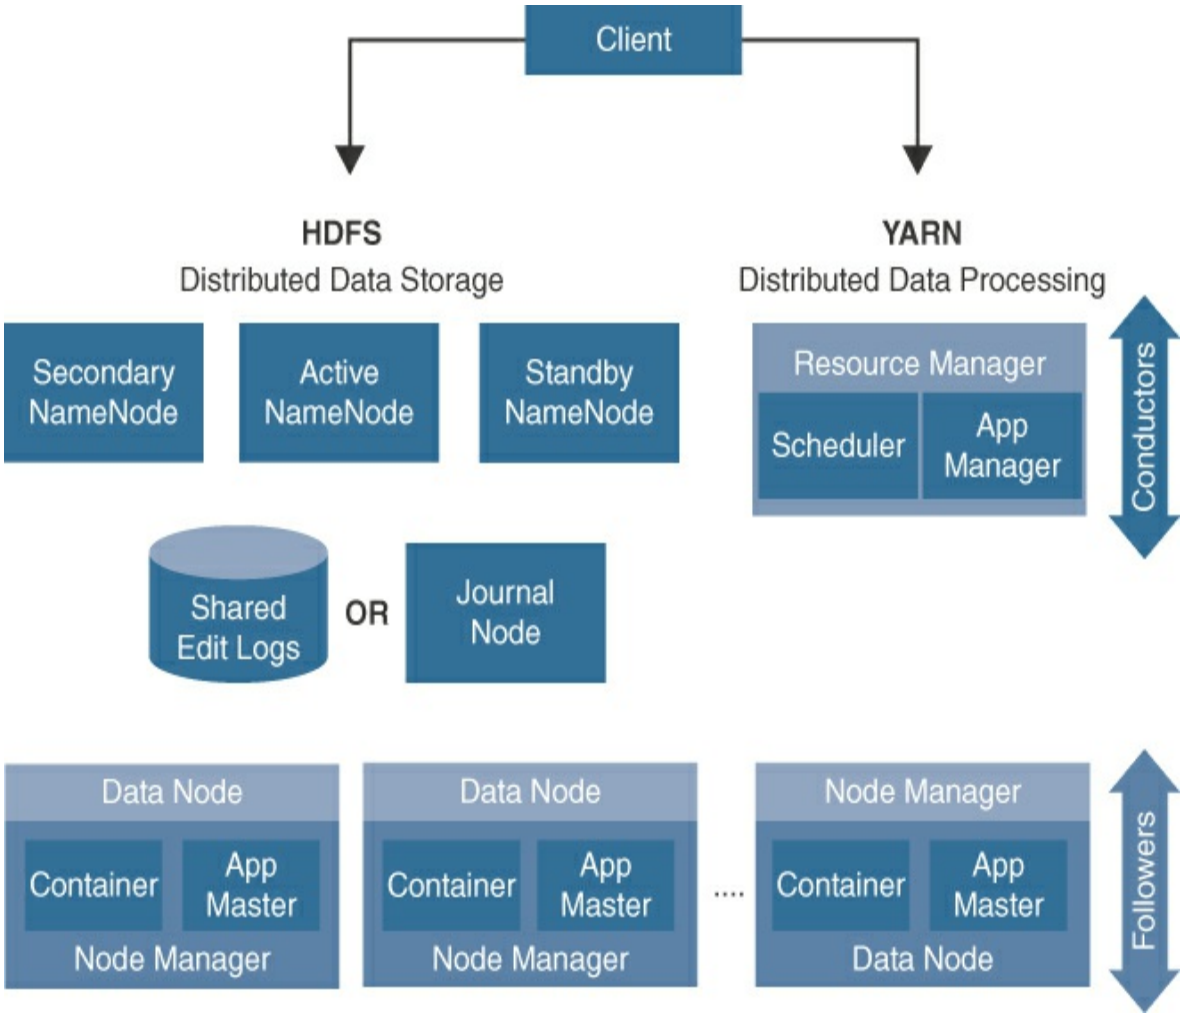

FIGURE 5-22 High-Level Diagram of a Hadoop Architecture

There are many other supplementary Apache components to Hadoop to provide value, such as Hive or Pig for data warehousing, Flume for data ingestion, and ZooKeeper for coordinating distributed applications. Apache Kafka and Apache Spark are options that can help Hadoop work faster by helping to accommodate real-time data since they can use in-memory data storage, which is a limitation of MapReduce. The open-source community surrounding Hadoop is huge and very supportive, enabling you to identify many use cases through research.

Hadoop Challenges

One of the bigger challenges regarding using Hadoop that requires more

work than what can be provided by a supplementary add-on is addressing security concerns. Hadoop has some security options, but they are fragmented because vendors and the open-source software community have had to retrofit security features into an Apache framework that wasn't developed with security in mind. Big data is becoming a critical resource for many organizations, bringing it into the spotlight as a high-value target for cyberattack. To better understand the security challenges with big data, the Cloud Security Alliance (CSA) Big Data Security Working Group listed the following top ten security and privacy challenges to overcome in CSA's 2012 publication "Top Ten Big Data Security and Privacy Challenges":

1. Secure computations in distributed programming frameworks
2. Security best practices for non-relational data stores
3. Secure data storage and transactions logs
4. End-point input validation/filtering
5. Real-time security/compliance monitoring
6. Scalable and composable privacy-preserving data mining and analytics
7. Cryptographically enforced access control and secure communication
8. Granular access control
9. Granular audits
10. Data provenance

Ajit Gaddam presented his whitepaper "Securing Your Big Data Environment" at Black Hat USA 2015 regarding recommendations to address the preceding list of concerns. Gaddam's whitepaper proposes a Big Data Security Framework, shown in Figure 5-23. I find that many of these security recommendations fall in line with topics found in popular guidelines such as NIST and ISO.

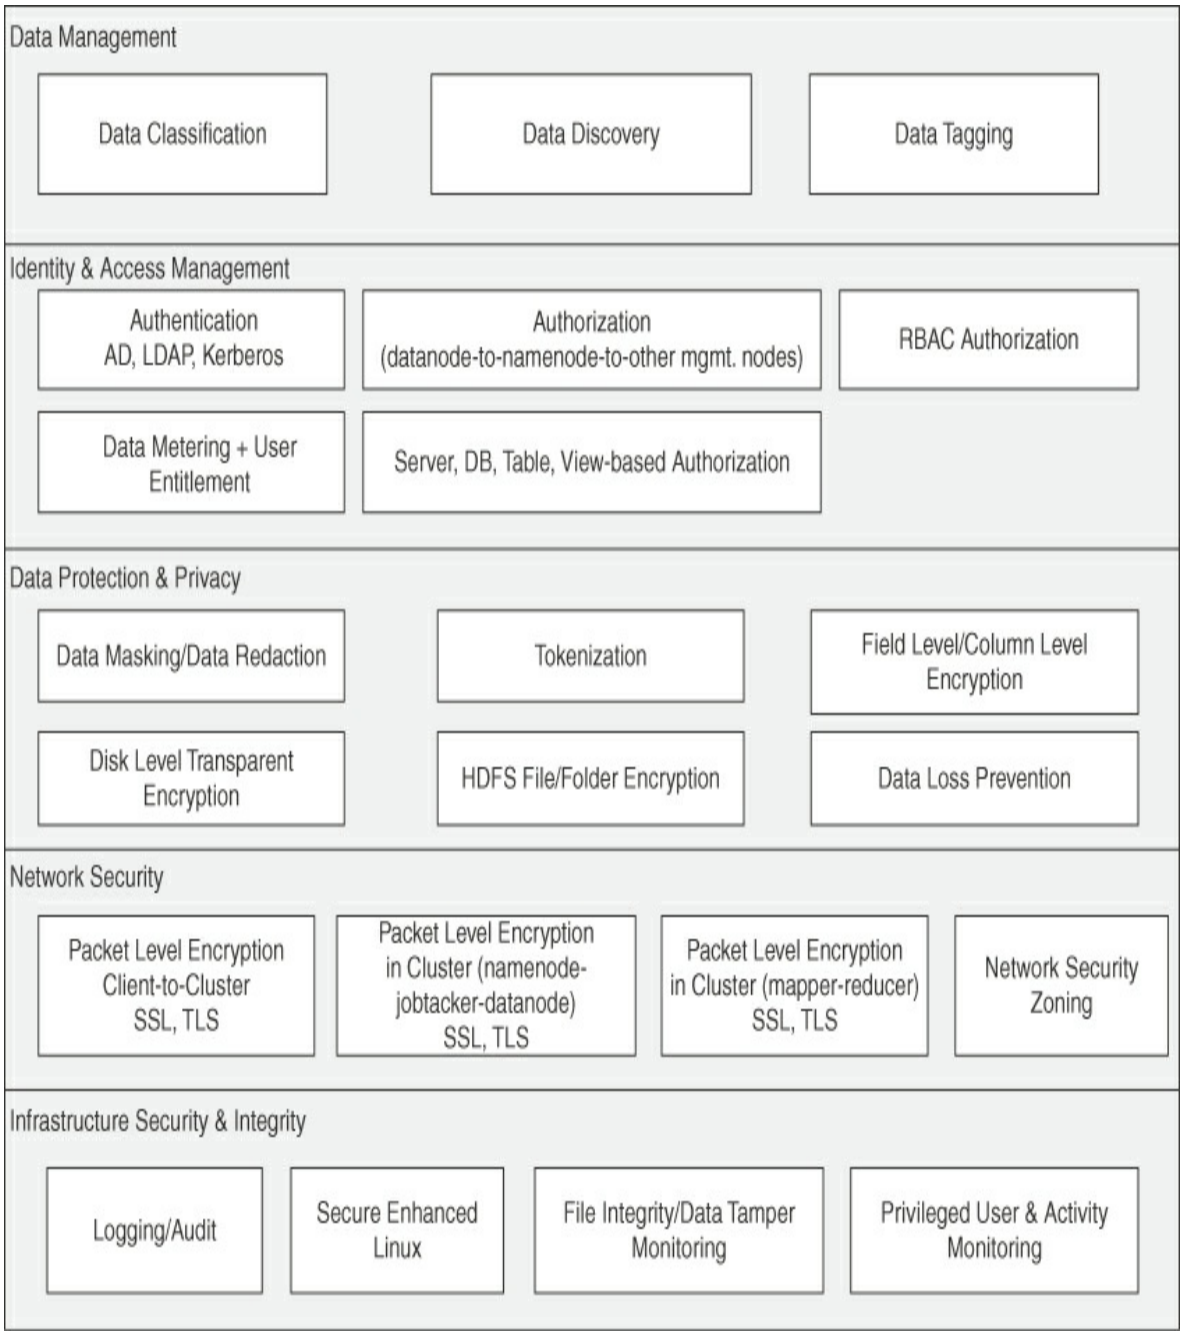

FIGURE 5-23 Big Data Security Framework

A summary of how this framework applies to securing Hadoop can be broken into five focus areas:

- **Data management:** Addresses how big data is collected and identified to ensure the scope of what to protect is properly assessed and understood.
- **Identity and access management:** Focuses on how to control who can

access what aspects of the big data being collected by your organization. This focus area also includes validating that big data is coming from an authorized source, to avoid including modified or unwanted data.

- **Data protection and privacy:** Protects what data is collected and ensures any privacy concerns are addressed.
- **Network security:** Targets securing data as it is collected and moved between systems that are part of or leveraging the Hadoop infrastructure.
- **Infrastructure security and integrity:** Addresses securing the Hadoop components within its framework.

Securing Hadoop

Hadoop includes different security options such as Kerberos, encryption within HDFS, HDFS file permissions, perimeter security using Apache Knox, and authorization applications such as Apache Ranger to address data protection, network security, and infrastructure security concerns. I have covered many of these security topics in this book, and my recommendation is to leverage the Big Data Security Framework as a point of reference, similarly to using any popular guideline, as you address the concerns presented within this framework. To assess a current Hadoop setup, you can perform a capabilities assessment to ensure a defense-in-depth approach is included, which was covered in Chapter 1.

Most modern SIEM solutions have configuration options for working with Hadoop or similar technology. As an example, Splunk uses a Splunk Hadoop Connect application to simplify the integration process between Hadoop and Splunk. With this app, Splunk can forward events to Hadoop for long-term archiving and additional batch analytics. A short summary of what Splunk Hadoop Connect offers is the ability to export, explore, and import big data. Hadoop collects and indexes massive streams of machine data in real time and allows Splunk to navigate and inspect HDFS directories from the Splunk Hadoop Connect user interface. The Splunk Hadoop Connect user interface can also import and index Hadoop data into Splunk, making that data available for searching, reporting, analysis, and visualization without the limitations of requiring MapReduce code. Figure 5-24 shows a diagram of the relationship between Splunk, Hadoop, and the Splunk Hadoop Connect

application.

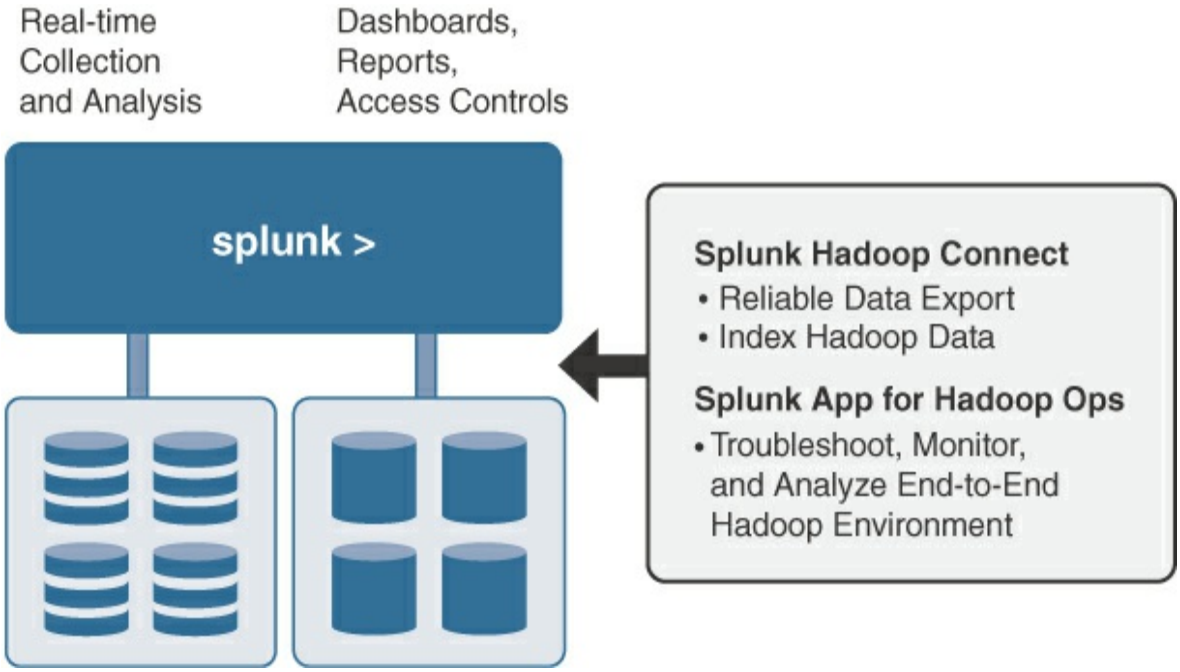

FIGURE 5-24 Splunk, Hadoop, and Splunk Hadoop Connect App Diagram

Big Data Threat Feeds

Another common use case is the threat feeds that vendors provide to their security tools. For this use case example, you as the customer do not deal with managing how big data is being used. Instead, the vendor has its own infrastructure to collect big data for the purpose of generating actionable intelligence for the vendor's products and customers. Looking back at the Cisco Talos example, the big data sources listed as resources for research translate into different types of intelligence that is funneled to Cisco products. As a research team, Talos is pulling huge amounts of data about real-time attack behavior; however, as a customer, you will only see the results of the Talos research, such as how an IP address is associated with malicious behavior. Figure 5-25 is an example of the result from this type of big data research displayed in Cisco Umbrella. Figure 5-25 shows a summary of why the 17ebook.com website is malicious based on big data research performed

by Cisco Talos, including how often the website is accessed, what occurs when systems access this website, who owns this website, if this website is associated with other malicious websites, and real-time updates regarding the risk associated with this website. Cisco Talos is the researchers that enable Cisco security tools such as Cisco Umbrella.

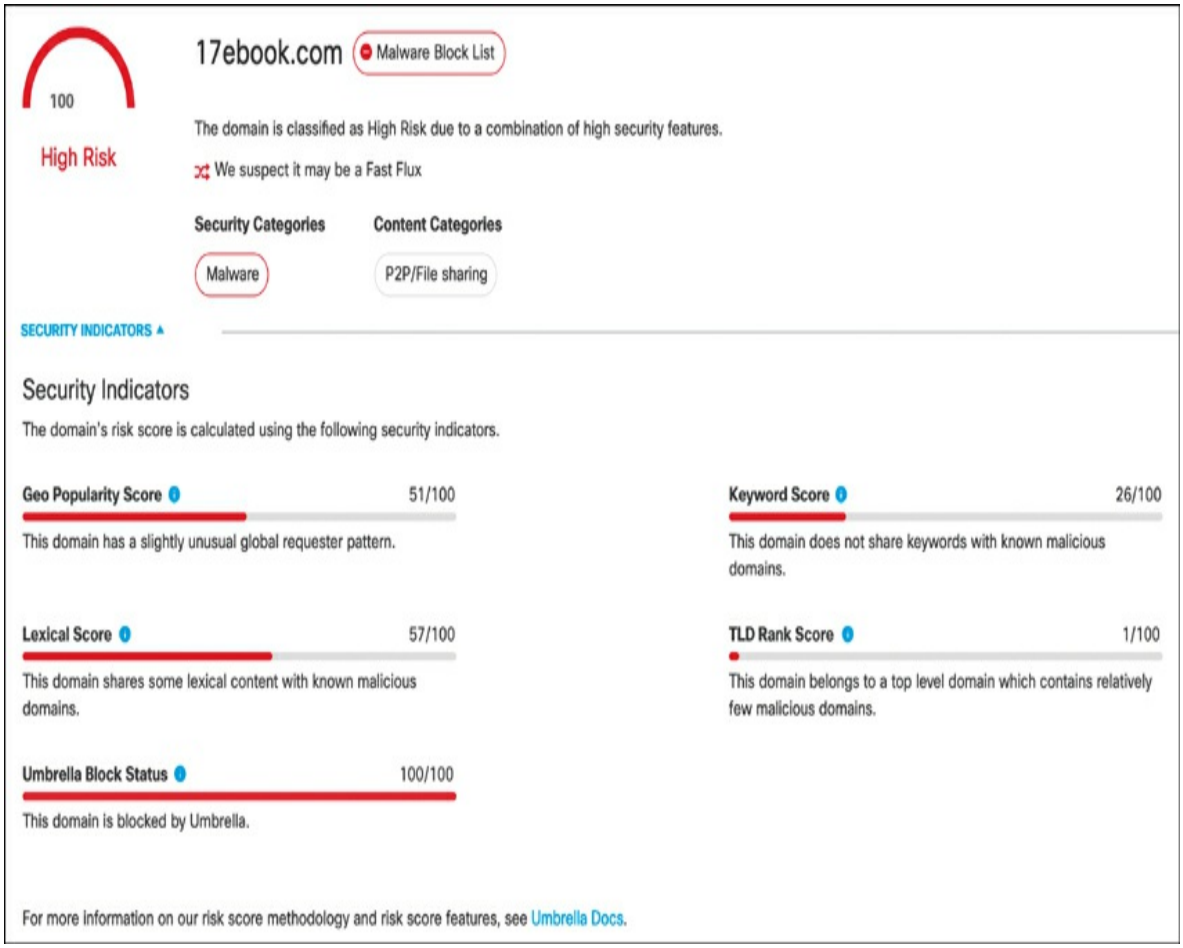

FIGURE 5-25 Cisco Umbrella Big Data Results Example

Most security vendors leverage their own research or a third-party big data research tool to accomplish the goal of delivering actionable intelligence within their products. VirusTotal is a popular tool that leverages big data focused on malware, which many tools will compare artifacts against to get real-time validation of the associated risk with a file. The key to success for this is the focus of using big data to generate actionable intelligence, which means translating the trends seen within big data to actions a SOC can take as

a response. How can this happen if there are so many different types of data that can be leveraged when viewing big data? There can't be a manual way to sort through big data—an analyst isn't viewing the big data or creating tools to look for specific things. If you are thinking that automating finding useful data is the correct answer, yes, that is sort of true; however, one of the biggest values of big data is identifying patterns regarding what you don't know rather than flagging things that are known. Programming automation to identify something must have a way to identify a known thing, which means you are by default ignoring whatever doesn't get identified as that known thing. The answer to this big data unknown-value problem is that the industry leans on machine learning.

Machine Learning

Machine learning uses statistics to find patterns in big data. When using the Internet, you might not know it, but you are seeing the impact of machine learning. Think back to the last time you used a search engine like Google. When you search for something, the search engine suggests what is the best response customized for you. If you and I both type the word "voting" in a search engine, we will receive differing results that are based on our respective previous browsing history. The same concept applies when you are on a social media website and the advertisements are catering to your taste. Social media sites posting advertisements that spark your interest isn't happening by accident. Social media services are constantly collecting data about you and comparing you against everybody else in their massive big data datasets to see what is unique about you to accommodate your specific tastes. This is all done using machine learning—their algorithms are generating content customized for you based on what they have learned about you. If you want to test this concept, try going to one of these same sources using an online proxy such as ProxyScrape (https://proxyscrape.com). You will find the advertisements and recommendations are different than what you see using your trackable access. This is due to the proxy being used by multiple people, causing random data to be seen by the tracking algorithms.

Machine Learning in Cybersecurity

Machine learning (ML) can be used to benefit your SOC as well. Machine learning techniques can help your security tools identify patterns and threats with no prior definition, rules, or attack signatures, and with higher accuracy. Just as social media platforms use ML to identify you uniquely from all other users, security tools use ML to identify malicious activity uniquely from all the other data that represents normal behavior. This is why big data is a key component to ML's success. Without enough data on what is considered normal and what is considered malicious, it would be hard for ML to accurately flag something as being a potential threat.

Why does ML matter in regard to cybersecurity? Traditional security tools are limited based on having to be fed information, typically in the form of attack signatures, about the threats they are designed to defend against. One approach is having attack signatures of known threats more commonly called signatures used to compare things against with the hopes of finding a match. If a threat doesn't match a signature, the threat goes undetected. Data management tools such as a SIEM solution use correlation rules; however, those are also predefined rules looking for a sequence of events, meaning you have to know what you are looking for. An unknown threat, also called a zero day, can bypass traditional approaches to security since there isn't a way to set up a system to look for something that isn't known.

Machine learning can step in to fill the gaps in traditional security capabilities by addressing the unknown. The ML model is designed to take a different approach by starting with identifying anomalies rather than considering and building defenses against what is bad. Machine learning works by first creating a model of normal behavior of your users and network traffic. Threats are detected based on a combination of comparing collected data about a potential threat against threat data pulled from external big data resources (hashes of malicious binaries, known malicious URLs, etc.), threat characteristics pulled from external big data resources (attack behavior seen on other networks), and unusual behavior as related to what is considered normal traffic or behavior. In concept, to evade ML, a threat would have to not be known via the big data resource, not act as a known threat as seen by the big data resource, and act like normal users and network traffic within the environment it has breached in order to remain undetected. If the threat

attempts to communicate outside of the network, ML would evaluate against known behavior the outgoing communication, any files being uploaded or downloaded, and other factors, and this would be continuously occurring as new data is being collected. As you can see, it would be extremely hard for a malicious party to remain undetected since all actions are scrutinized against normal behavior along with threat data obtained from big data resources. This approach is extremely effective against large-scale exploitation such as a world-wide phishing campaign or an exploit targeting various organizations.

Artificial Intelligence

Why hasn't machine learning been heavily used in the security market the last 5, 10, or 20 years? Large data sources have always been available; however, the challenge is how a machine engine consumes and analyzes large datasets. The term artificial intelligence (AI), which means having the cognitive ability to automate tasks, is one of the developments that have made ML possible. ML needs to continuously get "smarter" as it is provided data—it must learn from experience or it won't benefit from seeing large datasets. The AI component of ML allows security to keep up with the bad guys.

I find that the term artificial intelligence is overused by product marketing teams, who commonly use the term to market a tool that doesn't actually have any AI capabilities. As Stuart Russell and Peter Norvig state in their book *Artificial Intelligence: A Modern Approach*, "AI needs to be able to deal with unknown environments/circumstances in order to achieve its objective/goal, and render knowledge in a manner that provides for new learning/information to be added easily." A good way to evaluate whether a vendor is using AI in a particular solution is to ask the vendor how their solution identifies a positive and negative match without knowing what to look for, and then follow up by asking what the solution does with that new information. If their answer is related to monitoring for anything predefined, then their solution is looking for already established items, which isn't AI. Instead, they are referencing advanced pattern machining designed to look for known behavior.

Machine Learning Models

There are different models used by machine learning as in regard to cybersecurity. First is a *supervised machine learning* approach. Supervised learning means the tool learns from a dataset that contains inputs and known outputs. This translates to a security tool analyzing new behavior and determining if it is similar to previous known good or bad behavior. Another model is *unsupervised machine learning*, which consumes data that contains only input variables. This approach means there is no preset correct answer leading to discovering new patterns in data. Unsupervised ML is important in the security world since it allows for catching unknown/zero-day threats. A third approach is *deep learning*, which targets unstructured data, picking out smaller pieces of data to construct a better understanding of the bigger picture. Hidden insights can be established between various data points based on how frequent things are seen, which are grouped into categories of interest. A good way to think of deep learning is in the context of how large amounts of social media are evaluated. An example is seeing patterns of speech from hundreds of Twitter accounts associated with threat researchers covering a specific term, which leads to the discovery that research is being done on a specific threat that hasn't been fully disclosed to the public outside of the bits and pieces posted on social media.

Machine learning is not a product you can go purchase for your SOC. It is, however, a capability that you can seek out within your security tools. If you are considering adding big data through the use of Hadoop or a similar tool, you will find there are options such as Spark that are better suited than MapReduce when running ML applications such as Naive Bayes and k-means. You will need to do research to ensure the applications are actually meeting the goals for your product as well as performing ML capabilities. A similar approach will apply regarding any vendor's solution that claims it offers leveraging big data and ML functionality. There are two models you can use to evaluate ML: hold-out and cross-validation.

Hold-Out Model

The hold-out model tests different data than the solution was trained on. To perform this type of evaluation, you need three elements:

- Provide a training set representing the dataset used to build predictive models.
- Use a validation set of data with which the system builds its understanding/method of learning during the training phase. Validation data may not be needed, meaning the system may not need to be trained.
- Provide test data that the solution hasn't encountered to evaluate how the solution will respond.

The hold-out approach is a quick and simple way to develop a testing environment targeting machine learning. An example of this approach is giving a security tool PCAP files of attack data to allow it to learn about an environment. After the learning process completes, you provide different attack data and see if the tool can address threats with different characteristics than what was provided during the learning process. If the learning process did not use fileless attacks, would those also be identified with the new attack dataset? What about changing up how attacks leverage stealth or timing? The hold-out approach works only if the validation and test datasets are completely different and no preset explanation between both datasets was provided to the technology using machine learning. Doing that would allow the solution to cheat by having some background knowledge of what it is expecting to see.

Cross-Validation Model

The cross-validation approach to modeling machine learning involves partitioning the original observation dataset into the training set with the purpose to train the tool. Then you use a completely independent set of data to evaluate the analysis. A common approach to delivering the cross-validation model is using the *k*-fold cross-validation. This approach works as follows:

1. Take *k* number of equal-size datasets.
2. Repeat the testing *k* number of times.
3. Each round, one of the *k* subsets is used as the test set and the other *k* subsets are combined as the training set.

An example of using this approach would be taking a PCAP and breaking it into smaller data subsets. During the first round, you would take the first PCAP as the validation set while the remaining PCAPs would be combined into a training set. You would repeat this process four more times, changing the validation subset PCAP until each subset gets a chance to act as the validation set. The result is every subset getting the chance to be the test once and part of the training set $k - 1$ times. The value of this approach is that it significantly reduces bias and variance because most of the data is being used as the training set as well as the test set.

The hold-out and cross-validation models are both useful if you want to perform a true evaluation of a solution's ML capabilities. If your SOC doesn't have the ability to perform thorough testing using these approaches, attempt to have the vendor answer the questions I covered regarding ML criteria as well as research if popular ML software is being leveraged by the solution. Examples of popular ML software include scikit-learn, PyTorch, TensorFlow, Weka, KNIME, Google Colab, Apache Mahout, and Accord.NET. When you identify the possible use of ML, focus on the value it provides and problems you hope it will solve. Using the capabilities assessment technique provided in Chapter 1 will help align the value of ML to your existing security capabilities.

Summary

Centralizing security data from a variety of sources has been the focus of this chapter. The first step is understanding the different types of data available and what tools can leverage that data to accomplish a specific goal. I reviewed how a data-focused assessment can help you scope available data and craft a plan to build a SOC that maximizes the value from available data. You learned how to leverage various types of logs, the most common formats for data, and how to leverage logs within security tools.

This chapter focused on using a SIEM solution to centralize security data. I reviewed how modern SIEMs collect data and provided a review of these steps, including how to troubleshoot when a SIEM solution does not produce your desired results. I also covered how SIEM solutions and other security tools work with third-party tools, by leveraging APIs. The final topics

addressed current trends with big data and machine learning and how to leverage huge amounts of data for security value.

At this point, I have covered many angles regarding how to collect data and convert it into actionable intelligence, but I have barely scratched the surface on how big data relates to threat intelligence, which is the focus of Chapter 7. I also have not yet covered what actions can be taken on data collected, which is the focus of Chapter 10. If you are wondering where topics such as SOAR and XDR come into play, they will be addressed in Chapter 10. Keep in mind that topics covered in Chapter 10 are dependent on this chapter, meaning you can't take proper action if the data you are acting upon is not good.

Before taking any action on data, you must understand what risks and controls you need to address, which is the focus of the next chapter.

References

Cisco Talos. (2018.) Talos Intelligence. Cisco Talos. https://www.cisco.com/c/dam/m/en_us/offers/pdfs/talos-group-whitepaper.pdf

Cloud Security Alliance. (2012, November). Top Ten Big Data Security and Privacy Challenges. Cloud Security Alliance. https://downloads.cloudsecurityalliance.org/initiatives/bdwg/Big_Data_Top_Ten_v1.pdf

Exabeam. (n.d.). Log Aggregation, Processing and Analysis for Security. In *The Essential Guide to SIEM*. Exabeam. https://www.exabeam.com/siem-guide/events-and-logs/

Gaddam, A. (2015). Securing Your Big Data Environment. Black Hat. https://www.blackhat.com/docs/us-15/materials/us-15-Gaddam-Securing-Your-Big-Data-Environment-wp.pdf

Hao, K. (2018, November 17). What Is Machine Learning? MIT Technology Review. https://www.technologyreview.com/2018/11/17/103781/what-is-machine-learning-we-drew-you-another-flowchart/

Marr, B. (2018, May 21). How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read. Forbes. https://www.forbes.com/sites/bernardmarr/2018/05/21/how-much-data-do-we-create-every-day-the-mind-blowing-stats-everyone-should-read/#6b00544160ba

MuleSoft. (n.d.). Types of APIs and How to Determine Which to Build. MuleSoft. https://www.mulesoft.com/resources/api/types-of-apis

Pearlman, S. (2016, September 7). What Are APIs and How Do They Work? MuleSoft Blog. https://blogs.mulesoft.com/biz/tech-ramblings-biz/what-are-apis-how-do-apis-work/

Rapid7. (n.d.). Leverage Rapid7 Vulnerability Intelligence to Add Deep Security Context to IBM's QRadar SIEM. Rapid7. https://www.rapid7.com/globalassets/_pdfs/product-and-service-briefs/rapid7-nexpose-ibmqradar-solution-brief.pdf

Rockafella. (2018, September 25). Overcoming Common Causes for SIEM Solution Deployment Failures. WordPress. https://virtualizationandstorage.wordpress.com/2018/09/25/overcoming-common-causes-for-siem-solution-deployment-failures/

Russell, S., & Norvig, P. (2015). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson.

Talend. (n.d.). What Is Hadoop? Talend. https://www.talend.com/resources/what-is-hadoop/

YoLinux.com. (n.d.). CVS Intro, Commands and Examples. YoLinux.com. http://www.yolinux.com/TUTORIALS/LinuxTutorialCVSintro.html

Chapter 6

Reducing Risk and Exceeding Compliance

All I want is compliance with my wishes, after reasonable discussion.

—Winston Churchill

This chapter focuses on various forms of compliance. *Compliance* by definition is to meet governing regulatory or contractual requirements. Requirements can come from an organization's leadership, such as a corporate-mandated policy, which would be considered a policy based on corporate compliance. Another possible requirement is meeting a legal obligation, which would be government-based regulatory or statutory compliance. Lastly, a compliance requirement can be industry compliance, meaning leadership sets the goal to meet a general recommendation. Not being compliant with a government-required policy will lead to fines and potentially jailtime. Not being compliant with industry recommendations will lead to gaps in security practices. Gaps will become failures in security. This is why many organizations include industry-recommended guidelines and government-required policies within their mandated corporate policy. Enforcing good security practices leads to a reduction of risk.

Note

Compliance is not security. Security is the practice of implementing effective technical controls to protect assets, while compliance is focused *only* on meeting governing, regulatory, or contractual requirements. Anything outside of those requirements is not considered, which is why SOCs need to look *beyond* compliance

when setting their baseline for security standards—and is why this chapter's title states "Exceeding Compliance." To exceed compliance, your security baseline must focus on being secure!

Chapter 1, "Introducing Security Operations and the SOC," emphasized that security best practices are made up of the right combination of *people*, *process*, and *technology*. An example of a people-focused best practice would be found within industry certification programs, such as the CISSP teaching proper security practices. An example of a process-oriented best practice resource would be a guideline such as those released by NIST. An example of a technical-focused resource would be Common Criteria, an international standard for computer security certification. When I consult with customers, I typically find that they have security technology in place as well as people with responsibility over that technology, meaning the customers have some level of *people* and *technology*. A common area of weakness in many organizations is a lack of processes or very weak processes made up of dated procedures or policies. Without up-to-date policies and procedures, situations develop where security events are captured by existing technology, but procedures are not followed regarding proper use of technology, causing events to be overlooked. Outdated policies and procedures cause vulnerabilities in an organization's security practice and become the weakest link leading to exploitation.

This chapter describes how to develop and maintain strong policies and procedures and introduces industry recommendations for policies and procedures found in standards, frameworks, and guidelines. Policies and procedures are requirements developed and enforced internally by an organization. Standards, frameworks, and guidelines are not required by law to be followed and usually are referenced externally by organization leadership, such as an industry best practice for a topic, but can be a best practice developed in-house as well. Many organizations use standards, frameworks, and guidelines as templates for creating policies and procedures, which I'll cover how to do in this chapter. The successful use of policies and procedures will depend on your organization's business objectives, what is required by local and federal government agencies, and what threats could impact the business. This chapter also covers common required compliance

based on government or service requirements.

Why Exceeding Compliance

I use the term "Exceeding Compliance" in the title of this chapter even though compliance can only be met or not met. My intent for this chapter is to cover how to develop strong policies and procedures using industry standards, frameworks, and guidelines to establish part of a strong baseline for security. The remaining parts of your security baseline will be addressed using security-focused services such as tabletop exercises, assessments, and penetration testing, which are also topics that I will cover. The combination of building compliance and security-focused capabilities within your security practice will dramatically reduce the risk of future exploitation of your people, process, and technology as far as the entity enforcing compliance is concerned.

> **Note**
>
> Many organizations can operate while not compliant with internal policies or external laws or regulations as long as they have a plan to eventually become compliant. Sometimes the allowed time to remediate is months or even years!

Policies

Policies are high-level mandatory rules that an organization sets in place. Think of a policy as the objective for a security goal. Common language for a policy is using a term such as "Acceptable Use" meaning what actions are and are not accepted. Because a policy is an objective, it is developed with the intent to not change often. Any specific details should be left out, such as naming a specific technology or person or specifying how tasks are performed, because they can change as the organization evolves over time.

Details for a policy are ideal for being developed into procedures, as covered later in this chapter.

An example policy statement is “Only employee-issued devices will be used within the employee network.” In this example, the vision is stated; however, details such as how this policy is enforced are left out. This policy can survive changes in people, process, and technology because it’s a goal lacking specific details. Another example is “All employees must use multifactor authentication to access the corporate network.” Again, the details are left out regarding which authentication factors are eligible, which technologies are to be used, and how this policy will be enforced, all of which are details better addressed in procedures.

Note

SANs provides an excellent resource for different policy templates at https://www.sans.org/security-resources/policies.

Let’s break down how to properly construct a policy by looking at an example. The different parts of this policy example will also be similar to what is included in developing other formal documents such as standards and procedures. None of these steps are required, but including them in your policy documentation is highly recommended and encouraged. Some sections are not needed if the details they cover are found in other sections. I, however, recommended using dedicated concept sections (more commonly called indexing) to simplify finding the location of each topic covered in a policy. This allows a reader to quickly identify areas of interest in the policy, which is particularly important when introducing new policies with major impact that will cause concern for employees. Employees will need a way to reference the policy and better understand how it impacts them or they will not accept the change. I will talk more about this concept when I address tasks that must be completed as you deploy new policies or make changes to existing policies.

Policy Overview

A policy starts with an overview. An overview can explain the history of the policy or situation prior to the policy being established, the scope of the policy—how and to whom or what it applies, when the policy applies, the intended audience that is expected to review the policy, and prerequisites or other details needed to comprehend the policy's meaning. Some of these items can be omitted if other sections of the policy cover the topic, such as a section dedicated to describing the history of the document. However, a concept may be stated more than once in a policy—the overview may include items that are repeated later in the policy because the overview is intended to be an introduction to the policy topic, which might require the reader to understand certain details before proceeding to any other topics within the policy. There is a lot of flexibility for developing an overview, but keep in mind that as the first part of the policy, the overview sets the tone for the rest of the document. Also keep in mind that the overview may be the only part of the policy that some people review.

Policy Overview Example

The following is a sample policy overview.

Overview

<COMPANY or **SECURITY TEAM>**'s intentions for publishing **<POLICY>** are not to impose restrictions that are contrary to **<COMPANY>**'s established culture of openness, trust, and integrity. **<COMPANY** or **SECURITY TEAM>** is committed to protecting employees, partners, and **<COMPANY>** from illegal or damaging actions by individuals, either knowingly or unknowingly.

Internet/Intranet/Extranet-related systems, including but not limited to computer equipment, mobile devices, software, operating systems, storage media, network accounts providing electronic mail, usage of WWW, and file transfer, are the property of **<COMPANY**. These systems are to be used for business purposes in serving the interests of the company and customers in the course of normal operations. Please

review Human Resources policies for further details.

Effective security is a team effort involving the participation and support of every **<COMPANY** employee and affiliate who deals with information and/or information systems. It is the responsibility of every **<COMPANY** user to know these guidelines, and to conduct their activities accordingly.

The first part of this example policy overview highlights the intent of the policy. In this case, the policy is designed to protect employees, partners, and the company from illegal or damaging actions. The first part of this overview also emphasizes the security team has not designed this policy to impact the culture of the organization. These two points establish a baseline for the thought process during the development of this policy.

The second part of this policy overview defines the scope related to various types of technology and usage of technology. Rather than explaining the details of each technology, this section refers to another human resources policy for additional details regarding what is defined as “business purpose in serving the interest of the company and customers.” It is common practice to cross reference other policies in this manner.

The final part of this policy overview explains not only who is required to follow the policy but also how following this policy impacts every individual at this organization. The policy requires all users to understand guidelines associated with the impacted systems of this policy. Rather than referring to another policy, this last statement refers to guidelines, which are different than a policy and contain more details. Referring to industry guidelines is also common practice, as is referencing other more detailed sources such as procedures. For example, a policy could require two-factor authentication and reference a procedure that provides details on how the policy should be executed in terms of which technologies, people, and steps should be taken to meet the policy requirements.

Other overview items that could have been included in this example policy are statements about the creation or revisions of the policy, who was involved with creating it, who is sponsoring this policy, who to contact with questions, and other items that help the reader better understand the policy as they

review the other sections that follow the policy overview. For this example policy, you will find many of the statements that were left out of the overview are addressed in other sections of the policy.

Policy Purpose

The purpose of a policy explains why the policy exists. Sometimes, a policy exists to meet a mandatory requirement. Some examples could be a policy enforcing HIPAA requirements to protect user privacy, meeting business goals such as reducing operational costs, or even meeting another policy requirement. Other times a policy's purpose is to reduce risk, meaning that conforming to the policy will reduce the chance of some unwanted action(s) from occurring. An example is banning the use of USB drives in order to prevent the release of confidential data of the company, its partners, and customers. The policy purpose is critical because it answers questions regarding why this policy must be met.

Policy Purpose Example

The following statement is an example of a policy purpose.

Purpose

The purpose of this policy is to outline the acceptable usage of **<SOMETHING>** at **<COMPANY>**. These rules have been established to protect the employee, customers, and **<COMPANY>**. Inappropriate use exposes **<COMPANY>** to risks including cyberattacks, malware, and compromise of systems and services, any of which can result in legal issues.

This example purpose explains that the reason for this policy is to protect users, the customer, and company from various types of threats. A different tactic could take the approach of meeting a compliance requirement, which would also result in similar goals of risk reduction. For example, the purpose

could state the intent is to be compliant with a guideline such as the Center for Internet Security (CIS) Top 20 Critical Security Controls. The purpose of the CIS Top 20 Critical Security Controls is protecting systems from cyberthreats, which has the same purpose as protecting users from various types of threats as stated in this policy purpose example. Specifying a purpose or referencing another source such as a guideline that has the same purpose are both common practices for language included within a policy's purpose section.

Policy Scope

A policy's scope defines what is and what is not included in coverage of the policy. The scope can include types of people, technology, and behavior as well as what is associated with each of these items. An example is listing employees and those that interact with employees. For this example scope statement, coverage would not include guests unless the guests have interacted with an employee but it would include employees and contractors. A scope should be developed in a whitelist format, meaning the policy must specify only what is covered and assume anything else is not covered. Using a blacklist approach such as "all nonemployee assets" would not be recommended because a nonemployee asset could be almost anything. It is better to specify more specific categories whenever possible, such as "nonemployee mobile devices, tablets, or IT equipment."

Policy Scope Example

The following is an example policy scope.

Scope

This policy applies to employees, contractors, consultants, guests, and other users within **<COMPANY>**, including all personnel affiliated with third parties. This policy applies to all equipment used by such people that is owned or leased by **<COMPANY>**.

Policy Statement

The policy statement is where the details (the meat and potatoes) of the policy are explained. Policy statements are written in a high-level format and can reference other sources, such as procedures, as a way to explain what is required to properly meet the goals of the policy. A policy can state multiple items and group content based on similar concepts. For example, all topics that cover protecting information could fall under a section titled "Security and Proprietary Information." Grouping concepts within a policy works well; however, recommended practice dictates that you split up a policy into more-focused policies if too many categories are needed to properly cover the policy concept.

Policy Statement Example

An example policy for "Security and Proprietary Information" follows. This example shows only one group within the policy, with sections 1.1.0–1.1.4. If the policy has other groups, they could be presented as sections 1.2.0–1.2*X*, 1.3.0–1.3.*X*, and so forth.

Policy: Security and Proprietary Information

1.1.0. All computing devices that connect to the **<COMPANY>** managed internal network must comply with **<COMPANY>** access policy.

1.1.1. System-level and user passwords must comply with **<COMPANY>** password policy. Providing access to passwords to another individual is prohibited.

1.1.2. All computing systems must be secured using a password-protected screensaver with automatic activation feature set to 5 minutes or less.

1.1.3. Any postings by employees of **<COMPANY>** to external or internal media must contain a disclaimer stating the opinions expressed are strictly their own and do not represent **<COMPANY>**, unless otherwise permitted to do so on behalf of

<COMPANY>.

1.1.4. Employees must use caution when opening email attachments from unknown senders.

This policy example has one topic group focused on information security, which could by itself be the complete policy or could be a topic group within a larger policy. Notice that none of the parts of this policy provide details on how a requirement is met. Also, notice that none of these statements would likely need to be modified as the organization changes, because the concepts are broad enough to adapt to change. Lastly, notice that consequences for not meeting this policy are not listed in this section. Consequences come in the next section, policy compliance.

Unlike the preceding policy statement example, which covers what is permitted, a policy statement can also be written regarding actions that are not authorized. For example, a policy for “unacceptable usage of company assets” would state various activities that are a violation of the company policy. Part of a policy covering unwanted behavior could state that no employee is ever permitted to engage in illegal activities. The options for what can exist in a policy are endless and should be based on a specific overall objective that is backed by leadership.

Policy statements can be very broad or very specific. Broad policies are ideal to give guidance and accommodate many different possible situations, while a specific policy is more direct, but anything outside of the specific defined criteria would not be a violation. I recommend layering both types so that when a specific policy doesn’t catch a violation, a broader policy is put in place to step in. For example, a specific policy could state “Users are required to change their passwords every 180 days.” A broad policy could state “It is the policy of this organization to secure credential information for all assets from unauthorized use or disclosure.” The specific policy requires users to take a specific action to reduce the risk of password compromise, but it doesn’t cover all actions users must or must not take to secure passwords. For example, if a user shares their password with an outside party, that would not be a violation of the specific policy example, but it would violate the broad policy that covers anything that endangers the security of credential

information, including the action of sharing a password. Prohibiting the sharing of passwords would likely be covered in separate specific policy, but if the policy writers overlook including that provision, the broad policy a covers the void.

> **Note**
>
> I recommend reviewing policy statements with a legal professional to ensure that they don't include any language that could be construed as illegal in the country in which your organization operates. Certain policy restrictions or requirements that are intended to secure your organization might also be considered, for example, a violation of employee privacy rights.

Policy Compliance

The policy compliance section states how the policy will be validated, what are the potential exceptions to the policy and the possible outcomes if someone fails to comply with the policy. This section is important not only to explain the validation process of a policy, but also to identify who is responsible for validating compliance and where to seek requests for exceptions. Will the InfoSec team enforce validation of the policy using detection tools, monitoring techniques, etc., or is it up to the user? Will a specific group have the authority to grant an exception to the policy? This section should answer these questions, though it shouldn't specify individuals who are enforcing the policy or how to contact them. Those details can change often and thus should be included in a procedure that can be referenced by the policy. An example resource could be "IT Security" or "The Security Operations Center," which both are generic enough to accommodate change within those groups. If a specific contact must be included, it is best to use a generic contact method, such as itsupport@company.com, that is tied to the current team within the IT security team.

Policy compliance must also state what will occur if the policy is not met. An

example of good use of language for noncompliant behavior could be “Any employee found to have violated this policy may be subject to disciplinary action, up to and including termination of employment.” This example provides the opportunity for the policy violator’s manager and/or human resources to determine the specific outcome if a policy is not met.

Related Standards, Policies, Guidelines, and Processes

This section is ideal for referencing other sources related to the policy. Sources can include other policies that align with this policy’s goals, sources that are referenced within this policy such as another larger policy, related procedures that would define in more detail what would be required to enforce this policy, related disciplinary and legal details, or anything else that would make sense to associate with this policy.

Definitions and Terms

For the definitions and terms section, include any definitions and references that will help the reader understand the policy. For example, if the topic of phishing is covered within the policy, a reference to the definition of phishing should be included in this section, not only to explain what phishing is, but also to explain how the term is used within this policy. Some concepts can be interpreted in different ways, and this section can help clarify which definition is being referenced. For example, if the policy covers a denial-of-service (DoS) concept, it might be appropriate to explain in this section if the denial-of-service example leverages a large volume of systems (known as a volumetric DoS attack) or exploits a vulnerability within a protocol to cause the system to become unavailable, which is a different type of DoS attack. Best practice is to consider the audience of a policy and include definitions for any terms that a reader might be unfamiliar with or terms that would need more clarity to understand how they are being used. A reference symbol such as * could be placed next to any term in the policy that the reader could seek out in this section for more clarification.

History

This section lists when the document was created and when any updates have been made. It is ideal to include the date of the change, who made the change, and a summary of the change. All three of these pieces of information are important to help validate how often changes are made and who to question about any changes. Recommended practice dictates changing the version of the policy and tracking version updates within this section, including the changes that were made to the new version.

Note

I recommend seeking templates and consulting when developing policies rather than developing policy details from scratch. Using trusted references not only provides external checks and balances to what is being enforced but also reduces the likelihood of developing ineffective processes that will not be well accepted by the impacted parties. I also recommend crossreferencing sources if they are available, such as “see <Guideline> for details” or “in compliance with <Guideline>,” regarding the purpose of the statement being made.

Launching a New Policy

The previous eight sections covered a practical layout for developing a policy. The specific sections that you include in any particular policy will depend on the type of policy and the intended audience. Regardless of the chosen format, it is critical that all policies are backed by leadership to ensure they are properly supported and enforced. Leadership that sponsors a policy is ideally part of the C-level team, meaning the CEO, CSO, or another position that has enough authority to provide required support to properly execute the policy as well as enforce any disciplinary actions that occur if the policy is not followed. Support can include budget for people, process, and technology.

Prior to launching a new policy, make sure to identify prerequisites to sustain the new policy, such as budget, to avoid a situation where expected results are unattainable. For example, if a policy requires multifactor authentication for physical access to doors but the existing doors do not have a method to provide multifactor authentication, that would be an obvious policy failure that will continue to be a failure until budget and resources are made available to acquire and deliver authentication technology for each physical door. It is common for policies to be launched with the expectation that the current state of the organization does not meet the requirements; however, provide a set time to correct the missing policy elements before any violation is considered. I have seen many requests for proposals from customers based on policies they are being asked to enforce before a specific date such as the end of the year.

Steps for Launching a New Policy

Launching a new policy should include the following three general steps to inform the organization about the changes associated with the new policy:

1. **Distribute the policy.** For a policy to be legally enforceable in relation to employees, the employer must bring it to the attention of all employees. Each employee that is impacted by the policy should be provided a copy of the policy, either in electronic form, hard copy, or both. Employees need to have easy access to any policy that impacts them, again either electronically, such as an intranet resource, or in hard copy, such as in a department policy manual.
2. **Obtain acknowledgment of receipt.** It is important to have each impacted employee sign and date an acknowledgment stating they received, read, and understood the policy and were afforded an opportunity to ask questions about it. File a copy of the acknowledgment in the employee's personnel file in the event you require it at a future date. There will always be a handful of employees who don't (or won't) sign such acknowledgments; in those cases, two management members can themselves sign an acknowledgment stating that they gave "the person" a copy of the policy but "the person"

refused to sign the acknowledgment.

3. **Provide advance notice.** If the new or changed policy significantly affects or changes any employment rights or benefits, such as vacation rights, you need to give employees advance notice of the policy implementation and a grace period for the change to take effect. Otherwise, the change can amount to "constructive dismissal." Constructive dismissal occurs when the employer unilaterally changes a fundamental term(s) or condition(s) of employment, effectively breaching the employment contract. The employee can either accept the change, essentially agreeing to the new contract, or refuse to accept the change, accept the employer's termination of the contract, and sue the employer for wrongful dismissal. Just what amounts to a fundamental change and the length of the required advance notice depends on the circumstances. In general, however, the more significant the change, or the longer the employee's service, the longer the advance notice of the change that's required.

How each of these three actions is carried out will depend on the impact the policy has on the associated parties. Minor policy changes or new policies with minor impact require far less advanced notice than a major policy change or a new policy with major impact. I find it common for the distribution of electronic or physical documentation for minor policy changes or new policies with minor impact to occur during an annual policy reminder session, training, or notification that includes the combination of various minor changes, rather than bothering employees with documentation for every minor policy change or new policy with minor impact. Acknowledgment of receipt can be a required action item following the policy reminder session, training, or notification, allowing for the acceptance of all policy changes at once as well as any new policies.

Policy Enforcement

Another important policy concept is that a policy must be monitored to validate that the policy is being followed; otherwise, policy violations will go unnoticed. I remember working for a government agency that had a policy

stating, “All contractor laptops must be scanned before accessing the network.” That policy was displayed on physical signs around the office; however, it was up to the contractor or sponsor of the contractor to notify security that a contractor’s device needed to be scanned. This poor practice led to many contractor laptops connecting to the network without being scanned because the organization loosely monitored and enforced the policy. I also don’t recall if any repercussion existed if a contractor was caught not having their laptop scanned outside of it being considered “frowned upon,” which further drove the behavior to ignore this policy. This example shows why it is highly recommended to include in your policy template a policy compliance section that states who is responsible for monitoring and enforcing the policy as well as the consequences of a failure to comply with the policy.

Although a policy compliance section states how the policy will be validated, this doesn’t solve the problem of ensuring the policy is being enforced once it goes live. Who is responsible for looking for violations? When a violation is identified, who makes the determination that the action is or is not a clear violation of the policy? Looking back at my example of a broad policy, “It is the policy of this organization to secure credential information for all assets from unauthorized use or disclosure,” I gave the example of an employee sharing his or her password as a violation. Suppose the employee gave the password to a contractor who was authorized to access the system on the date of the alleged policy violation. The employee could argue that providing the password to the contractor on that day was authorized and therefore no policy violation occurred. The employer could argue that it is a violation because the contractor continues to know the password after the date the contractor was authorized to access the system. If there isn’t another policy that requires contractors to be provided a temporary password, it is possible an employee would share a password while a contractor is authorized and technically would not be violating the policy. Which argument is correct? It depends on how your organization enforces policies.

Enforcing a policy starts with enabling the people responsible for validating a policy with the proper authority and training. These people are typically the managers and supervisors who manage the affected employees, but can also be a third-party validator. Training needs to include how to handle violations. When a violation occurs, the appropriate disciplinary penalty needs to be

imposed consistently and warnings of consequences need to be provided regarding future violations.

Policy violations need to be tracked. Chapter 8, "Threat Hunting and Incident Response," covers case management concepts. Case management also applies to tracking policy violations. The goal is to document when a violation occurs to gain an understanding of how successful the policy is, validate if changes need to be made due to excessive violations, and track parties involved for possible disciplinary actions. If tracking shows a number of requested exceptions for a policy, the risk management team will need to adjust the policy based on the recorded feedback. If an employee is found to have a history of policy violations, action will need to be taken to reduce future violations. All of these outcomes require a tracking system to document and adapt to policy violations.

Another, more formal approach to enforcing and adapting policies is to use an official certification and accreditation program.

Certification and Accreditation

A formal way to enforce policies is through a certification and accreditation process. The certifier is the party responsible for assessing if policy elements are being met. This practice is common for meeting formal certification programs such as PCI DSS (discussed later in the chapter), but can also be part of the process to enforce organization-driven policies. For certification programs, organizations will use a third-party certifier to validate that the processes, systems, products, events, or skills meet some existing standard. For general polices, an organization can authorize one or more people within the organization to certify if a policy is or is not met. This team is typically part of the compliance team.

Accreditation means a formal declaration by the certifier that what the certifier was reviewing meets all requirements. If an organization is looking to meet a formal certification program, once the organization passes all checks by the certifier, the organization can be accredited as being compliant. The entity that declares accreditation not only is confirming that the organization has met all requirements, but also is owning all responsibilities

for everything that it has accredited. If a violation is found post accreditation, the accreditor might be liable, depending on the situation. The same concept applies to an accreditation given by the compliance team within an organization.

Regarding general policies, somebody with authority needs to be able to declare an accreditation when polices are met. For example, if there are minimal security requirements that must be put in place before a server is allowed to connect to the network, somebody needs to know how to validate that those security controls are in place, representing the certifier, and somebody needs to formally approve the system is ready to go live, representing the accreditor. Without these two processes, polices will not be enforced because no one will know how to validate against a policy and no one will be concerned about breaking a policy because no formal validation is put in place.

Note

One extremely important value obtained from using a certification and accreditation program is that it addresses risk. Security policies are designed to reduce the risk of threats. If a certifier finds that a policy is not being followed, the certifier can identify why and recommend controls that would enforce policy compliance. Even though compliance doesn't take into consideration all risks, enforcing certification and accreditation of your policies will accommodate all risks associated with any policies being validated.

When it comes to the success of a policy, determining *how* the policy is executed by the organization is critical. *Procedures* fill in the missing details for policies. This brings us to the topic of procedures.

Procedures

Procedures are the details that support what a policy is designed to

accomplish. Procedures are more specific than policies and will change as people, process, and technology adapt to changes within the business. Think of procedures as the step-by-step requirements to accomplish a policy. Looking back at the previous policy scenarios, an example procedure to match the policy “Employees are permitted to use only company-issued devices on the company network” would explain the details of what is considered an employee-issued device. For this example, let’s say an employee-issued device is only one that is provided by the employer and running approved software. The procedure document would define this as well as explain how validation of an employee-issued device could be achieved either by verifying a hidden certificate or by adding the MAC address of the device to a whitelist that includes the MAC addresses of all corporate-issued assets. The procedure document would also need to explain details such as how acceptable devices are distributed, how they are validated as they connect to and are used within the network, who should receive requests for exceptions, and what will occur if a violation is identified. Notice that many of the details in this example procedure document will likely change, making the maintenance process of the procedure documentation just as critical as its original creation. Changes to procedures should be captured in updated versions of the associated procedures, while the policy being supported should remain untouched. I find that when an organization has failures in processes, the cause typically is not how the documents were created but rather a failure in documentation maintenance and training.

Procedure Document

The format of a procedure document can vary as long as the required steps are included to properly execute all required tasks. Topics included in a procedure document should cover the purpose of the procedure, what is required to perform the desired tasks, when tasks should be performed, who is responsible for the tasks, and any other required steps that must be performed to properly execute the procedure. Let’s look at an example procedure document for registering a website.

Procedure for Website Registrar at <COMPANY>:

The website register must record the following information and provide the results to register@Company.org within 30 days of launching a new website. A notification from the service provider will be provided regarding the status of the request to the email provided.

- List of all domain names that are registered
- The renewal dates for each domain name
- Names of all hosting services
- Any services that expire with expiration dates
- Associated costs and the responsible department
- Name and job title of the requestor
- Email of the requestor
- Purpose of the website

It is the responsibility of the registration requestor for all maintenance of the registered website. Network services at **<COMPANY>** shall provide notification of the expiration date of the website registration to the requestor 90 days prior to the expiration, after which it is the responsibility of the requestor to renew the website prior to the expiration date. Any expired websites must have the owners submit a new request for website registry.

This example procedure document lists the specific details regarding what is required to register a website within the company. All details are required to be sent to register@company.org, which will respond with status updates as the request is being processed. Responsibilities for both the requestor and organization are listed, including expectations for notification of the website expiring as well as who should renew the website. The steps provided in this document are clear enough that any employee who wants to register a website within the organization can follow the procedure to properly submit a request. Other procedure sections would likely be included with this example procedure section, including a section covering procedures for deploying content on the website, securing the website, and performing other website-related tasks. These topics could all be grouped under one large procedure

document for deploying websites at the company, or each topic could be a separate procedure document that references other documents as users need to request, stand up, and maintain a website at the company. Also, all of these procedures could fall under one or more policies, such as a policy for "deploying websites at <COMPANY>."

One theme I have continued to emphasize is that there isn't a set rule for developing policies and procedures. While this flexibility is great for accommodating creativity and the many dynamics found within an organization, having an open-ended format for building process documentation can cause confusion when an organization attempts to build and validate well-thought-out policies and procedures. I find that many organizations don't know how to validate their new or existing policies and procedures or identify gaps and are therefore unsure if their policies and procedures are effective. One recommended approach to testing an organization's policies and procedures is to perform a tabletop exercise. In the next section I will walk you through developing and executing a successful tabletop exercise. This will allow you to validate and improve your processes.

Tabletop Exercise

A tabletop exercise is a hypothetical scenario–based walkthrough of how an organization would respond to a security incident. The purpose of a tabletop exercise is to test the people, process, and technology conceptually without having these resources perform the actual response. Other services such as penetration testing are ideal when real-world testing of a policy is desired.

An example of a tabletop exercise is one that tests how an organization would respond to a ransomware pandemic. Ransomware doesn't just appear on an endpoint, meaning a tabletop exercise would consider all aspects of this situation from response to recovery. For this scenario, the following are sample questions that the team performing the tabletop exercise could be asked:

- Who is responsible for receiving the notification from the end user that malware has been identified on the user's system?

- What endpoint applications exist that could identify and alert IT of malware like ransomware?
- What tools exist to determine the type of malware being identified?
- Which team monitors endpoints and how does that team respond to events?
- When malware is identified, how is the situation scoped and contained?
- What remediation actions would take place and who would perform those steps?
- What forensic steps would be taken to identify how the malware was installed on the endpoint?
- What vulnerability and security assessment actions would follow this situation to ensure that other systems are not affected by the malware or vulnerability that led to the malware infection?
- Who would be responsible for any public or internal messaging explaining the impact from this malware?
- Which specific team members would be part of the response team? Who are their backups?
- What response time and result would be considered a success or a failure?
- How does the organization's current expected response align with desired successful responses?
- What gaps in people, process, and technology exist that would impact a successful response to this threat?
- Which members of the HR, management, and legal teams would be involved for any required disciplinary actions resulting from this incident?

As you can see from the list of example questions, a tabletop exercise can be very involved and require many hours to properly execute. You should also notice that members from different parts of the organization are required to thoroughly validate the people, process, and technology that would be evaluated during a tabletop exercise. This brings up the question of who should be involved in a tabletop exercise. The short answer is, it depends. The following section outlines different options for delivering tabletop

exercises, which require different people to be involved with the exercises.

Tabletop Exercise Options

There are different approaches to executing a tabletop exercise. One approach is to run the same scenario for different groups within the organization, which means performing multiple smaller tabletop exercises. For the ransomware example, the organization could first work with desktop support to capture how they would respond to a ransomware situation. Next, the organization could meet with the network security team and run the same ransomware scenario, but with a focus on their impact on the overall incident response. Using this smaller-group approach is ideal for keeping the exercise focused with fewer people involved and possibly reducing concern that other teams would influence or impact how a specific team would respond during the tabletop exercise.

Another approach to delivering a tabletop exercise is to include managers from each department and obtain input from all departments during the same meeting. This approach allows each department to contribute while the scenario is being covered, rather than requiring the team delivering the tabletop exercise to combine the responses of all the individual group exercises into a single tabletop exercise results document. This also allows for real-time debate as current and potential future responses based on changes occurring post table top exercise are analyzed. Some concerns when including all teams within an organization are the size of the meeting, the time required to hear feedback from all included members within the tabletop exercise, and the potential for conflict between departments. Limiting the tabletop exercise to only department managers can mitigate these and other problems; however, department managers might not know the details required to answer a question during the tabletop exercise. For example, a manager might not know the specific technology that would be used, but would know which member of his or her team would be able to answer such questions. Using the manager-led tabletop exercise format would require managers to take any unanswered questions back to their team and later provide a response.

Tabletop Exercise Execution

The best format for a tabletop exercise typically depends on the scenario that is being tested. However, the following tips apply generally to executing a successful tabletop exercise, paraphrased from the CSO article, “6 Tips for Effective Security Tabletop Testing” by Bob Violino:

- **Take time to prepare for the exercise:** Invest time in developing the objectives and the scope and selecting the proper participants. Planning is often the most time-consuming phase but will pay off regarding the success of the exercise. Having the wrong people involved or a scope that is too narrow or too broad will not only produce poor results but also cause those involved to believe their time was wasted, leading to a lack of support for any recommended adjustments following the tabletop exercise.
- **Involve multiple parties from different departments:** Security is the responsibility of every person in an organization. That means the response to an incident will impact many different departments. The best tabletop exercise results come from involving all parties that would be impacted and could provide value or cause complications during the incident response.
- **Establish the ground rules up-front:** A tabletop exercise represents walking through stressful situations. Some people might want to speak openly regardless of whether the topic is positive or negative, while others might not want to participate unless called upon. Not addressing the different personalities in the room will lead to frustration and loss in support for the results of the exercise. Make sure to set expectations up-front for required participation, what is and is not permitted to be discussed during the exercise, what can and cannot be talked about after the meeting, and any other ground rules to ensure a successful use of everybody’s time.
- **Leverage external resources if available:** Tabletop exercise services have been used by many organizations, and best practice documentation is available. Hiring a consultant that provides tabletop exercise services can be valuable because the consultant offers an external, unbiased voice and has experience executing successful tabletop exercises. If your team has

never delivered a tabletop exercise, it is recommended to seek external support and learn from the best practices of others before performing your own internal exercises.

- **Broader is likely better:** Discussion of some tabletop exercise topics can either go very deep into details or remain broad. Keeping the discussion broad likely is better unless the team involved in the tabletop exercise believes the topic being covered requires specific details. For most tabletop exercises, details can be captured by the expert after the tabletop exercise, allowing for all topics to remain broad and for all team members to be engaged (versus the risk of only hearing from specific experts). Keeping conversations broad also shortens the duration of the tabletop exercise because details are not addressed until after the meeting.
- **Use realistic scenarios:** The purpose of a tabletop exercise is to evaluate the response to a potential future situation. All scenarios need to be as realistic as possible in order to produce a realistic result.

Tabletop Exercise Format

It is recommended to include certain topics within the tabletop exercise document that is used during and referenced after the exercise is performed. First, an introduction should explain the purpose of the exercise. A goals section should be included (or goals should also be listed in the introduction section) so that the expected outcome is clear. All details about the meeting need to be documented, including the date, facilitator, and all participants. All processes that are evaluated should also be listed, including desired policies and procedures that might not exist but are planned to be developed, with the current status noted. The tabletop exercise scribes need to document each scenario, including its target capabilities being tested, objectives for the exercise, and details regarding how the organization would respond if the situation occurred. Documentation of recommended remediation may or may not be included depending on whether the purpose of the tabletop exercise is to just evaluate processes or to also include corrective actions. All of this needs to be included in the tabletop exercise post event documentation.

I am often asked which tabletop scenarios are best for an organization to run

through. Again, my answer is it depends on your business, as no single threat list fits all organizations. The tabletop exercise scenario list for an organization that runs an online business will be very different than one for a school. If you need a general guideline for scenarios, I highly suggest considering using playbook templates. Chapter 8, "Threat Hunting and Incident Response," and Chapter 10, "Data Orchestration," both provide a deep look at playbooks. In both chapters, I will run through different playbook examples provided by the Incident Response Consortium (IRC) at https://www.incidentresponse.com/playbooks/. Figure 6-1 shows the gallery of free playbooks available from IRC, which include recommended steps for different parts of responding to specific threats. Playbook templates are a great resource for building a template for your tabletop exercise, but I highly recommend prioritizing the threats you expect to impact your organization versus running through random playbook templates.

FIGURE 6-1 Incident Response Consortium Playbook Options

Tabletop Exercise Template Example

This section presents an example of a tabletop exercise template. If you decide to use a playbook template to develop your tabletop exercise template,

you can use this example as a way to structure what questions you will ask based on the steps recommended in the playbook template. For example, if a malware outbreak playbook suggests informing IT about the situation, you would include a question such as “Who would be informed when a malware outbreak event is identified?” The same process would apply as you evaluate your SOC service playbooks against a tabletop exercise. Playbooks will be covered in more detail in Chapters 8 and 10.

The following is a sample tabletop exercise template for disaster planning featuring many of these suggested sections.

Disaster Planning Guide for <COMPANY>

This forum is used to design and facilitate a tabletop exercise (TTX) as well as provide appropriate documentation of performance and findings during the exercise.

The TTX involves administrative staff, management, and other key personnel in an informal group discussion focused on a hypothetical situation.

The purpose of the TTX is to evaluate existing plans, policies, and procedures without incurring signification costs and time commitment required to deploy and test actual resources. A TTX allows participants to work through a problem in a controlled environment in compressed or simulated time without the pressures of an operations-based exercise.

It is recommended that the TTX be completed on a regular basis for potential threats that have been identified by the SOC and leadership.

Goals:

Participants in a TTX will:

- Identify strengths and opportunities for improvement
- Enhance understanding of new concepts
- Change attitudes and perspectives

Conduct Characteristics:

- Requires an experienced person to facilitate the TTX
- Promotes in-depth discussions
- Involves slow-paced problem solving in simulated/compressed time

Date/Facilitator/Participants:

Plans, Policies, Procedures referenced during TTX:

Scenario 1:

- Overview
- The time of occurrence and impacted system
- Identified objectives for operational period
- Identified tasks that need to be performed to meet objectives
- Response from team members regarding current capabilities

Scenario 2 (Repeat):

Tabletop Exercise Evaluation:

- Identified areas of strength
- Identified opportunities for improvement
- Identified roles of CUSTOMER in the TTX
- Identify any changes that may result from the TTX

Note

The Center for Internet Security (CIS) offers tabletop exercise templates. One example can be found at https://www.cisecurity.org/wp-content/uploads/2018/10/Six-tabletop-exercises-FINAL.pdf.

Another approach to validating policies and procedures outside of performing a tabletop exercise is to compare what would be performed by an

organization against what is considered industry best practice. This leads us to the topic of standards, guidelines, and frameworks, all of which can be reference points for what should be included in your organization's policies and procedures. Standards, guidelines, and frameworks are also useful for referencing what steps should be taken during a tabletop exercise so that existing policy and procedures can be compared against what the industry would recommend should be done when addressing a scenario.

Note

Many organizations incorporate standards, guidelines, and frameworks into their policies and procedures rather than building policies and procedures from scratch.

Standards, Guidelines, and Frameworks

Chapter 1 introduced the concepts of standards, guidelines, and frameworks. Although organizations are not legally obligated to follow standards, guidelines, and frameworks, many are highly recommended as security best practices and often used as templates for developing policies and procedures as well as meeting various forms of compliance. There are many resources that provide their interpretation of best practices, which can cause confusion when seeking out advice for improving an organization's security. When considering the merits of a specific resource, it is recommended to evaluate the reputation of the organization that created the resource, what was considered when the standard, guideline, or framework was developed, the target audience for the resource, when the resource was created, and if any vulnerabilities or other issues have been documented in reference to the use of the resource. Common practices for verifying the usefulness of a resource include speaking with your counterparts at other organizations in the same field of business to get their opinions as well as researching the company that is offering the resource.

Note

In some cases, the leadership of an organization may adopt a certain standard, guideline, or framework as a policy, which means the standard, guideline, or framework must be followed regardless if it is dated or potentially no longer relevant. This could cause problems and introduce vulnerabilities depending on what is being required. For example, if a policy requires to follow a resource that approves only older software on a system, updating or patching the software on that system might violate the policy and put the system at risk! I recommend leadership of an organization to speak with technical experts regarding potential impact of change to reduce the risk of policies causing negative impact to the organization.

Next I will review some of the most commonly used security standards, guidelines, and frameworks, all of which were introduced in previous chapters. Many of these resources will have overlapping concepts, which should be expected because they are all attempting to provide industry best practices for various security topics. Keep in mind that it is common for people, process, and technology to change faster than external resources such as the following can adapt to, which could lead to introducing vulnerabilities even if the recommended resource is followed. This is why you should validate the resource's creation and maintenance date as well as how relevant the resource is to your organization's goals prior to leveraging the resource. The same concept applies when developing policies and procedures that reference external sources. With proper validation, standards, guidelines, and frameworks can be excellent resources for assisting with developing policies and procedures. Remember that meeting compliance doesn't mean you are meeting a best practice for a security baseline. You will need to exceed compliance requirements using a more holistic view of security.

NIST Cybersecurity Framework

Introduced in Chapter 1, the NIST Cybersecurity Framework (CSF) is one of

the most popular frameworks consisting of standards, guidelines, and best practices for dealing with cybersecurity-related risk. Chapter 1 also introduced the NIST Framework Core and five framework Functions, as shown in Figure 6-2.

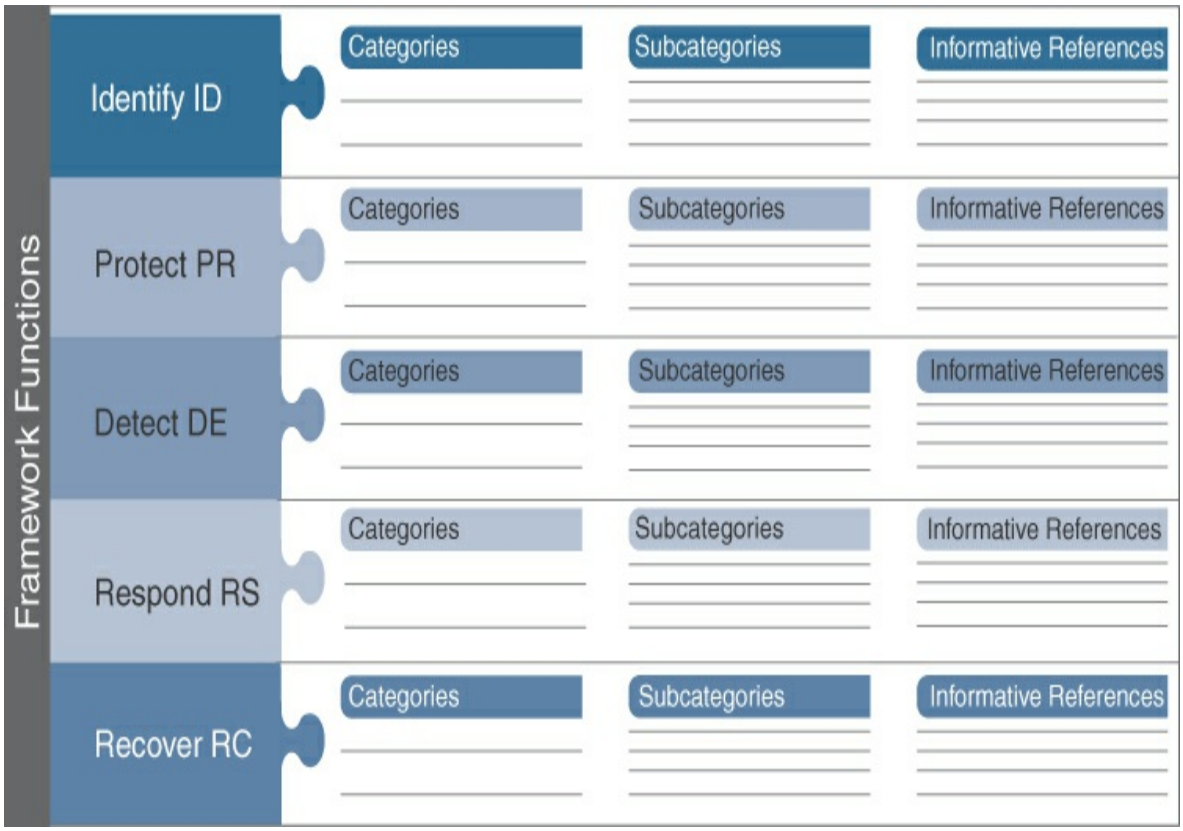

FIGURE 6-2 NIST Cybersecurity Framework Core Structure

As a refresher from Chapter 1, each of the five functions of the NIST Framework Core has a high-level view on one aspect of security. Anything under the Identify function would represent steps used to manage risk to systems, people, assets, data, and capabilities. The Protect function would cover topics that develop and implement people, process, and technology to reduce risk proactively and ensure delivery of critical services. Reactive risk recovery topics would fall under the Recover function since they cover actions that follow an incident, including topics covering how to restore and resist future compromise. Detect and Respond functions cover topics that involve how an organization identifies and responds to incidents.

Using the NIST Cybersecurity Framework

The CSF tiers provide context on how an organization views cybersecurity risk and the processes in place to manage that risk. The NIST CSF approach to grouping topics enables an organization to pick an area of interest and review targeted recommendations to best understand what NIST would consider best practices for people, process, and technology. Tiers cover an increasing degree of rigor and sophistication in cybersecurity risk management practices. The tiers approach helps determine the extent to which cybersecurity risk management is informed by business needs and is integrated into an organization's overall risk management practices. This approach also considers an organization's current risk management practices, threat environment, legal and regulatory requirements, information sharing practices, business/mission objectives, supply chain cybersecurity requirements, and organizational constraints. Understanding the maturity of a recommendation is critical not only to identity an organization's current capacities but also to set goals and milestones to improve the associated people, process, and technology against what is being recommended by NIST.

The NIST CSF can be used by first identifying the capability that you want to assess against what NIST views as industry best practices. You can use the NIST capability assessment process to understand where your capability stands today based on its current maturity level as well as identify goals for improvement based on gaps identified during the assessment process. The NIST CSF uses a tier model to categorize different maturity levels of capabilities, where the fourth tier represents the most mature capabilities. It is ideal to target all capabilities to reach NIST's highest tier in maturity, but some capabilities may not need to reach this level. For those situations, a lower tier can be set as the target profile during the evaluation process.

NIST Tiers

The following list breaks down each tier of the NIST Cybersecurity Framework. Each higher tier assumes capabilities from a previous tier plus additional improvements. Some of your capabilities might fall between tiers, making it a judgment call as to where your capabilities would be documented according to the NIST recommendations.

Note

It is left to the reviewer's discretion to determine at which tier an organization's current capabilities should be classified. This could lead to a false sense of security or false positives if poor judgment is used when determining current capability tier ranking.

- **Partial (Tier 1):** Organizational cybersecurity risk management process is not formalized and risk is managed in an ad hoc approach and sometimes reactive manner. There is very limited awareness of cybersecurity risk at the organization level. The organization does not understand its role in the larger ecosystem with respect to collaborating with or receiving information from external resources.
- **Risk Informed (Tier 2):** Risk management practices are approved by management but might not be established as organization-wide policy. There is some awareness of cybersecurity risk at the organizational level, but an organization-wide approach to managing cybersecurity risk has not been established. There is some collaboration with external sources via receiving and sharing security information; however, it is limited and not a repeatable collaboration.
- **Repeatable (Tier 3):** The organization's risk management practices are formally approved and expressed as policy. Cybersecurity practices are regularly updated based on business/mission requirements and the changing threat and technology landscape. There is an organization-wide approach to managing cybersecurity risk. This includes well-defined policies and procedures, which are periodically reviewed and updated. The organization leverages and contributes to external resources through established collaboration relationships. This leads to an awareness of the organization's understanding of existing risks based on industry standards, guidelines, and frameworks.
- **Adaptive (Tier 4):** The organization adapts its cybersecurity practices based on previous and current cybersecurity activities, including lessons learned and predictive indicators. Through a process of continuous improvement incorporating advanced cybersecurity technologies and

practices, the organization actively adapts to a changing threat and technology landscape and responds in a timely and effective manner to evolving, sophisticated threats. The relationship between the organization and cybersecurity goals is clearly defined and understood and considered when decisions are made. The organization has adopted external partnerships that are essential to decisions regarding understanding and responding to the current threat landscape.

An organization can use the NIST Cybersecurity Framework as part of its systematic process for identifying, assessing, and managing cybersecurity risk. The best approach is to match existing capabilities against what is recommended by NIST and identify at which tier your organization's current capabilities should be pegged. When evaluating a capability, you need to consider the entire lifecycle of the capability, including how the capability is planned, designed, deployed, run, and decommissioned. During the lifecycle of a capability, its maturity rating can be evaluated at any point against the NIST tier structure to better understand its current state of maturity. For example, the tier level of a capability should increase as you plan a new capability, deploy it, and eventually tune it so it provides maximum value. In this example, you would want to continuously update how the capability's tier or maturity is documented rather than document it during the planning phase and never adjust its perceived maturity as it goes live and provides more value to the organization.

NIST CSF Capability Assessment

NIST CSF tiers are a way to understand the maturity of a capability. A capability, however, must be evaluated to understand which NIST tier it falls within. You cannot assume all of your capabilities are at maximum maturity, or Tier 4 in the NIST CSF. You must continuously evaluate your capabilities with the goal to improve their maturity, which is represented as an increase in NIST CSF tier status. Using the NIST CSF to assess an organization's capabilities not only enables the organization to map each capability to a NIST Tier to better understand the organization's current strengths and weakness but enable the organization for a roadmap to improvement.

The following steps illustrate how an organization could use the Framework to create a new cybersecurity program or improve an existing program:

Step 1. Prioritize and Scope: In this step, the organization must identify organization or mission objectives along with high-level organizational priorities. NIST offers many recommendations for your organization to follow, some of which will be more relevant to your organization than others. It is up to your reviewer to determine the order of how the people, process, and technology are evaluated by NIST based on various business and technology factors related to the organization.

Step 2. Orient: Identify all related systems and assets, regulatory requirements, and overall risk approach to what is being assessed.

Step 3. Create a Current Profile: Document which NIST Framework Core Categories and Subcategories (introduced in Chapter 1) are currently being achieved.

Step 4. Conduct a Risk Assessment: Assess existing capabilities to determine gaps in what is identified as achieved and missing. The topic of performing an assessment will be covered later in this chapter in the section "Risk Assessment."

Step 5. Create a Target Profile: Create a target profile that focuses on the CSF Categories and Subcategories assessment and describes the desired cybersecurity outcomes.

Step 6. Determine, Analyze, and Prioritize Gaps: Compare the current profile against the target profile to determine what changes in people, process, and technology will help the organization achieve the desired target profile.

Step 7. Implement Action Plan: Prioritize and execute on actions needed to achieve the target profile.

Keep in mind that the NIST CSF tier approach is a way to measure a capability's maturity and that the assessment process places each capability within a tier. It is up to the organization to determine the desired tier the capability being evaluated should be functioning at as well as to set that tier as the target profile. All capabilities do not need a target profile of Tier 4, Adaptive, but any critical capability is ideal for Tier 4 performance.

An example of using the NIST CSF to validate industry best practices would

be to view the Identify (ID) function, Asset Management category, and ID.AM-2 subcategory that specifies “Software platforms and application within the organization are inventoried.” What also makes the NIST CSF a valuable option is that it references other industry guidelines to support its recommendations. For example, the “Informative References” column for the NIST CSF ID.AM-2 subcategory points to the following industry guidelines that align with this recommendation: CIS CSC 2; COBIT 5 BAI09.01, BAI09.02, BAI09.05; ISA 62443-2-1:2009 4.2.3.4; ISA 62443-3-3:2013 SR 7.8; ISO/IEC 27001:2013 A.8.1.1, A.8.1.2, A.12.5.1; and NIST SP 800-53 Rev. 4 CM-8, PM-5. I recommend viewing the other recommendations in an “Informative References” section of any NIST CSF best practice to better understand the security control based on seeing more than one guideline’s view of the same topic. During your assessment of NIST CSF ID.AM-2, you would identify which tier (1 through 4) represents your organization’s capability to manage the inventory of its software platforms and applications. Check out the full NIST CSF to review the various topics that you can use to assess your security capabilities for people, process, and technology.

Note

The NIST Cybersecurity Framework is available free of charge at https://doi.org/10.6028/NIST.CSWP.04162018.

ISO/IEC 27005

The International Organization for Standardization (ISO) creates documents that provide requirements, specifications, guidelines, or characteristics that can be used consistently to ensure that materials, products, processes, and services are fit for their purpose. Part of ISO’s value is its family of standards that provide best practices for various cybersecurity concepts.

One popular publication is ISO/IEC 27005, *Information technology – Security techniques – Information security risk management*, which offers guidelines for information security risk management. As shown in Figure 6-3, the ISO/IEC 27005 steps for managing information security risk include

establishing context; conducting a risk assessment by identifying, analyzing, and evaluating risk; and then treating the risk. Ongoing activities during these steps include communication, consultation, monitoring, and review. Any organization can use the details provided in ISO/IEC 27005 to help improve how it handles information security risk management through auditing against ISO recommendations or performing a tabletop exercise and using ISO recommendations as the source against which to compare existing practices.

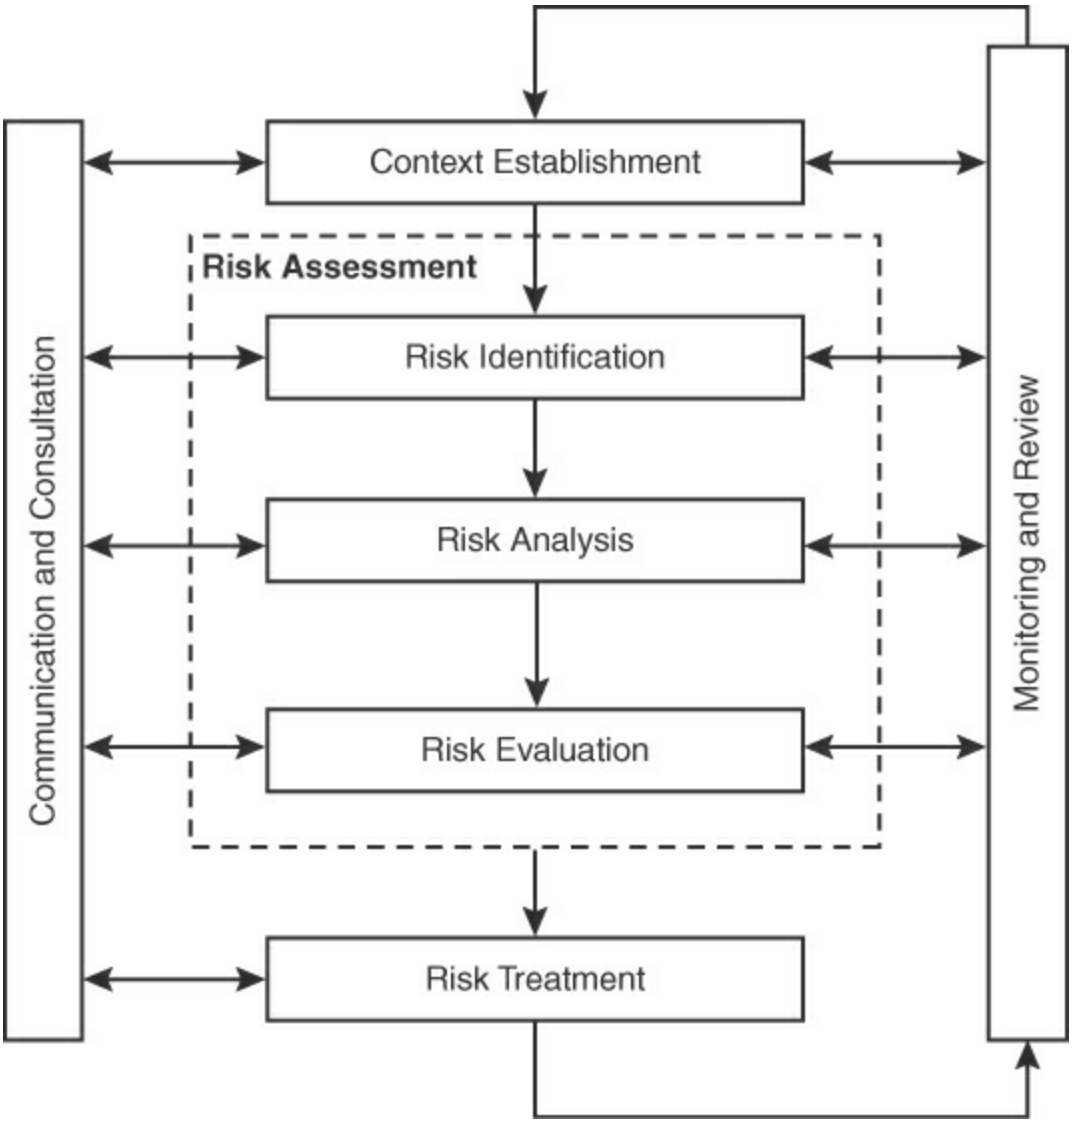

FIGURE 6-3 ISO/IEC 27005 Information Security Risk Management Process

An example use case from the ISO/IEC 27005 publication is section 7.2 under “Context establishment,” which focuses on a suggested information security risk management approach. ISO recommends using the appropriate risk management approach that addresses basic criteria such as risk evaluation criteria, impact criteria, and risk acceptance criteria. Additionally, ISO recommends that resources should be able to perform risk assessments as well as establish risk treatment plans, define and implement policies and procedures, monitor controls, and monitor the information security risk management process. Section 7.2 remains high level, making it an ideal resource for developing a risk management policy. Other sections of ISO/IEC 27005 provide more details that can help organizations develop supporting procedures for new or existing policies; however, procedures should always contain specific details that relate to your organization.

There are many other information security topics covered by ISO guidelines. For example, ISO/IEC 27034 provides guidance on information security to those specifying, designing, and programming or procuring, implementing, and using application systems. An application system normally consists of a user interface, business logic, and a database of some sort. ISO/IEC 27002 is extremely popular and an internationally recognized standard code of practice for information security controls. ISO offers hundreds of papers covering various IT topics, which you can search through at https://www.iso.org.

Note

Another great resource to learn more about ISO security guidelines is https://www.iso27001security.com.

CIS Controls

The Center for Internet Security (CIS) provides best practice solutions for cyber defense and builds and leads industry groups to enable an environment of trust within the cyber community. One CIS resource that is extremely popular is the CIS Controls, security best practices that provide specific and actionable ways to reduce the risk of exploitation. The CIS Controls are

derived from the most common attack patterns highlighted in leading threat reports and vetted across a broad community of industry security professionals. Members associated with vetting the CIS Controls include the NSA Red and Blue Teams, U.S. Department of Energy Nuclear Energy Labs, and various law enforcement agencies, which all give creditability to this resource.

The CIS Controls are broken down into the following three categories, depicted in Figure 6-4 (based on CIS Controls Version 7.1, the current version at the time of writing):

- **Basic CIS Controls:** Anything basic should be considered security table stakes, meaning all organizations should have some form of the recommendations in the basic category applied to their processes. Example basic recommendation categories include identifying what's on the network and dealing with vulnerabilities.
- **Foundational CIS Controls:** Defenses against threat categories such as malware defense or protecting an organization's DMZ or network boundary.
- **Organizational CIS Controls:** Controls that impact the entire organization, such as incident response.

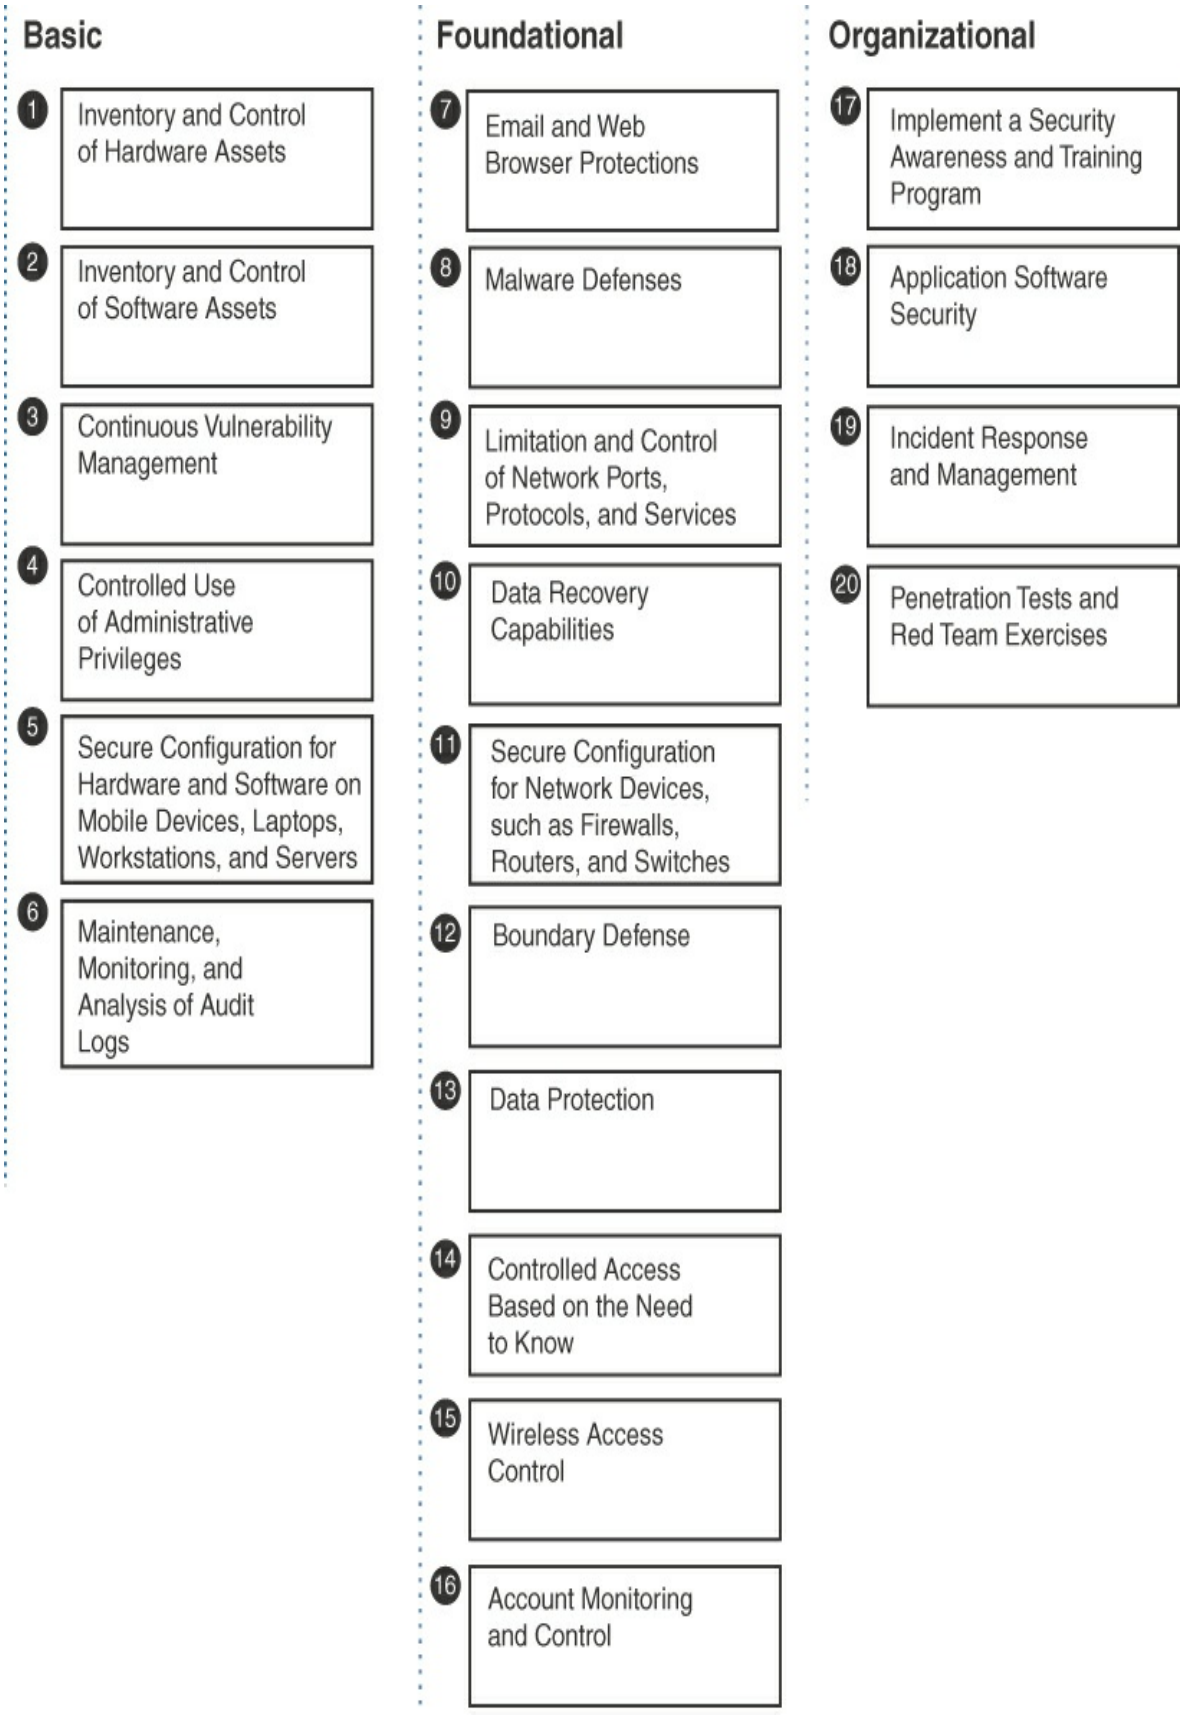
Basic
1
Inventory and Control of Hardware Assets
2
Inventory and Control of Software Assets
3
Continuous Vulnerability Management
4
Controlled Use of Administrative Privileges
5
Secure Configuration for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
6
Maintenance, Monitoring, and Analysis of Audit Logs
Foundational
7
Email and Web Browser Protections
8
Malware Defenses
9
Limitation and Control of Network Ports, Protocols, and Services
10
Data Recovery Capabilities
11
Secure Configuration for Network Devices, such as Firewalls, Routers, and Switches
12
Boundary Defense
13
Data Protection
14
Controlled Access Based on the Need to Know
15
Wireless Access Control
16
Account Monitoring and Control
Organizational
17
Implement a Security Awareness and Training Program
18
Application Software Security
19
Incident Response and Management
20
Penetration Tests and Red Team Exercises

FIGURE 6-4 CIS Controls Version 7.1

CIS takes into consideration the type of organization that would use its recommendations based on the organization's size and expected funding. For example, a small restaurant would likely not be able to invest in an expensive vulnerability management tool to evaluate a few endpoints used within its business. CIS uses a three-tier grouping system to identify which type of organization should consider the recommendations being offered. Figure 6-5 shows how CIS defines its system for grouping the different types of organizations expected to use the CIS Controls resource.

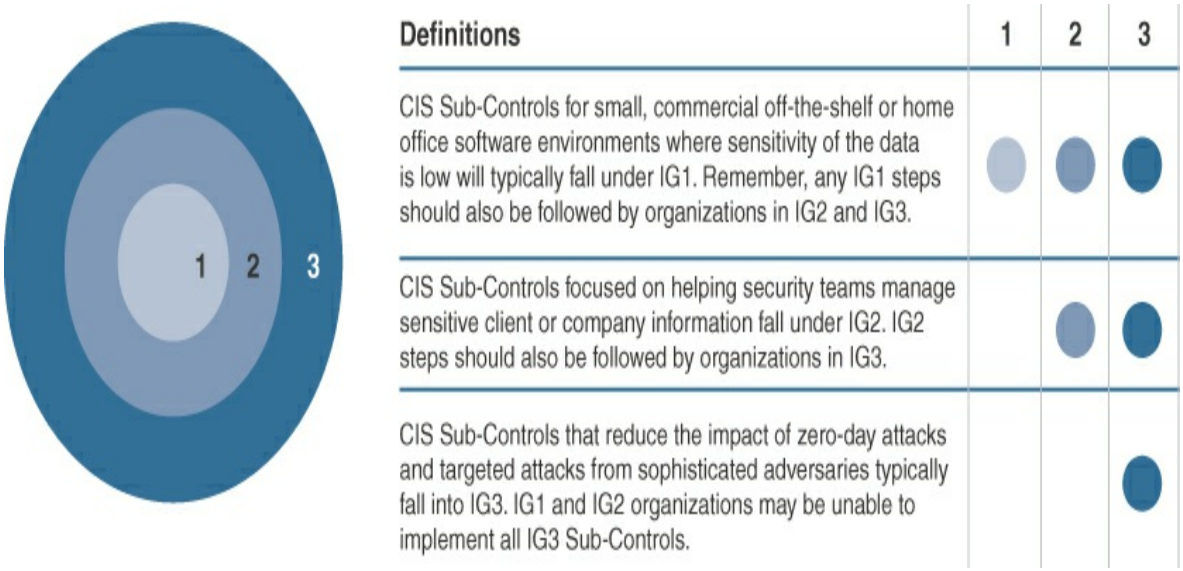

FIGURE 6-5 CIS Controls Organization Grouping

A sample use of the CIS Controls is referencing CIS Sub-Control 2.2, which states "Ensure Software is Supported by Vendor." This recommendation is classified as an "identify" security function and applies to all three organization groups, which means any size organization should follow this recommendation. The explanation section of this Sub-Control further defines this recommendation as only using authorized software according to the vendor and listing any unauthorized software as unsupported and a potential risk.

Any size organization can benefit from referencing the CIS Controls. They are short, tactical, and well-vetted by industry experts.

Note

You can download the CIS Controls at https://learn.cisecurity.org/cis-controls-download.

ISACA COBIT 2019

ISACA is an independent, nonprofit, global association that develops guidelines for information system best practices. One of its most popular offerings is the COBIT 2019 Business Framework for Governance and Management of Enterprise IT. COBIT 2019 offers detailed enabler guides that focus on governance and management best practices. COBIT 2019 also has professional guides targeting implementation, information security best practices, guidance on assurance, risk management, and other security topics. COBIT 2019 is available only to members, for a small fee. Learn more about COBIT 2019 and other ISACA offerings at https://www.isaca.org.

FIRST CSIRT Services Framework

The Forum of Incident Response and Security Teams (FIRST) is the global forum of incident response and security teams. FIRST offers various standards, publications, and best practice guides that are designed to provide best practices for various cybersecurity concepts. One popular offering from FIRST is the Computer Security Incident Response Team (CSIRT) Services Framework. The CSIRT Services Framework is a high-level document describing a collection of cybersecurity services and associated functions other security teams can learn from. The CSIRT Services Framework is essentially a guideline for best practices for incident response.

The CSIRT Services Framework Version 2.1 is composed of four elements: service areas, services, functions, and sub-functions: The five service areas are Information Security Event Management, Information Security Incident Management, Vulnerability Management, Situational Awareness, and Knowledge Transfer. Each service area has services, which in turn have

different functions and sub-functions. For example, the Information Security Incident Management service area has a service named Artifact and forensic evidence analysis, which contains four functions: media or surface analysis, reverse engineering, runtime and/or dynamic analysis, and comparative analysis. Each service, function, and sub-function has a section describing what it is, a section describing its purpose, and a section describing its outcome (measurable results of implementing it). Each function also has a list of sub-functions, if applicable. For example, function 6.3.2, reverse engineering, has four sub-functions, which are static analysis, code reverse engineering, potential behavior analysis and description, and potential signature design.

The CSIRT Services Framework Version 2.1 can provide value by enabling you to review your policies and procedures against what is listed as best practice for every service area and its associated topics. You can download the latest version of the framework from https://www.first.org/standards/frameworks/.

Note

FIRST also has a great framework resource for product incident response teams (PSIRT); see https://www.first.org/standards/frameworks/.

Exceeding Compliance

It is common practice to use any of the standards, guidelines, and frameworks previously discussed (and others) as reference points for developing and testing your people, process, and technology and as templates for building policies and procedures. Using resources such as a tabletop exercise or leveraging standards, guidelines, and frameworks is great but might not provide enough details to help you generate specific areas of improvement or best understand your specific needs.

As stated at the beginning of the chapter, it is important to understand that

compliance is something you either meet or do not meet. It is also important to understand that compliance does not consider all aspects of security, meaning any risk that falls outside of compliance is not taken into account in an audit for compliance. This is why you must establish a baseline for security that exceeds what is required for meeting compliance. To do this, you first need to meet compliance using a formal testing process known as an audit. Once your organization passes an audit, you identify where it may still have risk outside of compliance by using assessments and penetration testing. Next, let's look at how to meet compliance by using audits.

Audits

Audits are the process to validate that policies are properly enforced. Policies will include everything from regulations to what a organization's leadership deems is required. Audits alone are not designed to evaluate the security posture of the organization. Audits focus only on what is in scope for meeting compliance and are designed to be tactical rather than generic. An example would be auditing for compliance with a specific policy, such as a policy that prohibits unauthorized devices from being attached to the network. The audit would test only for unauthorized devices and not consider any other vulnerabilities. Keeping the audit focused simplifies execution and expected results, which leads to a quick benefit from the service. The downside of keeping the audit scope limited and focused would be overlooking other vulnerabilities. This can lead to a false sense of security even though the criteria for the audit were met.

Other services such as assessments and penetration testing should be used when evaluating the overall security of an organization. Reserve audits for more tactical evaluations such as meeting compliance for a specific policy. The format used for an audit depends on what is being assessed.

Audit Example

You should leverage any document that lists requirements to meet what is being audited, such as the targeted policy, and evaluate if capabilities are met.

In some cases, it might make sense to use a pass/fail approach, while other cases might require more detailed information. The following is an example template that could be used for auditing a firewall.

Firewall Audit

1. The organization should have a firewall or equivalent deployed to protect its internal network and devices against unauthorized access. [**Full / Partial / None**]
2. The password on the firewall should be changed from the default setting to an alternative strong password. [**Changed / Not Changed**]
3. The firewall password should meet the following:
   - 8 characters or longer [**Pass / Fail**]
   - Not the same as the username [**Pass / Fail**]
   - Does not contain any identical characters next to each other [**Pass / Fail**]
   - Is not made of up of only a dictionary word [**Pass / Fail**]
   - Includes upper- and lowercase letters and at least three numbers or special characters [**Pass / Fail**]
   - Has not been reused within a predetermined time period [**Pass / Fail**]
   - Does not contain the manufacture's name [**Pass / Fail**]
4. Each rule set on the firewall must be approved by an authorized individual and documented, including an explanation of the business need for the rule. [**Pass / Fail**]
5. Unapproved or vulnerable services should be blocked at the gateway firewall. [**Pass / Fail**]
6. Any permissive firewall rules that are no longer required should be disabled. [**Pass / Fail**]
7. All exceptions to firewall rules must have a documented expiration

date. All exceptions that exceed the expiration date must be removed. [**Pass / Fail**]

8. The firewall's administration settings should not be accessible from the Internet. [**Pass / Fail**]

The previous audit example shows a checklist approach to auditing a firewall. Each question provides answer options for the auditor to select to determine the appropriate response. Another approach to formatting the responses for this document would be to provide an open response section for each question, allowing the auditor to provide a much more detailed response about his or her findings, in addition to choosing from a predefined selection of choices. Another approach for this document would be to use a scoring system, where each question is worth a value, and then the auditor would add up all the values to generate a score. Any of these approaches work as long as the results successfully establish a repeatable evaluation of the people, process, and technology being audited.

Internal Audits

It is common practice to use audits to validate if specific policies or recommended standards, guidelines, or frameworks are being met within the organization. Audits are also used to validate if required industry regulations such as PCI DSS or HIPAA are being met, as described later in this chapter. The team responsible for performing an audit could be part of the SOC, part of a dedicated compliance team, or a partnership between one or more groups. Many organizations have a dedicated group that performs audits and is led by the chief compliance officer (CCO), who is responsible for overseeing compliance within an organization as well as ensuring compliance with laws, regulatory requirements, policies, and procedures. The CCO is considered the subject matter expert and has the responsibility of establishing standards and implementing procedures to ensure compliance programs are effective in identifying, detecting, and correcting noncompliance with laws and regulations as they apply to the organization. Because the CCO sits on the board, he or she must provide assurance to other senior management that

compliance is met as well as develop strategies and responses for any compliance-related complications. This is accomplished by using a combination of audit services and other assessment tools.

> **Note**
>
> It is highly recommended to seek the CCO's support for the SOC program and to incorporate the CCO's objectives within the SOC goals. CCO support will gain SOC visibility at the executive level as well as allow the SOC to leverage budget associated with the organization's compliance goals.

External Auditors

Most required industry compliance will have an external team that represents either the government or industry service that mandates the compliance and is responsible for validating if your organization falls within compliance. External auditors might charge a fee and require periodic auditing in order for your organization to be authorized to use a service or certain type of data. You can recruit your own external auditors to validate your current capabilities prior to the scheduled external audit or use internal services for audit preparation. For example, if your organization accepts credit cards as a form of payment, then your organization must comply with PCI DSS. There are companies that can assist with helping you prepare for the official PCI DSS audit and it is highly recommended that you use only external auditors that are approved by the PCI Security Standards Council to ensure a quality service when auditing for PCI DSS. The same concept would apply regarding seeking auditing services for any other type of compliance. You want to ensure the service provider is an authorized auditor for the subject you want to have evaluated.

When you consider external audit services, the following are some general questions you should ask when evaluating a potential audit service provider:

- How does your service validate that it meets the latest version of the

topic(s) being audited?

- What guarantees does your service provide if a topic considered compliant is found by another auditor to be noncompliant?
- What are the scope and costs for your services?
- What technology and steps are included within your services?
- What resources and access permissions are required for your service to be performed?
- What risk is associated with using your service?
- How long should the organization wait until performing another audit?
- Do you provide recommendations or remediation services for items that fail the audit?

Audit Tools

Another option you can use to audit for various types of compliance is to leverage cloud or on-premises auditing tools. Auditing tools can provide a template of questions your internal auditors can use to perform auditing services or provide some scanning capacities to gather and assess targets for compliance. An example is using LogicGate's risk management tools that offer compliance audit capabilities, including assessing for PCI DSS. Some network and security tools might also offer a compliance widget or report to help assist with gathering data for compliance; for example, a firewall might include compliance reports for NIST standards. You can also ask your trusted technology vendors if they offer mappings to popular standards, guidelines, and frameworks. Figure 6-6 shows an example of mapping the Cisco security products to the NIST Cybersecurity Framework. As with auditing services, recommended practice dictates that you validate the tool as an acceptable source for auditing the topic of interest to avoid having inaccurate results. Some vendors might claim to meet audit requirements but would actually only partially meet or outright fail if audited by an authorized service provider.

FIGURE 6-6 Mapping of Cisco Security Products to the NIST Cybersecurity Framework

Audits focus only on specific subtopics found within the topic being evaluated. You need to consider evaluating the entire organization for vulnerabilities regardless of whether the potential vulnerability is related to your organization's compliance requirements. This is how you exceed compliance: addressing risk beyond what is required to be compliant. The process you can use to evaluate for all vulnerabilities and risk is to perform an assessment and penetration test. First, let's review the concept of assessments.

Assessments

An assessment goes beyond what is evaluated by an audit by considering any vulnerability found within the scope of what is being assessed. For example, an audit against a policy for identifying unauthorized systems connected to the network would only validate what is connected to the network. An assessment of devices connected to the network would further evaluate the risk associated with each device such as if the device has potential vulnerabilities, what risk each device poses to the organization, and other vulnerabilities in the people, process, and technology associated with devices connecting to the network.

Note

Assessments are security-focused while audits are compliance-focused. Best practice is to use both to exceed compliance by ensuring you meet compliance as well as address risk outside of what is required to be compliant.

Assessment Types

Different types of assessments can be performed. A *threat assessment* looks at anything that could contribute to a disruption of normal services and operations. The following is a sample format for a threat assessment report.

Threat Assessment

Threat	Caused by Human/Nature	Strength of Threat	Motivation Factor
Disgruntled employee	Human	Medium	High

A *vulnerability assessment* focuses on identifying, quantifying, and rating weaknesses or gaps within your systems. The following is an example format for a vulnerability assessment report.

Vulnerability Assessment

Asset	Vulnerability	Severity	Exposure	Rating
Website	Incomplete SQL code	High	High	5

A *risk assessment* is focused on measuring the probability of a security breach occurring and the magnitude of the risk. This approach can provide results in a qualitative view (high, medium, low) or a quantitative view such as on a scale of 1 to 5. The following is an example format for a risk assessment report.

Risk Assessment

Risk	Likelihood of Occurrence	Existing Controls	Proposed Mitigation Measures
Virus attack	High	Antivirus	Improve Internet usage policy. Install anti-malware that is based on behavior and anomaly detection capabilities.

An *impact assessment* identifies how and the extent to which your business will be affected by a security breach. The following is an example format for an impact assessment report.

Impact Assessment

Risk	Business Impact	Customer Impact	Financial Impact	Regulatory Impact
Virus attack	High	Medium	High	Low

It is recommended to use all four assessment methods to evaluate your people, process, and technology. Assessment reports can break up data into sections that focus on these four assessment types or include different columns that cover results from each assessment method. The previous examples could be combined into one large report or be part of a spreadsheet that covers all related topics.

Note

The U.S. Federal Risk and Authorization Management Program (FedRAMP) offers a great resource for developing a security assessment report at https://www.fedramp.gov/developing-a-securityassessment-report/.

It is common for the scope of an assessment to be much broader than an audit and for assessments to be performed more frequently than audits. Assessments can include various types of people, process, and technology, and performing assessments is commonly one of the services offered by a SOC. Chapter 3, "SOC Services," covered different types of SOC services, including performing assessments.

Assessment Results

The results of an assessment are used to develop a vulnerability report that includes a list of potential vulnerabilities and details about each. A commonly included detail is the estimated risk associated with the potential vulnerability, to help the organization prioritize which risk to address first.

The most common method used in vulnerability reports to help organizations rank the associated risk with an identified vulnerability is the Common Vulnerability Scoring System (CVSS) from FIRST. Essentially, a CVSS score provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. The CVSS is widely accepted in the security industry, and popular security assessment solution providers such as Rapid7, Core Security, and Tenable offer the same CVSS scoring references within their products. You can calculate the associated risk of a vulnerability manually by referencing the CVSS; however, most security solutions will do the calculations for you and rank the vulnerabilities with a scoring or color-coding system. Table 6-1 provides a sample vulnerability report featuring the results from assessing a website, USB devices, and a server. Chapter 9, "Vulnerability Management," covers the CVSS in much more detail.

Note

You can learn more about how a CVSS score is generated at https://www.first.org/cvss/specificationdocument.

TABLE 6-1 Example Vulnerability Report

Asset	Vulnerability	Severity	Exposure	Rating
Website	Incomplete SQL codes left by a freelance web designer leading to SQL injection attacks	HIGH	HIGH	5
USB devices	Default operating system configuration allows all programs to run automatically	HIGH	MEDIUM	4
Server	Missing patch permits unauthenticated command prompt	HIGH	MEDIUM	4

Assessment Template

Different assessments will vary regarding the focus of what the outcome report will contain. A vulnerability assessment report will focus on the different vulnerabilities found and their associated severity, such as shown in Table 6-1. A risk assessment report can include vulnerability data, but the report will be more focused on general risk, including quantitative and qualitative computations. A threat assessment report can contain both risk and vulnerability data, but the focus will be on specific types of threats and how they could potentially impact the organization. An impact assessment report will be focused on your organization's assets and business and explain what to expect if certain types of events were to occur. There are many specific details that could be included depending on the desired outcome for the assessment service. My recommendation is to reference industry guidelines regarding what should be included in the type of assessment report you are looking to perform. An example of a guideline is the MITRE Threat Assessment and Remediation Analysis (TARA), found at https://www.mitre.org/sites/default/files/pdf/pr-11-4982-tara-methodology-description-version-1.pdf, which is a great reference point for a threat-focused assessment.

In addition to the specifics to the type of assessment being delivered, a final assessment report should include general topics, including an executive summary, purpose, list of technologies used, list of who was involved, and a summary of what was found. The next section is an example template for a risk-focused assessment. Many of the sections are common across all assessment reports regardless of the focus. Use this template as a reference for developing your assessment reports. I have left each section blank regarding example input and instead included descriptions about how to fill out each section.

Executive Summary

[Briefly summarize the scope and results of the risk assessment. Highlight high-risk findings and comment on required management actions.]

Detailed Assessment

1. Introduction

[Provide an overview of the project and what was done. This will prepare the reader for understanding everything to follow. This is not a complete summary, such as what was done in the executive summary.]

1.1. Purpose

[Describe the purpose of the risk assessment in context of the organization's overall security program.]

1.2. Scope of this risk assessment

[Describe the scope of the risk assessment, including system components, elements, users, field site locations (if any), and any other details about the system to be considered in the assessment.]

2. Risk Assessment Approach

[Describe the steps that were taken during the assessment.]

2.1. Participants

Role	Participant
System Owner	
System Custodian	
Security Administrator	
Database Administrator	
Network Manager	
Risk Assessment Team	

2.2. Techniques Used

Technique	Description
[List techniques used; e.g., questionnaires, tools.]	[Describe the technique used and how it assisted in performing the risk assessment.]

2.3. Risk Model

[Describe the risk model used in performing the risk assessment. For an example risk model, refer NIST SP 800-30 Rev 1.]

3. System Characterization

[Describe the technology that was assessed.]

3.1. Technology Components

Component	Description
Applications	[Describe key technology components, including commercial software.]
Databases	
Operating systems	
Networks	
Interconnections	
Protocols	

3.2. Physical Location(s)

Location	Description
[Include locations included in scope.]	

3.3. Data Used By System

Data	Description
[Detail data elements included in scope]	[Describe characteristics of data elements]

3.4. Users

Users	Description
[Detail categories of users.]	[Describe how users access the system and their intended use of the system.]

3.5. Flow Diagram

[Provide connectivity diagram or system input and output flowchart to delineate the scope of this risk assessment effort.]

4. Vulnerability Statement

[Compile and list potential vulnerabilities applicable to the system assessed.]

Vulnerability	Description
[List vulnerabilities.]	[Describe vulnerability and its impact.]

5. Threat Statement

[Compile and list the potential threat sources applicable to the system assessed.]

Threat Source	Threat Actions
[List threat sources.]	[List and/or describe actions that can be taken by threat source; e.g., identity theft, spoofing, system intrusion.]

6. Risk Assessment Results

[List the observations (vulnerability/threat source pairs). Each observation should include

- Observation number and brief description of observation (e.g., Observation 1: User system passwords can be guessed or cracked)
- Discussion of the threat source and vulnerability pair
- Identification of existing mitigating security controls
- Likelihood discussion and evaluation (e.g., High, Medium, or Low likelihood)
- Impact analysis discussion and evaluation (e.g., High, Medium, or Low impact)
- Risk rating based on the risk-level matrix (e.g., High, Medium, or Low risk level)
- Recommended controls or alternative options for reducing the risk]

Item Number	Observation	Threat-Source/Vulnerability	Existing Controls	Likelihood	Impact	Risk Rating	Recommended Controls

Vulnerability Scanners

One of the most common tools used for assessing vulnerabilities is a vulnerability scanner. Vulnerabilities will be included in all assessment report types since vulnerabilities represent potential attack vectors. You will need to include vulnerability scanning as part of your assessment process. Vulnerability scanners can be automated to perform periodic or on-demand assessments, depending on how they are used. The quality of the results will depend on many factors, including the data seen by the scanner and how the scanner evaluates a target. For example, a credential-based vulnerability scanner would be able to log into an endpoint device and assess the device without its results being impacted by the endpoint device's enabled security tools such as a host firewall. A non-credential-based scanner would not be able to log into the endpoint, limiting its results to what can be seen and not blocked by enabled security tools such as the host firewall. Chapter 9 covers vulnerability management best practices in more details, including how to use vulnerability scanners to address risk and how to understand the data the scanners produce. Use the steps from Chapter 9 to leverage a vulnerability scanner as part of your assessment process.

Note

Best practice dictates using both credentialed and non-credentialed scanning when assessing devices. Credentialed assessments provide more accurate results, but non-credentialed assessments are closer to what a potential attacker would have access to. Both approaches provide their own form of value and should be leveraged equally.

Assessment Program Weaknesses

One important concept regarding assessments is that they lack validation of the results. They show “potential” risk rather than actual validated risk. The results being considered potential risk can occur based on how the assessment is performed. Many times, the tools used for an assessment capture data and compare it against a list of known vulnerabilities rather than actually attempting to exploit the vulnerability. This approach can lead to false negatives due to other elements that are not evaluated but are essential to truly understanding the nature of the potential vulnerability. For example, a vulnerability scan may find that a server has many vulnerabilities; however, the server might exist within a contained lab environment and not be a real threat to the organization, or the vulnerabilities that were found might be based on limited data captured by the assessment tool due to the server’s existing enabled security that limits what can be evaluated by the assessment tool. The best way to test any potential vulnerability to ensure that it is really a risk is to attempt to exploit it. Exploiting potential vulnerabilities is part of a penetration testing service.

Note

Assessments and penetration testing are services that help your organization consider risk that extends beyond what is addressed by meeting compliance requirements.

Penetration Test

A penetration test, also referred to as a pentest, goes beyond an assessment by performing the same steps that an adversary would perform to exploit an identified vulnerability. It is common for an assessment to report many more potential vulnerabilities than a penetration test of the same scope actually discovers due to how both approaches validate a target. By following the steps used by an attacker, a pentest truly qualifies whether a target is vulnerable to a real-world attack and, if so, at what level. A pentest typically has a higher cost and more risk associated with its services than an assessment, but a pentest tends to produce more accurate results. I

specifically use the language “tends to produce more accurate results” because the actual results received will depend on the scoping and expected goals. I have seen organizations invest thousands of dollars in penetration testing services that result in very little actionable outcomes. My goal in this section is to address some of those pitfalls that lead to poor penetration testing results.

Note

You should first perform an assessment and apply risk reduction steps before performing a penetration test, which helps to maximize the value of the penetration test.

NIST Special Publication 800-115

NIST SP 800-115, introduced in Chapter 3, identifies four phases of a penetration test: planning, discovery, attack, and reporting. The following is a summary of those phases.

Note

NIST SP 800-115, *Technical Guide to Information Security Testing and Assessment*, is available at https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-115.pdf. I recommend at least perusing the table of contents to discover the full scope of what it has to offer for conducting penetration tests and assessments.

Phase 1: Planning

The planning phase involves establishing the rules and the scope of the penetration test and obtaining necessary approval for the penetration test

before executing it. It is highly recommended to obtain approval from the highest level of leadership to ensure that all steps are authorized, in the event that something goes wrong. I have seen situations where a manager approved a penetration test that executive leadership was unaware of and the pentesting caused system failures. Both the manager and service provider were blamed for the system outages even though the penetration testers performed services within the agreed-upon scope of work. In the industry, having this approval is known as the "get out of jail free card." Make sure your get out of jail free card is signed by a C-level executive!

Phase 2: Discovery

The discovery phase is made up of two parts. First, all network ports and services that fall within scope are identified to develop a list of potential targets. Other techniques might be used, such as dumpster diving, with the goal of gathering information and building a potential target list. These data collection steps fall under the category *target reconnaissance* and are the most time consuming but also the most valuable work. The more quality data that is discovered and collected about a target, the better the capability to identify potential targets to exploit. Some example tools you can use to conduct target reconnaissance are NMAP and masscan to perform port scanning; Shodan to research web-facing resources such as websites, web cameras, and network-enabled IoT devices; and searching social media sources such as LinkedIn and Facebook for details on people that work within the target network.

The second part of the discovery phase is analyzing the potential targets for vulnerabilities. Vulnerability analysis is commonly performed by comparing the services, applications, and operating systems for potential targets against a vulnerability database or using a tester's own knowledge of vulnerabilities. Tools such as Metasploit from Rapid7 provide a database of known exploits, where you simply type in an identified vulnerability and find various weaponized exploits tied to the weakness. Figure 6-7 shows an example of searching for Apache Struts vulnerabilities within Metasploit. Notice how the number of available exploits expands beyond what can be displayed on the screen!

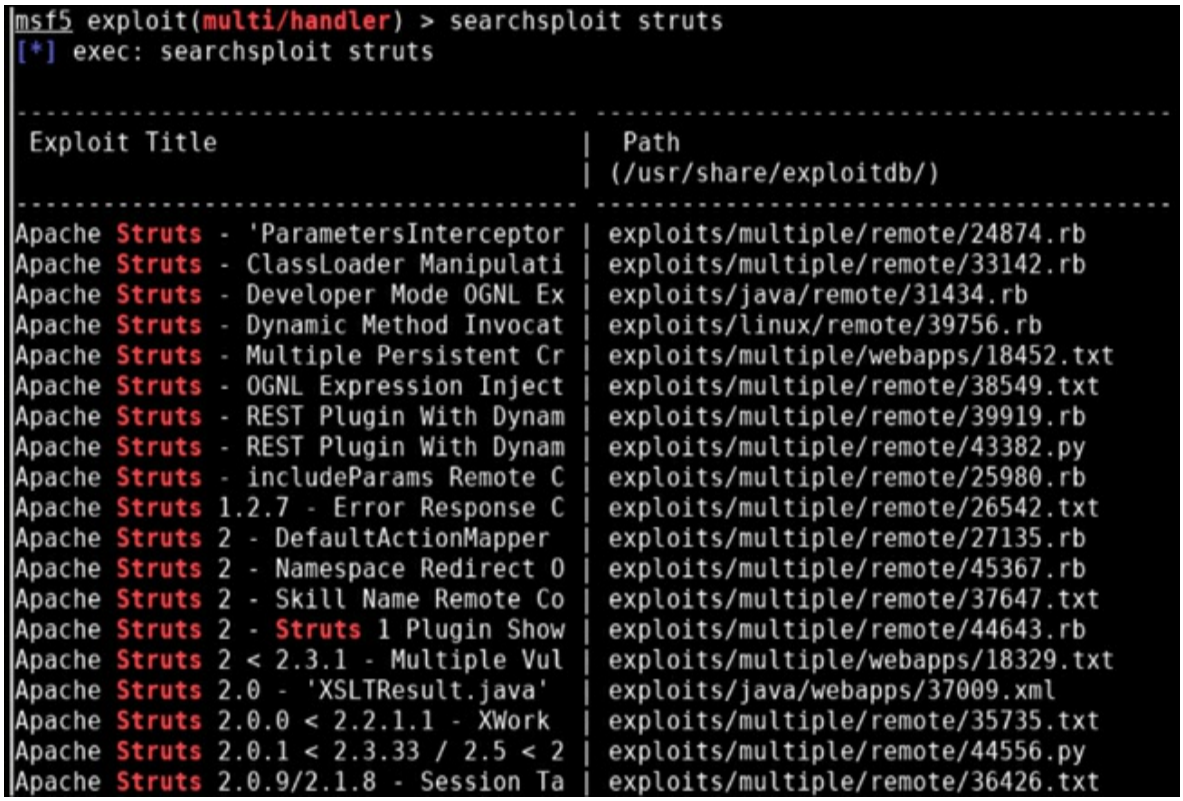

FIGURE 6-7 Various Apache Struts Exploits in Metasploit

Phase 3: Attack

Once targets are confirmed, the actual exploitation occurs during the attack phase. Remember that exploitation does not necessarily mean something bad has already occurred. Exploitation can open the path for malicious actions such as pushing a remote-access tool (RAT), ransomware, or other forms of malware to the system. An example of exploitation could be running an Armitage/Metasploit weaponized attack against a system that is vulnerable to an exploitation of Struts, as shown in Figure 6-8. The example shows a system has an open shell, meaning the attacker has full access to the compromised system. From this point, the attacker can perform a number of actions, including using this compromised system as a gateway into the network to perform a new attack against internal systems. In the example, the attacker showcases his or her level of access by running the **whoami** command.

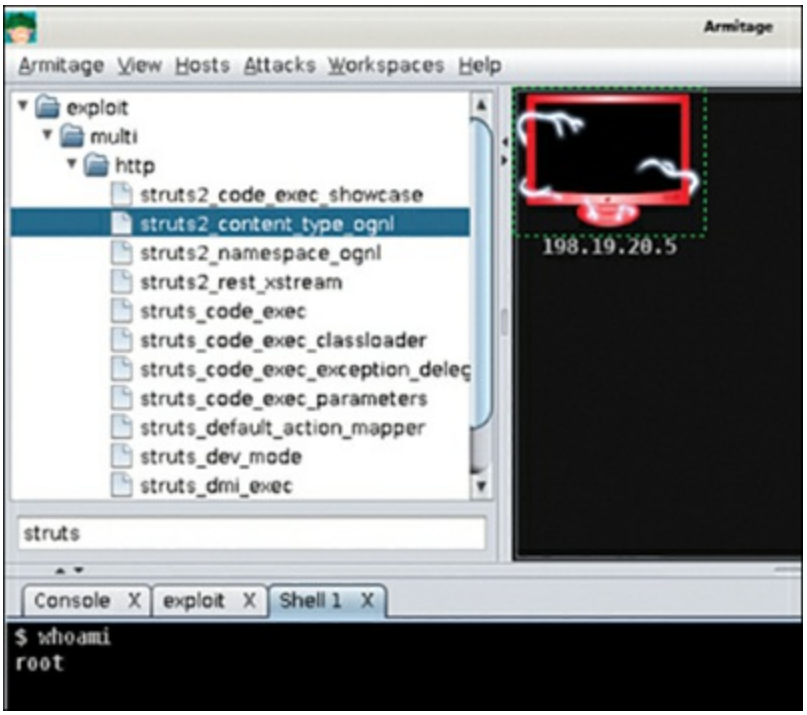

FIGURE 6-8 Using Armitage to Exploit a Struts Vulnerability

In other cases, the action of exploitation causes negative results, such as exploitation of a website to take it offline, known as a denial-of-service attack. Another example could be overloading the memory of a switch to cause it to fall back into a hub-like state, allowing an attacker to view all traffic within that device. From a penetration testing viewpoint, using the Cyber Kill Chain is a good way to view the steps that would be involved with exploiting a target and delivering something, which are the same steps penetration testers and real attackers would use against a target. As a refresher from Chapter 1, Figure 6-9 depicts the Cyber Kill Chain. The attack phase of the NIST penetration testing process corresponds to the Delivery, Exploitation, Installation, and Command and Control phases of the Cyber Kill Chain.

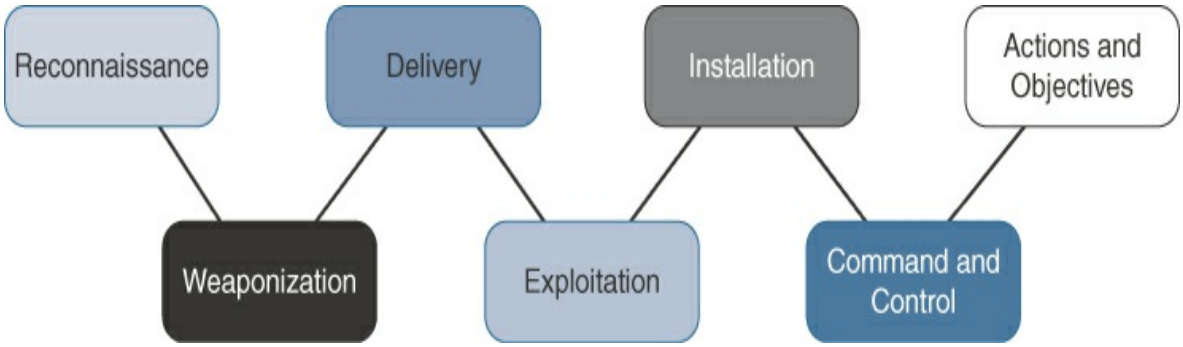

FIGURE 6-9 Cyber Kill Chain Attack Lifecycle

Note

Kali Linux and BackBox are great free penetration operating system builds you can download and use for penetration testing efforts. Metasploit is a great resource you can test your attacks against in a free and legal manner.

Phase 4: Reporting

The results of the attack phase are documented in the final phase, the reporting phase. Reporting should also include steps from the other phases so the entire penetration testing process is documented. The language used in a report is extremely important, meaning you must know your target audience and consider the impact that will result based on who reviews the report. I have seen people lose their jobs based on negative feedback about their work or how they were blamed for the results of the penetration testing exercise. I recommend at a minimum that your penetration testing report contain the following information:

- **Vulnerabilities:** What was found during the penetration test? List all vulnerabilities along with details about what they are.
- **Impact:** What is the potential negative impact that could occur if a real attacker were to take advantage of a vulnerability found during the penetration testing exercise?
- **Likelihood:** How hard is it to exploit an identified vulnerability? Do you

need root access or connectivity to the internal network, or can anybody from anywhere access and exploit the vulnerability?

- **Risk evaluation:** How do the identified vulnerabilities impact the overall business?
- **Recommendation:** What is recommended to reduce the risk of the identified vulnerabilities?
- **References:** Who was involved in the penetration testing exercise?
- **Additional details:** Make sure to include any appendices, glossary items, tools used, and other details that would help the reader follow the penetration test report.

Note

SANS offers great resources for templates for penetration testing reports, at https://pen-testing.sans.org/resources/downloads.

Figure 6-10 shows a diagram of the NIST SP 800-115 four-stage penetration testing methodology, which is a high-level view of performing a penetration testing exercise.

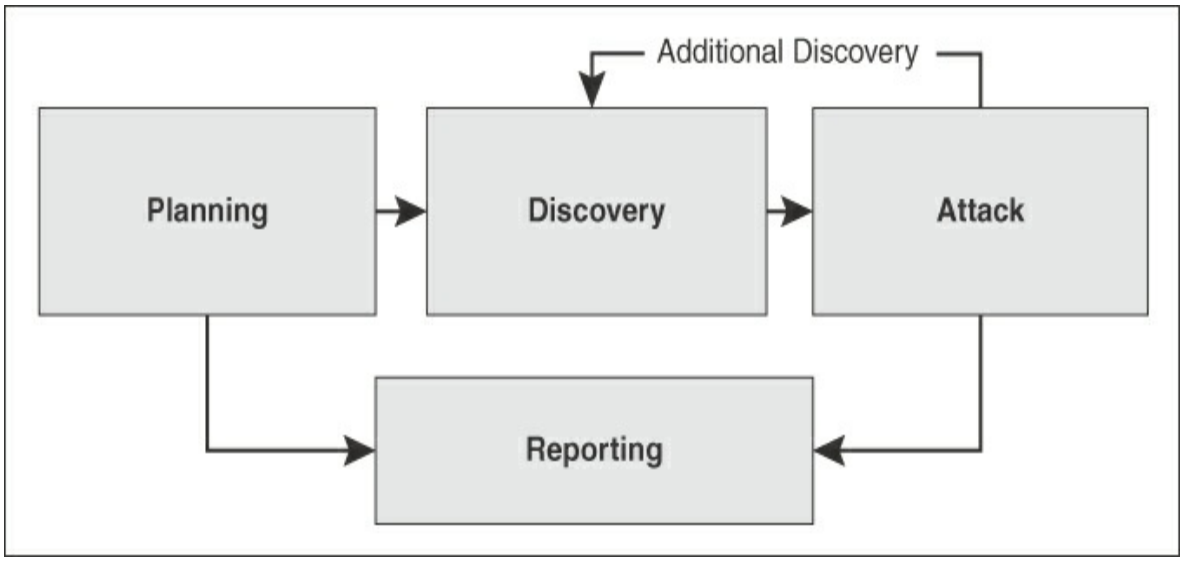

FIGURE 6-10 NIST 800-115 Four-Stage Penetration Testing Methodology

Additional NIST SP 800-115 Guidance

NIST SP 800-115 provides much more detail as you read further into the document. For example, the attack stage of the NIST four-stage penetration testing methodology is further broken down into four steps, as shown in Figure 6-11. The first step is to gain access to the target, also known as establishing a foothold. It is possible that the current privilege level will be limited based on how the access was accomplished, so the next step would be to attempt to escalate privileges to a higher level, such as administrator-level access. Once the highest level of system access is achieved, the third step is to browse the system to discover mechanisms to gain access to other systems. The final step represents the results of the exploit, which might be installing a backdoor, removing data (or documenting the potential to do so), or repeating the process against other internal systems. This cycle is repeated until enough results are generated that can be recorded during the reporting stage of the NIST penetration testing methodology.

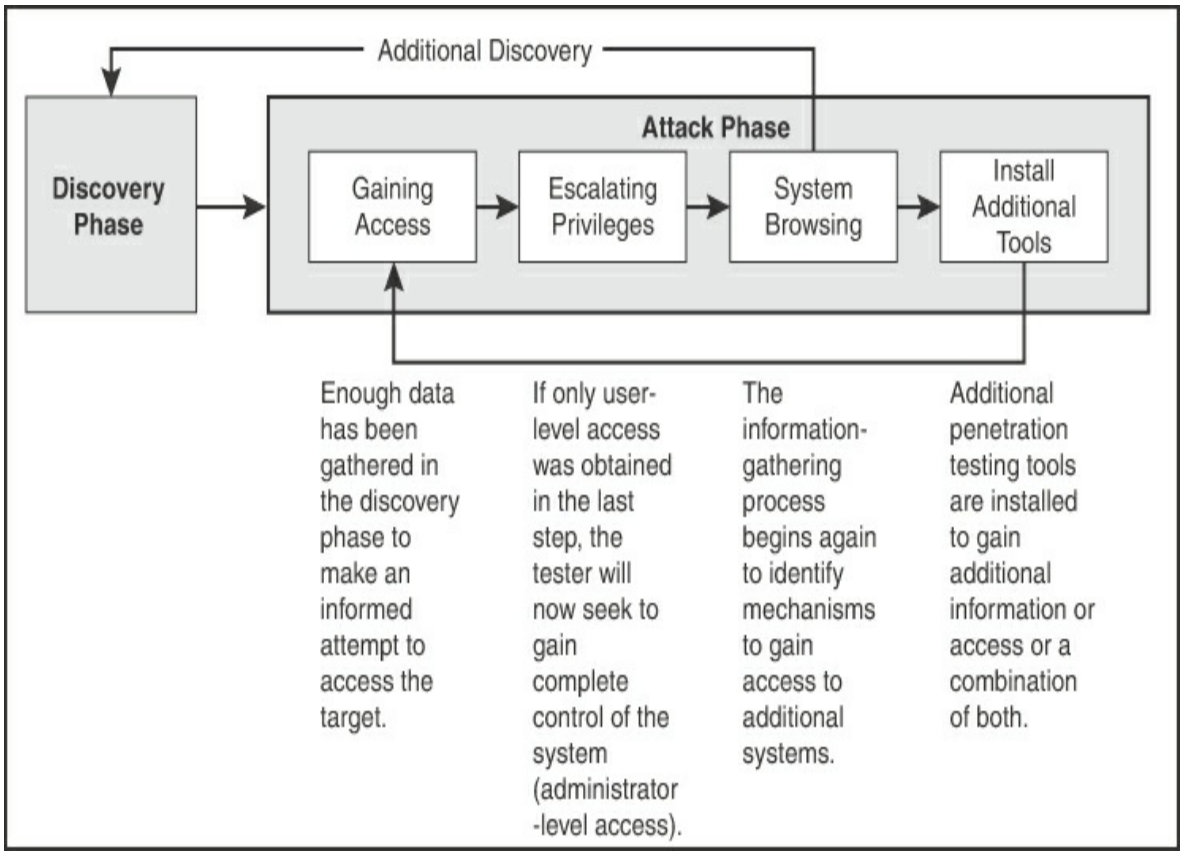

FIGURE 6-11 NIST 800-115 Attack Phase Details for Penetration Testing

Note

There are many other penetration testing resources available, including popular certification programs that train and test on core penetration testing skills, such as the Certified Ethical Hacker (CEH) certification program offered by EC-Council and the Offensive Security Certified Professional (OSCP) certification program.

Penetration Testing Types

Different types of penetration testing services are available and are differentiated based on the agreed-upon scope of the testing. It is common

practice to scope a penetration test in a known environment, unknown environment, or partially known environment (prehistorically called white, black, or gray format). The following are definitions of each approach:

- **Known Enviroment:** All information about the target is provided to the penetration tester. This includes its IP address, what software is running, who accesses the system, and other relevant information. A white-scoped penetration test is commonly performed by system owners responsible for the security of their systems.
- **Unknown Enviroment:** No information about the target is provided to the penetration tester. It is up to the penetration tester to research anything within scope and discover any potential target as well as attack the target using any method within scope of the service.
- **Partially Known Enviroment:** Some information about the target is provided to the penetration tester while other details are not. It is common to provide details that an attacker would find if they performed reconnaissance of a target, to reduce the time and cost of the penetration test. It is close to impossible to prevent adversaries from gathering intelligence about your organization using certain sources for their reconnaissance efforts, so some data exposure should be expected.

I recommend leveraging gray penetration testing services whenever using an external paid resource. You can assume and expect that adversaries are capable of researching your organization. Paying the penetration testing service to research your organization during a black-scoped penetration test can quickly increase the cost of the service while providing limited value.

Penetration Testing Planning

Beyond considering the type of penetration test to conduct, penetration testing planning includes other aspects to consider. The following list provides questions that need to be answered as you plan to perform a penetration test or request external penetration testing services. These questions would also need to be answered within any statement of work for a penetration testing service.

- What is the scope of the penetration test?

- What are the priorities for the penetration test?
- Who is authorized to conduct the penetration test?
- What are the data handling requirements?
- What are the penetration test's logistics?
- Will the penetration test be internal or external?
- How should sensitive data be handled?
- What should occur in the event of an incident?

My recommendation is to develop the requirements for your penetration test and identify potential service providers or develop an internal team to provide the service. To ensure the service will meet your scope, you need to have a scoping meeting, where you meet with the penetration testing team and review the scope of work to ensure all parties understand what needs to be done and what is considered a successful service engagement.

The most important aspect of the service is what is delivered, commonly referred to as the penetration testing report. To ensure that the final penetration testing report contains everything you desire from the service provider, you need to specify everything that needs to be included in your scope of work. The scope of work is essential for identifying what is expected from the penetration testing team.

Penetration Testing Scope Template

Much like an assessment report, a penetration testing report should be based on the focus of the service. A report for a web application–focused penetration test will look different from a report for a social engineering–focused penetration test. There are industry guidelines for the deliverable report for the various forms of penetration testing services. Offensive Security (creators of Kali Linux) provides a good example of a penetration testing report at https://www.offensive-security.com/reports/sample-penetration-testing-report.pdf.

Earlier in this chapter, I pointed out that the value of a penetration test is based on the quality of the scope of work. It is absolutely critical to ensure the scope matches the desired outcome or you will not receive what you

expect from the services. For example, I once delivered a penetration test in which I got access to the target's network using social engineering, meaning I attacked the people's trust and tricked them into giving me sensitive information. The organization, however, was expecting a technical penetration test, meaning they were only concerned about threats to their Internet-facing technology. The penetration testing scope of work did not specify the organization's desired outcome, which led to my service meeting the scope but the results not satisfying the customer.

To help you avoid receiving an undesired outcome from a penetration test, I have provided a template for a scope of work. Use this template as a method to summarize the questions that I provided in the "Penetration Testing Planning" section of this chapter; also pay attention to the additional information I have included in the penetration scope template. The more details you include regarding what you expect from the service, the better results you should receive from the service, assuming the penetration testing team meets what you request in your scope statement.

Penetration Testing Scope Statement

[Provide an overview of the purpose for the service request and expected outcome.]

Penetration Test Pre-Planning

[Provide details about what needs to be done prior to the penetration test. Who needs to be notified? What technologies need to be available? What other actions need to be taken prior to the start of the penetration test?]

Team Location(s)	Organization Location(s)	Client Personnel Aware of Testing	Resources Provided to Pentest Team	Pentest Technologies Used
[Response]	[Response]	[Response]	[Response]	[Response]
[Response]	[Response]	[Response]	[Response]	[Response]

High-Level Work Schedule: Project Scope

Description of Work/Pentest Boundaries	Assumptions and Constraints
What is tested? What are the social engineering test boundaries? What is acceptable? What are the boundaries of physical security tests? What are the restrictions on invasive pentest attacks? What type of corporate policies affect your test? [Response]	[Response]

Milestones	Due Dates
[Response]	[Response]

ID	Activity	Resource	Labor		Material		Total Cost
			Hours	Rate	Total Units	Cost	Total

Appropriate Authorization (Including Third-Party Authorization)

[List anybody who is authorized to perform the penetration test. Anybody outside of these people are considered not authorized and would be representing a real threat to the organization.]

Name	Title/Organization	Description of Authorization and Consent (Identify Reference Documents)
[Response]	[Response]	[Response]
[Response]	[Response]	[Response]
[Response]	[Response]	[Response]

Reconnaissance Pentest Activities

[Provide any requirements for reconnaissance activities. For gray or white penetration testing, certain reconnaissance action items can be removed, and the results can be provided to the penetration testing team. This will reduce the cost of the services based on the assumption that certain reconnaissance details are publicly available to real threats.]

Pentest Scanning Activities

[Provide any requirements for evaluating targets.]

Scanning Test Deliverable Name	Scanning Test Deliverable Description
[Response]	[Response]
[Response]	[Response]

Gaining Access Activities

[Provide what is permitted once access to a target is possible.]

Gaining Access Activity Name	Gaining Access Activity Description
[Response]	[Response]
[Response]	[Response]

Maintaining Access Activities

[Provide what is permitted to maintain access on a target, commonly called establishing a foothold.]

Maintaining Access Activity Name	Maintaining Access Activity Description
[Response]	[Response]
[Response]	[Response]

Covering Tracks Activities

[Provide what is within scope for removing the footprint caused during the previous exploitation steps.]

Covering Tracks Activity Name	Covering Tracks Activity Description
[Response]	[Response]
[Response]	[Response]

Pentest Analysis and Report Planning

[Provide an explanation regarding what outcome and details need to be included in the penetration test final report. Include whether you require recommendations for risk reduction, screenshots of the tasks performed, or any other details regarding what will enable your organization to use the report to improve its security.]

Describe Plan for Analyzing and Reporting Pentest Results
[Response]

All compliance concepts covered in this chapter thus far have been based on policy and procedures, also known as an organization's *process*. Processes are developed at the discretion of the organization and made into a required rule or policy that must be followed based on the details in any supporting procedure documentation. There are certain compliance requirements that are mandatory, regardless of whether the organization has a policy in place to enforce it. Not meeting certain types of required industry compliance can result in fines and other negative outcomes. The final topic for this chapter is reviewing required industry compliance.

Industry Compliance

Most industries are required to comply with specific laws, standards, and/or regulations mandated by legislatures, government agencies, or industry regulators. Violations of government requirements are considered breaking the law, with repercussions ranging from fines to jail time. The reason behind government-enforced laws is to protect those that are associated with what is being protected, which is a good thing. For example, HIPAA protects

people’s privacy, particularly their health information. If a company in the health care industry fails to protect its customers’ privacy, HIPAA requirements can be enforced by the government with strong repercussions if they are not followed, forcing the proper security controls to be put in place.

Industry services also have similar goals for enforcing requirements. Industry regulators don’t have the power to assess legal penalties; however, they have other methods to encourage following their policies, including reduction of allowed services as well as requirements for passing audits before services are reinstated. For example, most people have credit cards, which if stolen would mean access to their protected data. When people use credit cards to purchase something, they want to have confidence that their data will not be used inappropriately. If a company fails to protect a customer’s credit card data, it’s the buyer that is exposed to risk because their financial and other personal data is what is impacted.

Laws such as HIPAA as well as industry service requirements such as PCI DSS force any affected organization to enforce security best practices to protect the people that could be impacted by a breach, which in the example of PCI DSS are the customers using their credit cards with a vendor, and with HIPAA is the customers’ privacy. Without these requirements, organizations would use less secure methods to protect credit card and privacy data, causing an increase in data theft and a decrease in customer trust of the credit card network and health care services that contain private data.

Compliance Requirements

There are a handful of government-enforced laws and industry requirements that your organization will need to be aware of depending on your line of business. It is also good to know when other organizations must meet an industry compliance requirement as well. Looking back at the PCI DSS example, I once received my credit card receipt at a restaurant and saw that my entire credit card number was printed on it, a violation of PCI DSS. I could’ve reported the violation to my credit card company, but I chose to inform the manager about the issue. According to PCI DSS, the restaurant should have printed only the last four digits of my credit card number on the receipt, sufficient to enable me to confirm the credit card number charged is

the one I used; however, the remaining data is protected. I recommend including considerations of who you do business with based on if they meet required industry compliance.

> **Note**
>
> It is common practice for organizations to advertise that they meet various types of industry compliance, with the goal of gaining the confidence of potential customers. For example, an organization might advertise on its website that it follows secured supply chain or software development best practices when creating products.

The following are some of the most commonly required industry compliance standards and laws that your organization, or organizations you do business with, may be required to follow:

- **PCI:** The PCI Data Security Standards help protect the safety of that data. They set the operational and technical requirements for organizations accepting or processing payment transactions, and for software developers and manufacturers of applications and devices used in those transactions.

> **Note**
>
> Learn more about PCI at https://www.pcisecuritystandards.org/pci_security/maintaining_payment_security.

- **HIPAA:** The Health Insurance Portability and Accountability Act of 1996 developed regulations for protecting the privacy and security of health information. This rule applies to health plans, health care clearinghouses, and any health care provider who transmits health information in electronic form, including business associates. The HIPAA Security Rule is broken down into safeguard categories. Technical safeguards evaluate technology that is in contact with HIPAA-protected data. HIPAA also

references other guidelines within its technical safeguards, such as what can be found within various NIST standards for what it considers acceptable encryption standards. Physical safeguards target risk to physically accessing HIPAA-protected data and systems with HIPAA data. Topics include inventory of hardware and physical security concepts. Administrative safeguards address policies and procedures and bring privacy rules and security rules together. For example, conducting risk assessments allows an organization to address vulnerabilities in both physical and technical systems. Other administrative topics include developing policies and procedures according to industry-recommended best practices.

The U.S. Department of Health and Human Services offers various checklists at HHS.gov that you can use to audit your organization for HIPAA compliance. It is up to your organization to validate and meet HIPAA compliance requirements to avoid fines and other negative outcomes.

Note

You can download HHS's summary of the HIPAA Privacy Rule at https://www.hhs.gov/sites/default/files/ocr/privacy/hipaa/understanding/summary/privacysummary.pdf.

- **SOX:** Sarbanes-Oxley Act of 2002, better known as SOX, focuses on corporate fraud by requiring all publicly held companies to enact internal checks and balances and procedures for financial reporting. CEOs and CFOs are directly responsible for the accuracy, documentation, and submission of all financial reports and can be penalized for failure to comply regardless of whether it is intentional or not. SOX also has data security requirements, including use of SOX-compliant software. It is up to your organization to identify if you have SOX obligations and, if so, to meet compliance requirements.
- **FISMA:** The Federal Information Security Modernization Act of 2014 protects U.S. government data and compliance with it is mandatory for all

U.S. federal agencies as well as contractors that wish to do business with a U.S. federal agency. To comply with FISMA, organizations must follow FIPS 199, FIPS 200, and the NIST 800 series guidelines. The top areas of focus that must be met are as follows: maintain an information system inventory; categorization of systems based on risk; develop and maintain a security plan; adhere to NIST SP 800-53 security controls; conduct risk assessments according to NIST SP 800-30; and have annual certification and accreditation of systems. Penalties for failing compliance can include censure by Congress, a reduction in federal funding, and reputational damage. If you work for the U.S. government or plan to do business with a U.S. government entity, you must follow FISMA compliance requirements.

Note

Learn more about FISMA at https://www.dhs.gov/cisa/federal-information-security-modernization-act.

- **FedRAMP:** The Federal Risk and Authorization Management Program is another U.S. government-wide program and compliance with it is required for any government agency leveraging cloud technology. FedRAMP provides steps for performing security assessments, authorization, and continuous monitoring of cloud products and services. FedRAMP, like FISMA, uses the NIST SP 800-53 security controls for its blueprint. Cloud services must have their security assessed and be granted an Agency Authority to Operate (ATO) or Provisional Authority to Operate (P-ATO) in order to be accepted by a government agency, unless a special exception is granted. If you work for a U.S. government agency or plan to do business with a government agency using cloud-based technology, you need to follow FedRAMP requirements.

Note

You can learn more about FedRAMP at https://www.fedramp.gov/faqs/.

- **Data sovereignty laws:** Data sovereignty is a country-specific requirement mandating that data is subject to the laws of the country in which it is collected or processed as well as must remain within its borders. Many countries have had data sovereignty laws for decades; however, new privacy laws are making data sovereignty requirements more public. Countries such as Russia, China, Germany, France, Indonesia, and Vietnam require that their citizens' data must be stored on physical services within the country's borders, meaning data can't exist in a cloud service outside the physical country. The reason for these laws is that countries are concerned about data being misused when it leaves their borders because, technically, they are no longer legally able to protect such data. An example where data sovereignty will come into play is cloud services that have data warehouses located around the world. Organizations that must meet their government's data sovereignty requirements will need evidence that their data will be stored only within data warehouses residing within their country or they will not be permitted to use the cloud service. It is up to the organization to identify and adhere to any data sovereignty requirements, which will vary from country to country.
- **Industry-specific compliance requirements:** Your organization might be required to comply with one or more requirements specific to its industry. As an example of a compliance requirement specific to the energy industry, I'll briefly mention NERC CIP. The North American Electric Reliability Corporation (NERC) has been in charge of maintaining the operations and functions of bulk power systems, more commonly called the electric grid, since the early 1960s. In 2008, NERC developed the Critical Infrastructure Protection (CIP) compliance framework, made up of 11 control families, to mitigate cybersecurity attacks on the electric grid. There are useful general recommendations as well as those that are very specific to the security of bulk power systems. The framework can be found at https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx.

Note

I included NERC CIP as an example of a very specific industry compliance requirement. Industries such as manufacturing, education, and government also have industry-specific certifications. PCI DSS is a specific service requirement, meaning you only need to follow PCI DSS if you are involved with credit card services. I recommend consulting with a compliance specialist if you are unaware of all compliance requirements your organization is obligated to follow.

Summary

Meeting compliance is essential to any successful cybersecurity program. This chapter started by describing how to develop and maintain policies and procedures, which are the ingredients of an organization's process. Next, it covered how to evaluate existing policies and procedures through the use of a tabletop exercise. You then learned how standards, guidelines, and frameworks are resources that your organization can use to build or tune your policies and procedures; although following such resources is not required, they often represent industry best practices. Next, this chapter looked at how to meet compliance using audits as well as how to exceed compliance by focusing on all risk using assessments and penetration testing. Finally, this chapter surveyed several common legal and industry-based compliance requirements that might apply to your organization, depending on the nature of its operations.

Remember that compliance can only be met or not met in perspective of the entity enforcing the compliance; however, meeting compliance does not mean you are secure. Many organizations will continue to operate while they are not compliant as long as they have a plan in place to eventually become compliant, which not only increases the risk of compromise but also demonstrates a lack of focus for meeting what needs to be considered the bare minimal effort to be secure. Exceeding compliance means to not only meet compliance, but also to evaluate all risk, with a focus on actually securing the organization. Make sure to include both meeting compliance and security goals as part of how you establish your security baseline. This will

result in your organization not only exceeding compliance, but also reducing the risk of a future compromise within your people, process, and technology.

Chapter 7 dives into an extremely useful resource for securing your organization against attacks, known as threat intelligence.

References

Archieve360 Team. (2019, February 14). Data Sovereignty and the GDPR; Do You Know Where Your Data Is? Archive360. https://www.archive360.com/blog/data-sovereignty-and-the-gdpr-do-you-know-where-your-data-is

California Association of Health facilities. (n.d.). Discussion Based Tabletop Exercise: Design Template & Documentation. CAHF. http://www.ndltca.org/image/cache/Tabletop_Exercise_Template.pdf

CIS. (2018, October 18). Tabletop Exercises: Six Scenarios to Help Prepare Your Cybersecurity Team. CIS. https://www.cisecurity.org/wp-content/uploads/2018/10/Six-tabletop-exercises-FINAL.pdf

FIRST.org, Inc. (2019, November). Computer Security Incident Response Team (CSIRT) Services Framework Version 2.1.0. FIRST.org, Inc. https://www.first.org/standards/frameworks/csirts/FIRST_CSIRT_Services_Framework_v2.1.0.pdf

Gitanjali, M. (2018, May 2). A Security Assessment Template for Small Businesses: Evaluate Your IT Security. GetApp. https://lab.getapp.com/security-assessment-template-for-small-businesses/

LaBoissonnière, L. (2018, February 2). Be Diligent: 5 Practical Steps to Enforceable Workplace Policies. McInnes Cooper. https://www.mcinnescooper.com/publications/be-diligent-5-practical-steps-to-enforceable-workplace-policies/

National Institute of Standards and Technology. (2018, April 16) Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. NIST. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf

North American Electric Reliability Corporation. Critical Infrastructure

Protection (CIP) Standards. NERC. https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx

Oltsik, J. (2018, January 11). Research Suggests Cybersecurity Skills Shortage Is Getting Worse. CSO. https://www.csoonline.com/article/3247708/security/research-suggests-cybersecurity-skills-shortage-is-getting-worse.html

SANS. (2020). Security Policy Templates. SANS. https://www.sans.org/information-security-policy/

Vlallno, B. (2014, October 27). 6 Tips for Effective Security Tabletop Testing. CSO. https://www.csoonline.com/article/2838365/planning-for-a-security-emergency-from-the-tabletop-down.html

Chapter 7

Threat Intelligence

An intelligent hell would be better than a stupid paradise.

—Victor Hugo

Knowledge is one of the most powerful currencies. Some might even say it's the ultimate currency. How you react to a situation depends on how much knowledge you have of the situation. Imagine the difference between having knowledge that the stock market is going to crash before it happens versus having no knowledge until the event is well under way. If you're heavily invested in stocks, that piece of information could make the difference between securing your wealth and ending up in bankruptcy. A similar comparison can be applied to the world of cybersecurity. Having knowledge of a threat allows you to prepare a response instead of reacting to the impact of the threat, which at that point is too late. This is why threat intelligence has become and will continue to be a critical component of a successful security operations center.

This chapter dives into the world of threat intelligence. First, you will learn about general threat data and how it can be converted into threat intelligence. You will then learn about the different types of threat intelligence and how each type can be used within your SOC. This chapter organizes the different types of threat intelligence into four categories:

- Strategic threat intelligence
- Tactical threat intelligence
- Operational threat intelligence
- Technical threat intelligence

Intelligence is more than the data you collect. Intelligence is what you do

with that data or what awareness it provides. That is why the term *actionable intelligence* is used often when describing cybersecurity threat intelligence data. Actionable intelligence provides guidance for choosing the best actions. You will learn how to leverage external resources for intelligence, including how to evaluate the return on investment. You will find that all threat intelligence data isn't going to benefit your SOC and, in some cases, adding the wrong data can overwhelm your SOC analysts and negatively impact your security tools.

Let's start this chapter with an overview of what threat intelligence is and how to plan for a threat intelligence project.

Note

Every SOC practice eventually incorporates threat intelligence as it works toward maturing its services. I am a firm believer that threat intelligence is a necessary ingredient for improving how a SOC makes decisions and, if properly used, can dramatically improve how and when threats are detected.

Threat Intelligence Overview

Imagine having the knowledge that allows you to stop a cyberattack before it occurs. Some people might equate this to a pitch for some sort of fortune-telling scam, but actually it's a very realistic scenario if you have the right data. What exactly is the "right data"? In many cases, it is information about the mannerisms of the attacker, the network traffic and patterns created by the attacker, what capabilities the attacker has, and which common vulnerabilities the attacker exploits on a system. These types of information about the attack, the attacker, and the motives of the attacker provide security context. A simple definition of what threat intelligence is, is *context*. Context could include hash values to match files against, behavior patterns to look for, threat actors to be aware of, and many other items that will help you make informed decisions about your response to an attack. Gartner has a

slightly different definition for threat intelligence, which is the following:

> Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice, about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject's response to that menace or hazard.

Gartner summarizes threat intelligence as evidence-based knowledge, which means you can rely on it to make informed decisions about how to respond to a threat. This means if I give you a bunch of IP addresses with no context such as a warning list posted on a website, you won't understand what they mean—those IP addresses are just data. If I tell you that these IP addresses are bad but I do not explain why, you have only one specific use of the IP addresses, which is to block them. That would represent threat data, which can be a form of threat intelligence depending on how it is used, but by itself is not threat intelligence. Many people have a misconception about threat data, so let's address that up front before we move into the topic of threat intelligence.

Threat Data

All threat-related data is not threat intelligence. Many security tools are driven by a specific type of threat data, but that doesn't mean the data provided to the tool gives you, the user, any value. You can have the best packet parser and analyzer on the planet, but if you can't tell that tool what to look for, then that tool has very little use to your SOC. Many security vendors refer to data delivered to their tools as "threat intelligence," which may be true if you understand what that data is and how it applies to your goals. If you have no idea what is fed to a vendor's tool, then essentially the vendor's tool is receiving threat data, which has nothing to do with threat intelligence.

Threat Data Example

Here is a nontechnical example of comparing threat data and threat

intelligence. I have two kids. The younger is a one-year-old boy and the other is an older girl. The boy will smash anything he can get into contact with. If my daughter wants to build a LEGO castle, there is the threat that her brother will smash it. If I tell my daughter her castle can be smashed, that information doesn't inform her about the specific threat (her brother). It just points out there is a risk that her castle can be smashed. With this knowledge, she can't make any new decisions outside of maybe increasing the strength of her castle. But I didn't tell her about the threat, which is that her brother will destroy any castle she builds, regardless of how much effort she puts into reinforcing it. If I tell my daughter, on the other hand, that she needs to build her castle outside of her brother's reach or he will see it and come smash it, she not only understands the threat, she understands how to reduce the risk of her castle being destroyed. She understands what the threat is (her brother). She understands how the threat can identify her castle (he can see it). She understands that reinforcing the castle is not the proper risk reduction strategy. Instead, she understands her best next step is to move the castle out of the view of her brother, which is a response to the risk that her brother will see her castle and destroy it. This is an example of threat intelligence because she can understand the threat and take action based on what she learned from me informing her about the threat (her brother).

Threat Data Value

Threat data can be extremely useful for enhancing security tools. In the past, the native capabilities of security tools enabled them to thwart most tactics used by malicious parties; today, however, the threat landscape is rapidly evolving and expanding beyond any tool's local defense capabilities. Effective security tools now must include a way to *adapt to change*, which is based on continuous learning in the form of threat data. This is how tools such as IPSs, antivirus, and sandboxes stay relevant. The security vendors continuously update their tools with new detection capabilities in the form of signatures and behavior analysis to adjust to the results of threat research. This is why threat data is extremely relevant in the security vendor space.

Note

Some security vendors claim their tools use machine learning and artificial intelligence to stop attacks without signatures. However, those machine learning and artificial intelligence training sets still need to be constantly updated and trained to understand new logic and capabilities implemented by attackers as they update their techniques. Without these adjustments, the existing capabilities will quickly become dated and of little value. Chapter 5, "Centralizing Data," covered machine learning in more detail.

To be clear regarding threat data, getting updates to existing security products is not necessarily what the industry refers to as threat intelligence. Security tool updates only apply to the tool's capability of being updated, omitting many other details about the threat vector, which a technical threat intelligence feed would contain. This specific difference of what is omitted from the consumer's view separates a vendor provided threat data fed to a security tool versus you feeding a tool an external threat intelligence feed. When you provide the threat intelligence feed using threat intelligence, you understand what it contains and can control how the intelligence impacts the function of the tool. For example, a basic antivirus checksum update would have hash matches for only the latest threats used to pattern match against files. Any update to the antivirus product would not have any other context about the threat, such as who is a target, where the threat is originating from, and other details that could be used by your SOC to better understand, detect, and respond to the threat. Remember, a SOC has many roles, including nontechnical responsibilities that don't focus on details such as hash values, and those roles need the context associated with threat data in order to use it to make a decision outside of blocking what is matched. There are details that can be pulled from threat intelligence, such as the associated risk with a threat that has nothing to do with technical threat data feeds used by security products to block high risk resources. This missing data from many threat data feeds sent to security tools is a very useful type of threat intelligence that can help SOCs make decisions about how to respond to a situation. Threat intelligence is about using all of the context associated with data.

Threat Data Limitations

An important limiting factor that influences the quality of the intelligence provided by a vendor of a network or security product is that the *vendor sells its product to multiple customers*. Each customer's network is different, so how could a vendor provide a list that encompasses threats targeting small businesses as well as threats targeting enterprise organizations? Keep in mind that security tools can't just check for everything—that would mean billions of possible threats. A security vendor must consider all of its customers and develop a threat data feed that is generic enough to be useful to everyone, yet still able to prevent many of the common threats. Granted, many vendors have specialties and attempt to work with different industries and business segments, but they still must take the "accommodate all types of businesses" approach. This can lead to a threat data feed that not only doesn't cover all the threats that will impact *your* organization but also fills your tools with additional data that has nothing to do with your line of business. For example, a vendor's IPS default security feed will likely provide a list of vulnerabilities that could be exploited; however, your organization might not own the products associated with the vulnerability. Therefore, the IPS is wasting resources looking for an attack that could never happen on your systems. The same concept applies to the IPS not receiving signatures regarding tools specific to your organization that are not part of the default signature update. As a result, some resources would be exposed to attack until the IPS is manually adjusted to protect those devices. In both cases, the vendor's threat data updates leave gaps in your security capabilities and can also lead to a false sense of security.

The point of understanding the limitations of vendor threat data updates is to recognize the gap in threat data that needs to be filled. This is why it is critical to understand what threat intelligence is and how it can fill this gap found within many security tools. Filling this gap allows organizations and their tools to make faster and more informed decisions about security. *To be crystal clear, enabling vendor updates doesn't mean you are obtaining threat intelligence.* You are just receiving the vendor's generic threat data that all other customers are receiving. Also, it is important to point out that *threat intelligence isn't always technical details about a threat* such as a bad domain, IP address, or hash of a malicious artifact. Malicious domains, file hashes, and IP addresses are normally what the industry refers to as *indicators of compromise (IOCs)*. IOCs are considered threat intelligence only when combined with context. I find these are some of the most common

misunderstandings about threat intelligence in the industry.

Threat Intelligence Categories

Now that you understand what threat data is and how it can be confused with threat intelligence, it's time to move to the topic of threat intelligence. Referring again to the Gartner definition of threat intelligence, it is *evidence-based knowledge*, which means you are able to take action based on the conclusion you make with the evidence that is provided from the data. There are many forms of data that can provide threat intelligence to your SOC. Maybe it's social media. Maybe it's a vendor's threat data feed. Maybe it's another organization warning you about a threat. As mentioned in the chapter introduction, I organize all threat intelligence–related data into four categories, which is the focus of the remainder of this chapter.

Table 7-1 summarizes the four categories of threat intelligence for quick reference.

TABLE 7-1 Summary of Four Categories of Threat Intelligence

Strategic threat intelligence	Nontechnical threat intelligence that is heavily risk-based, used by high-level decision makers
Tactical threat intelligence	Provides details of threat actor tactics, techniques, and procedures (TTPs)
Operational threat intelligence	Reveals actionable information about specific incoming attacks
Technical threat intelligence	Technical details about threat indicators such as malicious IP addresses and hashes of malicious artifacts

Each category contains a different type of intelligence that can help a SOC better understand a current situation based on the context the threat intelligence feed provides. Threat intelligence can help leadership make decisions based on potential threats, but the data's context must be in a format they understand and can benefit from, or it will end up being useless

data. For example, if a CEO needs to decide if the organization should open a new office in another country, telling the CEO about specific technical attack techniques seen by attackers from that country will not help inform or influence the CEO's decision. What the CEO needs is strategic threat intelligence regarding the risks of opening the new office, the likelihood of the risks occurring, and the impact if an event should occur. This is why it is important to understand the type of threat intelligence you have access to and ensure that it is delivered to the appropriate audience so the receiver is able to benefit from its context.

Let's look closer at each threat intelligence category to better understand how the different data types can benefit different parts of your security operations center.

Strategic Threat Intelligence

Strategic threat intelligence views threat data from a high level rather than including technical details such as which threat actor is involved or specific hashes of malware. The purpose of strategic threat intelligence is to help executives make strategic decisions by giving them a broad understanding of threats that could impact their organization. The simplest way to think of this category of intelligence is to imagine a boardroom full of C-level leaders having a conversation about security. The conversation would be about risk to the organization and impact to operations, and would be outcome oriented. The specific details such as which tools are used, how the threat could be executed, and other "in the weeds" type of information would be handled at the operations level.

It is common for strategic threat intelligence to be provided in a report or briefing format (often summarized in the executive summary). The details can come from policy documents created by nation-states, various forms of media, recent publications, specialist activity, white papers, industry guidelines, and research reports. Organizations such as Gartner and Forrester can be contracted to develop on-demand strategic reports or more generic reports can be acquired.

One challenge with strategic threat intelligence is obtaining value. Many

sources are saturated with raw data and sometimes have hidden objectives or biased data. Filtering through strategic threat intelligence can be a manual, time-consuming process if the resources are not leveraged properly. Best practice is to view a strategic threat intelligence request as a project that contains a continuous feedback loop between the requestor and analyst to ensure that more accurate results are obtained. I look deeper into this concept later in this chapter.

Tactical Threat Intelligence

Tactical threat intelligence provides details about tactics, techniques, and procedures (TTPs) used by threat actors. The purpose of this category of intelligence is to better understand how threats will execute their attacks so that defenders can be better prepared to respond. Responses could include improving security tools, identifying gaps in capabilities, and modifying the people and process responsible for responding to attacks.

Tactical threat intelligence is intended to be used by technical roles responsible for an organization's defense. Job roles could include system architects, administrators, and other security staff. Executive roles tend to rely on their technical staff to leverage tactical threat intelligences, while the executives look for a nontechnical summary in the form of a strategic threat intelligence report. Sources of tactical threat intelligence include open-source tools, honeypots, data collectors on dark networks, scanning technology, malware analysts, closed-source networks, and technical experts. Expected details in tactical threat intelligence include potential targets of attack, attack vectors such as phishing or a malware type, and which tools or technical infrastructure are used by the attacker. An example could be data on how ransomware uses different types of vulnerabilities to infect hosts. Knowing these details will allow a SOC to evaluate its capability to defend against ransomware and be better prepared for a future attack. Tactical threat intelligence is not focused on a specific ransomware campaign, which could include a combination of tactics and customization unique to a specific threat actor. Tactical threat intelligence is broader in scope.

Operational Threat Intelligence

While tactical threat intelligence covers details on attack behavior, it isn't focused on a specific attack or campaign. Attackers tend to use a combination of exploits, known as *chained exploitation*, which viewed as a campaign can be fingerprinted, monitored, and sometimes linked to a threat actor. Knowing about a specific campaign allows defenders to track impact and risk associated with that campaign as well as validate if it changes its tactics or is likely to target their organization. As an example of the difference between tactical threat intelligence and operational threat intelligence, tactical threat intelligence of ransomware would look at a general ransomware category, while operational threat intelligence would focus on a current ransomware campaign targeting a specific Apache Struts vulnerability and possibly link it to a threat actor located in a specific part of the world.

Sources for operational threat intelligence include field-level resources that have access to threat campaign data. For example, a group of ethical hackers may have intercepted or compromised a threat group's communication or infrastructure. Examples include resources who have breached a chat room speaking about a threat campaign, identified a darknet forum selling victim information from a threat campaign, or maybe seen an advertisement or communication from a group promoting an attack for political or social reasons, such as hacktivists launching a targeted attack campaign. Social media sources such as Twitter and Facebook could also contain operational threat intelligence. Resources such as the FBI's InfraGard program can alert subscribers to current campaigns that are encountered by organizations that contribute to this threat intelligence program. Many vendor research departments and threat intelligence teams, such as Cisco Talos (https://talosintelligence.com/) and Fortinet's FortiGuard Labs (https://www.fortinet.com/fortiguard/labs), provide updates on blogs and other public media about active malicious campaigns.

Note

An important side note is that there might be questionable or illegal actions being taken to obtain certain forms of operational threat intelligence. You might want to ask the threat intelligence provider

legal and ethical questions regarding how operational threat intelligence is obtained when evaluating such resources.

Technical Threat Intelligence

Technical threat intelligence refers specifically to the threat actor's tools and infrastructure. This form of intelligence is more specific and detailed than tactical threat intelligence and focuses on indicators of compromise, or IOCs. The purpose of using focused technical threat intelligence, commonly fed into security tools, is to provide rapid distribution and response to threats. A simple way to think of technical threat intelligence is that it is any artifact or behavior that indicates a compromise, requiring immediate attention.

Technical threat intelligence can be part of a vendor's updates, the associated limitations of which I previously covered. Technical threat intelligence can also be general feeds that include lists of malware hashes, registration keys, or other file artifacts associated with malware, email characteristics associated with phishing campaigns, malicious URLs or domains, and IP addresses associated with malicious behavior such as known command and control (C&C) infrastructures or exploit kits. If these datapoints provide context that can help your SOC make better decisions, they would be considered threat intelligence; however, if these datapoints are just an unknown feed, such as data classified as bad without any context pumped into a tool, then they would represent threat data.

Technical threat intelligence typically has a shorter lifespan than other types of threat intelligence and is consumed in high volume by various types of security tools to help such tools be aware of recent attack indicators. It is common for a security tool to receive threat data from its vendor; however, the tool's capability is enhanced with a third-party threat intelligence feed more specific to the organization's need. For example, a school might purchase a threat intelligence feed that includes many specific technical details about IOCs targeting schools, which allows the school's SOC to know when the school is being targeted by a certain threat actor as well as helps tune their security tool beyond the generic threat data that the tool receives from its vendor. This data will need to be continuously updated, as IOCs are

constantly changing.

Threat Intelligence Context

Understanding the context associated with threat data is the key value that makes it threat intelligence. Threat intelligence can be obtained both internally and externally. Chapter 5 covered how security tools generate event data, which can be consolidated and converted into an action. I call this data *internal threat intelligence*. Combining the intelligence provided by vendor updates, events you see locally within your network, and threats seen through external threat intelligence feeds provides a balanced system of data allowing your organization to be best prepared for future events. The key to success is how to balance external and internal intelligence, which will vary based on the expected outcome for the use case being addressed. Figure 7-1 represents the concept of properly leveraging internal and external threat intelligence.

FIGURE 7-1 Balance Between Internal and External Threat Intelligence

External threat intelligence feeds come in many flavors. A threat intelligence feed that is free likely gathers data using only open sources. Commercial threat intelligence feeds requiring a fee, in contrast, generally provide more unique data that can be specific to a market segment and gathered from closed sources such as marketplaces in the criminal underground or from

security/data collection tools planted within various types of networks. It is critical to be aware that *just because a threat intelligence feed is commercial doesn't mean it's better than a free source*. In fact, some commercial sources are actually just aggregations of open-source feeds. For example, I have seen paid services that are composed completely of compiled open-source tools. I asked one such service, "What am I paying for if you only use open-source feeds"? I was told I'm paying for the effort to run a collector of open-source tools, giving me one spot to collect from multiple open-source feeds, as well as for the consolidated and "cleaned up" data. This example shows how important it is for you to question what you are paying for. For some organizations, having cleaned-up open-source data is worth paying for because the time and support to build such feeds has a cost, while many other organizations expect a paid service to include unique data and value that could not be obtained from free open-source services. Additionally, there are organizations that provide highly specialized threat intelligence feeds, such as feeds that identify threats from darknets and cybercriminal markets, feeds that specialize in information about zero-day vulnerabilities, and feeds that provide industry-specific context regarding threats against critical infrastructure, the healthcare industry, or some other sector.

Note

Unless you don't have the time to do the curation yourself, don't waste money on aggregated open-source feeds. There are better options available.

Threat Context

There are specific factors about threat context that can be extremely useful. One important concept regarding leveraging the context associated with threat intelligence is *prevalence*. Knowing how prevalent a threat is can help determine if a threat is widespread or an isolated incident. Imagine a situation where you identify malware within your organization. If research shows that particular malware has very low prevalence, you should be concerned that

you are dealing with a targeted attack or something new. If threat intelligence points out that the threat is very widespread, the context associated with the threat can help you learn from how the industry has responded to the threat. Many vendors' threat intelligence feeds provide metrics regarding how widespread an identified threat is to help customers understand the context pulled from threat prevalence. Figure 7-2 is an example of how Sophos uses the concept of threat prevalence in its products.

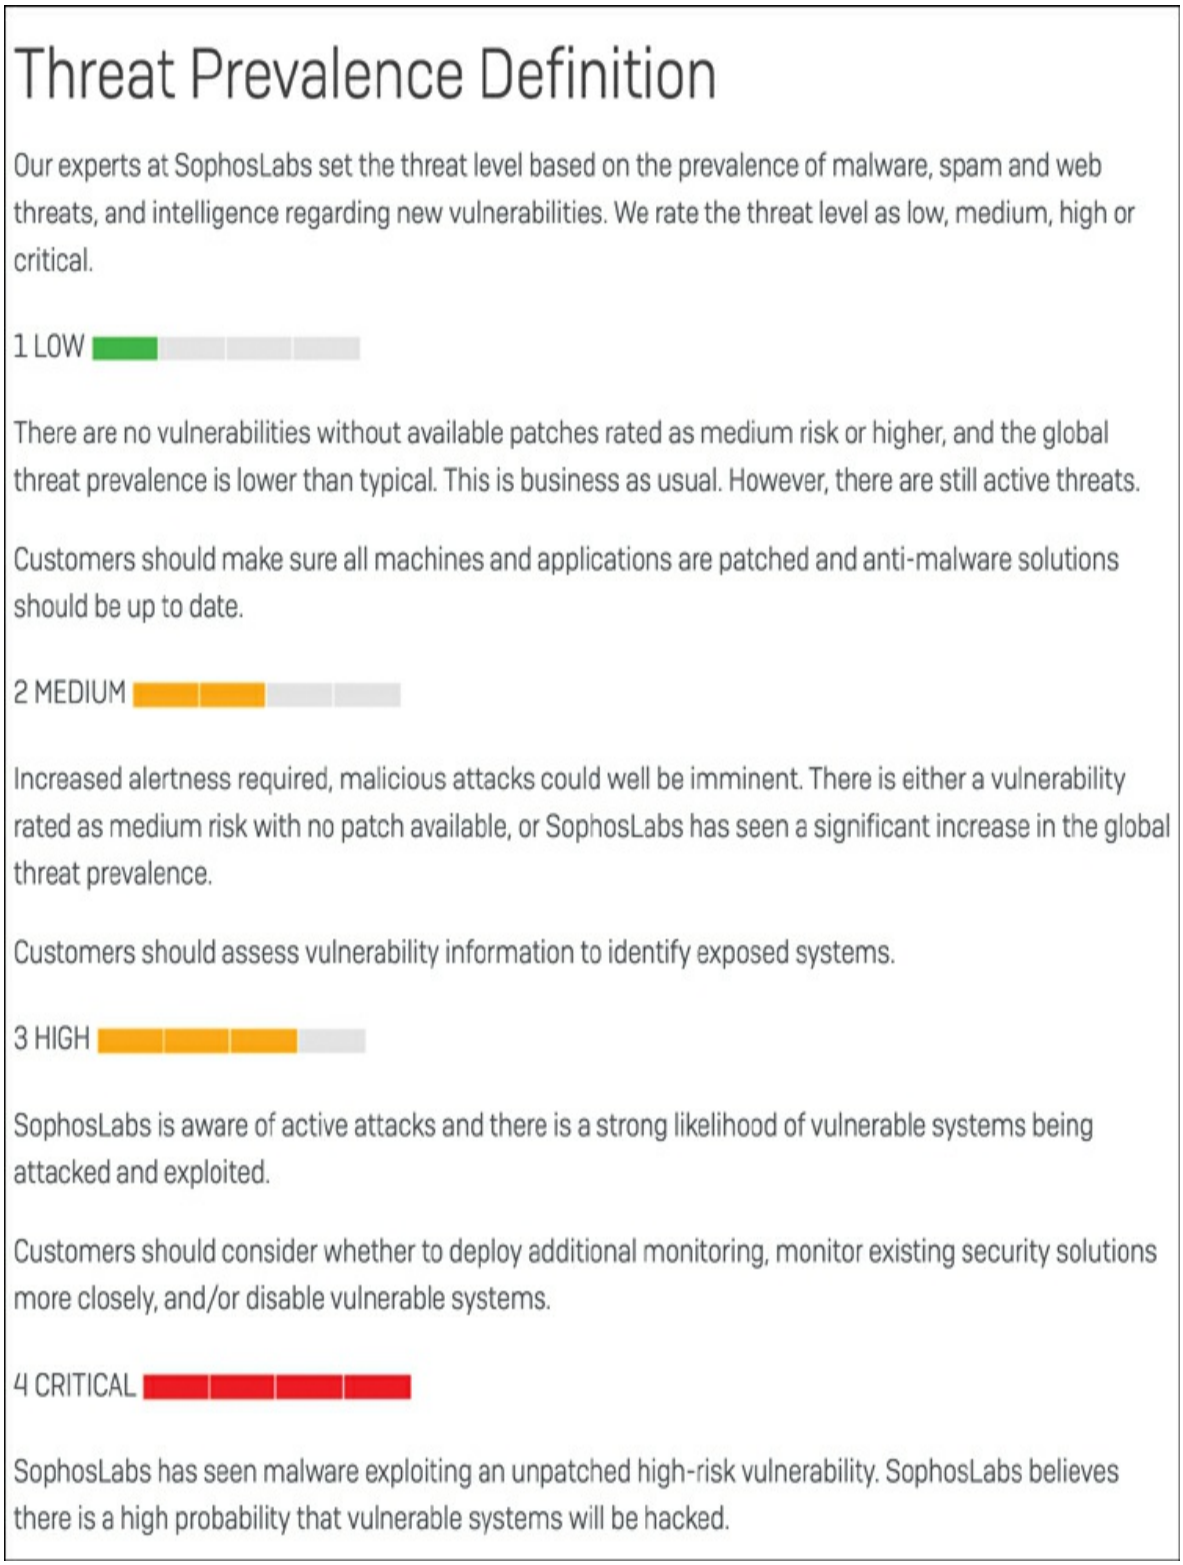

FIGURE 7-2 Sophos Threat Prevalence Usage Example

Another context datapoint is the age of the data. For example, a URL that was considered malicious a few months ago might be perfectly legitimate now. Maybe a trusted website was compromised and for a period of time was used to deliver malware. Once the website owners remediate the compromise, the URL is no longer a threat. The age of threat intelligence can help you adjust your response based on whether the threat is considered very new, possibly no longer a threat, or possibly a false positive. This is why I highly suggest you validate the quality of data associated with free threat intelligence feeds. I find that many contain very old data, which likely doesn't provide much value because the threats identified have changed.

Other important context items you want to gather from threat intelligence are potential impact, potential victims, attack trends, and anything else that can help the SOC better understand the threat that needs to be addressed. This brings us to the next topic, which is how to identify which context is ideal for your SOC using the proper threat intelligence evaluation process.

Evaluating Threat Intelligence

Regardless of the type of external threat intelligence feed you choose, it will provide a nonprioritized list of data that does not have any context regarding impact to your specific organization. It is up to your organization to pinpoint what you need so that you avoid being saturated with useless data that provides more harm than value. To get the best value, you must select feeds that can be used properly by the intended party looking to benefit from such data. You must also evaluate whether the source is *reliable*, meaning the threat intelligence provides data that is *accurate*, *relevant*, and *timely*. If any of these three key factors is not present, you should consider that source to be unreliable and you should not use it.

Before you consider evaluating threat intelligence feeds, you first need to understand your requirements for threat intelligence. You must evaluate the purpose and impact you hope to achieve from using external threat data. Consider the following series of questions as you evaluate an external threat intelligence feed:

- Who is the audience of the threat intelligence?

- What risks are unique to your organization's industry?
- What does your organization's network infrastructure and security capabilities look like, and how could they benefit from additional intelligence?
- What security capabilities or processes could benefit from threat intelligence?
- Is the threat intelligence being considered supported by your SOC's existing technologies? This includes the available format, how the data is delivered, how it is secured, and so on.
- Does the provider have a strong history of providing accurate and timely data?
- How often is the data updated? Data that gets updated only every 30 days will likely not be useful versus data that is frequently updated.
- Is there another option already available versus adding new or more threat intelligence?
- What budget and resources are available that could be used to process and apply threat intelligence to your practice?

There are some key elements you are looking to capture with these questions. First, you want to consider the beneficiary of the threat intelligence. Next, you want to collect data that is relevant to your industry. For example, if you work for a bank, you want to know about threats that impact the banking industry. Yes, it wouldn't hurt to know about a threat that has targeted other industries; however, it would be better if the focus of the threat intelligence feed was more specific to your line of business. Not using relevant data could also in some cases be more harmful than helpful.

Note

There are ways to tune threat intelligence feeds to cull irrelevant data; however, best practice is to focus your effort on threat intelligence feeds that are relevant to your line of business versus collecting all data sources. More is not always better in regard to threat intelligence feeds!

Another part of the evaluation being performed by asking the previously listed suggested questions is reviewing your network environment and security capabilities. The section, "Collecting and Processing Intelligence," later in this chapter, reviews different technology categories that are known for digesting and benefiting from external threat intelligence feeds as well as how different groups operationalize intelligence. For example, if you have a limited number of tools that can use technical threat data, then you are constrained by what could digest any additional data. It is best to invest in external intelligence once you have fully deployed existing tools and are ready to tune their capabilities.

The next part of the provided questions focuses on the cost to achieve the benefit. As mentioned earlier, some threat intelligence feeds are free while others require payment. You also learned that some threat intelligence feeds (technical or nontechnical threat data) could be mixed in with a lot of noise, making it difficult to abstract the intended benefit.

Threat Intelligence Checklist

Once you have answered questions regarding planning for threat intelligence, you can develop a pre-evaluation checklist to help you shop for the best threat intelligence resource for your organization's goals. Table 7-2 is an example checklist of what the results of this initial assessment could look like. By having this data, you will be able to quickly narrow down which type of threat intelligence feeds would be most ideal for you to evaluate.

TABLE 7-2 Example Checklist for a Request for Threat Intelligence

Type of intelligence	Tactical
Intended audience	Security tool administrators
Tools that support threat intelligence	SIEM, email security
Supported formats	STIX

Expected benefits	Improve SIEM event correlation by adding more context; improve malware detection in email
Industry focus	Banking, PCI DSS, finance,
Potential budget	$50,000 annually

Content Quality

Another important part of evaluating threat intelligence is understanding the quality of the external data. As you shop for a threat intelligence service, you will find vast differences in what is provided and costs ranging from free to thousands of dollars. By developing a threat intelligence checklist, you will have narrowed down the possible offerings that could work for your objectives, saving you time during the research and trial process.

Key Content Quality Factors

Your SOC must base the quality of any threat intelligence resource on the same key elements. You must judge whether the data is accurate, relevant, and timely to ensure it will provide value. *Accurate* means the data represents real threats rather than containing many false alerts that will cause more distraction than benefit. *Relevant* relates to how likely the threat will impact your organization. Acquiring threat data aligned with your market segment will make that data much more relevant than a generic source. *Timely* means the data is recent enough that responding to the potential threat would allow you to prepare for a real potential attack. Preparing for dated attacks provides little value and also gives you a false sense of security because you won't have data on current threats that are more likely to be seen attacking your resources. Use these three factors to measure the quality of the threat intelligence before starting your collection.

Content Quality Checklist

The best way to identify how accurate, relevant and timely a threat intelligence resource can be for SOC is to ask the service provider the right

questions. The following are some fundamental questions you should include in your threat intelligence feed evaluation process once you are ready to consider potential candidates:

- What are the data sources for the threat intelligence?
- What is the percentage of unique data?
- How long is the data relevant?
- How reliable is the threat intelligence?
- Is there a portal or other resource to gain more information about an event found within the feed?
- How accurate are the results from the threat intelligence?
- What is the return on investment?
- What is the total cost to use this service?
- Is this a subscription, and is there a trial period as well as minimal contract required to obtain the threat intelligence?
- Are there ethnical, legal, or compliance violations to consider?

The first question to ask a threat data provider is where it gets its data. This will quickly uncover whether the provider uses open-source tools, private data sources, or a combination of the two. Next, asking about the percent of unique data helps you determine how much overlap you will receive from a threat intelligence feed that draws upon data from multiple sources. For example, if a feed includes data from three banking sources providing similar data, the feed will seem like it is drawing on a single data source, not three unique data sources. Overlap has value insofar as it indicates a threat is common across multiple collection sources, but if you are regularly receiving a lot of overlapping data, you may be paying for what is advertised as a larger service than what you are actually receiving once you remove duplicate data.

I gave an example earlier in this chapter about how a technical threat intelligence feed can be used to update an IPS detection database. Using old or generic open-source technical threat intelligence resources will limit the tool by enabling it to look for threats that are not current. This can cause more harm than value considering the wasted resources. The same concern relates to the question of how reliable the threat intelligence resource is. This

is extremely important for operational threat intelligence because the expectation is that details about threat actors are included. Unreliable data can quickly lead to an overload of false positives, causing more trouble than value in the SOC as well as invoking the wrong actions. An example of this could be threat intelligence highlighting a specific domain as malicious, although later research might show that the domain was either spoofed or used as a proxy for the attack. Once again, it is important for an external threat data resource to be accurate, relevant, and timely.

In reality, many SOCs often do not evaluate the effectiveness of their threat intelligence after they have implemented it. Sometimes this effectiveness of the threat intelligence is not examined until it is time for contract renewals. I recommend a continuous evaluation of how effective and actionable threat intelligence is from a specific vendor or feed. If it is not effective or causes false positives or false negatives, it will waste the SOC's time and potentially put the SOC at a greater risk of missing an attack. The preceding list of questions also includes one about whether a service provides a portal to validate a finding, since it is common to ask why something has been flagged as bad. Identifying a threat is fine, but it's more important to know why what is being identified is a threat!

Testing Threat Intelligence

Outside of understanding details about the threat intelligence are the questions I suggested you ask regarding operationalizing a threat intelligence service. Ideally, an external threat intelligence service will allow you to use its feed for a trial period. You also want to consider if any additional data modification (such as custom filters) will be needed before the data can be used and how much effort is expected to abstract the value from the noise. Data formats for threat intelligence are covered later in this chapter in the section "Collecting and Processing Intelligence."

Using this approach to organize your requirements for a threat intelligence service will not only help keep your evaluation of security feeds organized, but also help keep your testing criteria clear, allowing anyone, regardless of technical level, to understand how well each offering satisfies your organization's goals. One of the biggest challenges in large organizations is

acquiring abstract services such as threat intelligence feeds using an external procurement organization. The requestor within the SOC knows what they want; however, the group responsible for selecting the service might not and could purchase the wrong service.

Note

It is common for a procurement team outside of the SOC to base decisions solely on price, which, in the world of threat intelligence, would mean selecting a bunch of open-source feeds or subscribe to a very low-price service. As you work through this chapter, you will find that relying on an open-source service (free or inexpensive) can end up costing more in time, money, and resources in the long run. In some cases, you will be better off not using a free resource as it won't deliver the desired impact.

To summarize threat intelligence value, consider a threat intelligence feed that can act as an ongoing stream of contextual information related to current threats. The feed must provide accurate, relevant, and timely data in order for it be valuable. The purpose of using threat intelligence will depend on the requestor's mission. The reason might be technical, such as enhancing the detection capabilities of security tools, or it might be nontechnical, such as aiding a leader to better understand the risk associated with a change or threat. Make sure to have a purpose for using threat data or it will just end up being additional data and never provide any value.

This chapter opened by stating that the best way to determine the value of threat intelligence is based not on the data *but on how it is used*. This important concept of *actionable intelligence* is tackled throughout this chapter. A threat intelligence feed could be rich with data; however, if your people, process, and technology don't understand how to digest and operationalize the data, then little to no value will be obtained. This brings us to the next topic, which is how to plan a threat intelligence project.

Planning a Threat Intelligence Project

I'll continue to say repeatedly in this chapter that the success of using any form of threat intelligence depends on how it is used and what you understand from it rather than what was collected. In order to use any form of external data properly, you must first collect and process it in a way that it can be operationalized for a specific goal. Before doing so, you need to decide what goal you are trying to accomplish. You can then match that goal with what should be collected and what should be filtered out and test to see if your goal is achieved. Once achieved, you can make adjustments as you become more comfortable with the new data resource.

Some tools that digest threat data support multiple formats, while other tools are limited in the data formats they can accept. Some tools have methods to manipulate nonconforming data into an acceptable format as it is being collected (known as parsing), while other tools require the data feed to meet specific format requirements in order to process it. Some sources will have a lot of useless data surrounding what your SOC team finds useful, making it difficult to find value between the noise. This section covers how to handle all of these scenarios.

Let's look at the data expectations your SOC needs to have for the four categories of threat intelligence, starting with evaluating strategic threat intelligence.

Data Expectations for Strategic Threat Intelligence

Strategic threat intelligence is different from the other forms of threat intelligence because it is typically requested by the executive level and is used by nontechnical people. This means the audience won't necessarily understand the details involving a request and could ask for data that doesn't exist or is not possible to obtain. An example would be a request for "the names of those responsible for the recent attack" or "the location where we are going to be hit next by the attackers." The details behind an attack are likely not available, and it is impossible to know where you will be hit next. It is also important to point out that having the names of a malicious party located somewhere on this planet provides absolutely no value in regard to

your organization's security posture. Attempting to predict a specific attack is also not useful because attacks can come from anywhere and at any point. You are better off assessing your vulnerabilities and addressing those weaknesses with the assumption that leaving a vulnerability exposed will be the next point of attack.

Planning a strategic threat intelligence project starts with *asking the right questions*. This conversation should take place between the requestor of the intelligence within leadership and a SOC analyst or researcher that will develop the final deliverable. An example conversation could be a request to understand the risk associated with a change within the organization, such as investing in cloud security services, expanding into a new line of business, or providing a new service, all of which would require a change in people, process, and technology and would have security risks associated with the change. Approaching your request with a topic that has a clear agenda but requires research to obtain the best answer will help the analyst understand the scope of the request and deliver a product that can be used to help solve the challenges outlined in the agenda. You need to maintain open communication and a strong feedback loop as a strategic threat intelligence project is being developed by your team, rather than providing all of the details about the project upfront and expect for the best results post research without any further communication or feedback to the analyst. There are many aspects that can change as the analyst performs research, which the analyst should be able to ask about to ensure their research is in line with the project's goals.

The results of a strategic threat intelligence report will not consist of a specific yes or no recommendation for every item addressed. Instead, it is common for *risk* or *confidence scores* to be used, because many factors that are unpredictable, such as economic or incidental risk, must be considered but can't be confirmed with 100% accuracy. There are exceptions to this rule when a specific yes or no recommendation can be provided. For example, asking if support for something can be validated would require a yes or no answer, such as "Does this tool function in a cloud environment or not?" It is, however, more common for topics in nontechnical reports to use answers based on the likelihood of something occurring using the risk or confidence score system. Certification programs such as the CISSP by $(ISC)^2$ recommend using similar language and provide formulas for calculating

potential regarding risk, likelihood, and potential impact, and these calculations are commonly included in nontechnical threat intelligence reports.

Once a draft report is received by the C-level or SOC member requesting the nontechnical threat intelligence, success should be measured based on a few factors:

- Was the report produced in line with the scope?
- Do the results help with the decision it was designed to assist with?
- Is the right level of technical and nontechnical details included?
- Is the data reliable, relevant, and timely?

Data Expectations for Tactical Threat Intelligence

Tactical threat intelligence is more specific than strategic threat intelligence regarding what you expect to collect and how it will be used. The purpose of the results is to understand the tactics, techniques, and procedures (TTPs) used by threat actors to improve an organization's defense. Tactical threat intelligence can be further refined into the following four categories:

- **Attack vectors:** The type of attacks being used. Examples include a phishing campaign or wrapping documents with malware. The goal of collecting this type of intelligence is to understand what vectors are being used by malicious parties so you can better prepare for an expected attack. If you find that phishing exploits are increasing in frequency within your market space, you can validate that your anti-phishing techniques are up to date, provide a company-wide warning, and provide training to better prepare the SOC team for the current threat.

  Attack vector elements you will want included within your tactical threat data feed include the following:

  - How are targets being selected?
  - Which vulnerabilities are being targeted?
  - What level of privilege is required, and could higher privileges be obtained?

    - What is the objective of the attack (financial, destruction, etc.)?
    - What are patterns to look for?
- **Tools:** Any specific tools being used by the attackers. For example, a specific type of exploit kit being used to attack a handful of vulnerabilities, or a specific type of distributed denial-of-service (DDoS) tool that exploits a communication protocol to take down a service. Having details about the attacker's tools will allow defenders to understand how the tools function and validate if the organization is vulnerable to a potential attack.
- **Infrastructure:** How the attackers communicate, where stolen data is transferred, how payments are made if any form of ransomware is being used, and what type of encryption is used during the exploitation and exfiltration process. Knowing these details allows defenders to monitor for similar communication, close ports or protocols associated with the threat that are not needed to run the business, and better prepare a response if malicious activity is identified. An example of closing ports could be filtering out communication to Tor networks if tactical data shows there is a high risk of a ransomware attack that communicates through Tor networks.
- **Stealth strategies:** How the threat avoids detection or bypasses security capabilities. Examples include using beaconing or other randomized communication or breaking up malware components and reassembling them post compromise. Attacks could also be identified as encrypting payloads or hiding data within certain files, such as embedding text in PDFs, allowing the SOC to be more aware and prepared to examine similar files that exist in the organization.

Having a tactical threat data resource providing elements from these four categories allows a SOC to have a broader understanding of potential threats, from vectors and tools to infrastructure and strategy used. The SOC can first validate whether existing security capabilities could detect and prevent the potential threat, and perform any adjustments if there is doubt about any tool's effectiveness. The SOC can then evaluate people and process to see if a proper response could be initiated in the event of an attack. An incident response plan could be developed and tested to ensure the organization is ready for the potential threat. Combining these efforts would allow the SOC

to be more effective against current threats.

> **Note**
>
> It is important that your tactical threat data be aligned with your market segment. You don’t want to waste efforts on threats that are not likely to impact your business.

Deciding what tactical data would be ideal for your project must be based on the same key elements covered earlier in this chapter. You must judge whether the data is accurate, relevant, and timely to ensure it will provide value. Remember to use these three factors to measure the quality of the threat data before starting your collection.

Data Expectations for Operational Threat Intelligence

Operational threat intelligence is focused on specific attacks, whereas tactical threat intelligence is broader. Expectations for operational threat intelligence projects are much more specific and detailed-oriented. Some threat data sources will offer a blend of tactical and operational data, meaning some TTPs will include very specific details about an attack campaign, allowing you to prepare for how a specific threat is being carried out. For example, you might find TTPs related to a phishing campaign that include operational data showing the emails are related to a specific group of attackers residing in a specific part of the world. The operational data could also speak to the latest trends of how the attackers are targeting victims, allowing you to not only prepare your tools, but also anticipate when to expect the attack, including how the phishing emails are formatted.

An example of operational threat intelligence is based on a real-world ransomware campaign a few years back that occurred in response to Microsoft pushing out updates to its operating system in a slow rollout fashion. Microsoft would contact organizations to notify them when it was

their turn to receive a new software update. A specific group of attackers got wind of this approach and developed a phishing campaign that targeted companies that hadn't received the update. The attackers pretended to be Microsoft providing the organizations their promised update, but the "update" was ransomware used to extort the victims of the attack. In this example, tactical threat data would identify the ransomware and the phishing attack that was used; operational threat data would provide context around who was delivering the attack and why the attack was being performed, better preparing an organization for the trick used by this campaign. In this example, you could alert staff to validate any Microsoft communication and evaluate any update files for malware before applying them to critical systems.

There are pitfalls regarding operational threat data that you must consider. The following common challenges are likely to occur as you evaluate operational data and must be considered before operational data is collected to avoid capturing data that won't be able to be converted into actional intelligence.

- Many threat actors are not based out of English-speaking countries, meaning communication will require translation. Without a translator, you will not be able to abstract value.
- Data could suddenly become unavailable if chat room or other resources providing operational threat intelligence is taken down or your access is revoked. Consider how reliable the resource is and the likelihood it will be available long enough to obtain value.
- Operational data will be raw and include a lot of noise if pulled from a chat room or a message board requiring manual monitoring and filtering to obtain value. You must ask yourself if it is worth the effort to collect the data that provides value even if it's free and available. Later in this chapter under the section titled Scrapers, I will provide an example of using operational data from Twitter feeds.
- Having detailed operational threat intelligence could be hard to use if obfuscation tactics are being used such as code names to refer to the attack. This is common when resources are publicly available such as social media commination channels.
- Operational data could be inaccurate by design and posted by bots to act

as a distraction while a threat is being executed. Validating bot traffic continues to increase in difficulty as bot technology matures including elements of artificial intelligence. Bots are also capable of quickly spreading across all social media channels smothering real data.

Many of these challenges can be overcome if the proper process is used to onboard the external threat data. After you overcome these challenges, you must once again take the same approach of judging the data as you do with tactical data. That means assessing whether the data is accurate, relevant, and timely. Because operational data will contain more context and overlap with tactical data, you must consider if there is a way to process the additional context outside of what is seen within tactical data to avoid missing context that many attack tools will not understand but that the people within the SOC will benefit from. I will look at this concept in the upcoming section "Collecting and Processing Intelligence."

Data Expectations for Technical Threat Intelligence

Technical threat data is specific to the threat actor's tools and infrastructure, while leaving out any additional context. The purpose of having this focus is to gather intelligence to improve detection and prevention capabilities. There are many technical threat intelligence artifacts to consider when evaluating technical threat data. At a bare minimum, I recommend looking for the following items when evaluating a technical threat data feed before deciding to include it as part of your project. These are the most common use cases I find within any SOC's requirements for technical threat data.

- Suspicious or known malicious domains/registered URLs
- IP addresses associated with malicious activity
- Latest malware hashes
- File artifacts from malware samples
- Subject lines or email content associated with phishing campaigns

Collecting a list of malicious domains is an easy method to develop a blacklist that can be used to block access to current remote threats. A similar approach is to use IP addresses associated with malicious activity. Host and

gateway security tools can quickly digest and adjust their blacklists, making the collection of data to operationalize a defense process simple. One challenge you might face is hitting IP address limits not allowing storage for all of the IP addresses or domains. Another challenge is false positives within the provided data. False positives can result due to many factors, including a malicious source cloning a safe website, the use of proxies by the attacker, Tor networks, or other deception and stealth tactics. IP addresses are even harder to validate because they can easily be cloned and manipulated. Testing and filtering might need to be applied to reduce the size of data that will be processed and to remove data that isn't reliable.

Capturing malware hashes and artifacts is another relatively easy approach to collect and operationalize threat data. Many pattern-matching tools such as antivirus and sandboxes will compare a suspicious artifact against a list of known malicious hash values. The vendor of a security tool will provide its own data feeds; however, complementing what is provided by the vendor's updates with industry-specific threat data can help make your security tools more effective at detecting threats that matter to your organization. I find many security tool vendors are adding this capability by supporting data feeds formatted in STIX and TAXII, which is a topic I will cover shortly in the section "Technical Processing."

Monitoring emails for content associated with phishing is yet another simple and effective method to improve detection against identified phishing campaigns. The value of this approach will depend on how often the phishing language and content are changed as well as whether the attackers use email security avoidance techniques. Using a combination of operational and technical data feeds will help your SOC adjust to the latest phishing trends.

Technical threat intelligence might seem like something you can just consume and benefit from, but once again I must remind you about the previous examples of filling your tools with threat and vulnerability details that are not related to your organization. Doing so will reduce the effectiveness of those tools rather than providing any form of benefit. To avoid this pitfall, you will need to ensure the technical data is related to your field of business and could be processed by your technology to ensure the data is relevant to your specific needs. Pushing a ton of random open-source technical threat intelligence feeds will cause more harm than value. Most organizations that are negatively impacting their tools with additional, yet not

valuable, threat intelligence feeds don’t know they are doing so. When asked about the value, the typical response is “More is better, right?” This means they have not considered what is being added and what purpose the feed is supposed to provide. In the upcoming section, I will explain how to avoid this mistake.

Outside of business relevance, how you collect the data will be important to how your tools will benefit from the details within the feed. This brings us to the next topic, which is how you will collect and process what you have identified as useful threat data resources. If you can’t properly collect and process threat data, you won’t be able to convert the threat data into threat intelligence.

Collecting and Processing Intelligence

As you plan your threat intelligence project, you will want specific results, which will be your success criteria. The threat intelligence service provider might have everything you need in its threat data offering; however, you could encounter problems with abstracting what you find valuable from the rest of the data. Problems could include what format your tools support, the quality of the data, or too much noise in the data. For example, suppose you obtain an operational-type data source such as a feed from a public forum that speaks about current cyberattack behavior; however, most of the communication is hidden in secret code and mixed in with non-security conversations. You could expend resources collecting this threat data, but if you can’t figure out a method to quickly abstract what you need, you will not be able to convert what you collected into actionable intelligence. Another example could be a technical threat intelligence feed that has what you want but delivers the feed in a format that isn’t accepted or understood by your security tools. I have had this occur when attempting to digest feeds within a security information and event management (SIEM) solution. Figure 7-3 shows an example of using Splunk to pull in data from a security tool; however, Splunk is not properly parsing the data being digested. When I attempt to see actions taken by the security tool within Splunk, Splunk can only identify events that were allowed. In its current state, the data from this security tool provides little value within Splunk; hence, it’s threat data but not

threat intelligence.

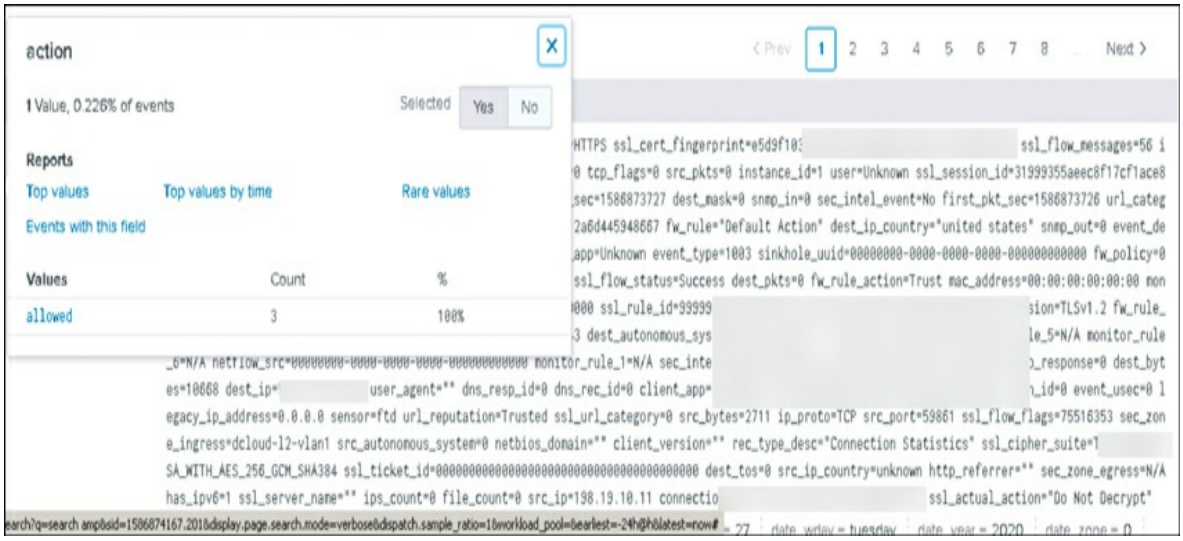

FIGURE 7-3 Example of Splunk Not Correctly Processing Threat Data

Note

The issue represented in Figure 7-3 could be caused by many factors, including the tool sending the data, the tool receiving the data, or how the data is transferred. Troubleshooting these situations will differ on a case-by-case basis and will depend on the tools and data type involved. The problem in Figure 7-3 was how the data was being sent, which when fixed, allowed Splunk to correctly display tons of details in a easily searchable format.

There is good news regarding overcoming situations in which a tool lacks support for the format of a threat data feed or in which a resource such as a report or public website is full of unwanted noise. There are ways to adjust data into an acceptable format and processing techniques that can filter out what is not wanted. (I find that most data format problems can be solved by doing a little research using Google or by speaking with the vendor processing the data.) This section first covers best practices for ensuring nontechnical threat data is converted into a useful threat intelligence format.

It then covers techniques for collecting and processing intelligence from social media and other web-related operational data resources, including using data scraping and using monitoring tools that can alert you when keywords are used so that you can find value between the noise. Finally, for technical resources, this section looks at various types of data formats and accommodating threat data that doesn't follow common data formats. Keep in mind there are many organizations that will automate collecting this information for your organization and provide it to you in a digestible format.

First, let's cover best practices for ensuring nontechnical threat data is converted into a useful threat intelligence format.

Processing Nontechnical Data

Earlier in this chapter, I explained the concept of strategic threat intelligence, aka nontechnical threat data. I pointed out how it is common for this form of intelligence to be delivered as a report or briefing, which is requested based on a project or decision that must be made. In order to receive the best results, I covered how the C-level or SOC member requesting the nonthreat intelligence must ask the right questions to the assigned analyst that allow the analyst to research the right topics leading to a useful report. If the C-level or SOC member asks too many open-ended questions, the analyst assigned to research your questions will return very generic answers, leading to a useless report. This process should not be a one-time assignment. The requestor should not dump a bunch of requests on the analyst and expect the final report to be delivered at a later time. A best practice for executing a strategic threat intelligence project is to function in a feedback-based workflow. This is the best way to ensure the analyst is understanding the C-level or SOC member's request and that adjustments can be made if the early rounds of results are not leading to a deliverable that will help with a future action.

Modified Waterfall Model

One popular engineering model that I modified for the purpose of understanding a feedback-based workflow is the waterfall model, as shown in Figure 7-4. In this example, I have adjusted the language of the waterfall

model to meet the needs of a typical strategic threat intelligence project lifecycle. Notice how each arrow represents a point of check-in, meaning a conversation between the stakeholder and analyst.

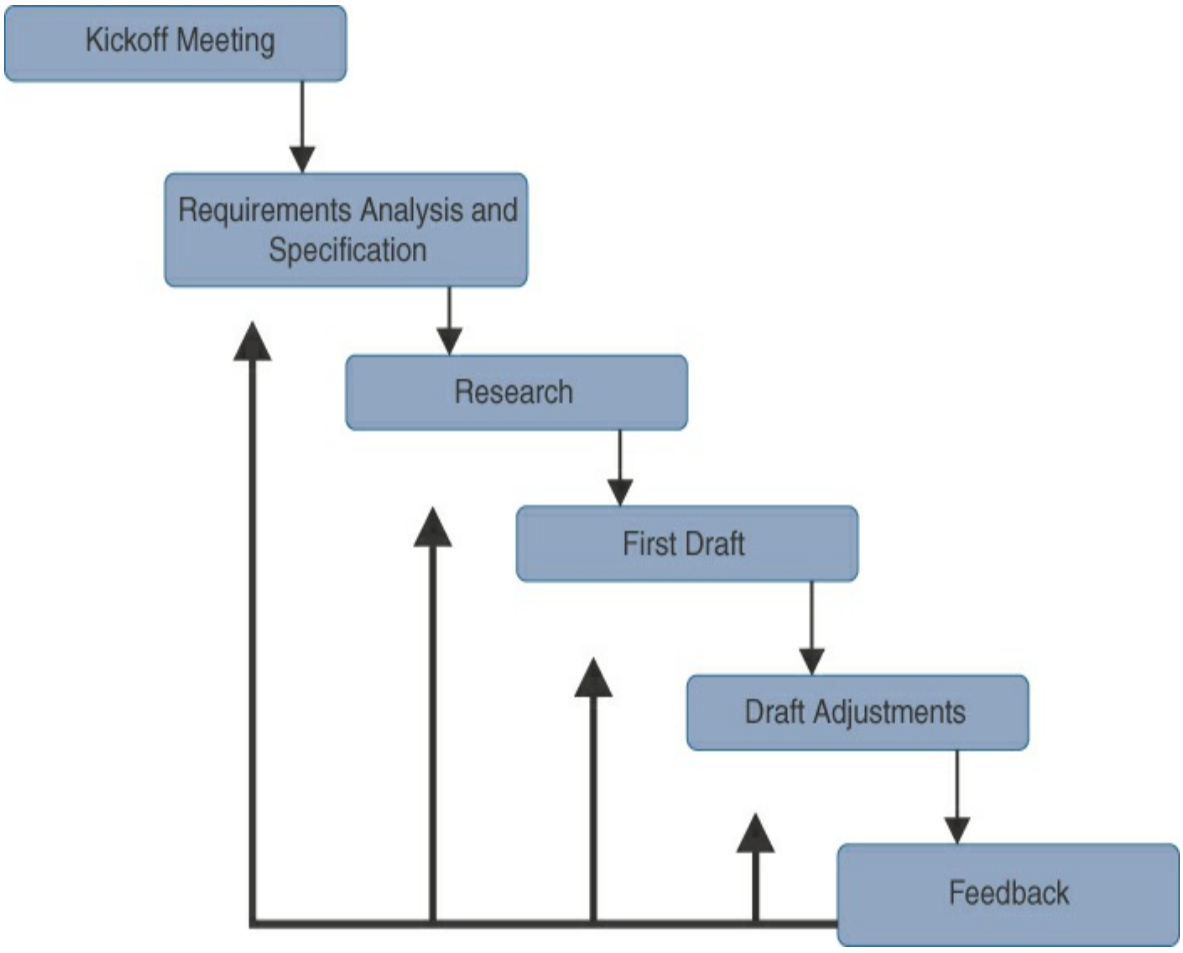

FIGURE 7-4 Modified Waterfall Model for Feedback-Based Workflow

The best way to explain the model I created is to walk through an example. Let's pretend that a C-level executive at a manufacturer is interested in opening a new datacenter in another country but is concerned about security risks and potential threats. The C-level exec requests a report that highlights potential risks, with corresponding concern rankings, and provides estimated costs for countermeasures. The process will work as follow:

1. The kickoff meeting for this project would lay out the scope of the type of data that is needed for the CEO, including where the datacenter could

be created, what type of systems could be used, and the type of data that could be seen within the datacenter.

2. The analyst would first validate each of these items to ensure that the analyst's understanding of the request is clear before research is performed.
3. While researching the requested items, the analyst will collect a lot of data; however, only certain topics would be useful to the C-level exec, while other content would not help the C-level exec with upcoming decisions about the datacenter.
4. The analyst would present a general review of the type of data available and confirm which data resources are valuable before proceeding with a draft.
5. Once the draft is complete, another checkpoint would occur to ensure that all topics needed by the C-level exec are being addressed.
6. A few more checkpoints might occur as the draft is converted into a final report, allowing for adjustments as the research and data are prepared for the C-level exec. These checkpoints allow for continuous tuning to the research and writing process, reducing the risk of the analyst including too much noise or not capturing what is required for the final report to be useful.

This same approach can be used for other forms of nontechnical report development. You might think that this process is overkill, with too many touchpoints, and you might be concerned that including these touchpoints will increase the cost and time to complete the work. I highly recommend spending the additional time and money to ensure the results provide the most value. Not doing so will likely cost you more in inaccurate results, making the entire project a waste of time and resources for all parties involved.

Note

You might be wondering why a SOC would be involved with helping a C-level exec understand the risk associated with a new datacenter. Remember that the risk management service covered in Chapter 3, "SOC Services," can be responsible for any risk to the organization. A

mature risk management team within the SOC will be seen by the organization as the trusted advisor for all security risks!

Operational Data and Web Processing

I previously explained how some threat data will be focused on a specific threat campaign. You will want your SOC to monitor for trends in attack behavior and monitor when active campaigns are occurring within your market space. Obtaining this data can require collecting data from various sources, including social media, blogs, and other forums that are mixed with a lot of noise made up of unrelated conversations and topics.

There is no simple way to find useful data from public resources because a huge part of threat intelligence is learning about concepts and events of which you are not aware—you don't know what to search for. However, you can take advantage of tactics that will alert you of potential threat data you believe will be useful and worth investigating. One approach is enabling web alerts when a specific type of conversation is being mentioned. Google Alerts is an example of a free tool that you can use to receive alerts when certain terms or conversations are taking place on public forums. Other examples include BuzzBundle, Mention, Talkwalker Alerts, and Hootsuite. Figure 7-5 is an example of setting up Google Alerts to be alerted regarding details on the Ryuk ransomware. The more details provided, such as a hash value or data about the threat actors, the better the results provided. Your SOC might get wind of an attack campaign and set up Google Alerts for specific characteristics about the campaign as a way to monitor the Internet for researchers posting updates about the campaign. This is a much more effective method than manually searching Google every time a topic of interest arises.

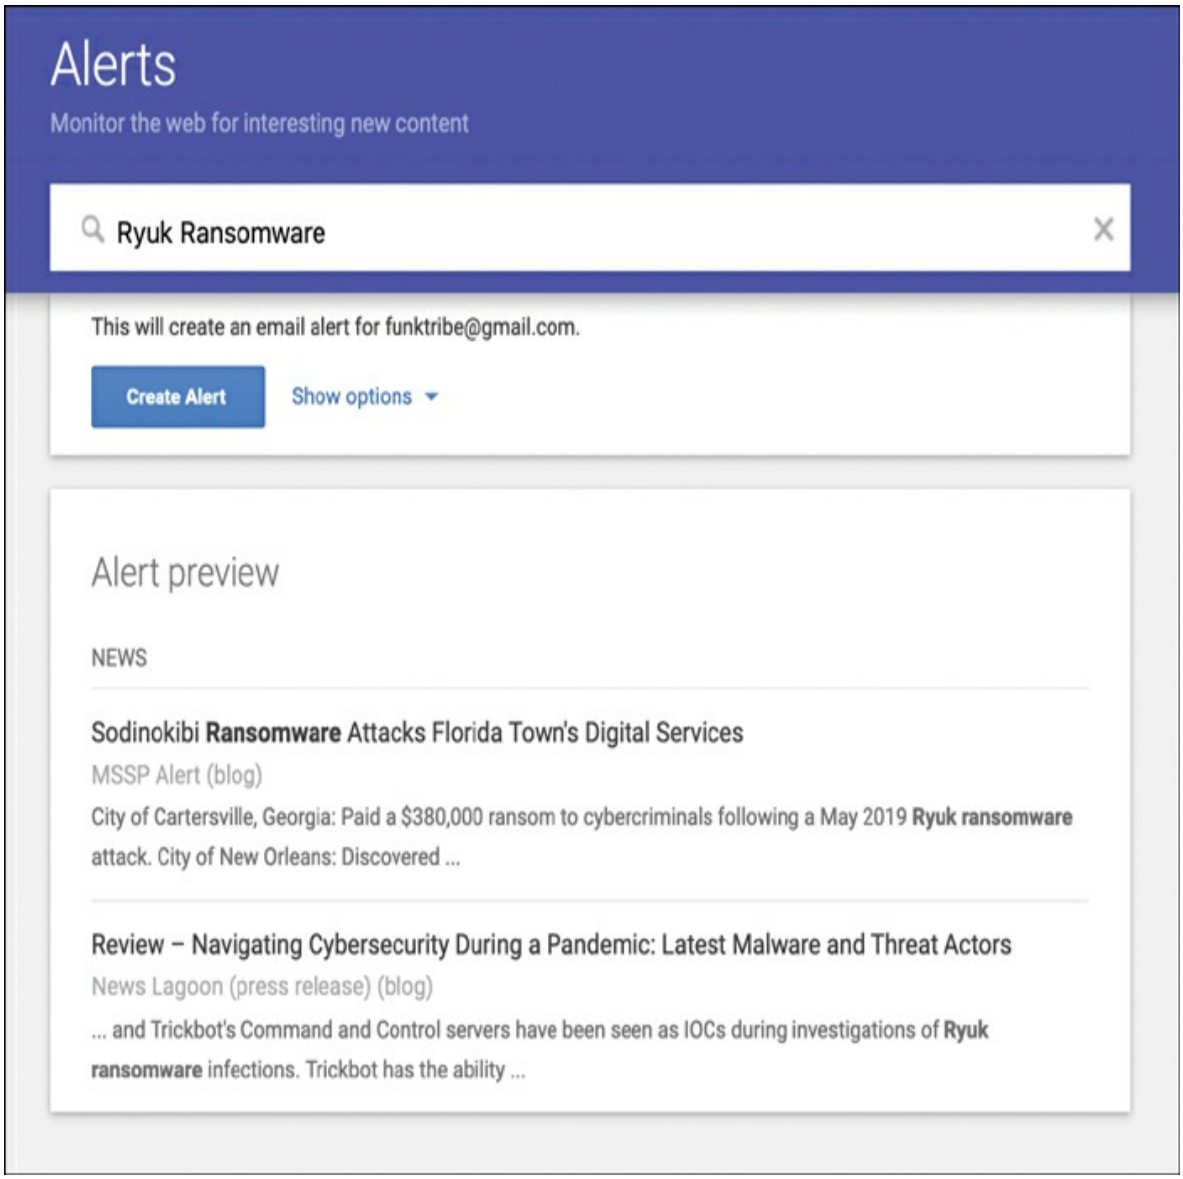

FIGURE 7-5 Google Alerts Example

Another option for identifying new threat data is to use social media and other resources to follow security professionals and influencers who post about current events that matter. Examples of important influencers in the cybersecurity industry include Bruce Schneier, Brian Krebs, Jeff Moss, and Dan Kaminsky. You can also follow vulnerability and risk advisory feeds/blogs for industry trends. Examples include the United States Computer Emergency Readiness Team (US-CERT), National Vulnerability Database (NVD), and SANS Internet Storm Center. Other feeds and reports to consider are analyst reports such as threat reports from Dragos and Cisco Talos and

the much respected research reports from Verizon (Data Breach Report), Gartner, and Forester. Keep in mind that the audience for many of these sources is very generic, so it might be worth investing in custom strategic threat intelligence versus manually processing what is publicly available.

Scrapers

Another challenge you are going to encounter as you are alerted to web-based resources with threat data is identifying what is useful from the rest of the noise. You also could have tools that are able to convert specific details such as IOCs into actions, but those details may be mixed in with a ton of useless details that need to be removed. This is where scrapers can become handy. A scraper can be configured to identify key terms, copy them, and paste them into a useful format such as a .csv file. An example could be setting up a scraper to collect hash values or IP addresses associated with malicious behavior, which can be parsed into cells that a security tool would be able to translate into a report or action. Examples of scrapers include Import.io, Dexi.io, Scrapinghub, and ParseHub.

One scraper you can use that is free but limited in capabilities is the Scraper Chrome extension. Simply install the Chrome Scraper extension, find a page that has data of interest, and you can highlight to scrape similar data. Figure 7-6 is an example of scraping hash values off of the Cisco Talos blog that speaks about various types of cyberthreats. These hashes could be fed into a SIEM solution or other security tool to scan for pattern matches.

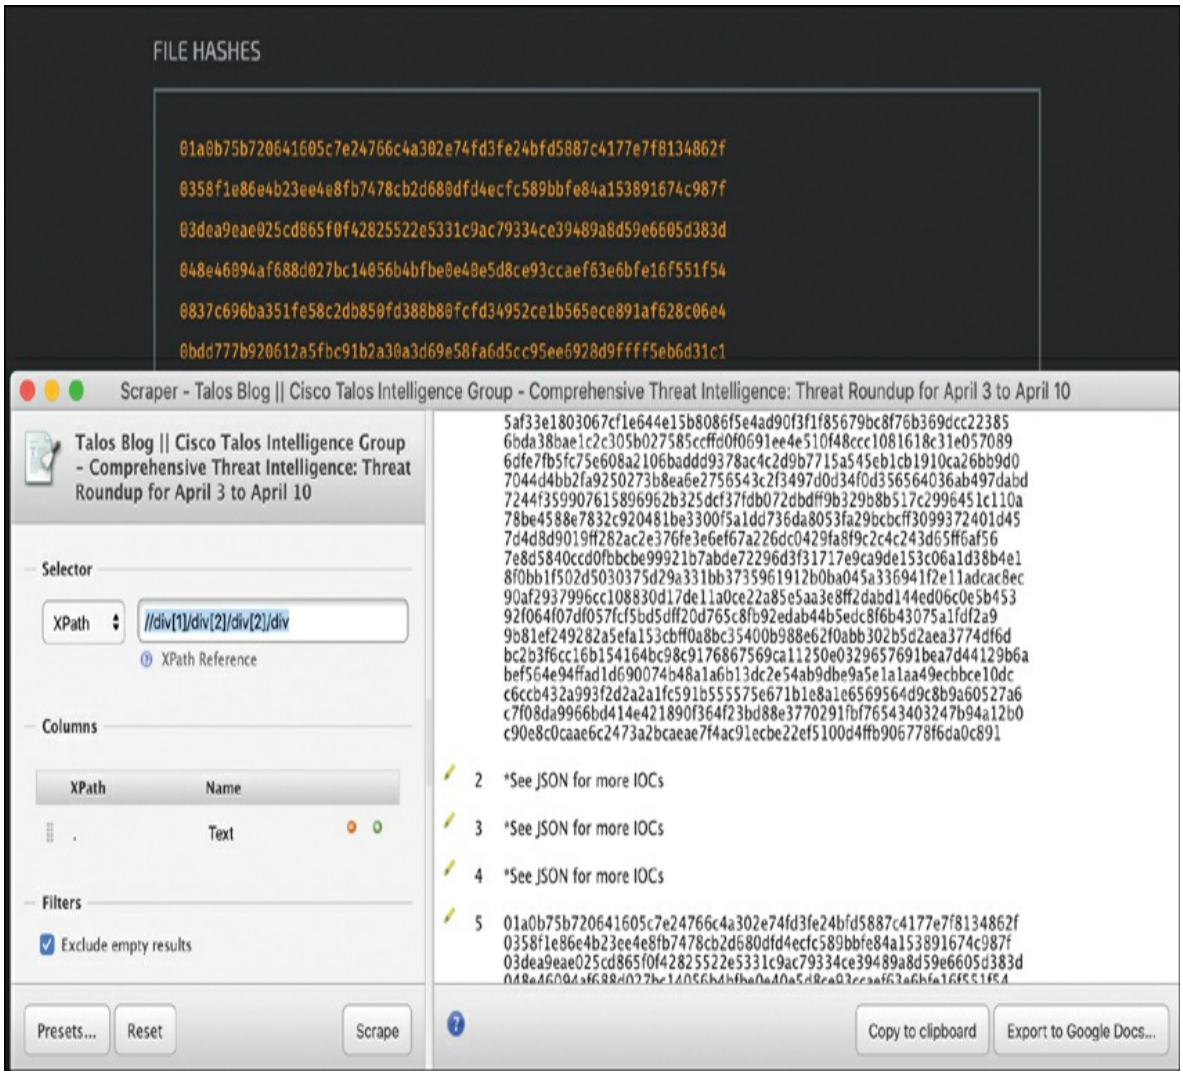

FIGURE 7-6 Google Chrome Scraper Collecting Hashes from the Cisco Talos Blog

Social Media

Social media can also be useful for threat data if leveraged correctly. Let's look at Twitter as an example use case. Twitter isn't just a personal platform for promoting or sharing personal events. There are bots that share tons of data, including recent IOCs and threat detection rules. There are Twitter bots that are configured to monitor Internet of Things (IoT) devices. There are open-source honeypots that log their findings on Twitter. Figure 7-7 shows some examples of this type of behavior pulled from Twitter.

FIGURE 7-7 Twitter Threat Data Behavior Examples

Why should you care about Twitter threat data? Think about the speed of information released on resources such as Twitter. One of the three key ingredients to valuable data is that it's timely, which is where Twitter can be a useful tool for notifications of new threats. The same value can be applied to identifying new vulnerabilities. You are more likely to catch a very new vulnerability within social media circles than waiting for a vendor to provide an update or for a vulnerability to make its way to national news. In regard to situational awareness best practices, you will want to be informed way before a weakness is broadcasted nationally, since that would mean the vulnerability has been out and caused enough problems with organizations to allow it to become a big news item. Twitter can help you become aware of emerging threats.

Your challenge to leverage a resource such as Twitter will be to pull relevant and reliable data. This can be accomplished, once again, using scraper technology. You won't be able to use the scrapers I previously covered because they are not capable of understanding how Twitter works. One Twitter-friendly scraper that you can use is Twint (https://github.com/twintproject/twint). Twint is different from other Twitter-friendly scrapers because it doesn't use Twitter's API, which would have a lot of limitations of what can be scraped. Instead, Twint uses Python, allowing you to scrape a user's followers, following, tweets, and more while evading API limitations. Using Twint, you can search for specific keywords such as "0-day," "CVE-," "CVE-2020-*," and "bugbounty." As data is collected, you can use additional CVE-related tags to look deeper into the results and weed out the noise. An example is searching for "Github" and "CVE" to identify GitHub repositories with proofs of concept of recent

vulnerabilities. Hashtags such as “ThreatHunting” can collect data and populate a widget within your SIEM solution to alert when recent cybercriminal campaigns are being socialized on Twitter. There are hundreds of use cases to leverage an always-flowing data feed such as Twitter for recent events.

Social Media Example

An example of using this approach in action was posted by the Trend Micro team monitoring the confluence vulnerability explained within CVE-2019-3396. By correlating Twitter conversations with the term “confluence” and CVE-2019-3396, the Trend Micro team was able to develop what is shown in Figure 7-8 representing different Twitter accounts speaking about the same DDoS bot that scripted an attack against the CVE mentioned.

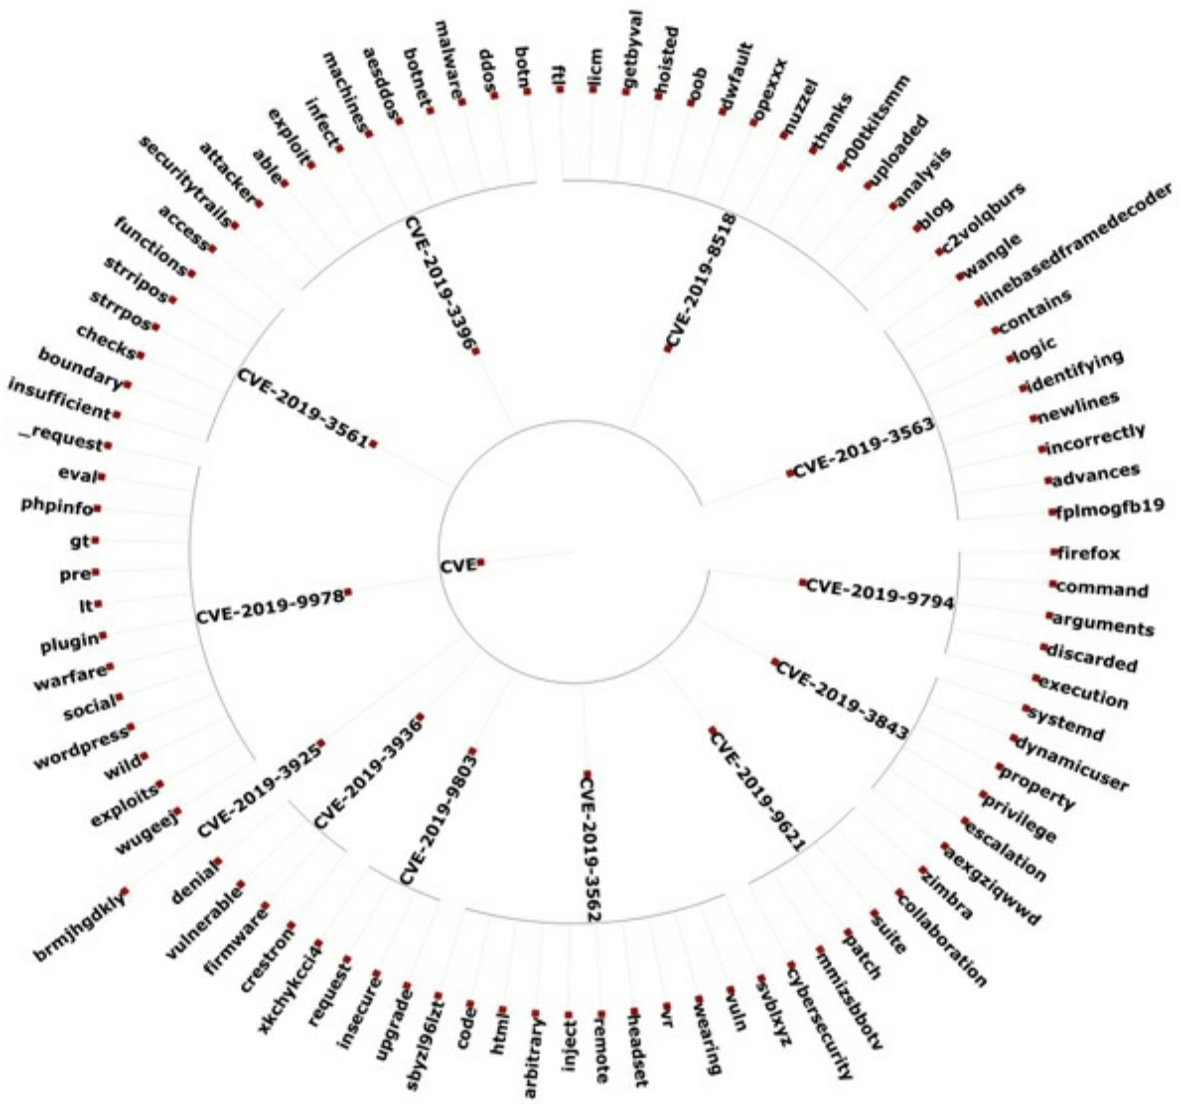

FIGURE 7-8 Example of Twitter Data Being Harvested into Threat Intelligence

These use cases and examples show a ton of value; however, you will run across the challenges I covered earlier in this chapter in regard to operational threat data. Twitter is brimming with bots, which can quickly spread fake news across multiple accounts, tainting your results from scraping Twitter. Unfortunately, there is no simple solution to overcoming fake news. The Trend Micro team that provided the previous example warned that bot traffic in certain cases will overwhelm real user traffic with the goal of creating panic over a fake event. The Trend Micro research also showed that if you attempt to validate events across multiple platforms such as Facebook,

YouTube, or Telegram, the same bots were also active on those and other social resources, further spreading the fake news and smothering articles that contain facts and other useful data.

Social media resources such as Twitter can also be used as a platform for delivering attacks, making the data a risk to your tools and environment. An example is the Anubis Android banking malware that used Twitter and Telegram as part of the C&C infrastructure. Also, malicious code and other elements don't have to be shared over text within Twitter. Data could be hidden in images, known as *steganography*, or links to malicious sources could be cloaked by using proxies or URL modification techniques such as Tinyurl.com or QR codes. Certain forms of malware will hide their final payload within images and could use social media as a way to get the payload captured within your scraping, allowing malware to bypass your security filters.

Cybersecurity is a continuously changing field. Resources such as Facebook and Twitter are part of many people's daily lives, making them ideal candidates for operational data sources if used properly and securely. With some research, you can find a scraper for every popular social media resource. For example, Ultimate Facebook Scraper (https://github.com/harismuneer/Ultimate-Facebook-Scraper) is a viable option for performing similar data abstraction as I covered using Twitter. I suggest trying one resource with specific goals and growing your social media threat data input as you discover value.

Technical Processing

The other side of collecting and processing data is supporting technical threat data. This is commonly used to enhance existing security tools by providing more IOCs that can be identified across all parts of the network and cloud. The challenge is having the ability to process useful data in an automated fashion that is acceptable across multiple vendor platforms. Each vendor has a set of acceptable formats and limitations. The good news is that there are industry standards that vendors are encouraged to support. The vendors that are forward thinking are including support for open standards, while the proprietary products are becoming more irreverent as the industry becomes

more dependent on threat data resources.

This section covers standards that can be used by various types of tools to leverage technical threat intelligence. Your goal for collecting and processing threat data is to spend the least amount of effort to capture and abstract key data elements. Then you want to deliver that data to a security tool to enhance its existing capability.

XML

Extensible Markup Language (XML) is an HTML-based language that defines a set of rules for encoding documents in a format that can be read by humans as well as machines. Simply put, XML is used to describe data. The XML standard is a flexible way to create information formats and electronically share structured data. XML is the foundation of popular formats used to share data, and many tools can accept data formatted in XML directly. A simple XML document could look like the following based on the common first program that displays “Hello, world!” to the screen:

Click here to view code image

```
<?xml version="1.0" standalone="yes"?>
<conversation>
<greeting>Hello, world!</greeting>
<response>Stop the planet, I want to get off!</response>
</conversation>
```

JSON

At its core, JavaScript Object Notation (JSON) is a lightweight format used for storing and transporting data. JSON is commonly used for server-to-webpage communication but is also the foundation of other data sharing standards such as MISP and STIX 2.x. JSON uses human-readable text to transmit data objects consisting of attribute-value pairs. For example, a JSON attribute-value pair could look like the following:

Click here to view code image

```
{Date:03/01/2020, DayofTheWeek:Tuesday, Event:Phishing}
```

JSON is a low-overhead format competitor to XML and does not use the end tags found within XML files. Since JSON is a simple text-only format, it can easily be manipulated, such as using a JavaScript script to format JSON into JavaScript objects. This allows security tools or other middleware applications to be used to automatically format and extract expressions of important.

OpenIOC

OpenIOC is an open framework for sharing threat intelligence in a machine-digestible format. OpenIOC is an extensible XML schema that allows you to describe the technical characteristics that identify a known threat, an attacker's methodology, or other evidence of compromise. OpenIOC ships with a base set of IOCs provided by Mandiant. The IOCs describe over 500 facets of environments that are useful to track down attackers, which have been vetted through Mandiant's research team. There are free open-source editors that allow you to create or add your own sets of IOCs and extend OpenIOC. SIEMs such as Splunk accept OpenIOC as a method to provide threat intelligence data.

STIX

Structured Threat Information Expression (STIX) is a language developed by MITRE that provides a way to represent structured information about cyberthreats. Think of STIX as a way to model threat intelligence so that tools understand it. It allows for automation, sharing, and analysis of data. STIX is an international standard in OASIS, open source, and extremely popular within the security industry.

The STIX language has a number of constructs or components known as STIX Domain Objects (SDOs), including the following:

- **Campaign:** Set of related TTPs, indicators, incidents, and exploit targets.
- **Course or Action (COA):** Defensive actions against a threat (prevention, remediation, mitigation).
- **Indicator:** An Observable with context. An Indicator can also contain a time range, information source, intrusion detection system rules, and other

details.

- **Observable:** Dynamic event or stateful property, represented in CybOX.
- **Threat Actor:** The cyber adversary.
- **Vulnerability:** What is subject to possible exploitation.
- Plus, more with the latest STIX 2.1 release!

STIX acts as a list of indicators about different aspects of an entire attack campaign depending on what details are available. Security tools can parse and abstract each category, making it simple to operationalize the threat intelligence. An example is how detection tools such as an IPS will abstract and monitor for characteristics of the attack. Figure 7-9 represents the concepts of the relationship of STIX.

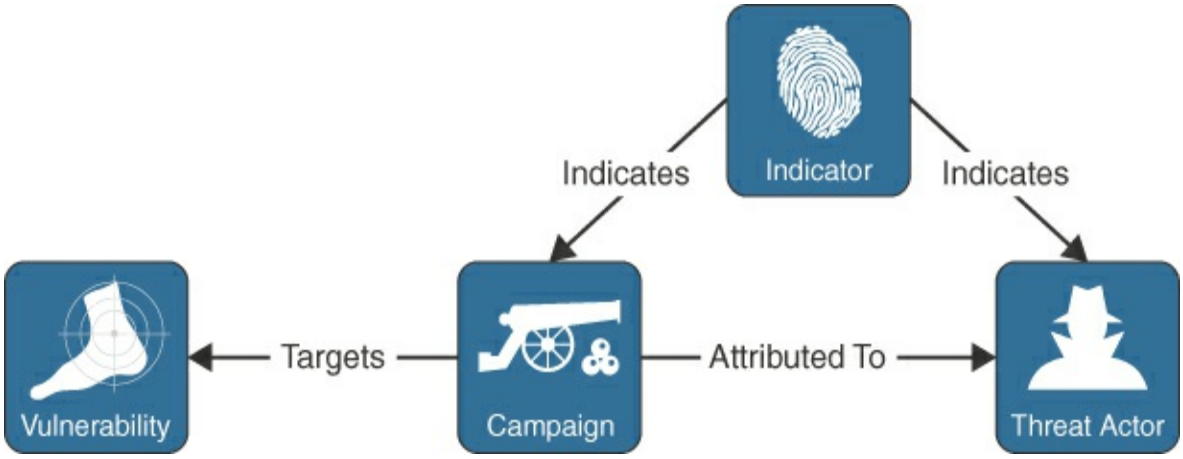

FIGURE 7-9 STIX Relationship Diagram

The next snippet is an example of a STIX 2.1 campaign object represented in JSON. Notice this example type shows it's a campaign, the spec_version representing STIX 2.1 and other constructs that are identified between each " " that provides the related data item.

Click here to view code image

```
{
   "type": "campaign",
   "id": "campaign--8e2e2d2b-17d4-4cbf-938f-98ee46b3cd3f",
   "spec_version": "2.1",
   "created": "2016-04-06T20:03:00.000Z",
```

```
   "modified": "2016-04-06T20:03:23.000Z",
   "name": "Green Group Attacks Against Finance",
   "description": "Campaign by Green Group against targets in the
      financial services sector."
}
```

TAXII

Trusted Automated Exchange of Intelligence Information (TAXII) defines a set of services and message exchanges that enable sharing of threat intelligence across organizations, products, and services. TAXII acts as an application layer protocol for the communication of cyberthreat information in a simple and scalable manner. Threat intelligence is exchanged over HTTPS using a RESTful API that aligns with popular sharing models. TAXII is specifically designed to exchange threat intelligence represented in STIX making it the common method to share STIX data.

TAXII functions by using TAXII clients and servers. There are two primary services that support different sharing models: Collection and Channel. A Collection is an interface to a logical repository of threat intelligence objects provided by a TAXII server that allows a producer to host a set of threat intelligence data that can be requested by consumers. This is how TAXII clients and servers exchange information in a request-response model. Figure 7-10 is an example of a TAXII Collection-sharing model.

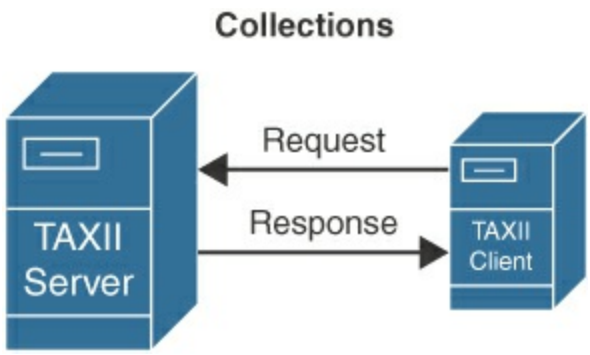

FIGURE 7-10 TAXII Collection Model

The other sharing model is the Channel model. A Channel allows producers to push data to many consumers and allows consumers to receive data from

many producers. TAXII clients exchange information with other TAXII clients in a publish-subscribe model. Figure 7-11 shows an example of the Channel model.

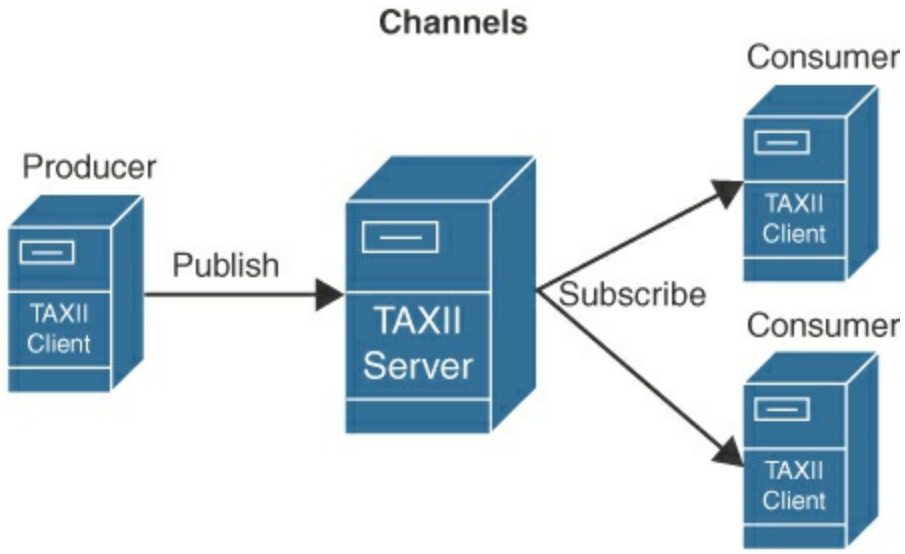

FIGURE 7-11 TAXII Channel Model

TAXII Collections and Channels can be organized in different ways. A TAXII server can support one or more API Roots, which are logical groupings of TAXII Channels and Collections. API Roots are instances of the TAXII API available at different URLs, and each API Root is the "root" URL of that particular instance of the TAXII API.

TAXII relies on existing protocols whenever possible. TAXII uses HTTPS as its transport for all communications outside of content negotiation and authentication, which happen over HTTP. TAXII was specifically designed to support the exchange of threat intelligence represented in STIX, but it can be used to transport other data formats as well. It is common to find security tools that support both TAXII and STIX.

The combination of STIX and TAXII is becoming a very popular format for sharing threat intelligence between large organizations and cybersecurity vendors.

Note

STIX and TAXII were developed to support each other, but they are independent standards. STIX is a language for modeling and representing cyberthreat intelligence, while TAXII is a protocol for exchanging cyberthreat intelligence. TAXII, as an example, can also be used to share data in formats other than STIX. To learn more about the specifications, visit https://oasis-open.github.io/cti-documentation/.

CSV

A comma-separated values (CSV) file is a delimited text file that uses a comma to separate values. Each line of the file is a data record. Each record in a CSV file consists of one or more fields, which are also separated by commas. A CSV file typically stores tabular data in plain text, in which case each line will have the same number of fields. A basic CSV file could look like the following example, which uses the data,time,engine,signature format:

Click here to view code image

```
2020-01-03,03.06.41,Ad-Aware,trojan.genericKD.3164632
2020-01-03,02:23:26,AegisLab,Uds.Dangerousobject
```

Regular Expression (Regex)

Regex is a text string format describing a search pattern. Regex is made of constants, which denote a set of strings, and operator symbols, which denote operations over these sets. Many programming languages incorporate regex search patterns, and many security tools can accept this format. A simple way to think of Regex is a much more feature-rich version of wildcards. Some examples of Regex characters include using the backslash (\) to give special meaning to a character that follows it, using a period (.) to translate to any single character (except line break characters) and using parentheses, (), to group characters together.

Regular expressions are powerful and have been used to advance computing concepts since the 1950s. Many routing protocols, search queries, and data loss prevention systems parse through data using regular expressions. Many

SOC analysts at some point in their career will need to be familiar with regular expressions to write a query for data they may not be able to find using any other method. Regular expressions can be processor and CPU intensive, and using them does take practice; however, as a security professional, you will be glad you made the investment in learning at least the basics of the structure and usage of regular expressions.

Technical Threat Intelligence Resources

There are many technical threat intelligence providers to consider that use variations of the formats previously covered. The following list covers a handful of popular free options. You can use these options to test how to collect and process different threat intelligence feeds and formats. I don't, however, suggest that you simply add all of these feeds for the sake of collecting data, as explained earlier in this chapter.

- **Emerging Threats Rule Server:** Popular open-source threat intelligence option that classifies emerging threats' IP addresses, tracks domain addresses associated with malicious activity, and tracks recent activity. This feed maintains 40 different categories for IPs and URLs as well as continuously updated confidence scores. See https://rules.emergingthreats.net/.
- **FBI InfraGard:** A collaboration between the FBI and private sector to share threat intelligence. There are 16 specific categories of infrastructure identified by the Cybersecurity and Infrastructure Security Agency (CISA), which include communications, energy and nuclear power, chemicals, healthcare, IT, transportation, emergency services, manufacturing, and financial. See https://www.infragard.org.
- **Dan.me.uk:** Offers a collection of tools. Some useful intelligence that can be pulled from this resource includes the TOR status of IP addresses, DNS blacklists, IP address checking for autonomous systems, and node lists. See https://www.dan.me.uk.
- **CINS Score:** Scores IP addresses according to their trustworthiness. One useful resource is the "CINS Army List" representing a list of malicious IP addresses that can easily be added to a blacklist within a security

product. See https://cinsscore.com/.

- **Blocklist.de:** Focuses on server attacks from SSH, FTP, email, and web server sources. Blocklist.de states “We report more than 70,000 attacks every 12 hours in real time using Whois (abuse-mailbox, abuse@, security@, email, remarks), the Ripe-Abuse-Finder, and the contact-database from abusix.org so we may find the abuse-address assigned to the offending host.” See https://www.blocklist.de/en/.
- **AlienVault Open Threat Exchange (OTX):** Free community-based project that monitors and ranks IPs by reputation. On this website you will find alert feeds called “browse pulses” that you can manually enter into a system to index attacks by various malware sources. Most pulses are automatically API-generated and submitted via the OTX Python SDK. Figure 7-12 is an example of a pulse covering SSH bruteforce logs. See https://otx.alienvault.com/browse/pulses.

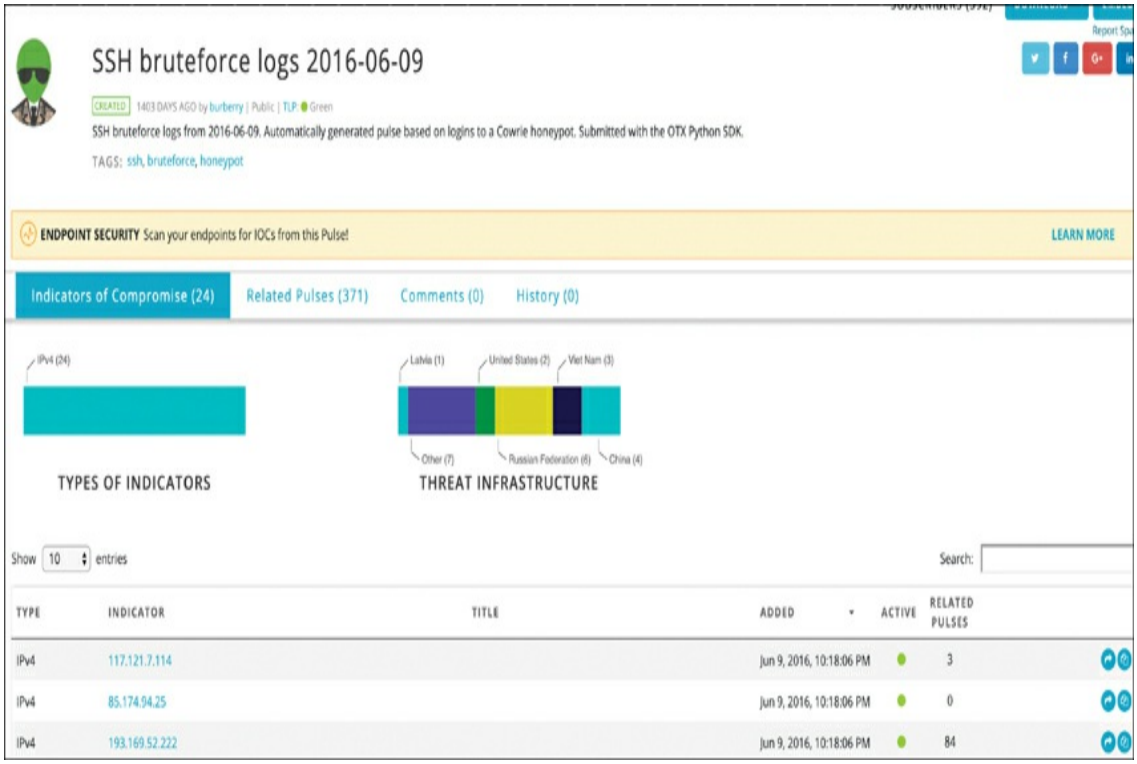

FIGURE 7-12 SSH Bruteforce Logs Example

- **Abuse.ch Feodo Tracker:** Focuses on botnets and command-and-control

infrastructure. The available blocklist is a combination of several blocklists, with specific attention paid to Heodo and Dridex malware bots (the name of the feed is a reflection of an older Trojan virus called Feodo). At the time of writing, there were 5,374 entries. See https://feodotracker.abuse.ch/.

- **Abuse.ch URLhaus:** Offers a depository of malicious domains tied to distributing malware. The database can be accessed via an API, allowing you to download CSV collections of flagged URLs, each website's respective status, type of threat associated with a website, and much more. Data goes back 30 days. See https://urlhaus.abuse.ch/.
- **IBM X-Force Exchange:** Has many features that are free to guests and registered users. It offers IOCs that are tagged by threats and threat actors. The site does require paid access to some premium features such threat reports and API access. See https://exchange.xforce.ibmcloud.com/.
- **Cyber Threat System from FortiGuard Labs:** Allows threat researchers to view and share threat intelligence and IOCs with Fortinet's FortiGuard Labs. The system is free to anyone who registers. It allows users to search via threat, threat actor, cryptocurrency access, and dozens of other IOCs. The API is provided to any user free of charge to help automate exchanging of information. See https://cts.fortiguard.com.

At this point, you are able to properly collect and process threat data. What that means is you now have a bunch of threat data, but that doesn't mean you have threat intelligence, because you have not set up any action to be taken or evaluated the data's context. Remember that reviewing the context of data so that action can be taken separates threat data from threat intelligence! The next step is figuring out how to convert data into action, which is the true judge of the value of a threat intelligence project. I call this creating *actionable intelligence*.

Actionable Intelligence

Once you have collected and processed threat data, the next phase is to transform that data into actions. As stated earlier in this chapter, the value of threat intelligence is not what is collected but how it is used and understood.

Converting captured data into actional intelligence can be broken into three steps:

Step 1. Perform analysis against your goals for adding the threat intelligence. Adjust or remove the intelligence if goals are not obtained.

Step 2. Summarize those results to stakeholders to ensure the results are acceptable, known as the dissemination phase.

Step 3. Collect any feedback and tune how the threat data is used, as these three steps are repeated to continuously improve the impact from leveraging the threat data resource.

Figure 7-13 provides a simple flowchart showing this process, starting with collecting data and ending with receiving feedback so that adjustments can be made to how the data is used.

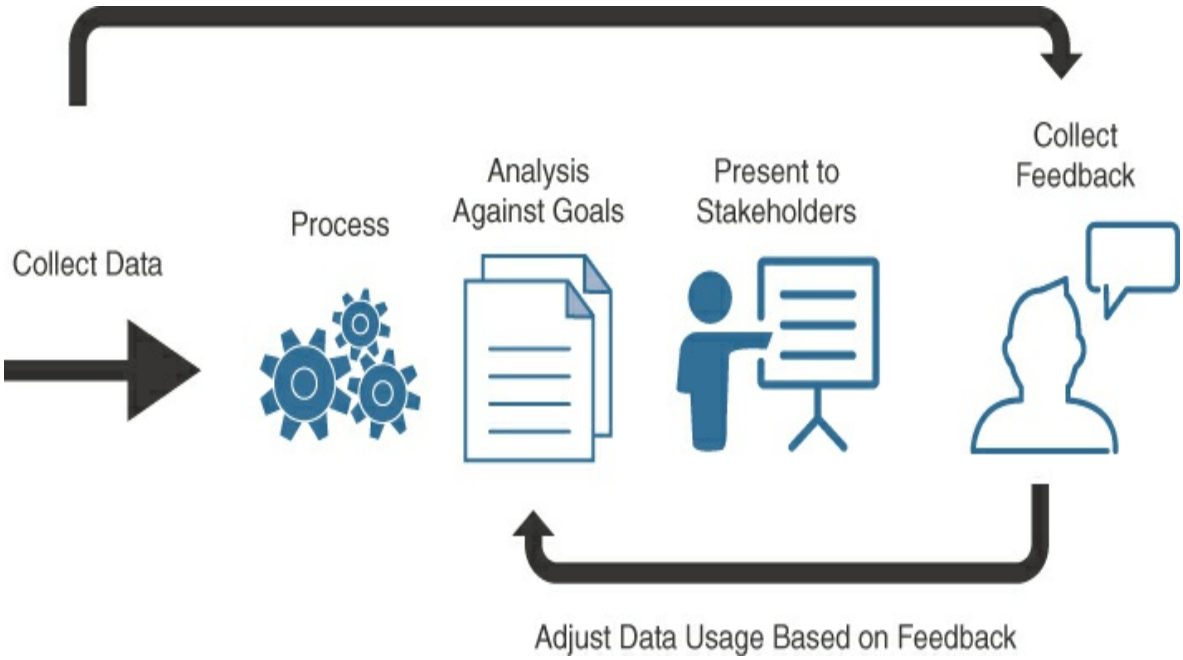

FIGURE 7-13 Actional Intelligence Flowchart

The following section introduces some popular security tools that leverage threat data. I will identify which capabilities can be impacted and what value can be obtained by leveraging external data and how that impact can occur

based on common attributes that can be pulled from threat sources. The specific configuration and tuning to make these actions work will vary on the type of data feed and technology being used.

Security Tools and Threat Intelligence

It is common for a SOC to leverage technical threat data to improve the detection and prevention capabilities of security tools. To better understand these use cases, I will review general categories of security tools and explain how threat intelligence can enhance their functionality. Most of these tools rely on technical threat intelligence, but some may also leverage elements of tactical and operational threat intelligence, which will also be highlighted.

In general, threat data will help security tools answer questions that can't be answered without an outside viewpoint of the threat landscape. Example questions that you hope to answer by using threat data include the following (most of these questions can't be answered using internally generated event data):

- Who is attacking the organization?
- What are the attacker's motives?
- What is their target?
- What tactics, techniques, and procedures (TTPs) are being used?
- What indicators of compromise should the SOC's tools look for?
- What actions can the SOC take to reduce the risk of exploitation?

As the SOC feeds threat data into security tools, the SOC will phase in how it is used regardless of the data that is contained within the threat intelligence feed. The first step is to monitor for mentions of the IOC or artifact associated with attackers. This will allow for testing to validate that the data is processed correctly and that the results meet the goals for adding the threat data. Earlier in this chapter, I provided an example of Splunk consuming data; however, that data was not properly processed, which means actions can't be taken. You will want to validate that the security tool is properly reading and using the data before you test any actions.

After the SOC validates that the security tool is properly consuming the threat data, the next step is to attempt to review results from before and after you added the threat data. Your SOC's goal is to confirm that the data is being used and is positively impacting how the system functions. Some security tools such as SIEMs have live widgets and reports that can be broken or tainted with the wrong data elements if external threat data is not properly utilized and filtered upon collection.

The SOC will want to adjust how the data is collected and make other adjustments according to the system's configuration options to modify results until the SOC has the best results for the SOC's and organization's business goals. Looking back at the Splunk example that didn't process the data from the security tool correctly, I had to adjust how the system formatted the data from that system until Splunk was able to properly categorize the events. At that point, I found Splunk would offer many new variables I could access based on various IOCs obtained from the threat data. Some items such as the timestamps were a bit off and required adjustments to how the collector within Splunk was configured, but eventually I got the data processed correctly so that I could quickly identify what I needed to find.

Once the integration of the threat data is configured correctly and tested, the final step is to add the results to the SOC's operation through analysis. Tools such as SIEMs will allow the SOC to create widgets that will continuously collect and display results pulled from threat intelligence feeds, such as current threats or top hashes of malware seen in the wild. The analysis stage is essentially when the SOC converts the collected threat data into threat intelligence, since actions are now being developed based on the results of the new data input. The SOC can also develop and test playbooks, which is a topic related to SOAR and will be addressed shortly. Figure 7-14 shows a diagram of steps to operationalize threat data into threat intelligence.

FIGURE 7-14 Steps to Operationalize Threat Data

To better understand how this transformation occurs, let's review how threat data could be used by popular security offerings. I start with a review of the impact on SIEM technology.

Security Information and Event Management

A common topic covered in the security industry is centralizing data from multiple tools. Sometimes, this is labeled as having a "a single pane of glass." The goal is to reduce the time to investigate an incident by having one place to perform the investigation and having the ability to correlate data from different tools to gain a better understanding of threats impacting your organization. SIEM technology attempts to solve this challenge by acting as that central point for event data, as I described earlier in this book. Market leaders for SIEM solutions include Splunk, QRadar, and LogRhythm according to sources such as Gartner. The goals of using a SIEM solution are to improve attack detection, speed up incident handling, centralize reporting, and provide a resource to measure compliance. With that in mind, where does threat intelligence fit in?

It is common for a SIEM solution to digest and correlate findings with internal telemetry, such as what you get from firewall and DNS logs. This allows you to match potential attacks to the external data that was collected. The value of a SIEM solution depends on the data it receives. If you send it limited data, you will get limited results. If the data sent is not good, the output will also not be good. Another way to say this is "garbage in, garbage out." If you don't configure it to present the information in your desired format, you will end up with dashboards alerting about hundreds of events

you will never be able to address, known in the industry as the “bug splat.” A SIEM dashboard throwing out thousands of alarms at your analysts will just overwhelm them, causing key alerts to be missed and the SIEM to provide little value. This brings us to the pros and cons of threat intelligence for a SIEM. I will start with the cons.

Adding threat data to a SIEM tool might seem like a simple process, but it is not. The purpose of the SIEM tool is to piece everything together (hence the single pane of glass, right?). If you mix in the wrong external threat data, you will cause a ton of misunderstandings within the SIEM tool’s built-in logic because the SIEM tool won’t be able to determine what is internal and external, forcing the SOC to either re-engineer everything or look at methods to isolate external data, leading to limitations in data correlation. I have found that some SOCs will light up threat intelligence, get trampled with tons of new alerts they can’t take action on, and soon after disable the threat data. This occurs due to an increase in false positives, contradictions in correlation data, and just more things to look at.

One other con I have seen is the complete opposite situation, where a SIEM solution is not generating any new alerts after the threat intelligence feed is added. This can occur if the threat data is not applied correctly and either forgotten about or the SOC didn’t find any major changes and felt the feed didn’t provide value. When this occurs, I tend to find that a scope or objective wasn’t set, meaning there wasn’t a set success criterion to measure against. The SOC simply lit up the feed and hoped for what would feel as “better” results. This is a common mistake when testing threat intelligence. The SOC should not just add a threat intelligence feed and generally look for some major incident they didn’t know about to pop up. In the real world, that doesn’t happen by automatically adding a threat intelligence feed to a SIEM.

Adding Threat Intelligence to a SIEM

Given that just dumping threat data into a SIEM tool will cause more problems than value, does that mean threat intelligence is not ideal to use in a SIEM tool? The answer is no—if you apply data with a specific goal and expected outcome. I covered how to evaluate threat data, with a focus on ensuring that the data is reliable, timely, and accurate and is relevant to your line of business. There are a handful of additional checkpoints that can help

with threat data integrations with a SIEM tool. Consider the following items as you evaluate threat data for your SIEM solution:

- Validate the threat data is in a format accepted by the SIEM solution of choice. If it isn't, how much effort is required to manipulate the feed so that it would be accepted?
- Will this new data source increase the monthly/annual SIEM bill? Many SIEM providers charge on a usage billing cycle, meaning the more data used, the higher the cost for using the SIEM solution. You could reduce the impact by filtering only what is relevant to your organization and using other tuning tricks. If you are billed by your SIEM provider, consider the impact of adding more data.
- What are the top threats you plan to address? Answering this allows you to focus the results of the threat data to specific goals that matter to your SOC. For example, are you a target for a nation state? Is phishing a top concern? You can use the SIEM tool to develop reports and live displays that digest threat data and answer these questions. In particular, you will find that some threats are better addressed by using external threat data while others are easier to detect using data from internal security and network tools. The two previous examples of phishing and nation state concerns are both better suited for using external threat data.
- Does the SIEM tool offer a way to capture additional context about events? This is where support for tactical and operational data will be extremely handy, allowing you to correlate attack data with additional details about the who, what, when, where, and how of the attack. Having context will make decisions regarding what action to take much simpler. This is especially true with SOARs and SOAR/SIEM integrations.
- What filters are available in the SIEM tool, and can they be applied against the threat data? The more you can focus on what is relevant to your specific needs, the more useful the threat data will be.
- Would the threat intelligence improve the confidence of existing detection capabilities? This is a huge question to answer since the SIEM tool is pulling in data from various internal security and network tools. As pointed out earlier, if adding external data weakens the SIEM tool's decision process, this will reduce the SIEM's confidence in alerts being generated, causing a breakdown of the value it provides. One common

method to overcome this is to be selective about which checkpoints/widgets within the SIEM tool are using the external threat data. The bad news is that it will require some re-engineering of certain widgets and reports if filters were not put in place regarding adding or removing external event data prior to when the reports and widgets were originally created.

Threat data can be converted to threat intelligence using a SIEM solution if it is added in a planned and meaningful manner. It is critical that you follow a solid rollout plan to ensure you maximize your value received while also avoiding any losses from adding the feed. If you just add threat data without any set goals or considerations for how the SIEM solution is currently being used prior to the threat data, you will run into problems with your deployment.

In summary, your rollout plan should include the following steps. Many of these steps follow the best practices I have covered in this chapter.

Step 1. Set objectives for using the threat data. What is your measurement for success?

Step 2. Configure the SIEM solution to accept the feed.

Step 3. Monitor that the data is being collected correctly.

Step 4. Identify if existing reports and live widgets are negatively impacted. If so, add filters to remove the external data from existing reports and widgets.

Step 5. Attempt to identify objectives using live feeds against filters that include the new threat data.

Step 6. Tune how the data is digested and troubleshoot any collection issues until you can identify your goals.

Step 7. Operationalize your findings in widgets and reports.

Step 8. Build the new use cases into your SOC practice and SOAR solution.

Security Orchestration, Automation, and Response

One major drawback of a SIEM solution is its limitations in what actions it

can take against an event. This is where a security orchestration, automation, and response (SOAR) solution steps in. A SOAR solution can provide case management, standardization, workflow, and analytics, all of which enable the SOC to be much more productive. Without a SOAR tool, a SOC would be left with a slew of alarms, leaving the SOC with the responsibility to manually investigate and track how events are being handled.

The benefits of using threat data for a SOAR tool are slightly different than those for a SIEM tool. One benefit is the impact on how playbooks are used. Having external data can be extremely useful for this purpose. Playbooks can include additional triggers that are impacted by threat data, allowing a SOC to take more proactive measures when attack campaigns are being seen in the wild. Many SOAR providers have playbook templates that leverage both internal and external threat data; not including a threat intelligence feed would limit usage of such playbooks. External threat data can also add confidence in when a playbook is triggered by adding additional context or checkpoints that must occur before the playbook is launched.

Other benefits from using external threat data with a SOAR tool are similar to those of using external threat data with a SIEM tool. Those benefits include improvements to dashboards and reporting and improvements to incident management and response. The assumption is, however, that the same considerations are made regarding choosing which data and how it is used as I covered in the SIEM section. As with a SIEM tool, you must follow a phased-in approach to adding threat intelligence to a SOAR tool or your capabilities will break, data within the SOAR tool will become tainted with false positives, and the overall value of the SOAR tool will be negatively impacted.

I highly suggest following the same rollout plan for adding threat intelligence to a SOAR tool as you would for a SIEM tool. The only difference will be identifying any default playbooks or other SOAR capabilities that are designed for leveraging threat intelligence and adding those to your evaluation plan to simplify your end results. Many SOARs likely have default playbooks that are similar to your goals, allowing you to have a starting point rather than developing each playbook from scratch. I recommend to first test default playbooks, which are built for leveraging external threat data, before creating your own. Once those default playbooks are capable of digesting the data, then attempt to modify the templates or

build your own, knowing that the threat data is captured correctly and available within the SOAR tool.

I recommend picking a handful of testing criteria as you add threat data to a SOAR solution. Those testing criteria should fall under three categories: playbooks, dashboards, and reporting. I also recommend to first test the impact of data within the SOAR dashboard before attempting projects associated with playbooks, reporting, or other automated tasks. The specifics of how to carry out testing will depend on the SOAR solution, other security tools, and your business objectives. I cover concepts such as SOAR technology and playbooks in much more detail in Chapter 10, "Data Orchestration."

Email Security

Email continues to be the number one platform used for delivering cyberattacks. Many of the threats such as spam are unique to email, and an email security gateway platform is essential to properly protect your email traffic. Solutions can be delivered as an appliance or as a cloud service and are designed to evaluate email traffic as it enters and leaves an organization.

The value of an email security solution is based on the capabilities it provides. Common features include the following:

- **Anti-spam:** Viewing email trends and the characteristics within the email to determine if it is spam
- **Antivirus:** Scanning attachments for malware
- **Outbreak prevention:** Identifying and preventing an outbreak of unwanted email
- **URL analysis:** Evaluating URLs that could send users to malicious web resources
- **Policy enforcement:** Evaluating the language and context of the email
- **Data loss:** Validating that compliance is not violated, such as by sending sensitive information or information that violates compliance

Many of these features could benefit from being provided external threat data properly. Antivirus could benefit from having additional artifacts and hashes

to validate within emails. Spam, URL, and outbreak prevention capabilities could benefit from obtaining blacklists of malicious web resources beyond what is provided by the vendor's updates. Any of the benefits will depend on the quality of threat data and how the vendor is able to process it.

Cisco Email Security Appliance (ESA)

An example use case would be using Cisco Email Security Appliance (ESA) based on AsyncOS 12, which supports external threat intelligence using STIX and TAXII. The benefits listed by Cisco include

- Proactively respond to cyberthreats such as malware, ransomware, phishing attacks, and targeted attacks
- Subscribe to local and third-party threat intelligence sources
- Improve the efficacy of the Cisco Email Security Appliance

These benefits are obtained by specific STIX indicators of compromise that are supported by Cisco ESA in this version of software. All other STIX data is not supported and used by this release. Those STIX IOCs include the following:

- File Hash Watchlist (describes a set of hashes for suspected malicious files)
- IP Watchlist (describes a set of suspected malicious IP addresses)
- Domain Watchlist (describes a set of suspected malicious domains)
- URL Watchlist (describes a set of suspected malicious URLs)

Similar support is provided in other market-leading email security technologies such as those from Proofpoint and Mimecast. How and what is supported will vary on the model and code version being used by the vendor of choice. Most capabilities that leverage threat data for email security are blacklist- and detection-oriented, meaning the enhancements will improve detection and blocking of threats.

Deploying Email Security Threat Intelligence

The best way to roll out a threat intelligence enhancement to an email

security tool is to focus specifically on which features you plan to impact and evaluate their effectiveness before and after the threat intelligence is added. You should also leverage sample threat data to test if alerts are triggered that match what is included in a threat intelligence feed, such as sending emails that match threat data email artifacts or attaching files that should trigger the anti-malware capabilities. Blacklisted URLs can be added to emails and sent to see if the threat intelligence feed has updated the email security technology to identify that a malicious URL has been included in the email. Without properly testing the impact of the threat intelligence feed, you will not know what impact the threat intelligence feed has to the email security technology.

Note

I find that many of my customers spend time and money adding threat intelligence to their email security solutions but then do not test the impact. They assume value is being obtained when the feedback is successfully uploaded. I always recommend testing that value exists.

Many other security tools such as anti-malware, intrusion prevention, and sandbox analysis solutions are all starting to add options for leveraging external threat data. The goal is to improve detection and prevention capabilities by adding additional context about threats, adding blacklisted external resources, and including more IOCs to look for when examining traffic for threats. Your SOC's goals for supporting threat intelligence should include offering the capability to digest an industry-recommended format such as STIX, the capability to learn more about an event if it is triggered by the security tool, and the capability to judge the impact the threat data has on the capability within the technology.

Feedback

The final step in a threat intelligence project is evaluating the results of adding the threat data. This is where the return on investment (ROI) is weighed against the cost to obtain threat intelligence resources to determine

whether value is seen or the threat intelligence project should be dissolved. Simply saying “The SOC sees more threats” will not be good enough to justify additional investments in dollars and human resources in the mind of a decision maker.

Looking back at the four types of threat intelligence (strategic, tactical, operational, and technical), different factors will apply to technical and nontechnical ROI. For nontechnical threat data, I pointed out how it is common for nontechnical threat intelligence to be delivered in a report or presentation. These reports and presentations are developed upon request, which the ROI can be weighed against if the decision pending on the research was impacted by the report. The feedback loop approach covered in this chapter dramatically increases the chance of having an impactful deliverable because the requestor has many opportunities to adjust the approach taken by the analyst as the strategic threat intelligence report is being researched and developed.

Technical threat intelligence is more outcome-oriented and can be tied to specific use cases. ROI can be judged on whether the use case is achieved with the new data, the amount of effort to maintain the feed, and value from having the use-case objectives met. Here are some checkpoints to use when judging the ROI of adding technical threat intelligence to your practice:

- Are you seeing alerts that include the new data? Are new threats now blocked? Have you tested to ensure such threats would be blocked? What would occur if such threats were not blocked and what would the cost be to the organization?
- What additional context is obtained outside of a block? Is this context obtained from the new threat intelligence feed and, if so, is it valuable?
- Does the new feed improve your incident response program? Could these improvements be achieved without the feed, and have you tested the response before and after the feed?
- Are the SOC and tools more proactive and aware of recent events?
- Are new capabilities such as new playbooks now available due to adding the threat intelligence?

These questions allow you to place a value on the new feed, which can measure its return on investment. You always want to question the accuracy,

relevance, and timeliness of the data to ensure you are adding quality data. You also want to identify overlap as you expand your threat intelligence usage, to ensure you maintain variety in data. As success is found within one team, more value can be obtained by meeting with other teams within the SOC and attempting to address their needs with the new threat data. For example, an analyst might initially acquire threat data for updating security tools, but the incident response team may find value by adding the threat intelligence to their tools as well. Expanding value allows for increasing the demand, which in turn increases the available funding for threat intelligence projects. This same concept holds true for enhancing reporting that is provided to nontechnical leadership. The more buy-in you can achieve, the more support you will obtain for your project.

Summary

The focus of this chapter was to understand the four categories of threat intelligence and how to benefit from each: strategic threat intelligence, tactical threat intelligence, operational threat intelligence, and technical threat intelligence. This chapter first explained the difference between threat data and threat intelligence, as this is a commonly misunderstood concept. Next, this chapter reviewed the lifecycle of technical and nontechnical threat intelligence, starting with how to plan and evaluate different threat intelligence options. You next learned how to collect and process threat intelligence from your chosen provider and then operationalize the results by improving the capabilities and data within popular security tools. The chapter wrapped up with methods to tune and improve the data collected to maximize the benefit of threat intelligence.

The next chapter looks at best practices for threat hunting and incident response.

References

Kropotov, V., Yarochkin, F. 2019, (July 30). Hunting Threats on Twitter:

How Social Media Can Be Used to Gather Actionable Threat Intelligence. Trend Micro. https://www.trendmicro.com/vinfo/us/security/news/cybercrime-and-digital-threats/hunting-threats-on-twitter

McMillan, R. (2013, May 16). Definition: Threat Intelligence. Gartner. https://www.gartner.com/en/documents/2487216/definition-threat-intelligence

MISP. (n.d.). MISP – Open Source Threat Intelligencer Platform & Open Standards for Threat Information Sharing. MISP Project. https://www.misp-project.org/datamodels/

OASIS Open. (2020). Introduction to TAXII. OASIS Open. https://oasis-open.github.io/cti-documentation/taxii/intro.html

OASIS Open. (2020). Sharing Threat Intelligences Just Got a Lot Easier! OASIS Open. https://oasis-open.github.io/cti-documentation/

Pokorny, Z. (2019, May 16). Cyber Threat Intelligence Feeds: Data Automation Solution. Recorded Future. https://www.recordedfuture.com/threat-intelligence-feeds/

SophosLabs. (n.d.). Threat Prevalence Definition. Sophos. https://www.sophos.com/en-us/threat-center/threat-monitoring/threat-prevalence-definition.aspx

Chapter 8

Threat Hunting and Incident Response

If some animals are good at hunting and others are suitable for hunting, then the gods must clearly smile on hunting.

—Aristotle

One of the staple services found in mature security operations centers around the world is incident response. Every organization has the expectation that the SOC will jump into action when a cybersecurity incident occurs. Who do you call when malware is found? The SOC! Who do you call when phishing attacks are identified? The SOC! Who do you call when data has been stolen? The SOC! The SOC is the organization's defense against cyberthreats. The key to the SOC's success when initiating its incident response service is ensuring the right processes are followed, as each incident will require different steps to provide the appropriate response.

This chapter looks at how mature SOCs provide incident response services. Responding to an incident includes the entire process of identifying the threat, confirming the threat is real, deciding which playbook to launch, containing and eradicating the threat, recovering from the incident, and following through with post-incident response steps such as hosting a lessons learned session. Responding to a malware infection, for example, does not just involve wiping the impacted systems; it requires determining how one or more of those systems were infected and whether other systems might also be infected. Responding to phishing emails doesn't just mean blocking the senders of the malicious emails—the next phishing campaign will come from another location, function differently, and can get through your defenses if you don't learn from the previous attack. Even when you respond strongly to

a security incident, that doesn't mean the same response will work the next time, which emphasizes why it is critical to always perform a lessons learned review to figure out how to improve after each incident.

This chapter covers how to deliver end-to-end incident response services, including how to identify, contain, and remediate the most advanced persistent threats. The chapter also covers what is involved when legal actions are expected to be taken as part of the response, which will require a digital forensic investigation. There is a lot to cover, so let's start this incident response journey by establishing a foundation of what the term security incident actually means to the SOC operation.

Security Incidents

A *security incident* is an event that leads to a violation of an organization's security policies and puts sensitive data at risk of exposure. A data breach is a type of security incident. A malware outbreak is a security incident. Destroying a server is a security incident. Preventing access to a system using a denial-of-service (DoS) attack is a security incident. A security incident isn't necessarily an intentional act. For example, an administrator misconfiguring a network device, causing part of the network to go down, is a security incident. Even privacy violations as defined by the U.S. Department of Homeland Security (DHS) are considered a security incident, making all privacy-related events security incidents. All of these situations would need to be reported to the SOC so that it can initiate the appropriate incident response.

The result of a security incident is *typically* (*but not always*) a security breach. By contrast, a security breach is *always* the result of a security incident. An example of a security incident that doesn't necessarily result in a security breach is a denial-of-service attack, because a DoS attack does not itself breach a system or network (although it could be part of a broader attack that exploits the service denial to breach a system or network). I point out the distinction between *security incident* and *security breach* to make clear that the SOC must treat all security incidents as potential security breaches but should not assume that a breach is always going to or has occurred. It is also critical to identify and understand all security incidents

that are associated with a security breach. The cause of a security breach is commonly multiple different security incidents chained together. Not understanding all of the security incidents that led to a breach opens the possibility that certain vulnerabilities that were exploited by a security incident were not identified during the incident response and are still exposed to potential future attack.

> **Note**
>
> Pay attention to the language used when you read about a security breach in the news. It is common that multiple security incidents related to different attacks or failures in security are discovered that led to the larger breach. If any of those security incidents had been prevented, the larger breach may not have occurred!

Incident Response Lifecycle

There are many guidelines for incident response best practices that present a similar lifecycle model for responding to an incident, but they tend to use different terminology to define the lifecycle phases. For purposes of this chapter, I'll present the popular, industry-respected incident response lifecycle offered in NIST SP 800-61 Rev. 2, *Computer Security Incident Handling Guide*. You should be able to adapt the material presented in this chapter to any other similar incident response lifecycle model that your SOC may choose as guidance. Figure 8-1 shows NIST's incident response lifecycle model, the four phases of which are preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity.

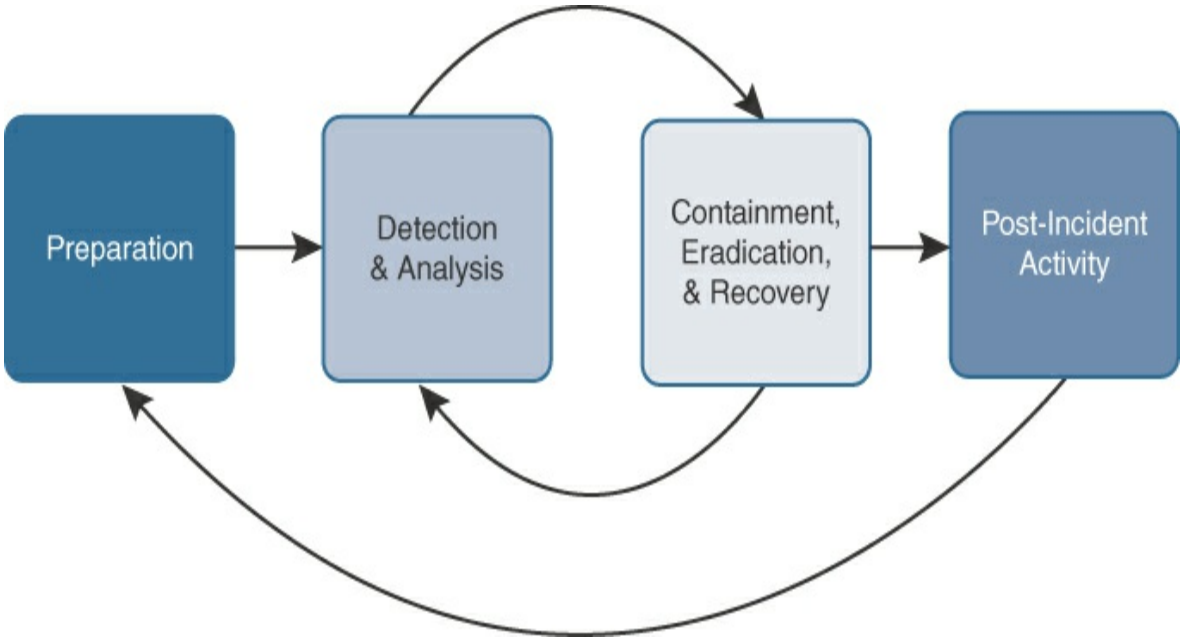

FIGURE 8-1 NIST SP 800-61 Rev. 2 Incident Response Lifecycle

The rest of this chapter guides you through the details associated with each of these four lifecycle phases. Notice in Figure 8-1 that these phases are interrelated, as indicated by the arrows showing how a SOC can move back and forth between phases. The preparation phase relates to all steps performed prior to launching the incident response service. The detection and analysis phase involves looking for security incidents, while the containment, eradication, and recovery phase involves the SOC responding to a threat. If there is a possibility that more threats exist, the SOC returns to conduct more detection and analysis. Once all incidents have been eradicated, post-incident activity occurs with the goal of learning from the experience and developing plans to improve the overall incident response program. Those recommendations are sent back to be incorporated into the preparation phase, as represented by the arrow pointing from the post-incident activity phase to the preparation phase.

The following sections take a closer look at the phases of the incident response lifecycle and how to incorporate the lifecycle into a SOC's incident response program.

Phase 1: Preparation

Reading this chapter would fall under preparing for an incident because you are learning about building an incident response program. Preparation is defined as developing the SOC's incident response capabilities, including identifying the required people, process, and technologies and developing how they will operate during an incident. Preparation also involves all work designed to reduce the risk of a security incident, essentially proactively responding to potential security incidents. Examples of such work include what the SOC does regarding vulnerability management, implementing security tools, enforcing strong security polices, and taking any other measures to reduce the chance the SOC will need to respond to an incident. As the incident response program matures, more time can be spent in the preparation phase, allowing for training as well as proactive security measures to be deployed and tuned. Preparation is an ongoing task and takes into consideration what is obtained from a lessons learned report developed after an incident occurs, in the post-incident activity phase.

Assigning Tasks with Playbooks

The first step for a SOC to prepare for an incident is to establish who is responsible for what part of the service. As introduced in Chapter 3, "SOC Services," the Incident Response Consortium (IRC) offers free playbook templates (https://www.incidentresponse.com/playbooks/), which are a great baseline for understanding different steps involved in responding to different types of incidents. Playbooks represent each step that is executed during an incident. Playbooks are also known as a SOC's procedure documents. I will be referencing the IRC's templates throughout this chapter.

The playbooks provided by the IRC are structured in the following different phases of an incident response service:

Note

The IRC's playbook phases are based on the NIST incident response

process

1. **Prepare:** How to prepare an incident response service
2. **Detect:** How to detect the incident
3. **Analyze:** How to analyze the threat and validate if it is indeed a threat
4. **Contain:** How to prevent further spread of the incident
5. **Eradicate:** How to remove the threat
6. **Recover:** How to return systems back to operational state
7. **Post-Incident Handling:** What was learned from the incident and how to improve future incident response to similar incident

Chapter 10, "Data Orchestration," includes more details about enhancing playbooks, including how to apply automation and other scripting capabilities to improve how playbooks are executed. Figure 8-2 shows the IRC Malware Outbreak playbook template for the prepare phase of responding to a malware outbreak.

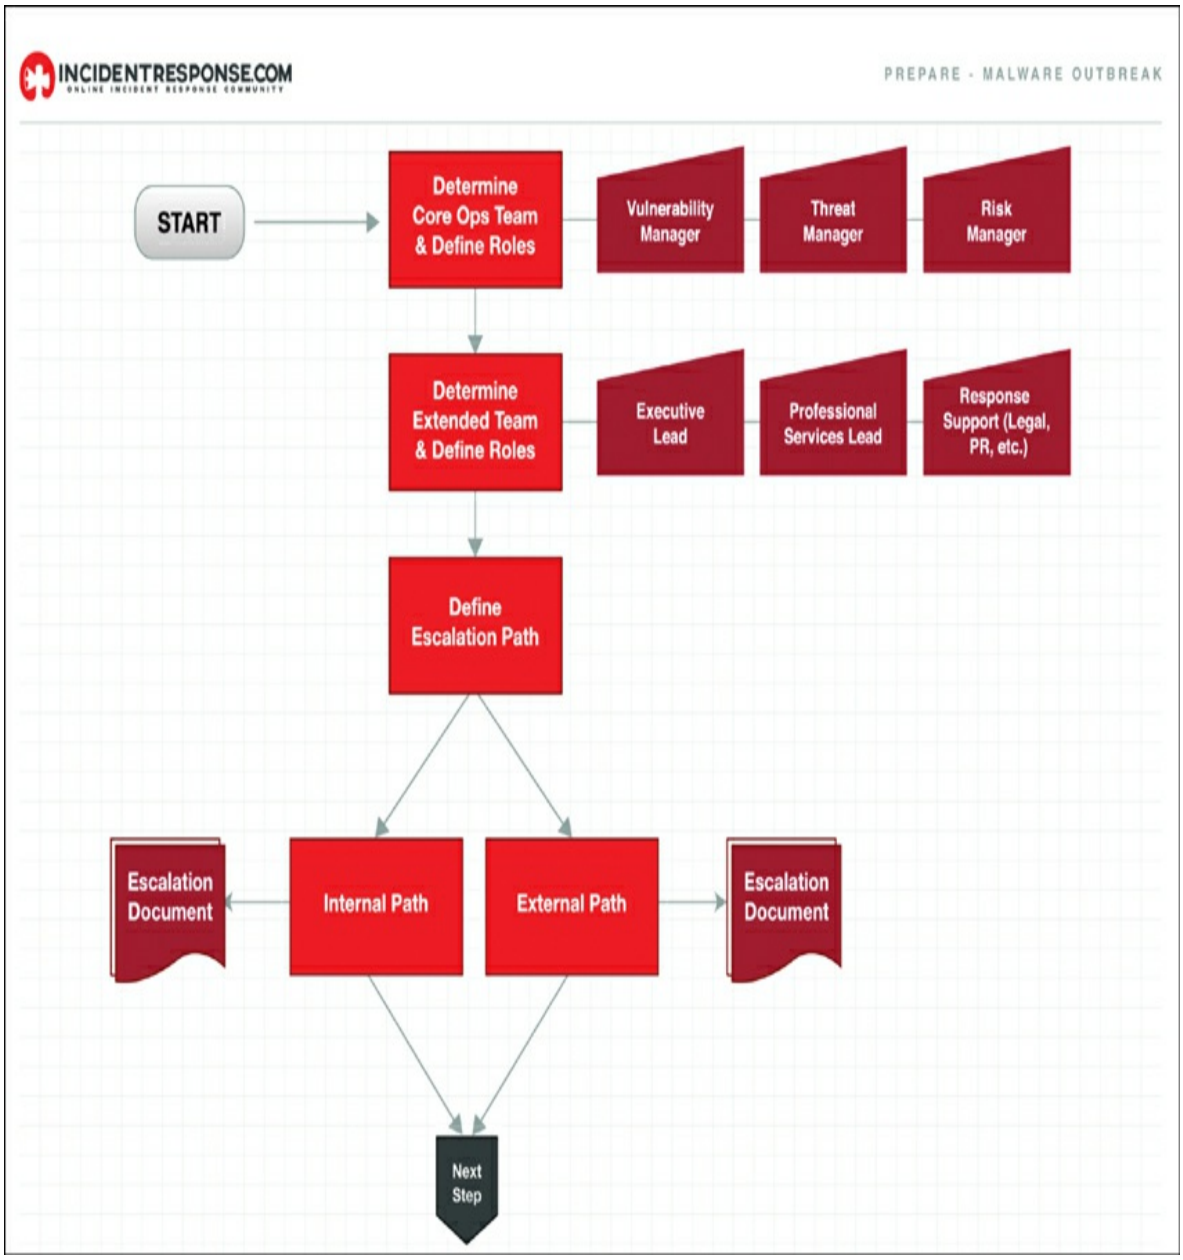

FIGURE 8-2 Incident Response Consortium Malware Outbreak Playbook Example

The Malware Outbreak playbook summarizes the prepare phase for responding to malware as doing the following:

1. Determine the core operations teams and define their roles.
2. Determine all extended team members and define their roles.

3. Define the escalation paths, meaning how the different tiers of support will operate.
4. Build polices regarding how internal and external escalation of an incident will flow.
5. Document all policies.

A quick summary of the prepare phase as described by the IRC is to first establish who will lead the incident response practice based on filling critical job roles related to management. Once those roles are filled, team lead roles need to be filled, rounding out the entire incident response leadership team. After the team is established, escalation paths for both internal and external resources need to be established so everybody working within or with the SOC incident response service knows their responsibilities and how to move a case along as the entire lifecycle of the incident is handled.

It is common for job roles to break up the order of responsibilities for reacting to an incident into a tier format. The industry calls this a *tier support model*.

Tier Support

The Malware Outbreak prepare phase template provides recommended job titles; refer to Chapter 4, "People and Process," to better understand what skills and experience are associated with those job titles. The escalation paths represent how a case will be passed between different teams that could have various job titles and responsibilities depending on what service is required to handle the incident. Some incidents will require digital forensic services, while others will require advanced malware analysis. This is why I recommend to always have some form of all eight of the critical SOC services outlined in Chapter 3, regardless of whether those are internal SOC services or third-party services, to ensure your SOC's incident response service has options available for handling any type of security incident.

Four-Tier Model

It is common for SOCs to use either a three- or four-tier model for

responding to incidents. Certain job roles will fall within each tier. Those tiers can be defined as the following:

- **Tier 1 (triage):** Tier 1 represents the front line when reporting security incidents. This level of support focuses on reviewing and assigning urgency to potential threats. It is common for entry- to low-level SOC members to start their careers in the SOC supporting tier 1. To help adjust for the lack of skills and experience, playbooks and other orchestration concepts are enforced.
- **Tier 2 (incident response):** Tier 2 analysts address the tickets created by tier 1; hence, tier 2 can be the beginning of the investigation depending on whether tier 1 has any responsibilities outside of opening cases. Some SOCs will have tier 1 qualify tickets before escalating them, while others have tier 2 review everything, meaning the starting point for tier 2 will be different depending on the SOC's incident response processes. Tier 2 analysts leverage security controls, policies, and intelligence to determine the scope and origin of the attack. It is common practice for tier 2 to execute mitigation, recovery, and remediation tasks.
- **Tier 3 (proactive cyber defense):** Tier 3 takes on tasks related to proactive measures and is also called on when the need for advanced threat hunting is required. Proactive measures include looking for vulnerabilities and access points into the network with the goal of preventing future attacks. Tier 3 also validates whether regulatory compliance, governance, and auditing is met before, during, and after a security incident. It is common for tier 3 members to be part of multiple SOC services, including compliance and vulnerability management teams. Advanced threat hunting can include the need for higher access to certain tools, the need for specialized experience, or additional services such as advanced analysis or forensic skillsets.
- **Tier 4 (operations, controls, and management):** The top tier represents those responsible for the incident response service as defined in the pre-incident documentation as management and team lead roles. These members oversee all aspects of proactive, active, and post-incident processes. Major crises will be brought to this level's attention, including events that could impact the entire business or a large group of customers. For example, if a major data breach is identified, tier 4 will be engaged to

handle everything from the internal/external messaging to how the entire incident response program will adjust its processes to reduce the risk of a future similar incident.

The incident response team will also engage other teams based on what resources are required for a specific incident and how the escalation processes are written up, as alluded to in the preceding description of tier 3 support and how members may be part of other SOC services. An incident can involve the need for a forensic specialist, which might exist within the SOC's digital forensics service team or could be an external party that is brought in during the incident. You will need to account for all of these relationships as you build the processes for your overarching incident response plan and, later, specific playbooks based on different types of incidents. This is why every SOC must have some form of the SOC services described in Chapter 3!

Communication

As you can see from the previous playbook example, an incident response service involves many parties as different types of incidents are properly addressed. Organizations need multiple lines of communication not only for engaging the incident response team but also to communicate with internal and external resources during the incident response process. That communication must be secure because many of the topics discussed will involve weaknesses within the organization that could be exploited if their existence were exposed to unintended parties. Chapter 2, "Developing a Security Operations Center," covered different forms of developing secure forms of communication

Another communication strategy you must develop is how to contact the incident response team when anybody internally or externally identifies a potential incident. The SOC can't assume that in the midst of an incident someone affected will be able to leverage computer technology to contact the SOC, making traditional forms of contact such as a phone number or publicly available message boards just as useful as internally accessible web pages. Imagine that an employee's laptop was destroyed in an incident and the employee doesn't have any digital method to alert the SOC of the incident.

Alternate options for communicating with the SOC can include phone numbers, email addresses, online forms, and secure instant messaging systems.

Key point

At least one communication method must allow for anonymous alerting. Without such an option, some people involved in an incident might not report the incident, leading to missed events. For example, the SOC could offer a website for submitting potential incidents and include the option to remain anonymous.

One proactive security-focused line of communication your business might want to consider is a specific place to point out flaws and bugs within products and services offered by your organization. Bug bounty programs encourage external parties to discover bugs before the general public is aware of them, preventing incidents of widespread abuse as well as rewarding external researchers for performing the ethical approach to responding to a potential vulnerability. Not offering bug bounty programs not only encourages those that find flaws to share them without your knowledge, but also limits those researching and improving your security to internal resources. I am a strong believer in the value of allowing outsiders to assist with your security capabilities. In order to encourage this behavior, I highly recommend offering a bug bounty program that includes how bugs or vulnerabilities can be submitted to your SOC for review.

Third-Party Interaction

Your SOC's incident response team should be prepared to interact with external third parties during the incident response lifecycle. Those outside parties include customers, media, other incident response teams (non-SOC organization resources), incident reporters, law enforcement agencies, Internet service providers, and vendors, as depicted in Figure 8-3 from NIST SP 800-61 Rev. 2.

FIGURE 8-3 Potential External Third-Party Communications During Incident Response

The incident response procedure documents need to account for all of these third-party resources. Only authorized staff should engage with external resources, to avoid communicating an incorrect or inconsistent message about an incident. It is common for those responsible for interacting with incident reporters or any news media channels to first attend training that focuses on how to limit what information is disclosed to only what is necessary and best for the organization. You do not want a random employee's statement "Things are falling apart around here!" appearing in the national news.

External support from vendors can be extremely useful but can also have an associated cost. You should identify which support services are included with the technology you own and how you can quickly obtain additional services when needed. Many security service providers offer packages that include proactive security assessment services mixed with reactive incident response services. I highly recommend seeking reactive services that if not used, can be transferred into proactive services such as assessment penetration testing or tuning services.

Lines of responsibility between the SOC and non-SOC teams need to be established up-front during the pre-incident response preparation phase. You do not want to have confusion regarding responsibilities during an incident, and problems can occur if non-SOC teams are interacting with impacted parties during a live incident. An example of a huge problem would be if the SOC determines an incident could include legal action and, before proper forensic steps can occur, a non-SOC desktop support representative re-images the impacted systems. These types of disconnections of responsibilities will lead to a breakdown of the quality of your incident response service.

Law Enforcement

Regarding law enforcement, I highly recommend your SOC to coordinate with all agencies that you might need to engage with now rather than post preparation phase planning. If you do not know how to contact law enforcement agencies that support your SOC, you need to establish those lines of communication before a real incident occurs. Call your local law enforcement and ask the following questions:

- Who should my organization contact if we experience a cybersecurity incident?
- What is the best way for our team to contact that person?
- What legal obligations are we committing to if we engage your agency?
- Are there any associated costs we should expect if we engage your agency?

- What type of events are escalated to higher-tier support within your agency?
- How are our customers and other associated parties engaged or involved in any incident we bring to you?
- What compliance or other rules of engagement should we be aware of?

Take the same approach with your service providers, vendor support, and other teams within your own organization. These details can be developed into a playbook, a topic you will learn more about in Chapter 10.

Law Enforcement Risk

On the topic of law enforcement, your SOC should address specific considerations before it uses this resource. Let's look closer at the value and risk associated with involving law enforcement in incident response.

As mentioned, one critical decision that the incident response team will make regarding certain types of incidents is whether and when to engage law enforcement. There are different reasons why you should or should not do so. There are situations in which you have to engage law enforcement, such as when you are legally required to do so, situations in which you can engage law enforcement but don't have to, and situations in which you do not want to engage law enforcement due to associated risks.

Why would you not engage law enforcement regarding an incident? The following are some things to consider:

- Law enforcement cares more about the law than your business.
- Law enforcement might seize computer equipment, for example, as evidence regardless of whether your business needs it to operate.
- The incident might not involve a potential crime and thus not warrant engagement of law enforcement.
- Engaging law enforcement can slow down the incident response process.
- Engaging law enforcement brings events into the public's view.

On the other hand, there are benefits to engaging law enforcement, such as

the following:

- Law enforcement can demand access to restricted resources if a warrant can be issued.
- Law enforcement has additional tools and resources to assist with your investigation.
- If the parties causing the crime are identified, they can be brought to justice, preventing future crime.
- Certain insurance programs and legal requirements state law enforcement must be engaged depending on the type of incident.
- If everybody engages law enforcement, criminals may be deterred from committing crimes based on the risk of being caught.

Incidents and Crime

One key thing that must be established before your SOC considers contacting law enforcement is that a crime has occurred. Laws being broken *usually* (but not always) stem from a security breach rather than an incident. Remember that a denial-of-service attack is not a breach, but it is a security incident and illegal to perform against an organization.

The U.S. Department of Justice (DOJ) differentiates the concepts of incident and breach as follows:

- **Incident:** An occurrence that actually or potentially jeopardizes the confidentiality, integrity, or availability of an information system or the information the system processes, stores or transmits, or that constitutes a violation or imminent threat of violation of security policies, security procedures, acceptable use policies or standard computer security practices.
- **Breach:** The loss of control, compromise, unauthorized disclosure, unauthorized acquisition, unauthorized access, or any similar term referring to situations where persons other than authorized users and for any other than authorized purpose have access or potential access to information, whether physical or electronic. It includes both intrusions (from outside the organization) and misuse (from within the organization).

Shawn Tuma, an attorney with the Texas law firm Scheef & Stone, LLP, who specializes in cybersecurity law, provides the following guidance on when to engage law enforcement based on the type of breach:

> *There are two primary causes of breach events: (1) Intentional wrong doing such as when an outside "hacker" penetrates the network and steals information or . . . when an employee intentionally accesses and takes forbidden information for his own purposes, both of which are generally considered criminal act; or (2) Carelessness or negligence such as when a company insider misplaces an unencrypted USB thumb drive containing PII information.*
>
> *Whether an event is considered to be an incident or a breach is determined by the nature of the event, not what caused it. Whether an event is considered to be criminal or negligence is determined by the actions that caused the event. Some incidents will be criminal but not a breach; some breaches will be negligence but not criminal. Both of the situations described above are breaches though the first was caused by a criminal act and the second was a result of negligence.*
>
> *Criminal actions should be reported to law enforcement. There may be situations where a negligence-based situation should be reported to law enforcement, but this will be determined on a case by case basis.*

Essentially, what Tuma is recommending is that organizations need to report cybersecurity incidents or breaches that are criminal acts to law enforcement as soon as possible, outside of very unique situations. Criminal actions that might have caused a cybersecurity incident are crimes that are no different than if someone were robbed on the street, and they should be treated as such. I partly agree with this thought process and personally believe you might find situations in which the risks of engaging law enforcement might make for a case to not do so. I have spoken with organizations that chose not to engage law enforcement after a network breach based on what systems were impacted and their belief that the threat was contained using their internal incident response service. They felt the additional value that law enforcement could provide, such as identifying more details about the attacker, did not outweigh the additional cycles and resources that law enforcement would need for the investigation. Unlike Tuma, I am not a lawyer, so I will leave

that decision to you and your organization's legal counsel. I also highly encourage you to evaluate the potential risks highlighted previously regarding using your local law enforcement and consider those risks prior to any law enforcement engagement.

Regardless of if and when you plan to engage law enforcement, it is important to have a plan for how to engage law enforcement, and that plan should be included in your incident response process documentation.

> **Note**
>
> I highly recommend engaging law enforcement proactively to identify all the points of contact and better understand the rules for engagement, rather than making your first call to law enforcement one that involves a real incident!

Ticketing Systems

Mature incident response teams use a ticket tracking system as part of the communication process. This feature is commonly available in case management software. Chapter 10 will show an example of case management using the popular security orchestration, automation, and response (SOAR) tool by Splunk called Phantom. Another example of SOAR capabilities is using Cisco SecureX case management features, which I'll use for the next few examples. Regardless of the case management tool that you use, you need to be able to create a report that is trackable for every incident. Figure 8-4 shows creating a new incident ticket for a potential malware outbreak.

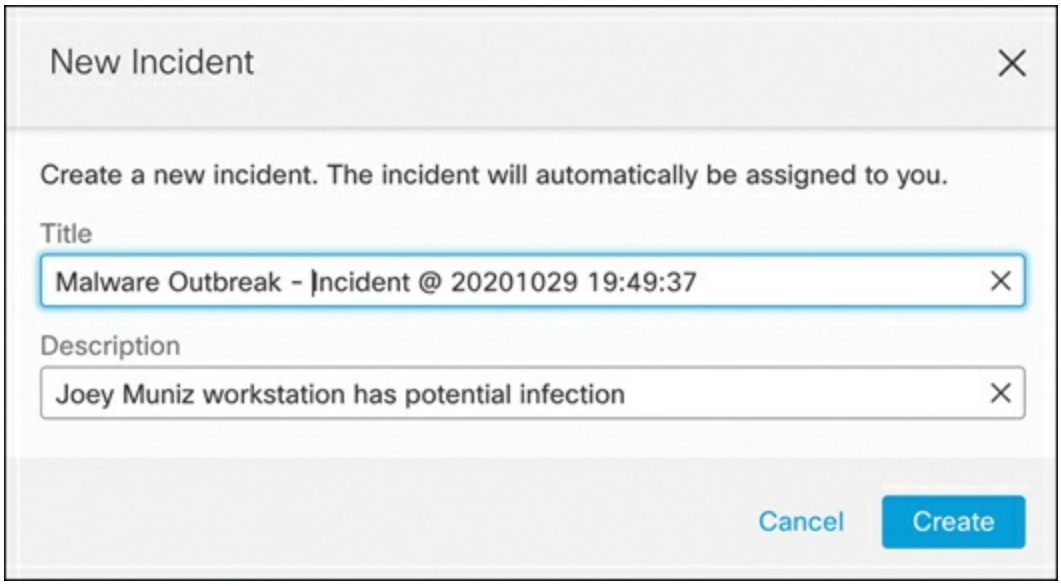

FIGURE 8-4 Opening a New Incident Ticket in Cisco SecureX

Once a case is opened (typically done by tier 1), a documented response is followed, commonly called a playbook (an example of which was presented earlier in this chapter). Playbooks include instructions such as who to alert, what systems to review, what actions to take, and anything else that is a manual or automated effort. As steps are performed, the case management system needs to be updated so that everything is tracked. This allows anybody involved with the case to know what has already occurred and what action items are outstanding. Figure 8-5 shows the new case created for the malware outbreak example.

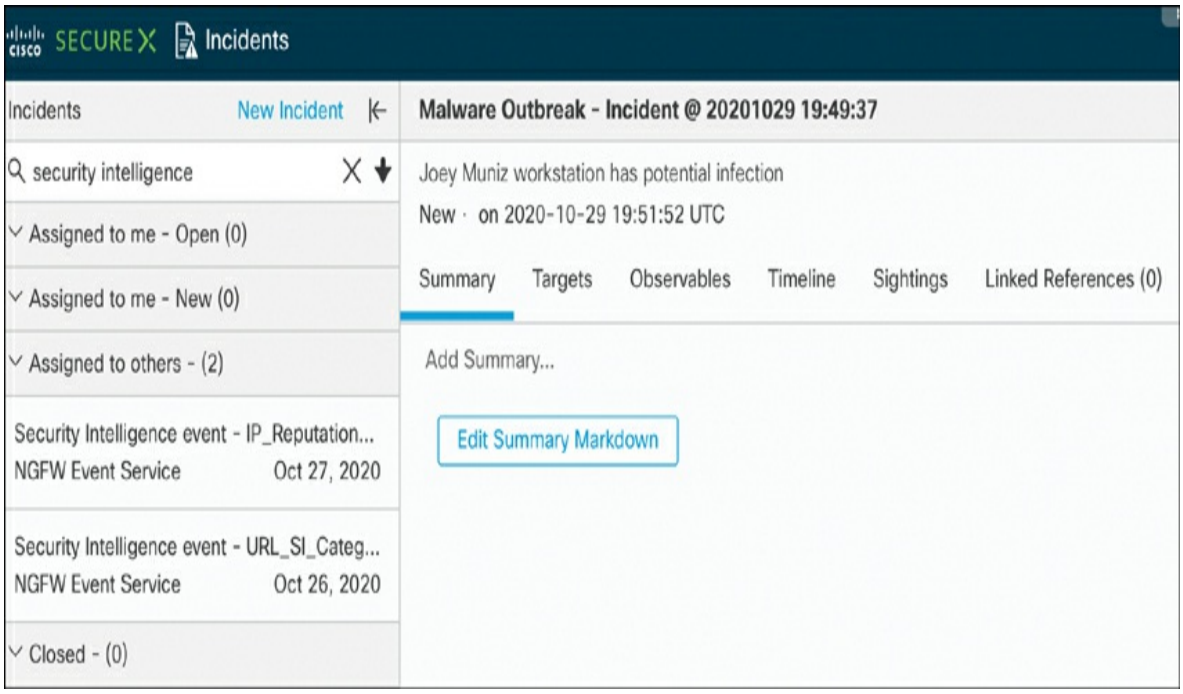

FIGURE 8-5 Example of New Case

As possible impacted targets are identified, you will need to list them to help scope the situation. You will also need to include the time and date that findings occurred, as things will change during the investigation. As other people are engaged, you will add them to the case. If you have notes or diagrams, most case management technology will allow you to upload or link that data. All of this data not only is useful for keeping the team working the case aware of the current status, but also allows the incident response management team to have quick visibility of all active cases so that they can make adjustments to workload and other resources. Chapter 10 dives deeper into how playbooks and orchestration work.

Note

If an incident has potential legal implications, a different case management tool that has forensic-focused tracking might be needed.

Other Incident Response Planning Templates

You can find many examples of incident response templates by searching online. Seeing another organization's view of how to plan an incident response program is a great way to get ideas and validate that you have not missed any critical points.

The following are some examples of incident response plan templates I found by searching online:

- **TechTarget's incident response plan template** (14 pages) includes scope, planning scenarios, and recovery objectives; a logical sequence of events for incident response and team roles and responsibilities; notification, escalation, and declaration procedures; and incident response checklists.

  https://cdn.ttgtmedia.com/searchDisasterRecovery/downloads/SearchDisasterRecovery_Incident_Response_Plan_Template.doc (automatically downloads)
- **Thycotic's incident response template** (19 pages) includes roles, responsibilities, and contact information; threat classification; actions to be taken during incident response; industry-specific and geographic-dependent regulations; a response process; and instructions on how to customize the template to your specific needs.

  https://thycotic.com/solutions/free-it-tools/free-privileged-account-incident-response-policy-template/ (requires registration to download)
- **Sysnet's security incident response plan template** (11 pages) includes how to recognize an incident; roles and responsibilities; external contacts; initial response steps; and instructions for responding to several common incident types, such as malware and unauthorized wireless access.

  https://sysnetgs.com/security-incident-response-plan-template/ (requires registration to download)
- **California Government Department of Technology incident response plan example** (4 pages) includes a 17-step checklist for incident team members to follow, with reference to more detailed procedures for specific types of incidents.

  https://cdt.ca.gov/wp-

content/uploads/2017/03/templates_incident_response_plan.doc (automatically downloads)

Phase 1: Preparation Summary

The list that follows summarizes the key points to keep in mind for the preparation phase of the incident response lifecycle:

- Identify incident response managers and team leads that will be responsible for the service
- Establish how communication will occur with internal and external parties
- Research all existing support contacts and establish any needed pre- and post-incident response service contracts with vendors and service providers
- Provide any training required for the incident response service and communication for third-party resources
- Identify all third-party resources and establish communication including law enforcement and other third-party resources
- Develop procedures for responsibilities for each tier supporting different types of incidents as well as how incident escalation will occur
- Acquire any missing technology or people needed to launch the incident response program or enforce changes based on lessons learned feedback

Phase 2: Detection and Analysis

Once you have built the incident response program based on all preparation tasks, your incident response service is ready to "go live," meaning the team is ready to start detecting and responding to security incidents. NIST SP 800-61 Rev. 2 calls this phase detection and analysis; *detection* refers to the SOC monitoring for potential incidents, and *analysis* refers to the SOC evaluating anything that could be an incident, meaning the work done by tier 1 or tier 2 to determine whether a case needs to be opened or whether the situation is

not an incident the SOC should be concerned about.

There are a few ways an incident will be detected. One way is that an employee or outside party will use one of the communication methods previously covered to alert the SOC of a potential incident. Another way an incident is created in the SOC is when the SOC detects a possible threat using detection techniques.

> **Note**
>
> There are other incidents that could occur outside of cyberattacks, such as the network going down due to a misconfiguration; however, this chapter focuses on incidents based on cyberthreats.

Incident Detection

Before the SOC can hunt for a threat, it must have a way to detect threat activity using different detection techniques or be able to develop a hypothesis of what to look for and monitor against that hypothesis. The hypothesis is based on developing a plan to detect different types of attack vectors. In Chapter 1, I covered how to develop security capability maps that are based on using industry guidelines that cover the most common attack vectors, such as email, gateway, and hosts.

Detecting threats within a capability map occurs by using detection technology. Chapter 1 introduced the concept of detection technologies as tools used to detect malicious behavior within network and systems. The industry commonly refers to this area of defense as *breach defense tools*. Breach defense tools are not focused on preventing threats but rather assume your prevention technologies such as firewalls have failed and now a threat has breached your network. Remember that all incidents are not breaches, but all breaches are incidents!

Core Security Capabilities

Chapter 1 introduced the three core security capabilities: signature, behavior, and anomaly detection, as shown in Figure 8-6. I highly recommend focusing on a combination of behavior and anomaly detection capabilities for your breach detection functions. Signature detection can help; however, signature detection through antivirus measures or an edge intrusion protection system (IPS) is more ideal for prevention. By combining behavior and anomaly detection capabilities, an attacker/malware would have to hide any malicious behavior but at the same time look like normal behavior. If the threat attempts to hide very well, it will look like a huge anomaly compared to everything else that is not hiding. If the threat doesn't hide that well, the malicious behavior detection engines within your breach detection tools should identify it. See Chapter 1 regarding security capabilities for more details on these concepts.

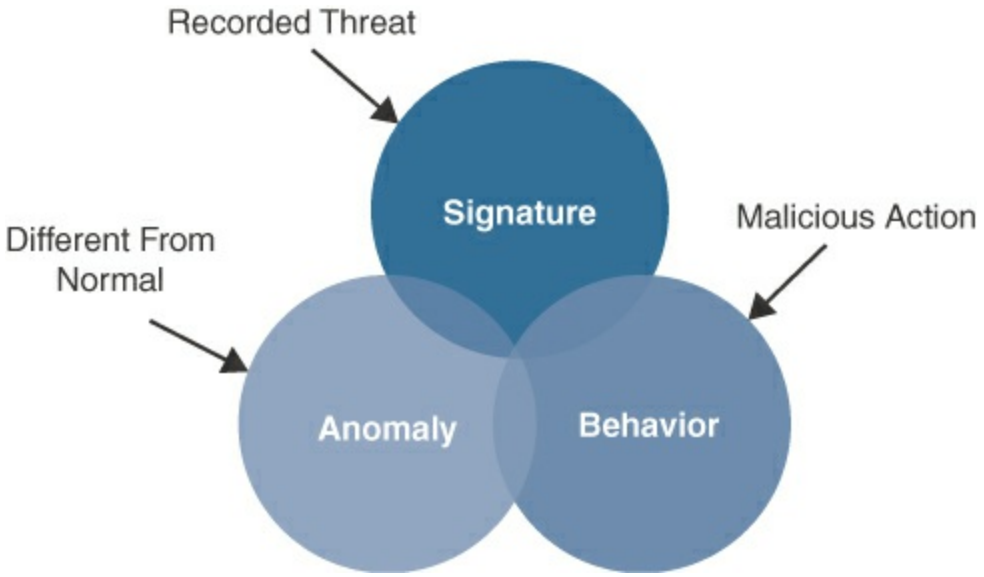

FIGURE 8-6 Core Security Capabilities

The following are popular options for breach defense:

- Connecting an internal intrusion detection system (IDS) to a SPAN/port mirror to monitor internal traffic
- Monitoring data within the network devices such as NetFlow
- Setting up a honeypot to lure threats to attack the honeypot rather than real targets

- Using tools that monitor outbound traffic looking for communication to high-risk resources such as the darknet, Tor, or known malicious resources
- Capturing and analyzing packets from within the network
- Your employees and their systems alerting you of a possible event

Perform a SOC security capabilities assessment, as described in Chapter 1, to ensure that you have one or more of the breach detection capabilities covered in the previous list. (When I say one or more capabilities, however, I don't mean that your end users should be your only detection capability!)

Chapter 5, "Centralizing Data," covered how to pull logs from security tools into a centralized tool and convert that data into actionable intelligence. Use those concepts to develop an incident response procedure document that covers who should monitor the centralized data tool, how events are converted into cases when applicable, and how this all works within your SOC tier support model.

As possible incidents are detected by your SOC or external parties, the second half of the detection and analysis phase kicks in, in which your incident response program analyzes what has been detected. Your incident response team's goal at this point is to determine if the threat is real and, if so, what is the proper response. This practice is commonly referred to as *threat analysis*.

Threat Analysis

The SOC incident response team will encounter situations where a potential threat exists but further analysis is needed to determine if the threat is real. For example, an employee might report "Since I ran this file, my computer has acted funny" and ask whether the file contains malware. Or a user could complain that an email received was flagged and an attached file was removed and stored in a sandbox. That user might ask why the file was removed, and the threat analysis team must be able to prove the file is a risk. Users could complain that external resources are being blocked and they want the SOC to approve allowing access to those resources. The threat analysis team will need to explain the risk behind the blocked resources through threat

analysis research based on why the resource was blocked. In all of these cases, the SOC needs to determine the risk associated with an artifact and resource before any actions can be taken, such as performing an incident response or alerting the end user that a risk doesn't exist.

The analysis service involves specific skillsets related to investigating artifacts and researching malware. As pointed out in Chapter 3, the analysis service answers questions about what an artifact is and its intended purpose (malicious or safe). A general comparison of skillsets between an incident response analyst and malware analyst is how each role works with malware. An incident response analyst wants to know the general risk associated with an artifact and is responsible for responding to the threat. A malware analyst will know how to identify if an artifact is malicious, isolate the artifact, and research its characteristics to better understand how it functions, what to look for to find all variations of the threat, and even possibly develop steps to reduce the risk of future attacks from similar threats. The average incident response analyst will not know how to disassemble malware or be able to read assembly language, which is why the analysis service contains SOC employees with very specific and specialized analysis skillsets. See Chapter 4 to learn more about skillsets associated with analysis specialists.

Detecting Malware Behavior

Malware developers know security tools and SOCs like yours exist. Teams I have worked with and friends in the security industry have taken down perpetrators of attack campaigns and discovered that they had research labs containing all of the security tools used in the industry. Malware developers typically have popular vendor firewall models, IPS tools, various flavors of sandboxes, most modern antivirus offerings, and anything else your organization and other organizations use to defend against cyberthreats. Keep this in mind as you analyze artifacts, and expect to encounter anti-detection, anti-disassembly, and anti-debugging tactics intended to confuse you and the tools you use to learn about threats.

Infected Systems

There are many things you should expect malware will do once installed on an impacted system. The following is a partial list of general malware behavior that you might encounter:

- Create new files and/or modify existing files
- Move or copy malicious files to a known location such as Windows system folder; malware here can run with user permissions
- Malware might not produce an executable but instead produce a DLL or batch script
- Delete temporary files that it had created
- Create registry entries (might also control network adaptors)
- Create unique identifiers that are benign but great for detecting malware families
- Disable security tools such as the Windows Firewall
- Phone out to malicious sources

With experience, your SOC analyst team will learn to identify these malware behaviors and many others; however, you do not have to have experience looking for malware behavior to analyze malware. The good news is that many of the industry tools I'll cover do the detection and analysis work for you! Figure 8-7 shows an analysis report produced by Cisco Advanced Malware Protection (AMP), an endpoint detection and response (EDR) program, displaying all of the indicators that resulted in an artifact being deemed malware. Analysts of any skill level can perform malware analysis, but the more skilled professionals are capable of working outside of the data provided by security tools. I highly recommend your SOC use industry certification and training programs to improve the skills of those responsible for analyzing artifacts, as this space is very specialized and hard to recruit for.

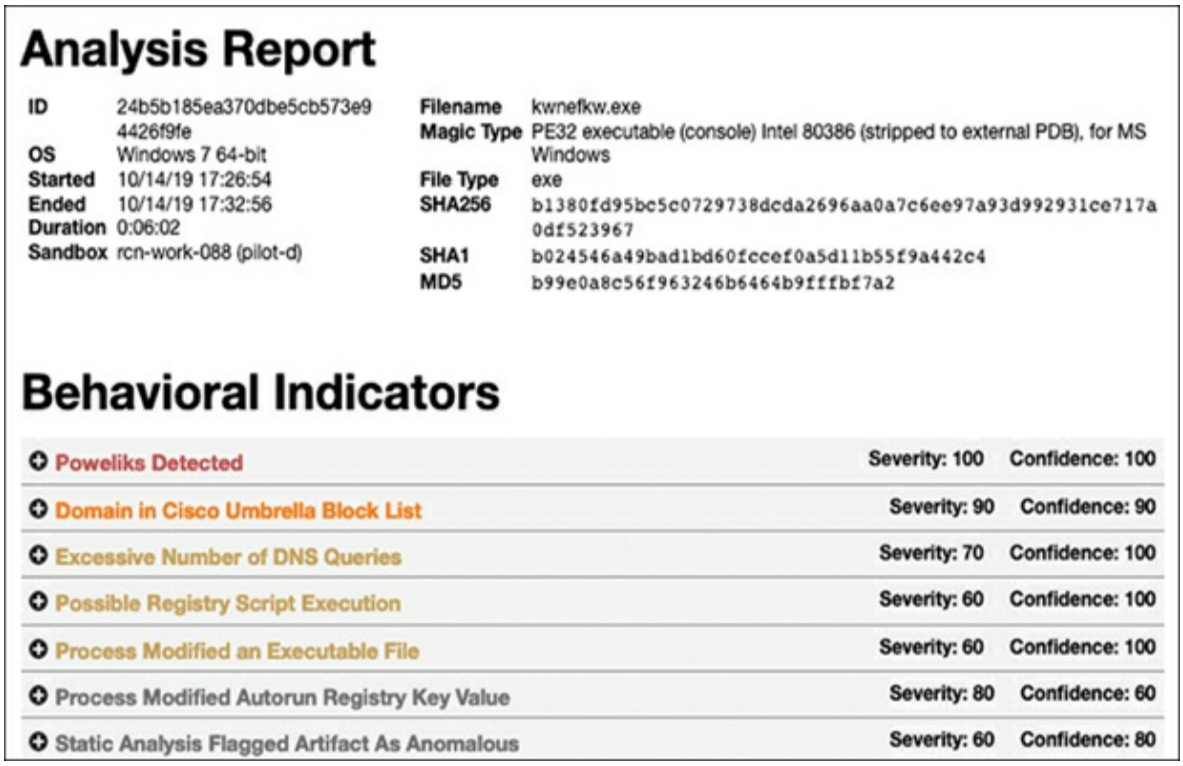

Analysis Report

ID	24b5b185ea370dbe5cb573e94426f9fe	**Filename**	kwnefkw.exe
OS	Windows 7 64-bit	**Magic Type**	PE32 executable (console) Intel 80386 (stripped to external PDB), for MS Windows
Started	10/14/19 17:26:54	**File Type**	exe
Ended	10/14/19 17:32:56	**SHA256**	b1380fd95bc5c0729738dcda2696aa0a7c6ee97a93d992931ce717a0df523967
Duration	0:06:02		
Sandbox	rcn-work-088 (pilot-d)	**SHA1**	b024546a49bad1bd60fccef0a5d11b55f9a442c4
		MD5	b99e0a8c56f963246b6464b9fffbf7a2

Behavioral Indicators

Poweliks Detected	Severity: 100	Confidence: 100
Domain in Cisco Umbrella Block List	Severity: 90	Confidence: 90
Excessive Number of DNS Queries	Severity: 70	Confidence: 100
Possible Registry Script Execution	Severity: 60	Confidence: 100
Process Modified an Executable File	Severity: 60	Confidence: 100
Process Modified Autorun Registry Key Value	Severity: 80	Confidence: 60
Static Analysis Flagged Artifact As Anomalous	Severity: 60	Confidence: 80

FIGURE 8-7 Cisco AMP Indicators of Compromise Example

Analyzing Artifacts

The first important point about analyzing artifacts is that there is no *single way of doing it*. I highly recommend keeping this concept in mind any time you investigate artifacts. If you find one way doesn't work, try something else. Do not dwell on one tactic failing, as there will likely be another way to accomplish the same outcome, which is determining if the artifact is a threat and what it does.

Different tools are great for different jobs and will offer different views of the same thing. A simple example is using different web browsers to analyze an external resource. I have Firefox, Chrome, and Safari installed on my Mac for this purpose. Sometimes things won't work within one browser but magically function properly when I switch browsers. Remember, if one tool fails, try another.

The most important recommendation I can make is to know your goal for performing an investigation. If your goal is to determine if something is a risk

to the organization, then research until you know that answer and stop any extra research. There is no reason to burn hours performing advanced analysis on something you know is bad when all you plan to do is remove it. If your goal is to understand how a threat could impact other systems, you will want to analyze how the artifact functions, requiring more research. If your goal is to identify characteristics that could lead to blocking artifacts that use the same logic as the malware, you may have to perform a complete disassembly and debugging of the artifact, requiring even more time and research. Take into consideration the time and resources needed as you choose what outcome you want from the analysis part of your incident response.

I recommend using the following approach to analyzing artifacts. Your end goal will determine when you stop working down this list, since each further step will require additional time and resources:

Step 1. Review existing tools, logs, and other sources related to the artifact.

Step 2. Use quick and simple static analysis techniques.

Step 3. Use quick and simple dynamic analysis techniques.

Step 4. Perform deeper static and dynamic analysis.

Step 5. Perform deep analysis, including a full disassembly and debug.

For most incidents, you will stop after the first step because existing tools will provide enough information to inform you that the artifact is bad. The first step might not produce enough evidence to make a decision, in which case you have to move on to doing some additional quick static analysis. Static analysis might not provide enough detail, forcing you to switch to running the malware using dynamic analysis techniques. It is possible that as you determine something is bad, the incident response process will require you to determine how the threat spreads, leading to more research being needed. Know that if you need this level of detail, you might have to invest a lot more time in research, including developing a secure place to interact with and disassemble the artifact. I believe the first three steps of my analysis list are ideal for an incident response analyst to perform; however, the steps requiring deeper analysis should be escalated to an analyst that specializes in analyzing artifacts.

Identifying Artifact Types

An artifact might look like one thing but really be something else. Malware creators disguise artifacts in this manner as a way to trick a user or system into performing an action. For example, a file that appears to be a Word file named evil.doc could actually be named evil.doc.exe and represent an executable that does bad things if run on a target system. Figure 8-8 shows the Windows settings that could cause a user to be unaware of hidden extensions. A file called “evil.doc” could really be “evil.doc.exe” representing an executable that does bad things if run on a target system. It is common to disguise the true nature of malware by using encryption, compression, and other techniques to trick users and security tools into not recognizing the threat’s file format.

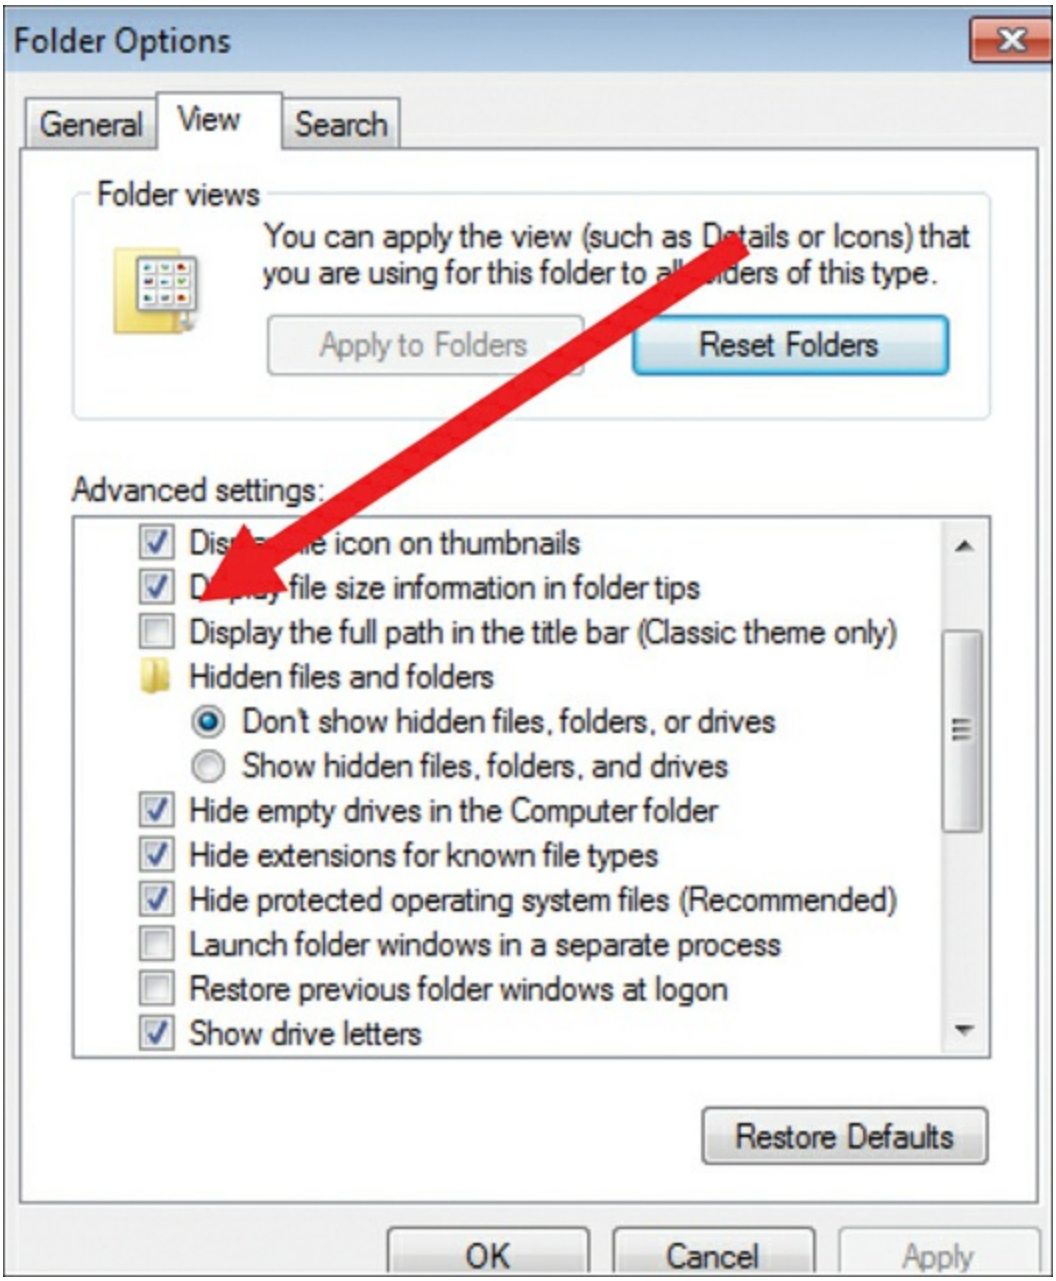

FIGURE 8-8 Hidden Extension Example

Magic Numbers

Every file has a magic number that represents what type of file it is. The absolute way to validate a file is to view this number. For example, to identify a GIF file, you can analyze the file at the packet layer to see if 47 49 46 39 is at the beginning of the file or if it contains GIF89. Many security

tools use this technique, which is how they can identify file types such as a PDF, GIF, or JPEG. Figure 8-9 shows an IDS rule that I wrote for Snort that looks for 47 49 46 39 or GIF8 as all traffic is analyzed. If either pattern is found, a rule is triggered that generates a log message stating "GIF file detected in network." I ran this experiment while teaching a digital forensics class at an industry conference. I went online, searched for GIF files, and downloaded them while analyzing everything with my Snort IDS detection rules. What was interesting was every few files that had the .gif file extension did not trigger my IDS alarm. When I analyzed the traffic using Wireshark, I found that those files that didn't trigger my rule but had the .gif extension were actually other types of files whose file extension someone changed without converting the actual file to the GIF format. Once again, this is why it is important to validate files rather than assume they are what they appear to be.

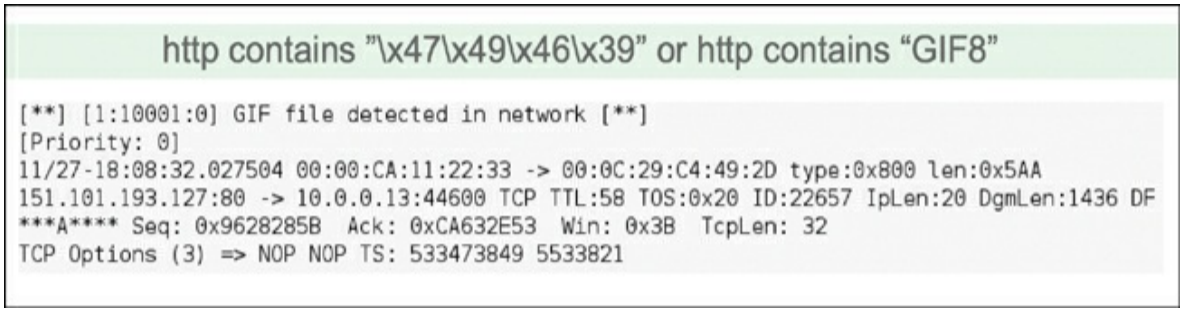

FIGURE 8-9 Magic Number Example Identifying GIF File Type

Note

Learn more about magic numbers for files at https://asecuritysite.com/forensics/magic.

File Identification Tools

An easier method for validating files is to use a file identification tool that essentially does the magic number analysis for you. Many file identification tools also provide additional details such as if the file has been packed, which

is a common characteristic of malware. Figure 8-10 is an example of TrIDNET, which is a free open-source file analysis tool. Simply click Browse, upload the file, and let TrIDNET analyze the artifact. Another option that is similar to TrIDNET is PEiD.

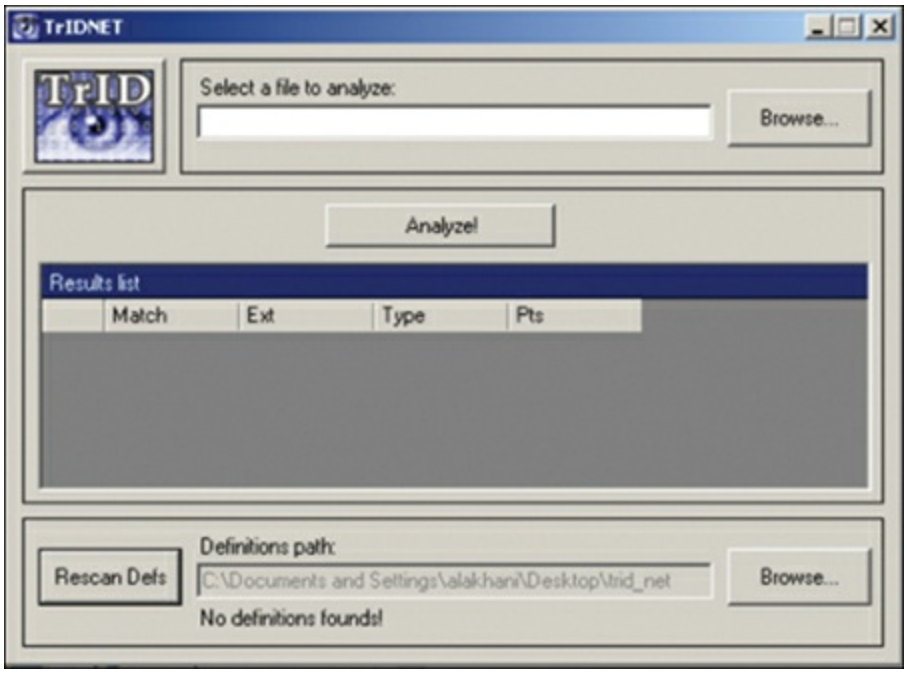

FIGURE 8-10 TrIDNET Free File Analysis Tool

Note

You can find TriDNET at https://mark0.net/soft-tridnet-e.html and PEiD at https://www.softpedia.com/get/Programming/Packers-Crypters-Protectors/PEiD-updated.shtml.

Packing Files

Most modern malware will be in the form of packed files. Malware developers use this technique to defend the malware against disassembly (and eventual eradication), a capability that many security tools are able to use on

unpacked malware to learn fundamental characteristics of how it works, which defenders can use within their tools to identify and prevent the malware. Essentially, packing means *executable compression* because it occurs after the malware program is created and converted into an executable.

Figure 8-11 represents the difference between an unpacked program (left) and a packed program (right). Within the header are details about the program. Programs call other programs and resources, which are labeled so it is clear what additional resources are needed for the program to function. Finally, you have the actual program. When a program is converted into an executable (known as *compiling* the program), the program is converted into machine-readable code. With the right tools and techniques, a researcher can reverse engineer the program and learn how it functions. To prevent this from occurring, packing the executable compresses and encrypts the code associated with the program, preventing research of how the malware functions.

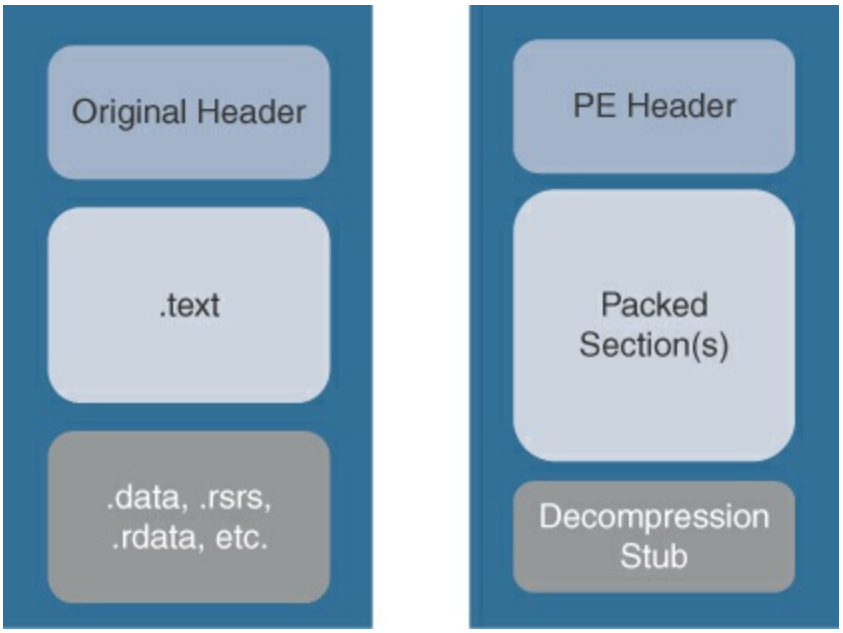

FIGURE 8-11 Unpacked vs. Packed Executable Concept

Packing occurs by using a packing algorithm. If an analyst can determine which algorithm was used to pack the executable, there will be a strong chance that the analyst can use the same algorithm to unpack the file. This is why using a file-identifying program like TrIDNET can be beneficial, as this

and many similar tools include the ability to fingerprint packing algorithms. Common commercial packing programs include UPX, PECompact, ASPack, Petite, WinUnpack, and Themida; however, many malware authors use customized packers to avoid the possibility of the packing algorithm being detected and used to unpack the malware file. If the algorithm used to pack the file can't be identified, another option to unpack the executable is to locate the stub file. Doing so can be extremely tedious and requires a specific skillset common only to experienced reverse malware analysts.

All is not lost if artifacts are packed. Useful datapoints can be learned from what is not encrypted, such as data within the header of the file, as well as by analyzing how the packed file looks. These analysis steps are performed through static analysis.

Basic Static Analysis

Static analysis involves investigating threats without executing them. Your goal is to learn about the artifact based on identification of various characteristics and matching those characteristics to what you or the industry knows about existing threats. The most obvious characteristic for matching an artifact is a known threat signature. Antivirus is a common example of a static analysis tool that attempts to match thousands of known threat signatures to any artifact on a system with the goal of quickly identifying and preventing a malicious file from being run.

There are other characteristics that can be quickly scanned for to get an idea of the risk. Examples of characteristics of interest that can be obtained using static analysis techniques include the following:

- Whether the artifact is packed and, if so, the type of algorithm used to pack the file
- Portable executable (PE) data
- String extraction and analysis
- Import table data
- Resource data

- Embedded objects

Each of these datapoints can lead to a better understanding of what the artifact is designed to do. A string extraction could lead to identifying a list of SMTP servers, supporting the conclusion that the artifact is designed to send spam or other useless messages. The PE header can show which libraries are required for the artifact to execute, where the routine code is located (commonly referred to as the entry point), or even when the artifact was created. Resource data can show if additional configuration data or files are supposed to be dropped on the system post execution. Embedded objects can represent additional artifacts that might not be packed or offer additional details through their own static analysis to lead to an overall conclusion about the risk associated with an artifact.

Common Packed File Attributes

The following list presents some common characteristics to look for when analyzing packed artifacts. Any one of these means you are likely dealing with a piece of malicious software.

- Packed and obfuscated code will at least include functions such as LoadLibrary and Get ProcAddress, which are used by malware to load and gain access to additional functions. This is a giveaway that what you're dealing with is malware.
- A packer signature that can lead to the conclusion of malware based on the type of packer that is detected.
- Abnormal entry point/section (example .test section in PE).
- Section or memory with read and write permissions. This is an obvious characteristic of malware because compilers don't do that, for security reasons.
- Large difference between physical size and run size of program.
- Data size is too large based on what is expected from similar file types (likely contains an executable or the file extension isn't accurate).
- Too few import functions based on what is expected from a similar artifact.

Peframe

Peframe is one of the many open-source tools that can be used to perform quick static analysis of an artifact of interest. Peframe can pull many of the data elements I previously covered as well as create different hash values of the artifact, make assumptions about how the file has been modified, and even query VirusTotal to see if the artifact has been flagged as malicious. Figure 8-12 is a screenshot of some of the results from running peframe against an artifact of interest. Notice there are different sections of results, such as what behavior might have been used to modify the file as well as a sha256 hash value. Peframe typically takes a few minutes to run, allowing for quick analysis of even the most complex artifacts.

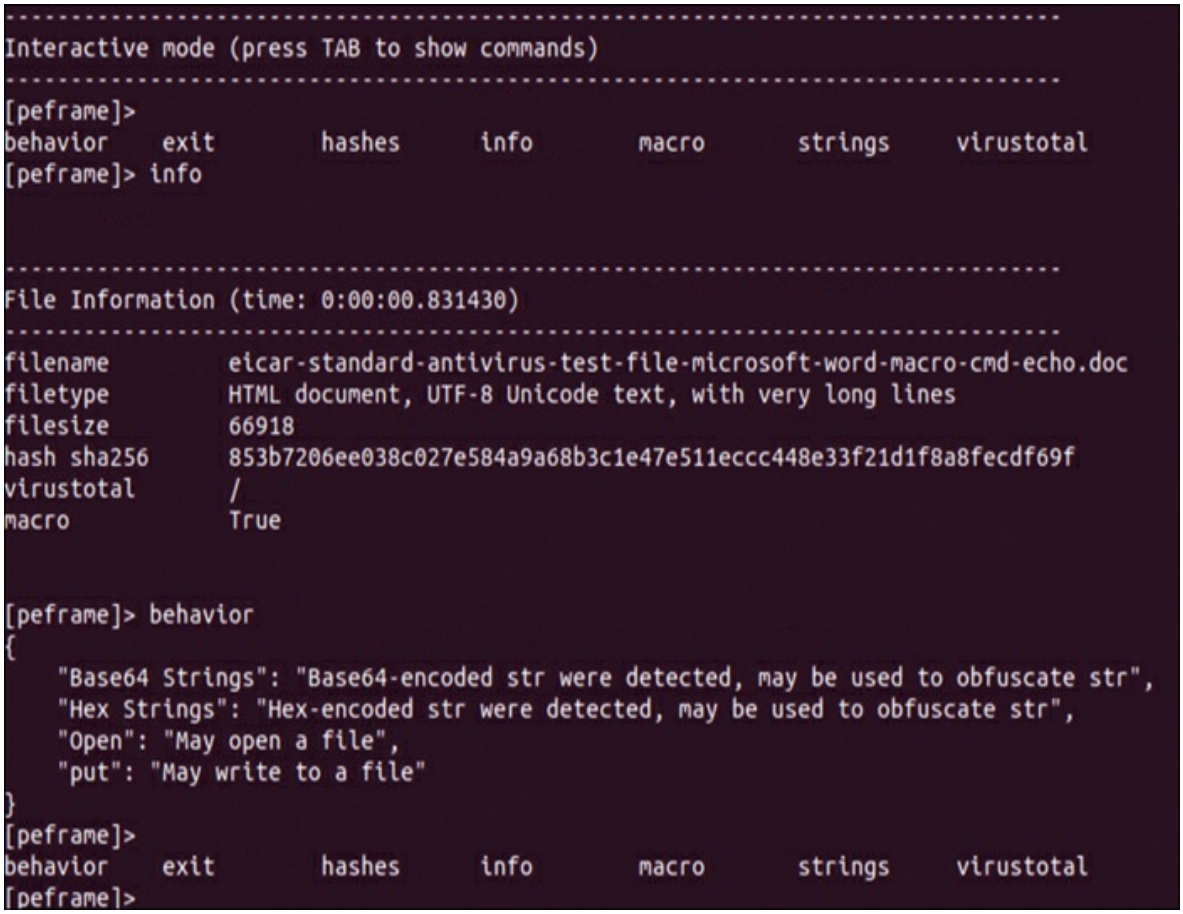

FIGURE 8-12 Example of Peframe Analyzing an Artifact

Advanced Static Analysis

If basic static analysis can't generate what is needed to make a decision or if additional information is needed about an artifact, your SOC can escalate the investigation to the analyst team that is capable of performing advanced static analysis techniques. Advanced static analysis involves steps used to understand how the malware was designed. Steps can include disassembling the malware down to its fundamental code and reverse engineering how different parts of the code relate.

As explained earlier in this section, malware developers are programmers. They use programming languages such as C++ or Python to develop malicious programs. Once a program is ready for use, it is compiled into machine code, commonly called an *executable*. Executables in machine code are extremely difficult for humans to understand. Although security researchers don't have access to the source code of an executable, they can convert an executable into a low-level language, known as *assembly language*, using a disassembler. Figure 8-13 represents this lifecycle of malware being developed to being analyzed by a security researcher.

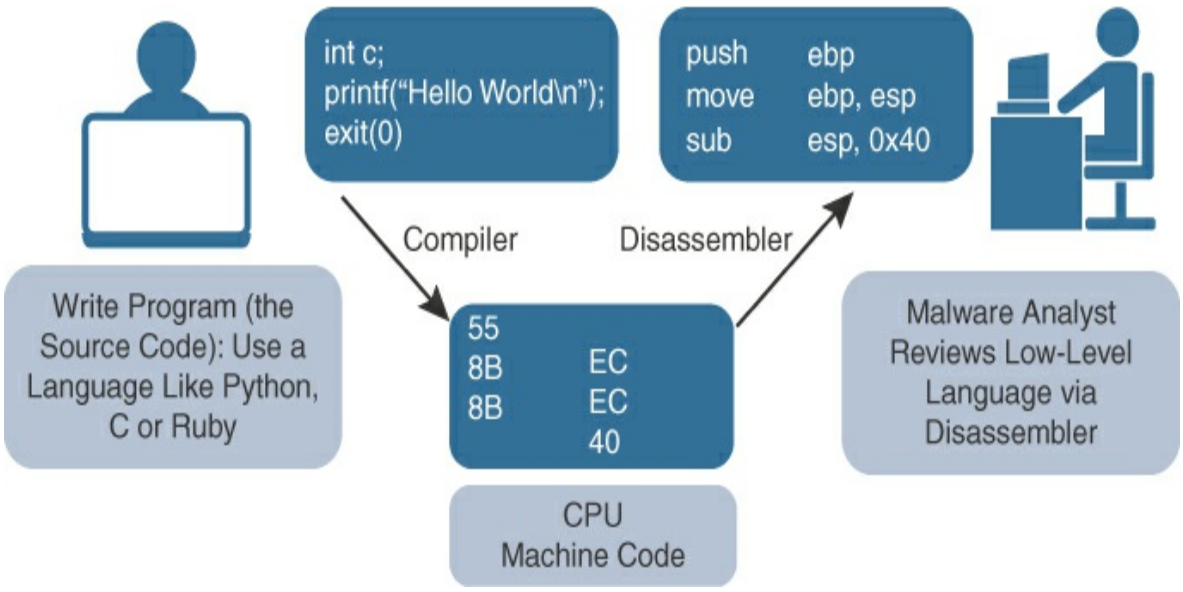

FIGURE 8-13 Programming to Disassembly Concept

One of the most basic programs new software developers learn is one that

prints to the computer screen “Hello, world!” When compiled, printing a simple “Hello, world!” program will look like random machine code. Figure 8-14 shows what a basic “Hello, world!” program written in C looks like when converted into assembly language. Notice that there are **mov** statements representing data being moved between computer storage sections such as registries and RAM. Understanding assembly language requires knowledge of both how assembly language is structured and how computers move data when programs are executed. It is important to reiterate that only analysts with specialized training in malware analysis will have these required skillsets.

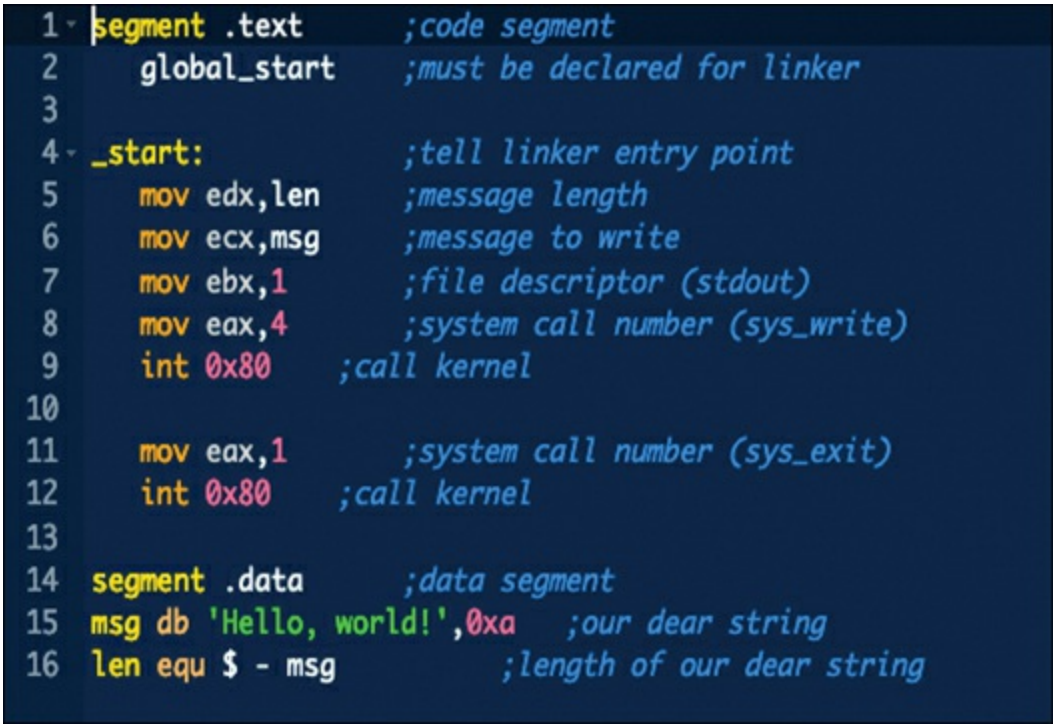

```
1  segment .text        ;code segment
2     global_start      ;must be declared for linker
3
4  _start:              ;tell linker entry point
5     mov edx,len       ;message length
6     mov ecx,msg       ;message to write
7     mov ebx,1         ;file descriptor (stdout)
8     mov eax,4         ;system call number (sys_write)
9     int 0x80      ;call kernel
10
11    mov eax,1         ;system call number (sys_exit)
12    int 0x80      ;call kernel
13
14 segment .data        ;data segment
15 msg db 'Hello, world!',0xa   ;our dear string
16 len equ $ - msg             ;length of our dear string
```

FIGURE 8-14 Unpacked vs. Packed File Concept

IDA Pro

One of the most popular disassemblers used by malware analyst around the world is IDA Pro. Figure 8-15 shows an example of using IDA Pro to disassemble a program. This example shows one of the most common training exercises performed when learning how to disassemble programs. In this example, my goal is to force a program to lead to a function that says I

have entered a correct password. The function shown in the lower-left command-line interface displays “Congratulations” and the function shown in the lower-right command-line interface displays an error message. My goal in this exercise is to figure out which function is used as the decision point to choose to either move to the correct message function or error message function. By identifying that decision function, I can modify things such as allow for a password or 1=1, meaning one will always equal one, forcing the program to always move towards the “Congratulations” function. This is common behavior found when a researcher removes password requirements from software, also known as cracking software, which is likely illegal to do.

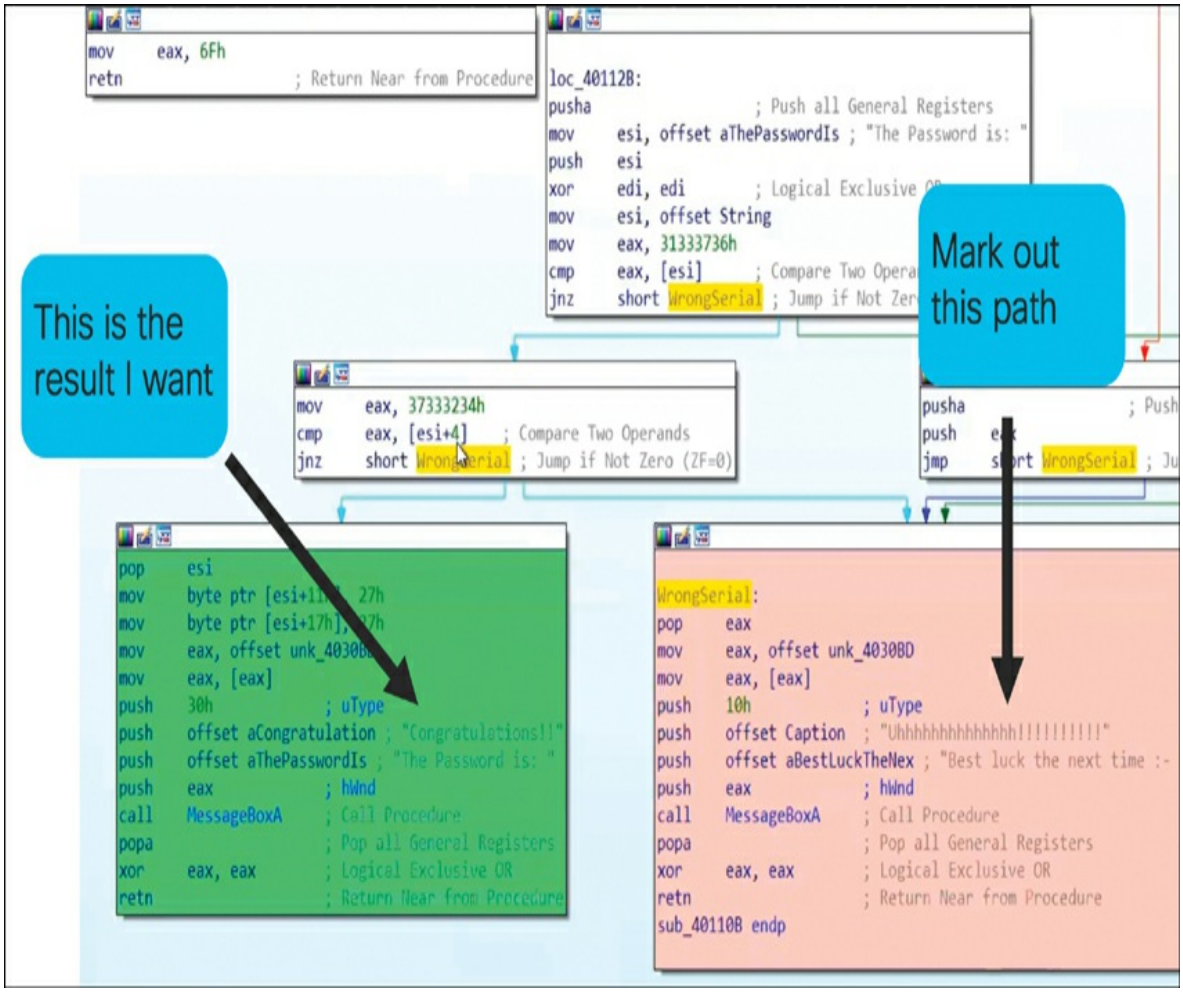

FIGURE 8-15 Understanding Disassembled Code Concept Example

I point out this example to help demonstrate the thought process behind

reverse engineering artifacts. The process can be extremely time-consuming to perform because you will likely need to review hundreds of lines of code, overcome anti-disassembly techniques, and debug aspects of the program that are not understood or are functioning in strange manners. This all assumes you are able to view the contents correctly. Figure 8-16 shows what a packed executable using UPX looks like. Notice on the left that the UPX label is displayed and very little data is displayed. There is little a researcher can learn from disassembling packed files. If the file can't be unpacked, basic static analysis is going to be the best approach for collecting information about the artifact.

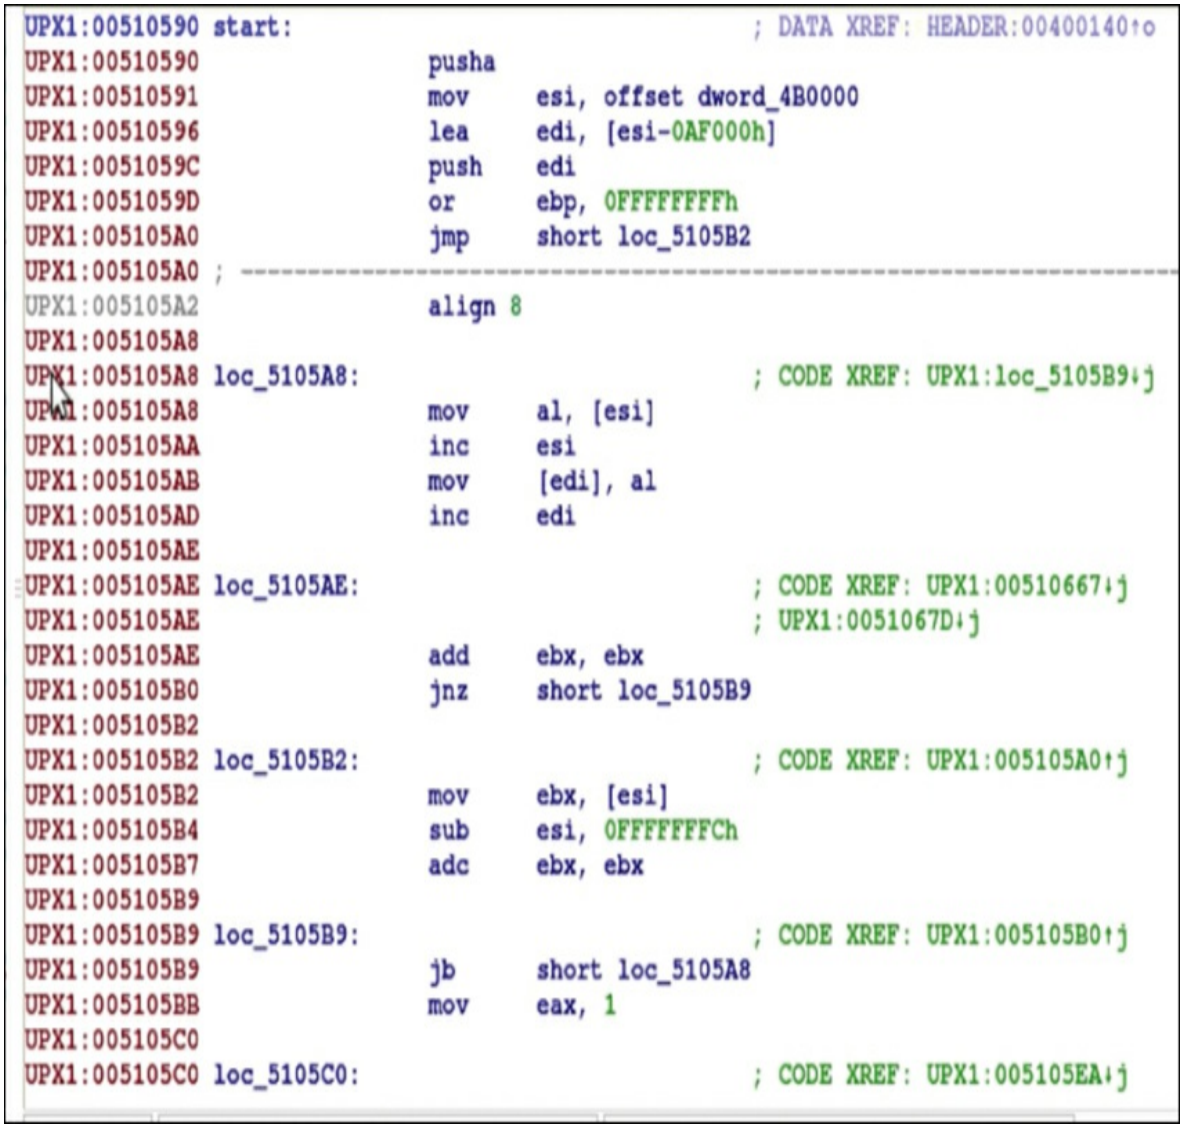

FIGURE 8-16 Executable Packed with UPX and Disassembled

Malware Analysis Story: WannaCry Kill Switch

WannaCry is a ransomware attack that occurred in May 2017 and hit systems globally. One malware researcher, Marcus Hutchins, who goes by the name MalwareTech online, disassembled WannaCry and found an interesting URL. He researched the URL and noticed that it was not registered. Out of curiosity, he registered the URL to see how the WannaCry malware would react to the URL becoming available. To his surprise, the malware reached out to the URL and, upon verification of its existence, shut down the WannaCry ransomware process. Essentially, MalwareTech had stumbled upon a built-in kill switch leading to the global takedown of WannaCry. This action instantly put MalwareTech in the spotlight as a cyber hero (though that later got him into trouble for some previous work he performed). Figure 8-17 shows an example of using a disassembler known as Ghidra to disassemble and review the infamous kill switch found within the WannaCry ransomware. The kill switch is the code with www in it.

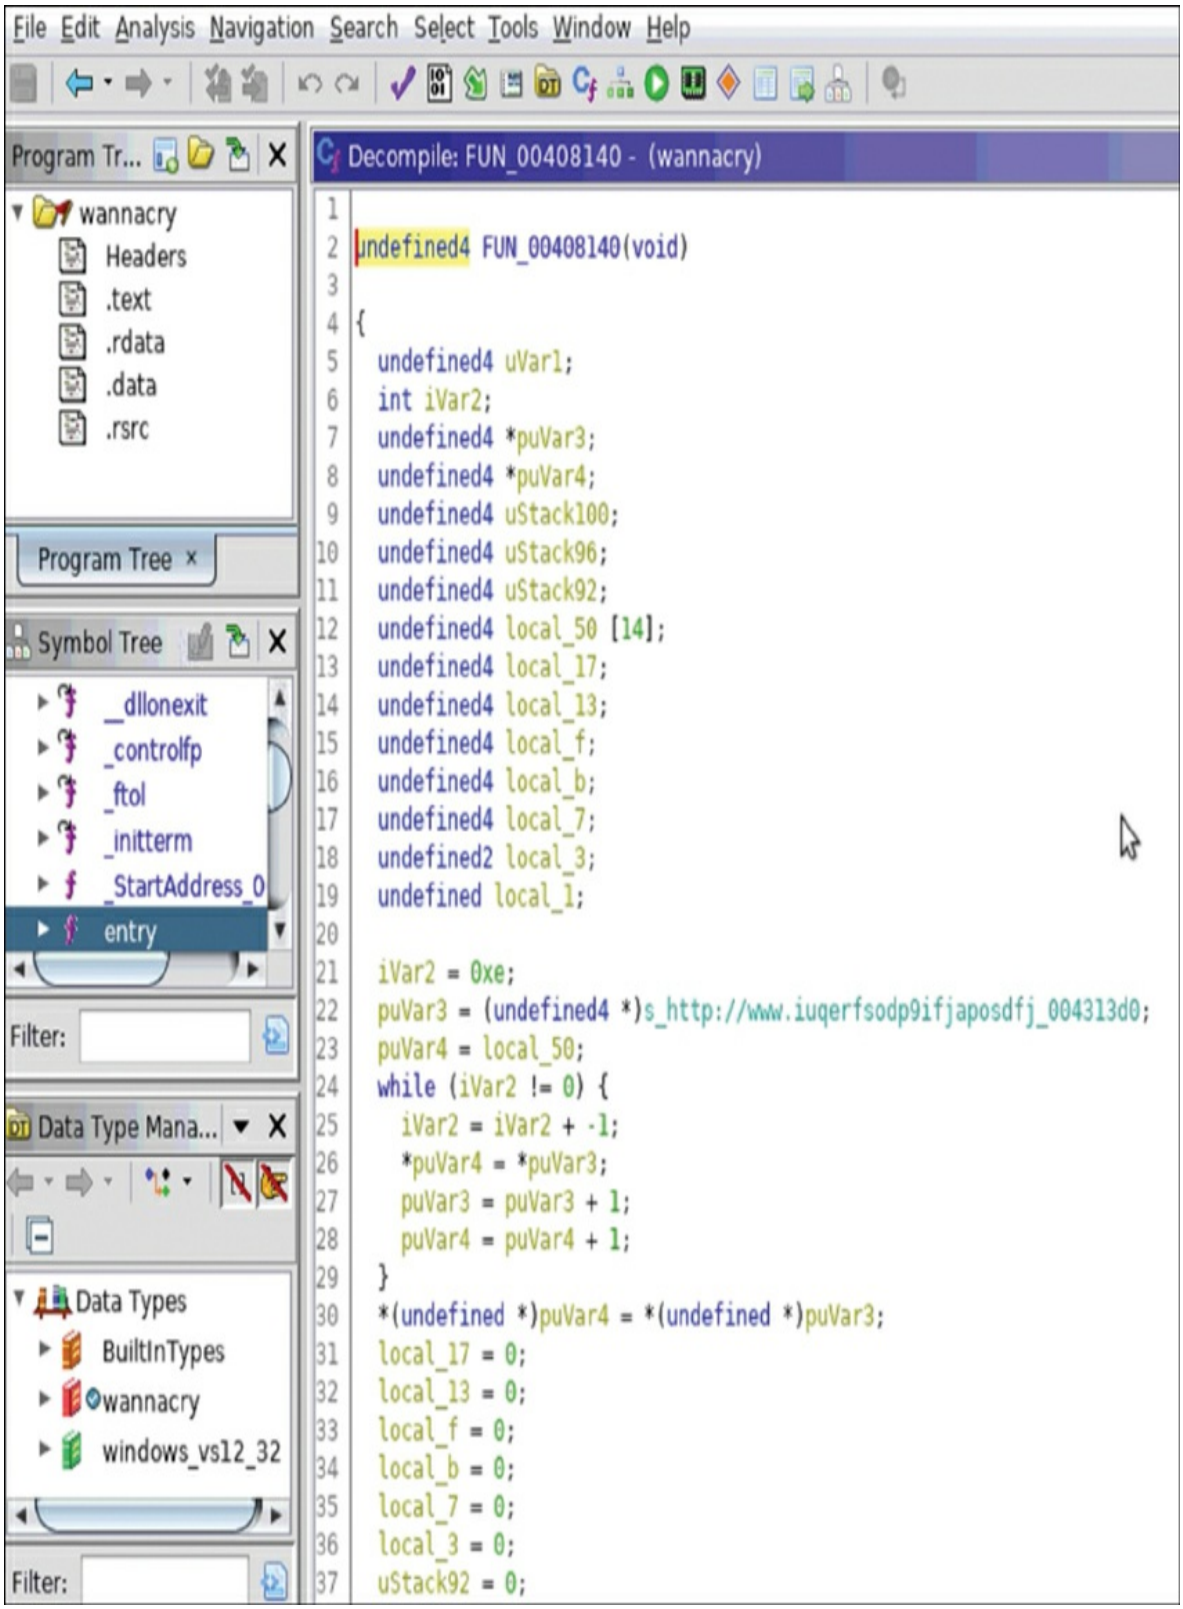

FIGURE 8-17 Ghidra Disassembler Viewing WannaCry Ransomware Kill Switch

Dynamic Analysis

Another approach to learning about an artifact is to execute it in a controlled environment and study its behavior. Doing so can show what the program does when executed, where it communicates to, and its intent upon execution. Do not test malware on a system with access to your network or the public Internet, as that exposes your organization and other organizations to the threat of being compromised by the artifact you are researching. Instead, you need a container to work within, which can be an isolated system or a sandbox.

Isolated Systems vs. Sandboxes

There is a difference between using an isolated system and using a sandbox. Sandboxes are designed to research artifacts, and most include many tools and detection mechanisms that can be used to learn about artifacts as they execute. Most modern sandboxes also include anti-detection tools that can be used to trick the malware into believing it has found a vulnerable host. Anti-detection capability is important because many malware authors design malware to only execute when specific human characteristics and behaviors are identified, with the goal of avoiding sandboxes and other security tools. Sandboxes can simulate mouse movement and provide fake network resources, all with the intent of convincing the artifact that it has found a real target to attack. Sandboxes are built around common operating systems, meaning you won't find sandboxes built for older or customized systems. If a sandbox option is available for your use case, I highly recommend leveraging it.

An isolated system is just a safe place to test artifacts. This means you must perform all testing as well as be able to recognize when malicious actions are occurring, because isolated systems don't have tools and detection mechanisms like those offered by sandboxes. This allows for deeper customization of the testing environment; however, it also requires a lot more work and skill to build a system that can properly identify all behavior displayed by an artifact. One example of this work is how a buddy of mine and I created isolated virtual images for Windows 3 and OS2 with the goal of seeing how current ransomware would execute on older systems. There are

no Windows 3 or OS2 sandboxes we could find publicly available and we had to use very limited monitoring tools to try to understand what the ransomware samples would do when run in these environments. The same results we obtained after weeks of research could have been obtained within a few minutes if we were able to use a modern sandbox.

Joe Sandbox

One free online sandbox you can use for dynamic analysis of artifacts is Joe Sandbox (https://www.joesandbox.com). After registering to Joesandbox.com, you will be able to upload any artifact to have the researchers at Joe Sandbox run that artifact through their sandbox (limited to 15 analyses a month/5 a day). Within a few minutes, you will receive an email with the results of their research, including if the artifact is malicious and even a video of what the artifact looks like if you run it on a system. By including Joe Sandbox in your SOC practice, you can claim to have basic dynamic analysis capabilities at no cost to your organization! Figure 8-18 shows an example of running an artifact in Joe Sandbox and receiving a verdict, including a video supporting why that verdict was made.

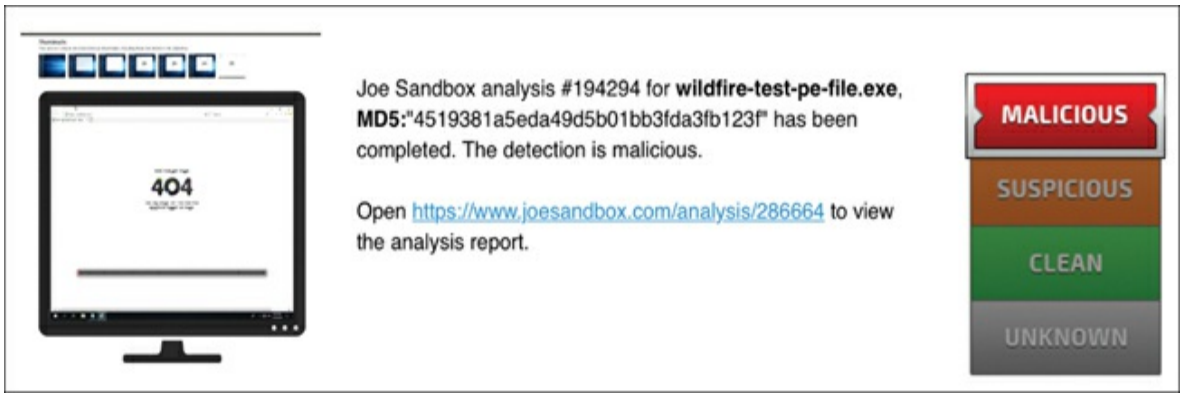

FIGURE 8-18 Example of Using Joe Sandbox for Malware Analysis

Note

The free version of Joe Sandbox publishes the analysis of an artifact for anybody to access. If you want privacy, you will need to upgrade

your Joe Sandbox account.

Other Sandboxes

There are many more advanced sandboxes available that you can run in the cloud or in your local environment. I highly recommend using an enterprise version of a sandbox if you are new to this field of study. There is always the risk that malware can breach your sandbox and lead to a security incident. Enterprise options tend to include support and training designed to help you be better prepared to research malware.

Open-source sandboxes also are available. One of the most popular options is Cuckoo (https://cuckoosandbox.org/), which is great but also complex. I personally did not find a great resource for building a Cuckoo sandbox and also found many YouTube videos misconfiguring Cuckoo sandboxes, which would lead to the potential of threats escaping the sandbox. For this reason, I highly encourage you to be careful when working with sandboxes and err on the side of caution, such as using a cloud option like Joe Sandbox that moves the risk from your organization to the cloud whenever possible.

Phase 2: Detection and Analysis Summary

The list that follows summarizes the key points to keep in mind for the detection and analysis phase of the incident response lifecycle:

- Incident detection is based on either reported incidents or alerts from breach detection capabilities made up of tools that use signature, behavior, or anomaly features.
- Threat analysis determines if the artifacts involved with an incident are an actual threat and, if so, the type of threat.
- Static analysis involves learning about the artifact without running it.
- Dynamic analysis involves running the threat in a controlled environment to learn about its behavior.
- Most modern malware is packed and therefor difficult to disassemble.

- Artifacts might look like a specific format but really be something else (for example, an executable with a .pdf extension). Therefore, all artifact formats must be validated to confirm their true format type.

Phase 3: Containment, Eradication, and Recovery

Once an incident has been discovered and confirmed as being a real threat, the third phase of the incident response program is to launch steps to counter the threat. Those steps should be documented and converted into different playbooks to ensure a repeatable process is executed as your incident response service learns from different incidents that are resolved.

The type of response your incident response service will execute will depend on various factors. Usually, the first decision the team needs to make is whether legal action might be required post incident response. This decision is critical for determining all incident response actions that follow; incidents with potential legal action require a digital forensics–based response, while incidents that don't involve potential legal action can be addressed with a "fix it as fast as possible" response. A digital forensics–based response can also include fix-it steps, but the incident response team first must carefully consider how they handle potential evidence, which likely includes a different set of skillsets found within the digital forensics SOC service. A simple example is if a system infected with malware needs to be preserved in its current state for purposes of evidence in a potential legal action, the response might include isolating the system, not powering it off, and having its hard drive copied so the copy can be analyzed by a forensic professional. If legal action is not a concern, a desktop support member could be authorized to reimage the system. Reimaging the system can return the system to operational state; however, that action also removes any digital evidence that could have been used during a forensic investigation. You will learn more about digital forensics concepts later in this chapter.

If the incident response team determines that digital forensics is not needed, the next step is to focus on containing the threat. Containment means finding and isolating all systems impacted by the threat. If you can't contain the threat, you will not be able to resolve the impact because the threat can continue to spread and cause chaos. The process of identifying all systems

impacted by a threat is commonly referred to as threat hunting. This brings us to the first part of the third phase of the incident response process, which is containment.

Containment

Containment means to prevent the threat from impacting other systems. Containment requires a method to identify all threats or containment can't properly be implemented. To identify threats you will need to perform threat hunting. *Threat hunting*, sometimes referred to as *blue teaming* when conducted in a testing lab, means to search for attack behavior and malicious artifacts within the environment you are assigned to protect. I point out the term blue teaming because in an attack-and-defend lab, the red team would be delivering attacks while the blue team would be looking for the red team's behavior. In a real organization, the SOC would be responsible for hunting for threats within all areas that fall under the SOC's purview. This can include systems within the organization, remote users, and cloud services, depending on the scope of the SOC's incident response service.

Threat hunting must be very thorough in regard to identifying the scope of an incident. If you can't identify the scope of the incident, you will never be able to contain it. This holds especially true with advanced persistent threats (APTs). As described in Chapter 6, more advanced threat models include a step called establishing a foothold, which means to gain access to a network. It is common for malware to spread the number of footholds within a compromised network so that if one foothold is removed, the attacker still has access to the network using another foothold. As a threat hunter, you need to identify all footholds before you can conclude that you truly have scoped out the entire incident. Closing the incident response prior to identifying all impacted systems not only leaves the attacker with access to your network to continue the attack unnoticed, but also provides a false sense of security to your SOC and organization.

To understand how to scope an incident, you need to know how to look for threats. The next topic of focus is understanding malware behavior.

Responding to Malware

Earlier in this chapter, you learned about both static and dynamic analysis techniques, which you can use to glean what an artifact is intended to do when executed. Your goal from analyzing malware is to develop a fingerprint that can be used to hunt for other impacted systems containing the same signature. One approach to hunting for malware is to develop behavior signatures based either on the behavior patterns you learned while researching, signatures provided by another tool, threat intelligence, or performing dynamic analysis on the artifact of interest.

Note

There is no universal system for naming malware. Malware typically is named by whomever first discovers the malware and publishes its details. Names such as Code Red or WannaCry are made up and have no actual meaning. Code Red came from a flavor of Mountain Dew that the researcher who discovered it liked to drink.

Malware Categories

Malware is identified based on what vulnerabilities it exploits and the category of malware it falls within based on its behavior. The following are the most common categories of malware along with descriptions of their behavior signatures:

- **Downloaders and launchers:** Malware that downloads other software. This type of malware is typically used for initial compromise of a target.
- **Adware:** Malware that presents unwanted advertising to user. This type of malware is often included with free software and browser toolbars.
- **Backdoors:** Malware that allows an attacker to connect to a compromised system and sometimes even take control of the system remotely.
- **Rootkits:** Malware that is designed to conceal the existence of other malware. Often rootkits hide backdoors.

- **Botnet:** A network of remotely controlled private computers with backdoors controlled by a command and control (C&C) server. Botnets often send spam and perform distributed denial-of-service attacks against targets.
- **Ransomware and scareware:** Malware that scares the victim into purchasing something or paying a ransom. Scareware typically blackmails the victim, while ransomware typically encrypts the victim's data.
- **Worm:** Malware that replicates itself to spread to other systems
- **Virus:** Malware that replicates itself into other applications, files, or even the boot sector. Viruses can harm computers, steal data, or perform other nefarious actions.
- **Spam malware:** Malware that sends spam from compromised systems. It is common for spam malware to be part of a botnet that could also perform DDoS attacks or use the resources of infected systems to create cryptocurrency.
- **Keyloggers/data stealing malware:** Malware that is designed to capture specific data, such as keystrokes for a password, credit card numbers, bank account numbers, and other personal data.

Phoning Home

Each category of malware has specific behavior associated with how the malware functions. Earlier in this chapter, you learned about common malware behaviors. One example of a common behavior associated with ransomware is phoning back to a malicious resource. This occurs when the ransomware uses asymmetric encryption, also known as public/private key sharing. In order for the encryption process to occur, the ransomware needs to perform a key exchange with the attacker's system located outside of the target's network. By identifying how this communication occurs, you could use detection tools to look for any signs of similar communication as a way to identify systems that have been impacted by the ransomware. This ransomware example isn't an ideal approach to responding to ransomware, but the lesson is that most modern malware has some form of outbound communication, which could be a backdoor alerting the attacker that a port is open, a virus downloading part of its payload, or any one of many other

behaviors.

Port Scanning

Another common behavior of modern malware is port scanning. This occurs when the malware is attempting to learn about the environment it has compromised so that it can find other systems and networks to spread to. Many behavior-based security tools monitor for unauthorized port scanning with the goal of identifying systems that have been compromised by malware. Figure 8-19 shows an example of a NetFlow-based security tool, Cisco Stealthwatch, mapping out compromised systems based on their port scanning behavior. It is important to point out that the 198.19.10.3 IP address in this example is named “scanner” and should be removed from this report because it is a vulnerability scanner used by the SOC!

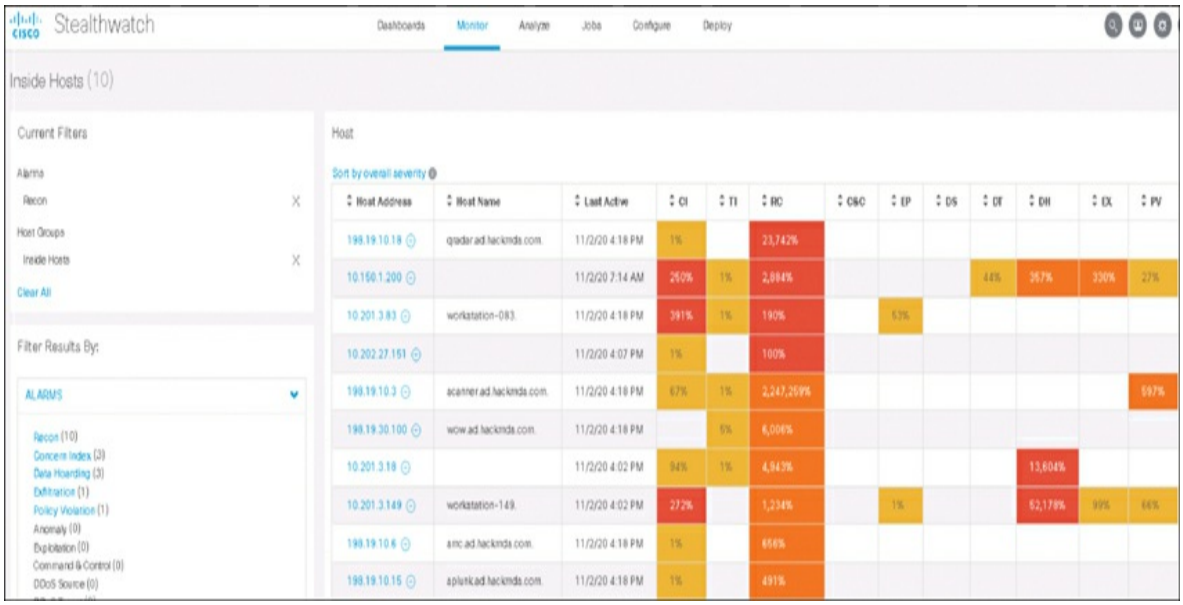

FIGURE 8-19 Cisco Stealthwatch Identifying Threats Based on NetFlow

Matching Hashes

A third very common approach to hunting for threats is matching a hash to a malicious artifact. A hash can be used as a digital fingerprint of a file, which can be searched for using an endpoint detection and response tool or other

endpoint tool to look for traces of the malicious artifact. Network tools such as application-layer firewalls can also be used to search for traces of artifact by scanning for the hash against network activity. It is very common practice for a threat hunter to pull hash values of threats found during the static or dynamic analysis of malware and hunt for a pattern match within various security tools. This is why many of the examples I demonstrated earlier in this chapter of artifact analysis included generating SHA hash files. The "Digital Forensics" section later in this chapter describes hashing in further detail.

Threat Hunting Techniques

There are many other specific techniques that can be used to identify threats based on behavior. A general list of things to look for regarding the detection details obtained prior to starting a threat hunt includes the following:

- Number of incidents associated with the incident, ranked by severity.
- Dwell time of any incident discovered, meaning how long the threat has been active.
- Number of detection gaps filled, meaning how many sources are stating the threat is real.
- What gaps in logging have been identified to understand what may not be known. This is important for understanding potential gaps in available data prior to starting research.
- Vulnerabilities identified.
- Insecure practices that were identified and should be corrected.
- Number of hunts that transitioned to new analytics as research was performed.
- The false positive rate of transitioned hunts.
- Any new visibility gained about the incident.

Performing Threat Hunting

As you perform threat hunting, you will want to include a few specific techniques as part of your process. The first technique is performing searching and mining of all possible data, known as data mining. The more data that is available, the better the results. As data size increases, you will want to lean on clustering techniques, which are often carried out by machine learning. Clustering looks within a large dataset for similar data points based on certain characteristics and eliminates redundancy to reduce the dataset to only what needs to be analyzed to understand an incident.

Grouping is another approach used to reduce what needs to be analyzed by a threat hunter. Grouping takes a set of multiple unique artifacts and links them together based on a specific criterion. The major difference between grouping and clustering is that in grouping, your input is an explicit set of items that are already of interest, while clustering is used when you don't know what to look for. Discovering groups within your research leads to identifying a tool or tactic, technique, and procedure (TTP) used by an adversary.

Stack counting is a common technique carried out by threat hunters investigating a hypothesis. This approach works by developing a hypothesis and counting how many occurrences are found as well as any outliers. This approach doesn't work well with large and/or diverse datasets but is extremely effective when proper filtering is applied to help zero in on the datapoints of interest.

By using the preceding techniques and available data, you will aim to answer the following questions about the incident:

- How is the attacker or malware moving around your network?
- How many systems have been impacted?
- What vulnerabilities need to be addressed to avoid future events?
- How is the attacker utilizing his or her tools?
- How is the attacker determining where he or she is going?
- Can the malware survive a reboot or simple remediation steps?
- Has the attacker stolen any data and, if so, what data was impacted?

The results of threat hunting will be based on the available data, how much experience the hunter has, and the type of tools the threat hunter can leverage

for the research. As all three of these improve, the maturity of your threat hunting program will increase. Figure 8-20 represents a conceptual maturity model mapping what is associated with maturing a threat hunting practice. By level 4, huge datasets can be quickly evaluated based on the techniques previously covered. This strong practice leads to strong verdicts based on many obtained factors that answered the questions threat hunters aim to answer.

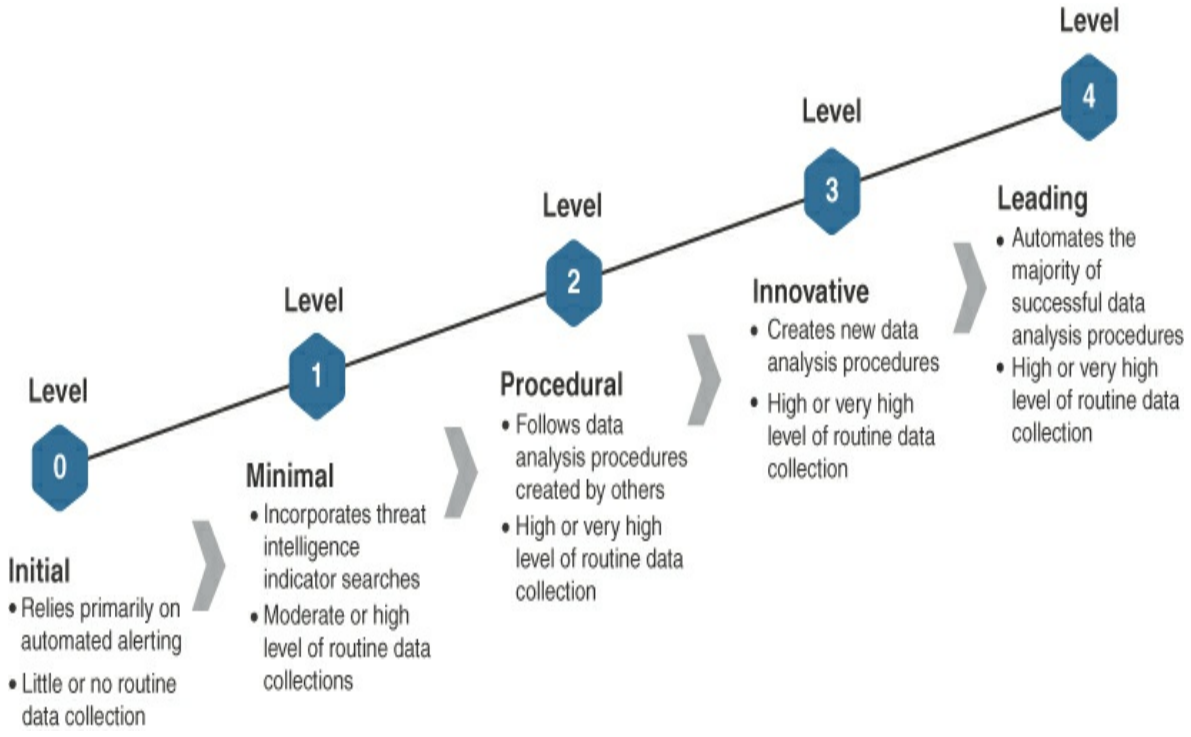

FIGURE 8-20 Threat Hunting Maturity Model

Threat Hunting Example

Let's walk through a short example of threat hunting using mining, clustering, grouping, and stack counting. For this example, I am tasked with investigating a potential malicious domain called success20.hopto.org (warning: this website is malicious, so don't try this unless you have the right tools for testing!). I'm using Cisco Threat Response to perform threat analysis of the suspicious domain in this example, but similar tools are available, such as Sqrrl. Figure 8-21 shows using Cisco Threat Response to

research success20.hopto.org.

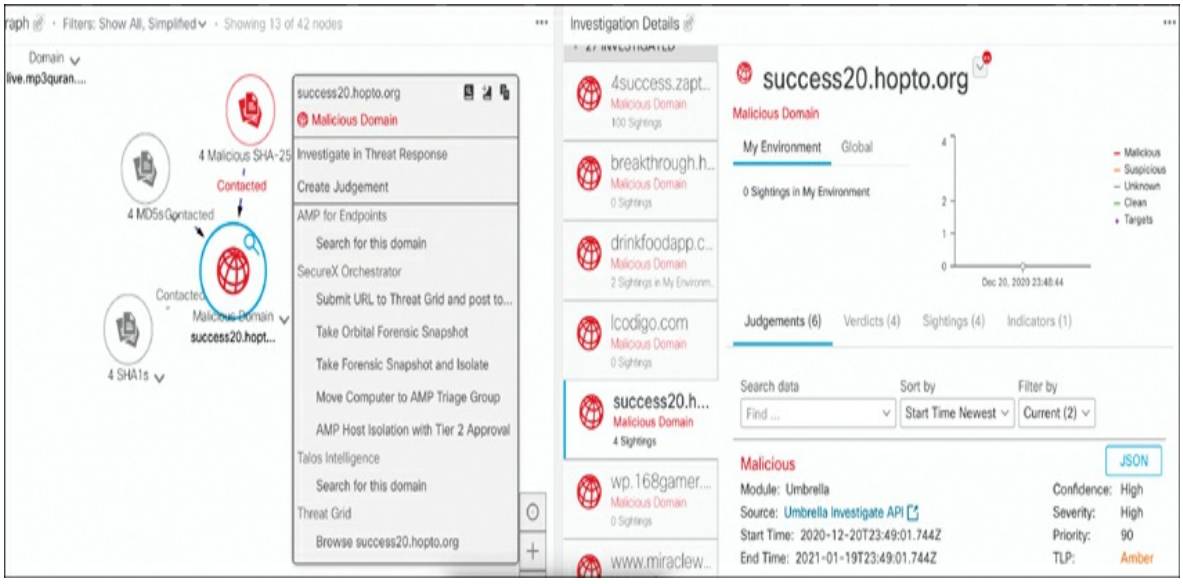

FIGURE 8-21 Cisco Threat Response Analyzing success20.hopto.org

Mining is represented by all of the datapoints that are pulled from investigating the suspicious domain. I can see different malicious files, domains, and SHAs that all have some relation to success20.hopto.org, represented as “Grouping” because each group is based on a shared datapoint. Figure 8-22 shows the grouping of other domains that are associated with success20.hopto.org via the same domain owner. This is useful for learning about all domains associated with a potentially malicious domain to block everything owned by the domain owner. It is very common for the domain owners of malicious domains to own multiple malicious domains. This allows them to cycle through domains as the domains are flagged as malicious and blocked by security tools.

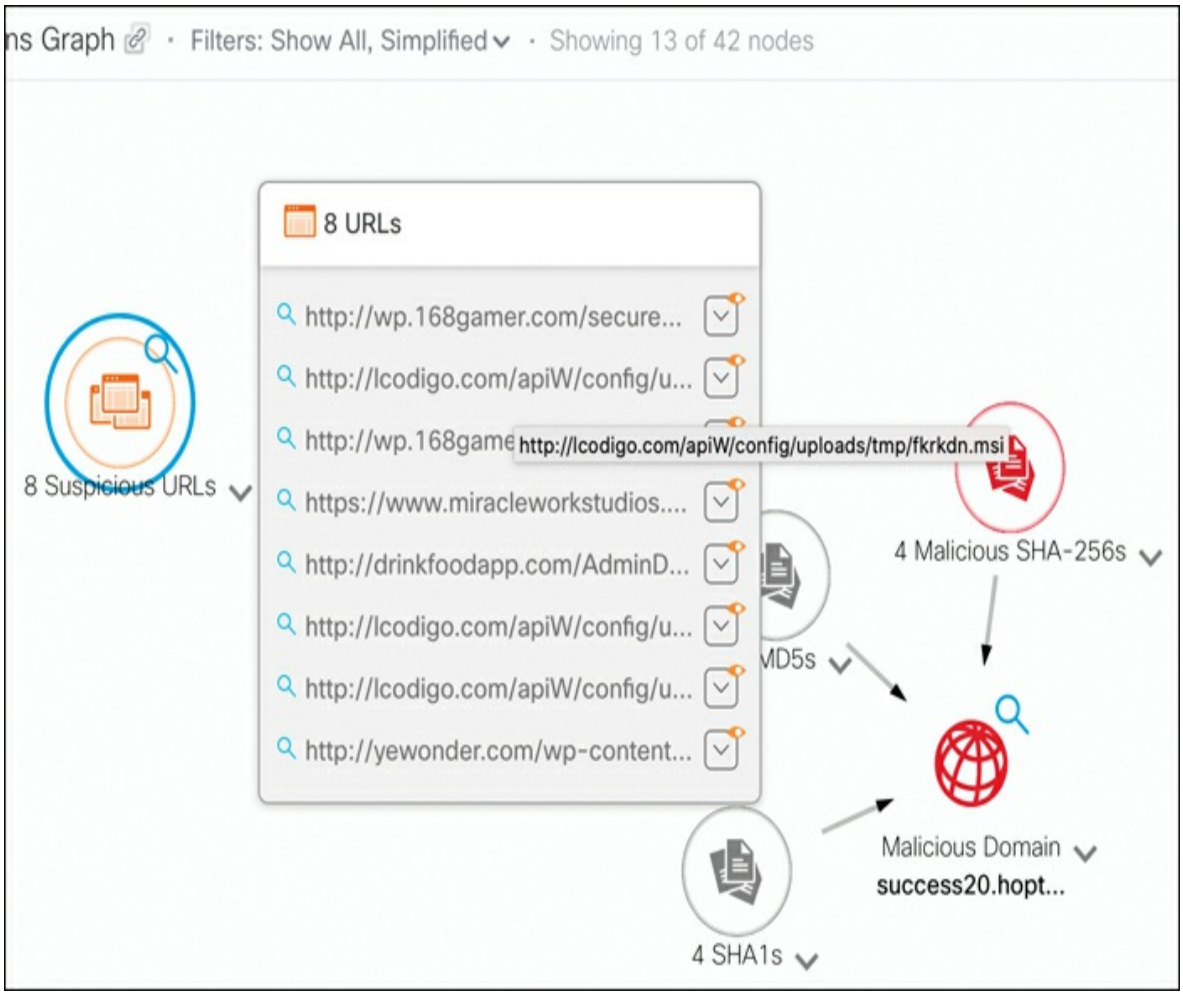

FIGURE 8-22 Other Domains with Association to success20.hopto.org

If I want to use stack counting, I can search for a specific thing, such as count how many domains are associated with the domain owner of success20.hopto.org. If I'm looking for other suspicious domains associated with the domain owner, I would create a stack, which would find eight other URLs as shown in Figure 8-22. Stacking is more useful when searching across a large dataset with expectations of having multiple hits.

Clustering would be represented by searching across a lager dataset for similar datapoints to see if a new relation could be established as well as to reduce what the threat hunter needs to look for. Figure 8-23 shows searching for interaction with success20.hopto.org within different security tools and threat intelligence data available within my SOC to see how many sightings have been identified. Based on the number of sightings and judgments (risk

scoring) that show success20.hopto.org is malicious, I can quickly determine the website and its associated artifacts are all malicious, allowing for a quick validation and summary of what the SOC needs to block.

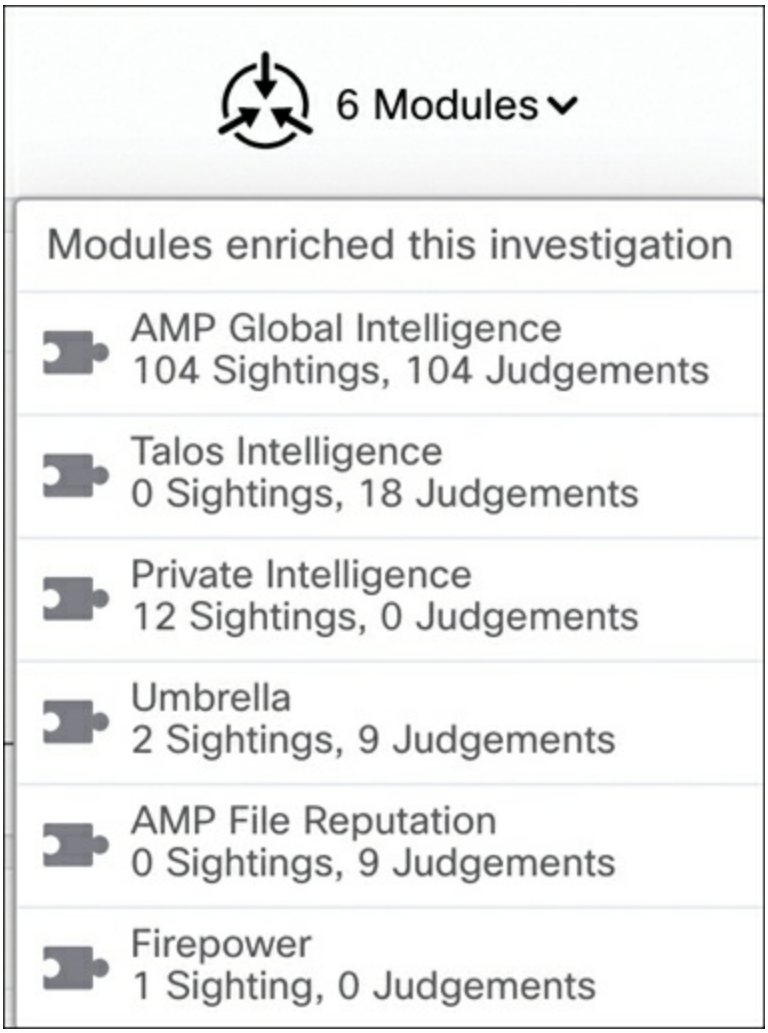

FIGURE 8-23 Searching success20.hopto.org Across Multiple Resources

This is just a simple example of how different datapoints about a suspicious domain can help a threat hunter lead to a conclusion about what is being researched. These techniques also help identify other associated risk, allowing for a complete analysis of all potential risk with an artifact, domain, or other topic that needs to be analyzed. Security tools can make research easy, and for very large datasets such as big data, data aggregation and correlation are required. Chapter 6 covered how to work with big data.

Note

Websites such as Malware Domain List (http://www.malwaredomainlist.com/mdl.php) and Internet Storm Center (https://isc.sans.edu/suspicious_domains.html) are resources that provide indicators of compromise (IOCs) you can use for research purposes. Be cautious when using these and similar resources as they contain real malicious domains!

Eradicate

Once your incident response team has contained the threat through threat hunting and segmentation/isolation of impacted systems, it is time to perform steps to eradicate all aspects of the threat. All aspects of a threat include all impacted systems and all vulnerabilities that were used to initiate the compromise and spread the threat. For example, a specific exploit could be used to gain access to one system, but another exploit could be used on that same system to get a domain login, allowing for malware to spread to other systems. In this example, eradication would include removing the malware, changing the domain passwords, and providing a response to the vulnerabilities that permitted the initial access to the first victim (commonly called “patient zero”) as well as the vulnerability that led to the larger compromise. The eradication step assumes you have contained the threat.

System Order

When performing eradication steps, you want to start with the directly impacted systems. These are the systems that are confirmed as compromised and must be addressed to reduce the risk of further spreading of the threat. Next, all systems that are or could be in contact with any impacted system must be evaluated to ensure they have not been compromised. By confirming all systems within reach of all infected systems are not compromised, your incident response team is validating the *possible* scope of the breach.

I specifically point out you are possibly scoping the breach based on how malware can jump between systems and not infect every system contacted during a major outbreak. Many forms of malware will look for specific indicators before launching an attack, the goal being to avoid deception technology such as honeypots and limit exposure to only systems that are a high percentage of successful compromise, meaning only attacking systems that are likely to be breached. I explained this concept earlier in the chapter when reviewing dynamic analysis techniques. With this concept in mind, after validating all systems within contact of any infected system, you will need to run a network-wide scan for any threat that was found during the incident response. You can do this using the hash search technique previously covered or behavior signatures developed as the associated malware is analyzed. Figure 8-24 represents the recommended approach to responding to impacted systems during an incident response.

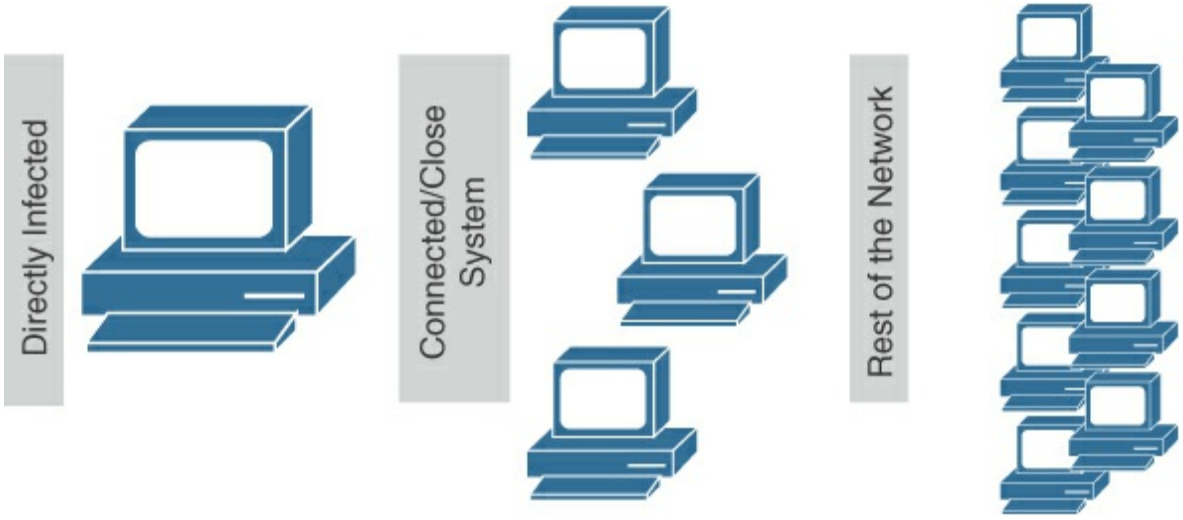

FIGURE 8-24 Order of Systems to Address

Containment needs to include a method to prevent the threat from spreading. Chapter 2 covered different types of segmentation techniques. I recommend complete isolation if possible, as systems on the same network segment can spread malware. Isolation of network resources should also be enforced to prevent the malware from sending data off the network. This action can alert the adversary that their malware has been identified, which means you want to execute swift containment of all threats to avoid the adversary adjusting

their attack plan based on being alerted that the malware is being contained.

Eradication Playbook

The steps involved to eradicate a threat will depend on the type of threat and damage that it caused. Returning to the IRC Malware Outbreak playbook template introduced earlier in the chapter, Figure 8-25 shows IRC's view of what eradication would look like for a malware outbreak. Preventing the spread includes disabling services, restricting access to resources, adjusting firewall rules, blocking services and files with antivirus, and patching systems. The Eradicate step of the Malware Outbreak playbook also includes various forms of communication that need to occur with all parties impacted by the threat, so that they know what to expect as remediation is being applied to their systems. The playbook has four specific steps to eradicate the malware, which include either to clean or quarantine the malware with antivirus, use a malware removal tool, or manually intervene.

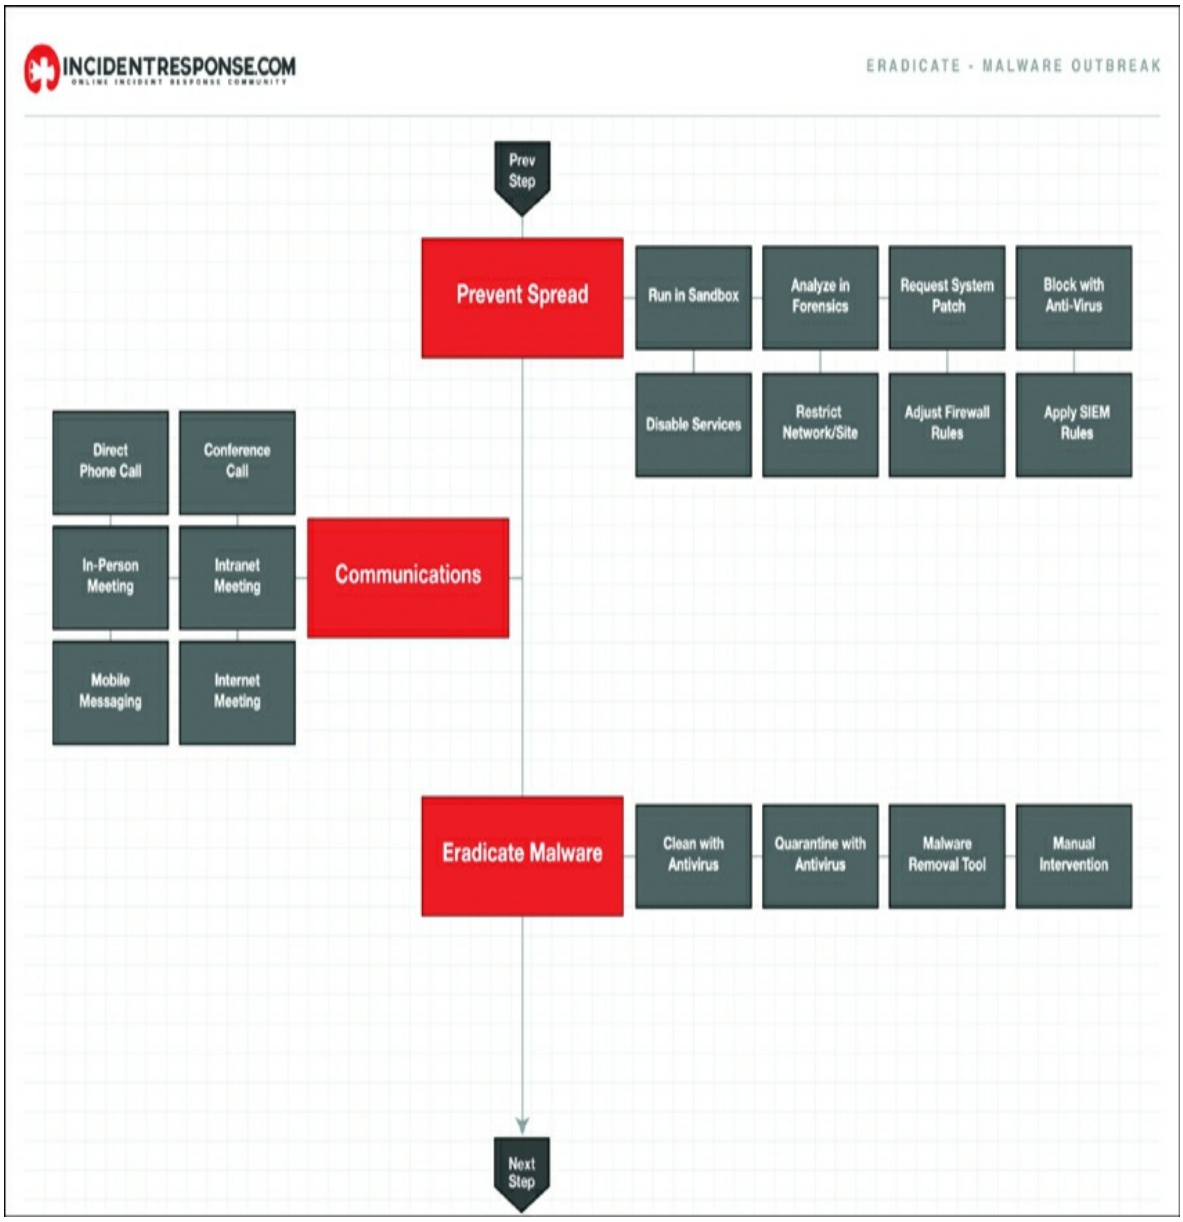

FIGURE 8-25 IRC Malware Outbreak Playbook Eradicate Step

Another IRC playbook example is the Eradicate step of the Data Theft playbook, as shown in Figure 8-26. This playbook example first focuses on learning how the incident occurred and confirming what is known. Communication once again is part of the playbook, informing those responsible or impacted by the data breach. Finally, any malware is eradicated, and the network is monitored for ongoing data loss to ensure no future data is being taken post remediation steps.

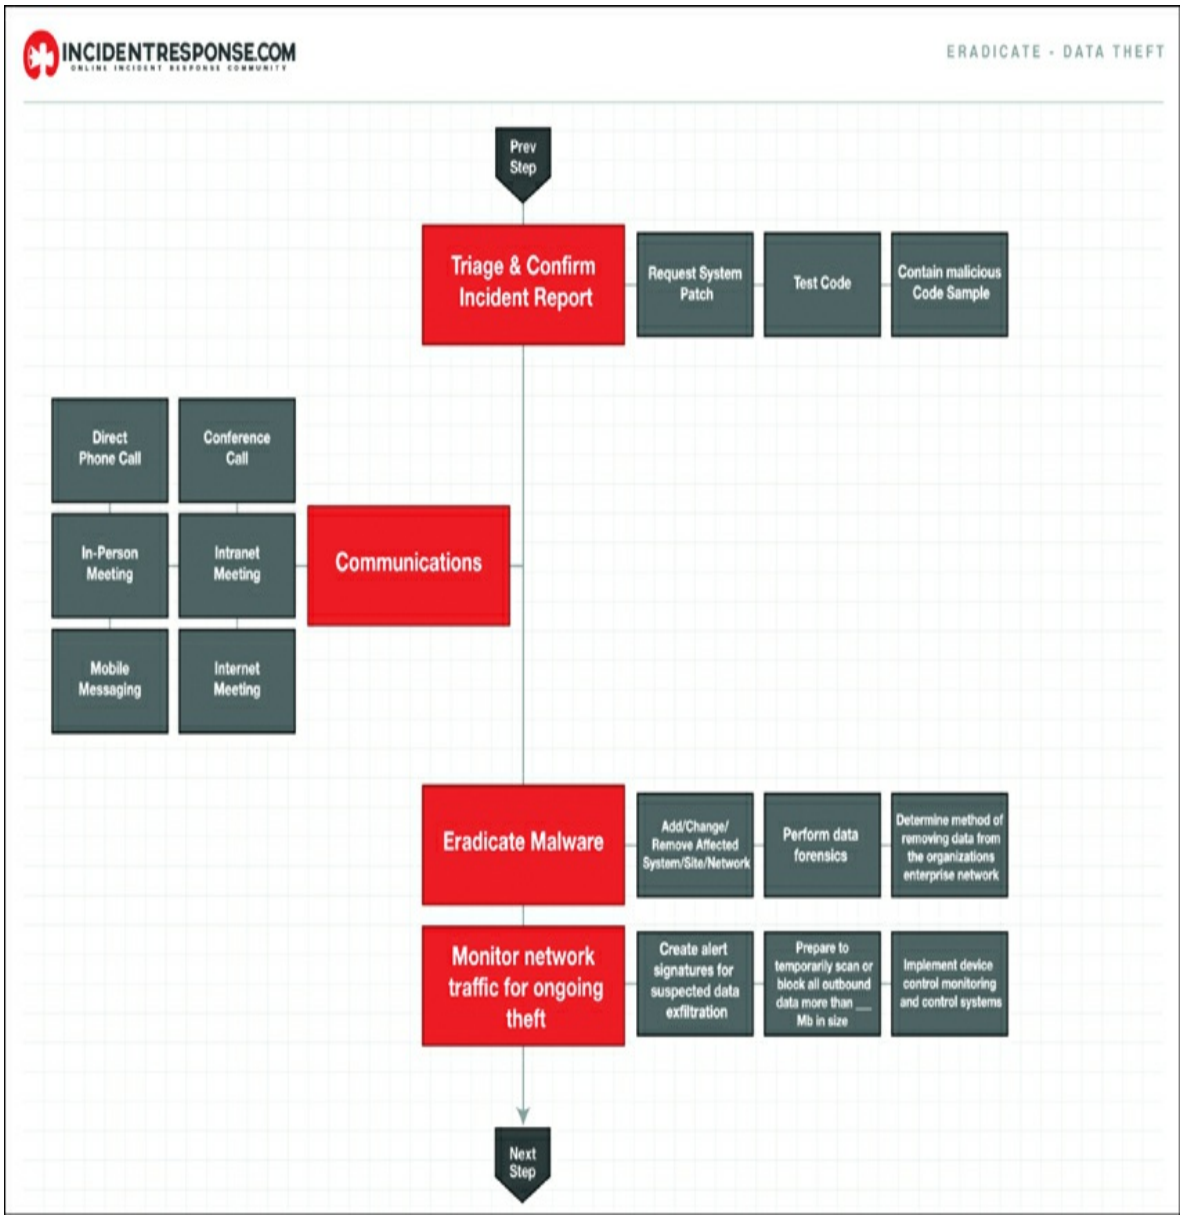

FIGURE 8-26 IRC Data Theft Playbook Eradicate Step

In general, eradication must include a way to remove the threat, contact all impacted parties, and ensure measures taken completely remove all aspects of the threat to avoid further compromise and damage. The SOC incident response team should expect to engage with other SOC teams to execute various remediation measures required by different types of incidents. Playbook templates such as those offered by the Incident Response Consortium are a useful starting point but must be customized based on real-world business usage within your organization. Other guidelines such as

NIST SP 800-83 Rev. 1 (covered in the next section) can be great resources for eradication templates.

Recovery

Recovering from a breach means to return all impacted systems back to an operational state and remove any temporary containment measures. This can include anything from reimaging systems that are corrupted by malware to rebuilding lost data and restoring data that was recovered from backups or other resources. Temporary security measures put in place during the containment and eradication phases can include isolating systems, shutting down services, and restricting privileges. As threats are eradicated, these security measures may or may not be removed depending on their impact to business and the likelihood of the threat not being fully contained.

Questions that need to be addressed as the recovery phase is executed are as follows:

- When can systems be returned to production?
- Have systems been patched, hardened, and tested?
- Can the system be restored from a trusted backup?
- How long will the affected systems be monitored and what will be the scope of the monitoring?
- What tools (e.g., file integrity monitoring, IDS/IPS, etc.) will ensure that similar attacks will not reoccur?

The decision to remove temporary containment measures, such as suspended services or connectivity, must be made only after containment and eradication have been confirmed. Additional validation, including time to monitor impacted systems, might be needed depending on the severity of the breach and sophistication of the threat. The decision to remove containment measures is a tough one because you do not want to reconnect systems and cause another outbreak. As a safety measure in the event that dormant threats are still present, I advise that the incident response team keep any containment measures in place as long as possible without causing major disruptions to operations outside of what has already occurred post incident.

Recovery Playbook

The steps involved with your recovery phase will depend on the type of event, which is the determining factor in which playbook you should follow. Returning again to the IRC Malware Outbreak playbook, Figure 8-27 shows the Recover step of the playbook. Notice this template includes a few recovery steps focused on IT systems impacted by malware as well as a step specific to removing temporary containment. Once the recovery process is completed, recovered data is returned to systems and the incident remediation steps are performed. These steps include some similar recommendations found within NIST SP 800-83, including updating security tools and security policies.

FIGURE 8-27 IRC Malware Outbreak Playbook Recover Step

This section focused on containing, eradicating, and recovering from a security incident. Steps such as wiping a system are ideal for removing malware; however, this would also delete any evidence of the malware.

When legal actions are possible, the previous incident response steps should not be taken. Instead, you will need to execute a digital forensics–based approach to your incident response.

Digital Forensics

As indicated earlier in this section, your approach to containment, eradication, and recovery will be completely different if the incident may involve potential legal actions, in which case a digital forensics–based approach is required. The specific approach will be based on the legal system of the jurisdiction in which your organization operates. The rules of evidence in many legal systems require digital evidence to be secured in a specific manner in order for it to be admissible. The rules usually are very stringent because digital evidence is considered extremely volatile, meaning it can easily be manipulated. Even if the *possibility* exists that digital evidence has been manipulated, most courts will deem the evidence contaminated and thus inadmissible in the case. For this reason, if your SOC decides that an incident might require legal action, it is critical to immediately implement a digital forensics–based approach to your incident response. This approach requires a different playbook than a standard incident response.

Note

You might be wondering how the SOC determines when to use digital forensics versus launching a standard incident response. Some situations will be obvious, such as when personally identifiable information (PII) has been stolen and legal authorities need to be notified. My suggestion for situations that are not obvious is to ask one question: Do you think legal action could be needed, which would require evidence that would be required for use in court? If the answer is yes, you need to use digital forensics.

One example of the difference between a standard response and a digital forensics–based response to an incident is the latter’s focus on chain of

custody (explained in detail a bit later in the chapter). If at any point custody of evidence is not documented, such as during transit or storage, a court will assume it was at risk of contamination and deem it inadmissible. I have been involved as an expert on both sides of cybercrime litigation, and when I'm working with the defense, I look for mistakes in chain of custody to get evidence removed from the case. I was involved in one case where the evidence was locked up, but the lock was a very basic door lock within an office building. My team successfully made the argument that a more secure method should have been used, such as a safe to secure the evidence, to prove the chain of custody. There was no proof that the office was locked while the evidence was stored, raising the possibility that somebody could have entered the office, modified the evidence, and locked the door behind them. There was also the possibility that somebody could have asked a cleaning person for a key to be let in or picked the lock with a basic lock pick kit. As you can see, if a strong argument can be made to question the security of evidence, critical evidence can become useless.

Chain of custody is just one of the critical requirements involved with performing a digital forensics–focused incident response. Next, I'll cover what I believe needs to be involved with all digital forensic efforts.

Digital Forensic Process

A digital forensic process is different than a traditional incident response process. Every step of the digital forensic process must be executed correctly, or the entire process can result in contaminated evidence. The following is my high-level overview of a properly executed digital forensic process:

Step 1. Identify that a crime might have been committed.

Step 2. Collect preliminary evidence.

Step 3. Obtain a court warrant for seizure of evidence if required.

Step 4. Perform first responder procedures.

Step 5. Seize evidence at the crime scene.

Step 6. Transport evidence to a forensic laboratory.

Step 7. Create copies of evidence.

Step 8. Generate images and examine them for evidence.

The first step is to identify a crime. Then you need to validate if the crime is real and worth investigating by collecting preliminary evidence. You also will determine what you can and can't access based on physical or legal restrictions. Many countries have specific laws that protect people's privacy, and if you violate those rights, not only will your evidence be useless, but you could also be charged with a crime for privacy violation. One example of a violation in the United States is going into a person's office or looking into a computer without authorization from the owner. These actions require a warrant unless the person has given up their rights to privacy. Why would somebody give up their privacy rights? Many organizations provide corporate-issued equipment to employees; however, the employee must agree to give up some privacy rights while using the equipment. This hypothetical situation would allow an organization to investigate a corporate-issued device without the employee's permission, assuming that is how the law works where that organization resides.

Note

I am not a lawyer and I recommend that you review your local and federal laws with an authorized legal professional before deciding to implement or take any action that could be a violation of privacy rights.

Preliminary evidence can be used to justify a warrant. There are different types of warrants and different reasons for a warrant being issued. Some warrants are issued out of concern that critical evidence will be destroyed before proper legal action can be taken. Some warrants are issued because privacy issues are blocking an investigation. The concept of legal warrants is out of scope for this book, but it is a concept you need to understand regarding your legal system and when a warrant should be used. My advice is if you believe a warrant might be needed, seek legal advice or engage law enforcement.

Once you are authorized to engage evidence, you perform first responder procedures, which protect any potential evidence from being contaminated and allow for the proper collection of evidence. Evidence is moved to a forensic lab where copies are made, and only the copies are investigated, not the original. Evidence is collected and a report is created determining what was found. That report can be used for legal action as well as for a better understanding of what occurred regarding the security incident and its impact.

All steps must be performed properly and in order for the evidence to be legally admissible. The most crucial step in my experience, and where I find organizations make the most mistakes, is the first responder step.

First Responder

The first responder is the person who has the first real engagement with the incident. The first responder's job is to determine whether a crime has been committed based on all preliminary evidence, decide whether in-house, third-party contracted experts or law enforcement should be engaged, and prevent any unauthorized access to the crime scene. This includes taping off the area associated with the crime and preventing anybody from touching anything. The first responder might be the same person that collected the preliminary evidence or might be engaged once another team was alerted about the crime and determined that there is the potential for legal action, which turned the situation into a forensic investigation.

Note

The first two steps of the forensic process I outlined are part of a general incident response plan rather than always considered a forensic–focused response. Once legal action is considered a possibility based on the preliminary evidence, the incident response process is flipped to a forensic process and a forensic specialist is engaged to perform the first responder steps.

The first responder documents the current state of the crime scene, representing the state of the crime scene prior to the start of the official forensic investigation. It is common to use a camera and photograph everything from the devices impacted by the threat to the area surrounding them. It is critical that this is performed to avoid a successful claim by the defendant's defense team that items were changed or removed during investigation outside of what was documented by the forensic investigator.

All evidence that is authorized to be collected is collected without changing its state, unless there is no other way to collect the evidence. It is critical that systems that are powered on remain powered on and that systems that are powered off are never turned on. Doing so changes the state of the system and is considered a form of contamination. There are professional tools that can be used to keep a system powered on as it's moved from a wall power source to one that is mobile. In the industry, the terminology for this part of the investigation is "bag and tag," meaning put what you find in a bag, secure that bag, and label it. At this point, anything that is collected is now the responsibility of the forensic team until it is returned to the owner. Everything must be documented, secured, and monitored at all times to maintain the chain of custody of evidence.

Chain of Custody

Chain of custody is a term used in forensic practices referring to the logical sequence that records the custody, control, transfer, analysis, and disposition of physical or electronic evidence in legal cases. Documenting each step in the chain of custody of evidence is essential because if uninterrupted custody cannot be proven, the evidence might be rendered inadmissible. This means chain of custody occurs at the point of collection and is maintained until the artifact is returned to the owner.

The first key part of the chain of custody process is identifying what can be collected. What should *not* be collected is any evidence that you are not authorized to collect or that is not relevant to the case you are building. For evidence that you can't collect but could apply to the case, you will want to seek a warrant from the proper legal authority to collect the artifact. For artifacts that are not relevant, you will need to mark them as processed and

not relevant in your Prepare/Extracted list. Doing so ensures they are not collected when the extraction process begins as well as includes their existence in the first responder artifact collection process. A prepare/extracted list is a document the first responder uses to report everything that exists within the crime scene to clearly define what does and does not need to be collected.

Whatever is relevant needs to be documented in a relevant data list. Documentation must include pictures, descriptions, and other recording allowing for a complete picture of the artifact that will be investigated. It is ideal to upload all of this documentation to a forensics case management tool if one is available. As data is gathered about an artifact, it is very possible that new artifacts or leads to new artifacts of interest will be discovered. For example, you can identify a computer and find other devices, such as a phone, are connected to it, leading to new potential evidence. Now you will need to repeat the documentation process and treat the phone as a new item with its own relevant data list. This entire process is repeated until no more artifacts need to be addressed. Figure 8-28 shows the steps taken prior to extracting artifacts from a crime scene.

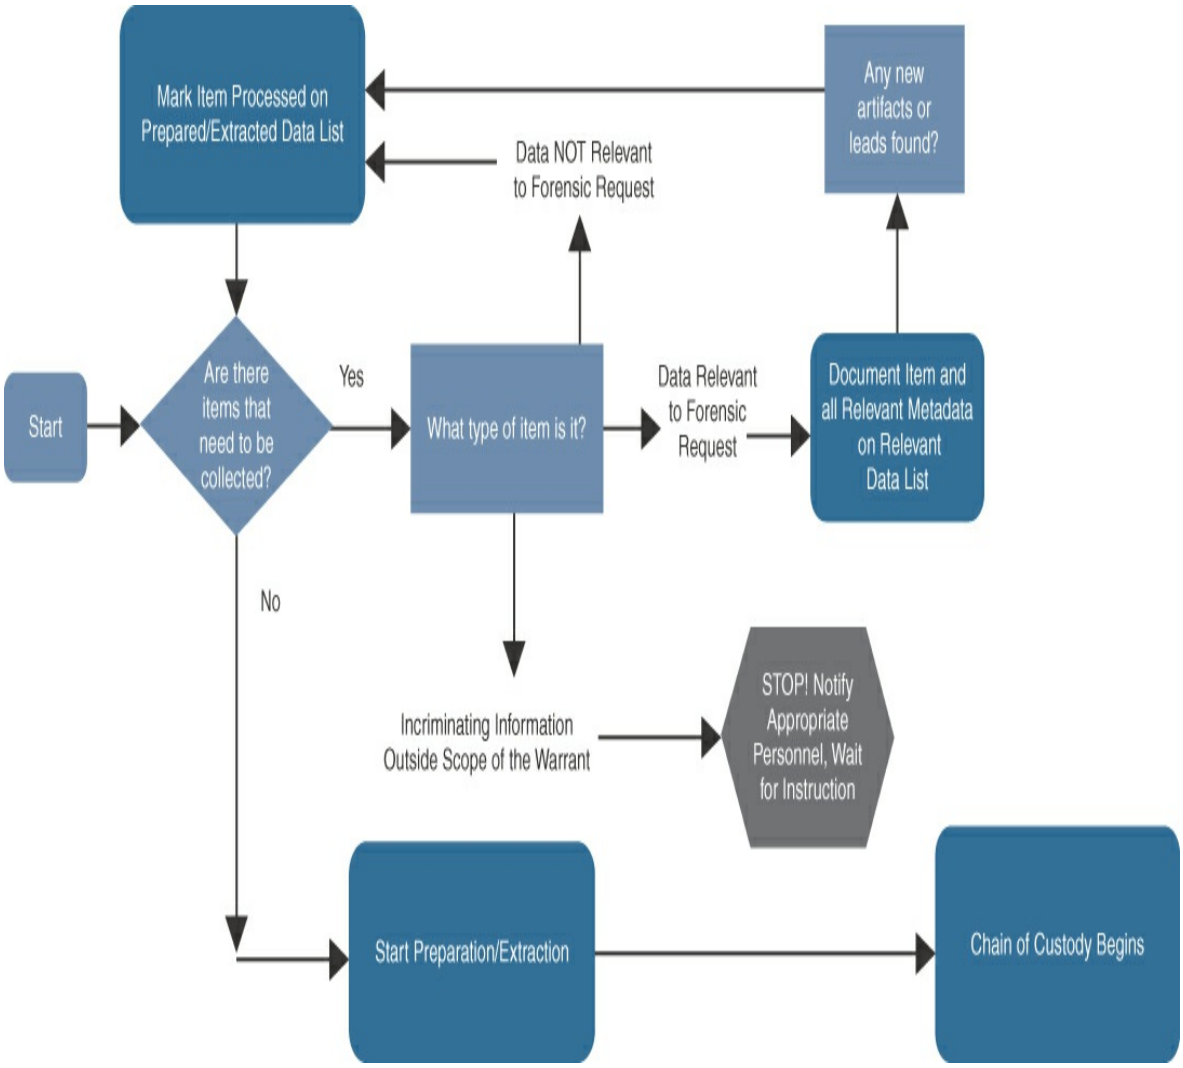

FIGURE 8-28 Chain of Custody Documentation Bag Example

Chain of Custody Process

Chain of custody occurs during the entire forensic investigation process, which can be broken down into the following four steps:

1. **Collection:** Artifacts are documented, collected, and secured until they can be transported to a place of examination.
2. **Examination:** There are different approaches that can be taken to examine artifacts. The entire examination process must be documented so that any questions about what was performed to find any evidence can be clearly answered.

3. **Analysis:** The analysis is the overall result of the examination. Analysis is designed to answer questions that are brought up in a specific case.
4. **Reporting:** The final step is to develop a summary of all documentation of the collection, examination, and analysis process. The reporting must include all tools used, who was involved and what was their role, and any issues that occurred during the entire process.

Securing artifacts during the collection phase is commonly done with a sealed bag that has a label designed for digital forensics, as shown in Figure 8-29. Any interaction with the sealed artifact is documented on the label, including transporting it to a new location or pulling it for examination. Sealed artifacts that are not being used must be stored in a secure location, such as a safe, and the secure location must have an access log to document all details of any access to the location.

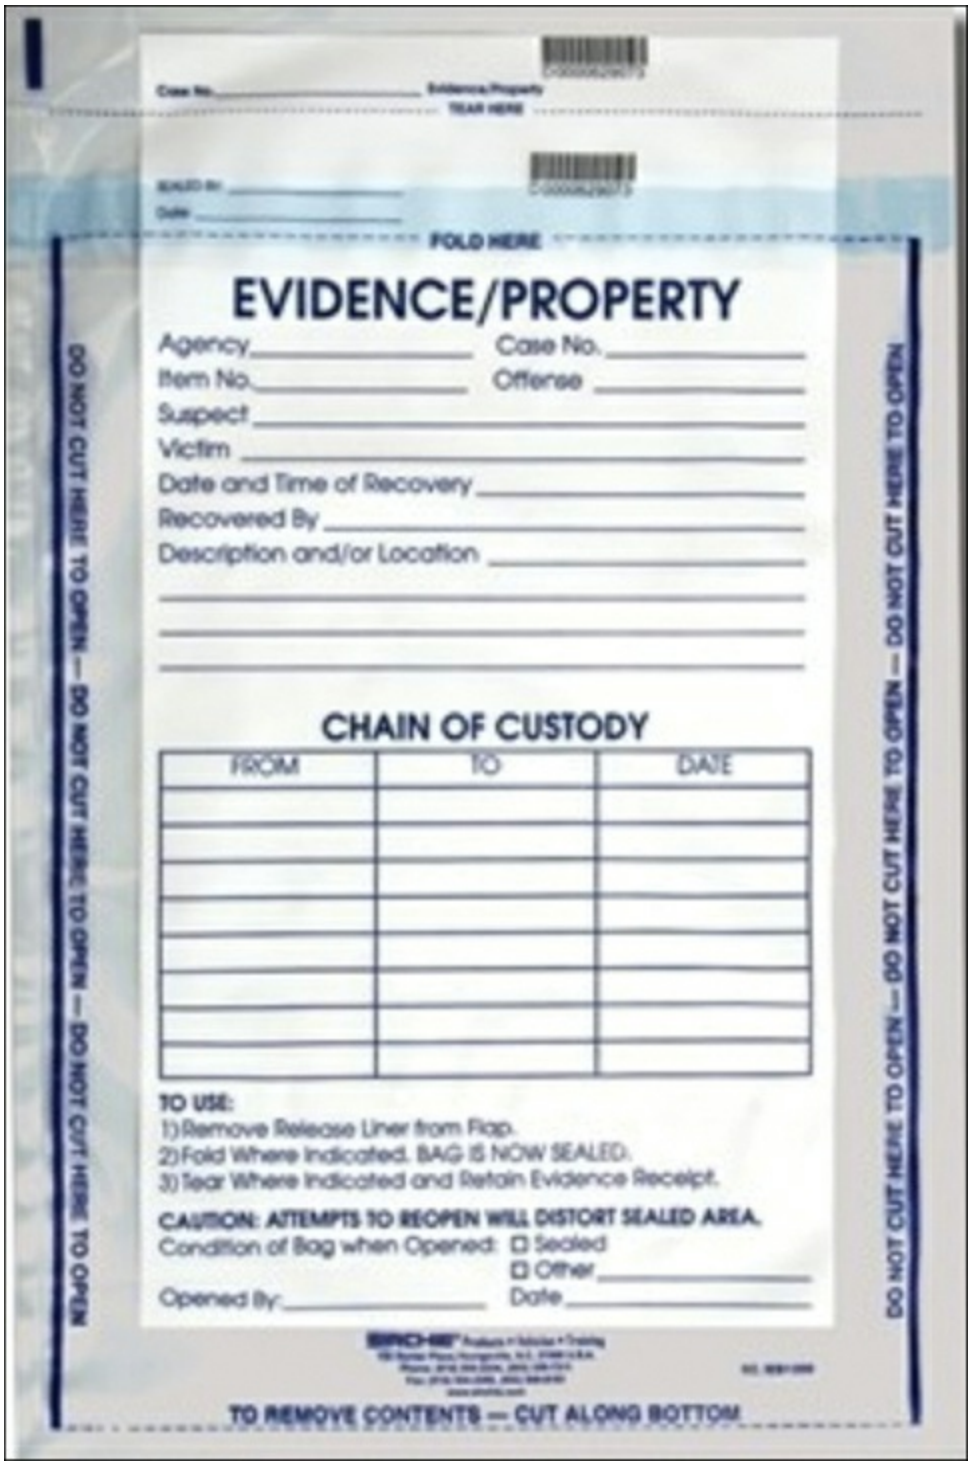

FIGURE 8-29 Chain of Custody Documentation Bag Example

Note

NIST SP 800-86, *Guide to Integrating Forensic Techniques into Incident Response*, provides a sample chain of custody tracking form you can download as your template for tracking artifacts.

The examination of artifacts needs to include various forms of documentation, including video recording, photography of steps taken, and a written report explaining exactly what was done and by whom. Anybody involved with the examination should be prepared to appear in a court proceeding if the only defense necessarily needs explanation about the examination process or how a conclusion was made during the analysis phase of the forensic process. All of these details will be included in the final report; however, live questioning might be required after the report is provided and used in a court of law.

To prove chain of custody, all examiners need to be prepared to answer the following questions:

- Where is the proof for the evidence found?
- How did you acquire the evidence?
- When was the evidence gathered?
- Who handled the evidence?
- How was the evidence secured when not handled?
- Why did that specific person handle the evidence?
- Are there any gaps in monitoring or securing the evidence?

More detailed questions will be asked about the specifics of the case; however, the previous questions are examples of the general types of questions that may be posed to a witness to establish whether a proper chain of custody was maintained. What is key is having proper chain of custody performed and documented to ensure all evidence is admissible in court.

Working with Evidence

Evidence can be anything that can help to prove a case. Technology that has

any form of data storge or memory can contain evidence. This includes computers, mobile devices, answering machines, memory drives, digital cameras, printers, networking hardware, servers, and IoT equipment.

There are a few fundamental principles that must always be followed when working with digital evidence. Make sure to follow all of these key points anytime you are working with evidence that could be used for legal action:

- Never work with the original
- Make one or more bit-level copies
- Validate the copies with the original using hash validation
- Secure the original and copies
- Enable write block (not allowing the ability to edit a file) when investigating a copy of the evidence
- Document everything

The last point is critical in rebutting any claim that the forensic team obtained evidence illegally or mishandled evidence during the investigation process. Most professionals use software to document the entire forensic process, allowing for simple access to any step that was taken, including a place to save notes, images, and other details obtained during that step of the forensic investigation.

One popular open-source option for documenting the forensic process is Autopsy. Autopsy is available in Kali Linux and can be accessed by going to **Applications > Forensics** and choosing **Autopsy**. Once Autopsy is running, you simply go to http://localhost:9999/autopsy and access the main GUI, shown in Figure 8-30.

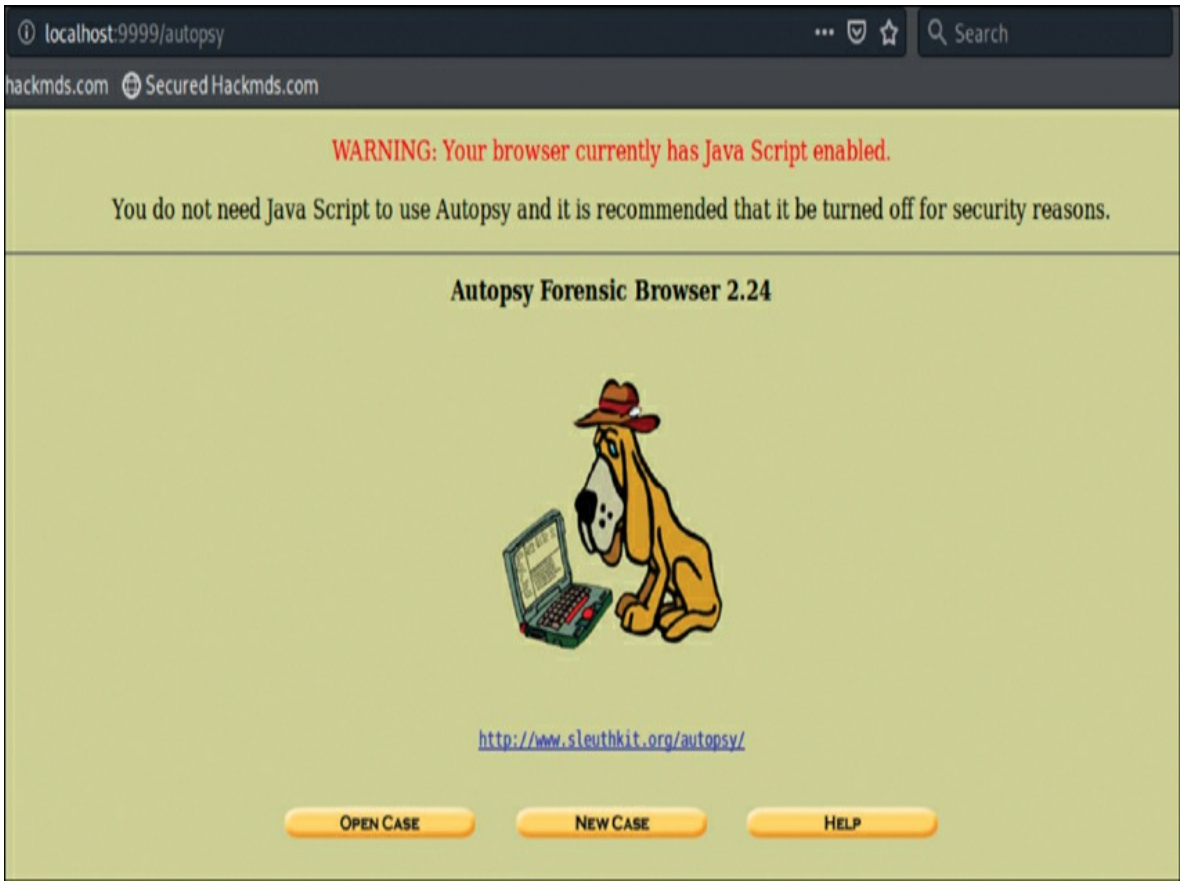

FIGURE 8-30 Autopsy Main Page

From the main page, you can open an existing case or create a new case. For this example, I'll create a new case and start investigating a computer. Once I create a host entry for the computer, I can upload images, make notes about what I have done with the computer, upload hashes to verify I have not changed anything as I investigate the copy of the computer, and add any other notes. Figure 8-31 shows some of these options in Autopsy.

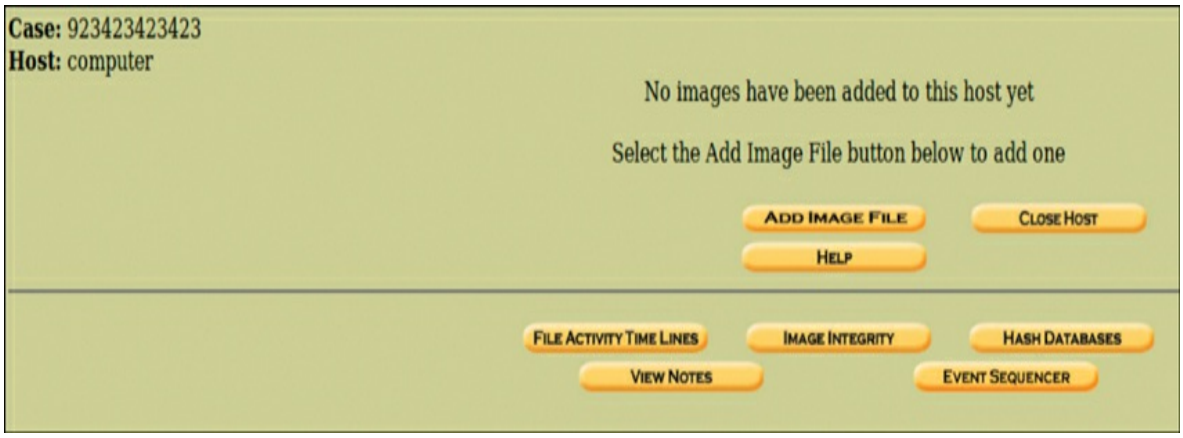

FIGURE 8-31 Autopsy Usage Example

Autopsy is a pretty basic forensic documentation program; however, it gets the job done if you do not have the budget to purchase an enterprise option. Enterprise options include many more features such as hashing, various bit-level data duplication capabilities, collaboration tools, and other useful tools that are not available with Autopsy. I highly recommend having some form of a forensic case management software application as part of your forensic process.

Duplicating Evidence

If you work with any original evidence, you are putting the entire investigation at risk of contamination. Instead, make bit-level copies, which are not backups. Backups copy only specific data and do not copy data stored, for example, in hidden clusters or space deemed by the system not useful. Malware authors design malware to leverage these hidden spaces to hide, which is why you need to copy everything using bit-level copies of evidence. Also, backups modify data, which means the possibility of change to the evidence potentially rendering it inadmissible.

There are dozens of free tools you can use to make bit-level copies of artifacts. The following are some popular options:

- **Linux dd command: dd if=** <*source*> **of=** <*destination*> **bs=** <*byte size*>
- **Linux dcfldd command: dcfldd if=/dev/sda hash=md5**

of=/media/diskimage.dd bs=512 noerror

- **GNU ddrescue: ddrescue** [**option(s)**] <*input file*> <*output file*> [**log file**]
- **Netcat (to copy to remote location): dd if=** <*source*> **of=** <*destination*> **bs=** <*byte size*> **| netcat** <*remote location IP*> <*port*>

Most modern digital forensic platforms used for investigating artifacts include the ability to duplicate, hash, and exam copies of evidence. Examples of popular platforms include EnCase from OpenText (formerly Guidance Software) and the open-source Digital Forensics Framework (DFF).

Hashes

You will need to validate any copies you make of artifacts using a hash to prove nothing has been modified during the duplication and examination process. *Hashing* means you have a one-way algorithm that doesn't produce the same results for two different things. Hashes also are randomized, meaning you can't predict the hash value before it is produced. If anything changes with the file, a completely new hash will be created within the same length. For example, the hash of "What's up" and the hash of "Hi, how are you" would be the same length hash result but would be different numbers. Even "What's up" and "What's Up" will be completely different hash numbers of the same length. If any of these rules are violated, the hash is no longer deemed safe to use.

These fundamental hash concepts can be summarized by the following three rules:

1. You can't predict the hash value of a file.
2. No two hash values can be the same, commonly referred to as a hash collision.
3. If anything changes to the source being hashed, the hash must change.

The two most common hash algorithms are Message Digest 5 (MD5) and the Secure Hashing Algorithm (SHA) family. Many professionals recommend using strong hash algorithms such as SHA-512 because older hash

algorithms, including MD5, have been proven to produce collisions, putting your case at risk if you validate your evidence with a vulnerable hash. Using a stronger hash will take a bit longer to compute the hash and produce longer hash values, but the effort is worth ensuring that your validation process is deemed secure.

One example of options you can use to create a hash of a file is the set of options built into Linux. After you complete a copy of a file, hard drive, or anything else you are investigating, you can validate it by using the **sha256sum** command to create an SHA-256 hash or the **sha512sum** command to create a SHA-512 hash, as shown in Figure 8-32. The figure also shows that I opened the data.txt file that was used to generate both hashes after running the hash commands, made a change to one character, saved the file, and ran both hashes again. Notice following the **nano** command that the hashes produced are completely different based on the change of one character in the entire file.

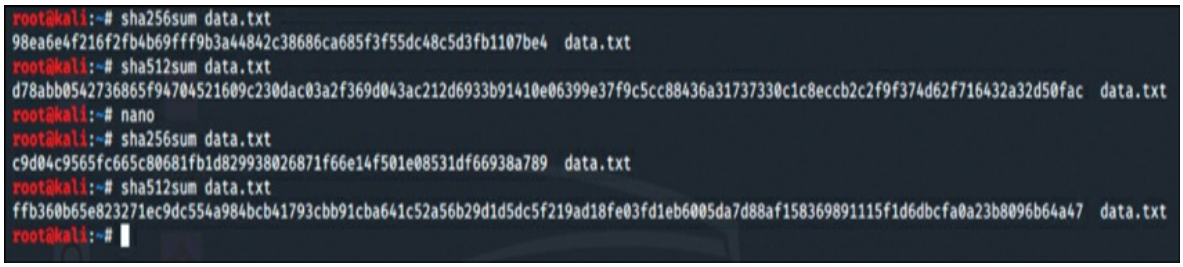

FIGURE 8-32 SHA-256 and SHA-512 Hash Examples

This example shows how hashing can be used to validate if any changes are made to a file, which is critical to prove that your investigation has not modified evidence as you use it to prove your case. If anybody questions your evidence, you need to be able to show the hash of the original evidence and the copy you used to find evidence both have a matching hash value. A forensic case management tool is a great place to store copies of the hash validation you should continuously perform as you conduct your investigation of the copy of the artifact in question.

Note

In some situations you will not be able to control changes to evidence. A prime example is when you encounter solid-state drives, which include firmware that is designed to increase the lifespan of the drive through data movement/cleanup. When hash validating a drive, these changes will appear as changes to the evidence even though they are just the result of the hard drive performing built-in processes. There are methods to work around this, including focusing on the part of the hard drive that contains the evidence; however, this topic is out of scope for this book. Check out this paper to learn more about dealing with solid-state drives: https://commons.erau.edu/cgi/viewcontent.cgi?article=1445&context=jdfsl.

Forensic Static Analysis

Static analysis in digital forensics refers to researching a target that is not powered on. This means all data has been saved. Do not power on a target system, because the system startup will change the state of the hard drive and system. If you make a hash-level copy of a system's hard drive, power it on, and make another hash, you will find that the hash values will not match. This minor change can cause evidence to be deemed inadmissible in court due to potential contamination.

You will want to research anything on the bit-level copy of the hard drive regardless of whether it was labeled as trash or data hidden within unusual places. For example, one place you don't want to overlook is *slack space*, the leftover storage that exists on a computer's hard disk when a computer file does not need all of the space it has been allocated by the operating system. That leftover space is available for malware writers to use to hide artifacts. Many forensic investigation applications such as EnCase include the ability to collect all slack space and examine it for additional hidden data. This only works if you have a complete bit-level copy of the system being investigated.

Another data point to look for when investigating systems is *metadata*, which is data about data. Metadata is important for identifying who created a file, when it was created, when it was last accessed, and many other details that allow the investigator to better understand how the evidence relates to a case.

I was once involved with investigating an employee's datacenter share drive that had inappropriate content. The employee claimed he didn't put the content in question on the drive, which was confirmed by looking at the metadata; it showed that some external party had gained access to the server and was using it to store inappropriate content. Without the details discovered in the metadata, the employee might have been in serious legal jeopardy because the content stored on his share drive was illegal and very troubling.

Key point

I highly recommend enabling write block whenever examining artifacts. This prevents you from making changes during your investigation. Remember that even the slightest change will change the hash value of the artifact you are investigating!

Now that you have an idea of what static analysis is, next let's look at how to recover data from a powered-off system.

Recovering Data

The most basic goal of a digital forensic investigation is to recover data. A common misconception among computer users is that when they delete data, it is erased from the computer and can't be retrieved; in fact, they are just telling the computer that the space occupied by the deleted data is now available, and the data is not removed until something is written over it. Many hard drives use a first-in, last-out approach, meaning deleted data won't be overwritten until all other space runs out. Consequently, old data can linger as long as a few years before new data is written over it. I have run recovery software on old USB drives of mine and recovered hundreds of files that I deleted years earlier.

I spoke at DEFCON about my research on this topic a few years ago. I had matched the publicly posted social media handles of HR employees in various organizations with the handles of sellers on eBay. My hypothesis was that HR employees were selling used computers holding extremely sensitive

data on eBay but were not following proper data deletion protocols, allowing buyers to acquire the systems and recover sensitive data. Any time I saw one of these eBay sellers list a hard drive or a printer with a hard drive, I purchased it with the intent to investigate whether I could recover extremely sensitive information from it. Sure enough, I collected driver's license information, passport details, documents labeled classified, and much more from printer hard drives, scanner hard drives, and external media, all sold on eBay. Think about devices such as scanners used for scanning sensitive data, IoT devices that collect data about users, cameras used for conference calls or security, USB drives that are shared with employees, and all of the other drives that have collected data that can reveal all sorts of details about their environment and associated users.

One simple way to recover data is to use a free open-source tool called Foremost. One great feature about Foremost is that it will automatically identify all files within a target location and organize them by file type. You can run Foremost by using the command **foremost -t all -t -I** <*location to scan*> **-o** <*where to send what is found*> against any hard drive that is connected to your forensic laptop. Figure 8-33 is an example of running Foremost on a USB drive found at /dev/sdb1 and outputting the results to a folder located at /root/Desktop/recover/.

```
foremost -t all -v -i /dev/sdb1 -o /root/Desktop/recover/
```

FIGURE 8-33 Example Running Foremost

Figure 8-34 shows the output from the previous run command. Notice all the nice folders; any matching file formats will be placed in their corresponding folder. Many commercial programs do the same thing but include fancy GUIs and are marketed as the solution for when you delete something by accident and need to recover it. My advice is to just use a free option such as Foremost that does a fantastic job of locating any lost data and collecting it in a nice file library.

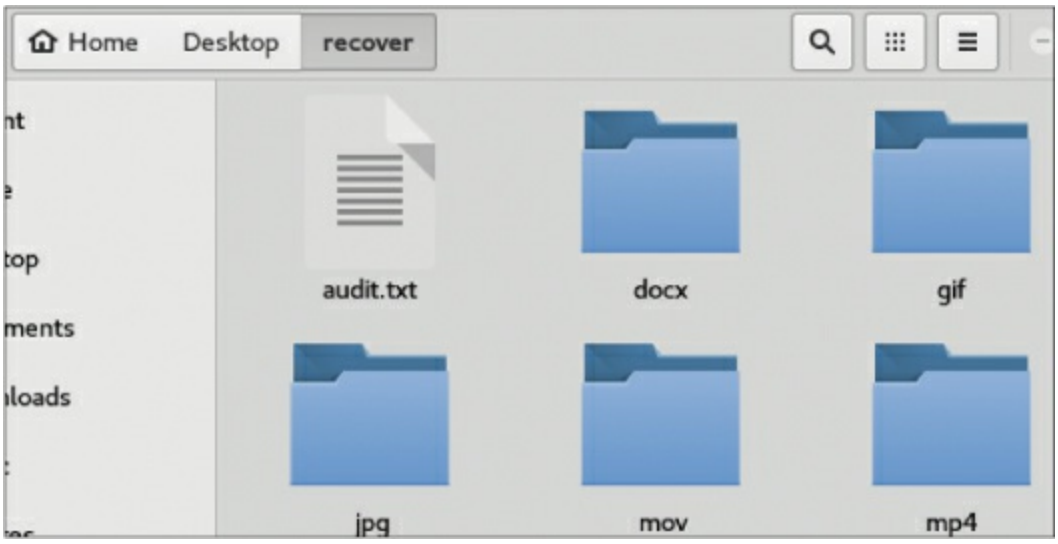

FIGURE 8-34 Example Foremost Output

Note

To properly delete or destroy the data on a hard drive, either physically destroy the hard drive (recommended approach) or use a write program that writes over the entire hard drive with 0s and 1s.

Forensic Dynamic Analysis

Dynamic analysis in digital forensics refers to analyzing a target system that is running (powered on). Do not shut off a target system that is powered on; not only because it changes the state of the system but also because you will lose any volatile data. *Volatile data* is data that is not saved but is used while the system is running. Common places to find volatile data include registries, cache, and RAM. Why care about volatile data? Remember that even encrypted data must at some point be unencrypted in order to be processed by the system. Also think about how your system remembers your user history to make life "easier" for you. This is the type of data that sits in a volatile state and can be collected with the right tools and techniques.

Some examples of volatile data you can find on running systems (and will

lose if you turn off the systems) include the following:

- Running processes
- Unencrypted data
- Passwords in clear text
- Who is logged in
- Instant messages
- Attached devices
- Executed console commands
- System information
- IP addresses
- Registry information
- Trojan horses
- Open ports and listening applications

Volatility

One tool you can use to view volatile data in RAM is Volatility. You will need to first dump RAM from a running system—a program such as DumpIT allows you to dump a copy of a running Windows RAM. If you don't know the type of RAM you are working with, you can use the command **volatility-2.5.standalone.exe –f** <*mem dump file.raw*> **imageinfo** to get a fingerprint of the type data you are working with. You will need to figure out the offset of the memory, which you can find by running the command **volatility-2.5.standalone.exe –f** <*mem dump file.raw*> **--profile=** <*suggested profile*> **kdbgscan**. With the offset, you can identify what processes are running on the system by using the command **volatility-2.5.standalone.exe –f** <*mem dump file.raw*> **--profile=** <*suggested profile*> **--kdbg=** <*off set*> **pslist**, as shown in Figure 8-35.

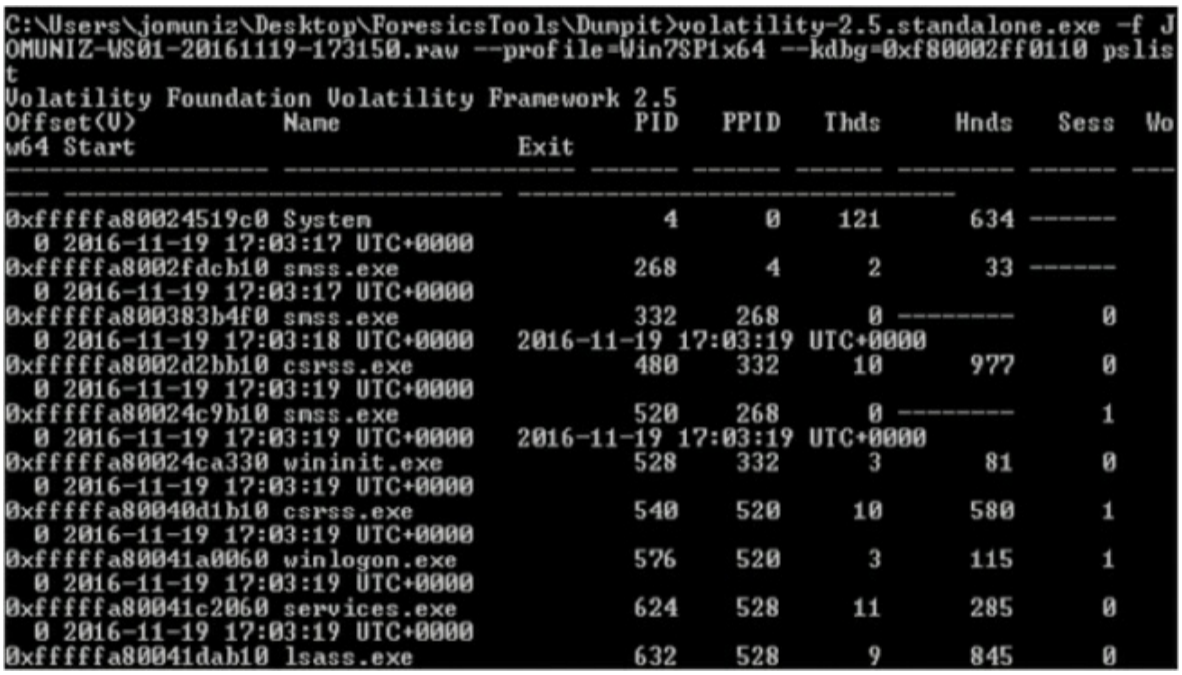

FIGURE 8-35 Finding Processes in RAM Using Volatility

You can also recover hash dumps with Volatility to see password hashes, as shown in Figure 8-36. This is useful for running against password rainbow tables to recover passwords.

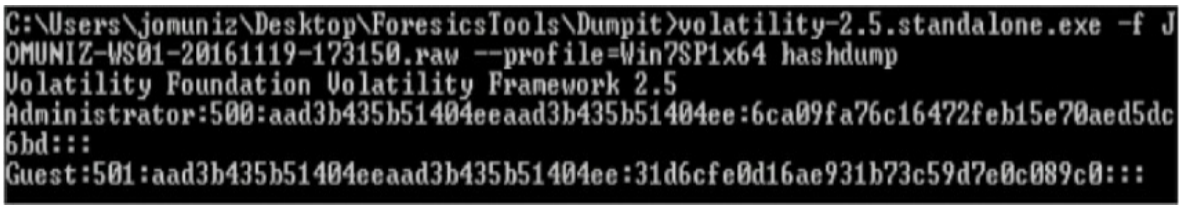

FIGURE 8-36 Dumping Hashes Using Volatility

If you can get access to the RAM of a terminal server, running the **cmdscan** option allows you to see all commands that have been run. Imagine all of the sensitive information you can recover with that history file, including admin passwords, critical IP addresses, and so much more! Figure 8-37 shows an example of running Volatility to access this info.

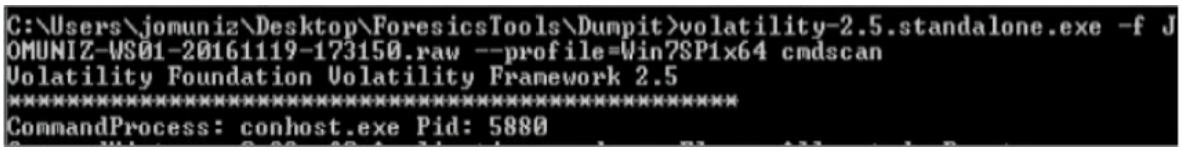

FIGURE 8-37 Dumping Command-Line History Using Volatility

One last Volatility example is pulling the history of what websites were accessed regardless of the browser being used. The command ends with **iehistory**, but it will dump the history of all browser usage, not just Internet Explorer. Figure 8-38 shows how that command looks in Volatility.

FIGURE 8-38 Dumping Browser History Using Volatility

Digital Forensics Summary

The examples presented in the previous sections provide just a glimpse into what can be pulled from a running system; active systems are a gold mine if you can keep them powered on while investigating them. You do not want to investigate the live system; however, you can pull dumps of the live data, mark the time of dump, and investigate the dump without impacting the original system.

Never power off a system or all volatile data will be lost and the state of the system will change as it is performing the shutdown process. If you must shut down a system, do a hard shutdown, meaning kill power without permitting the shutdown procedures to execute. Your goal is to reduce as much as possible the amount of change that is caused by your interaction with the system.

Digital forensics is a massive topic. Entire books are dedicated to this topic (including one I coauthored for Cisco Press called *Investigating the Cyber Breach*). The field has dedicated certifications, such as the EC-Council Certified Digital Forensics Investigator program. This section's brief

introduction to digital forensics is intended to serve as a primer of general topics you need to know to run a functioning digital forensics practice. Leverage other resources for more guidance on developing your digital forensic capabilities.

> **Note**
>
> NIST SP 800-83 (currently Rev. 1) is a great guide to conducting digital forensics in response to a malware incident: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-83r1.pdf.

Phase 3: Containment, Eradication, and Recovery Summary

The list that follows summarizes the key points to keep in mind for the containment, eradication, and recovery phase of the incident response lifecycle:

- The proper approach to incident response depends on whether or not the incident can lead to legal action.
- If an incident has potential legal ramifications, the investigation must immediately flip to a digital forensics–focused playbook.
- Threat hunting means to create a fingerprint of a threat and search all systems and networks for that fingerprint.
- Fingerprints used during the threat hunting phase are developed during the analysis phase, which include behavior fingerprints and hashes of malicious artifacts.
- An incident must be contained before it can be eradicated.
- Eradication includes handling all impacted systems. Some impacted systems might not show infection, as many forms of malware remain dormant until specific actions or processes are seen.

- Recovery includes all steps required to return all impacted systems back to a functional state.
- Digital forensics requires all steps to be performed correctly and a chain of custody to be maintained or evidence will be considered contaminated.
- Never work with original evidence during a digital forensic investigation. Work with copies and document everything.
- Never change the state of a system when performing a digital forensic investigation. If you must change the state, do so with the least amount of change, such as a hard shutdown.

Phase 4: Post-Incident Activity

The final phase of any form of incident response is to perform post-incident activity. This phase does not occur until all impacted systems have been returned to an operational state, meaning the entire incident has been handled and the case is ready to be closed. Before closing the case, the post-incident activities that follow the incident focus on reviewing and documenting how the SOC handled the incident, with the goal of learning and improving from the situation. Enforcing post-incident activity is characteristic of the higher maturity levels for SOC services. The discussion of how the SOC performed occurs during a lessons learned meeting.

When I speak about post-incident activities, I'm not just recommending that you host a meeting to debrief on the event and move on to other duties. Post-incident activities need to be a structured process in order for real benefits to be obtained. Otherwise, the lessons learned meeting will become a free-for-all filled with complaints and possible finger pointing, leaving the SOC worse off than if they had not met. A security incident translates to a failure in security, and the worst thing that a SOC can do is shame anybody involved with that failure.

Note

It is absolutely critical to enforce a mentality of improvement rather

than blame during a lessons learned meeting, as conversations about what occurred during an incident can easily get heated. I have personally sat in heated lessons learned meetings and witnessed negative conversations quickly become toxic, causing major disruption between employees and directly impacting how the SOC functions moving forward. I have seen people get so heated that they put in their notice to quit. My advice is to open the meeting by stating that the focus is improvement, not blame, and to quickly reiterate that point if conversation shows any sign of becoming toxic.

Post-Incident Response Process

The post-incident response process should include five key steps, identify, document, analyze, store, and retrieve, as depicted in Figure 8-39.

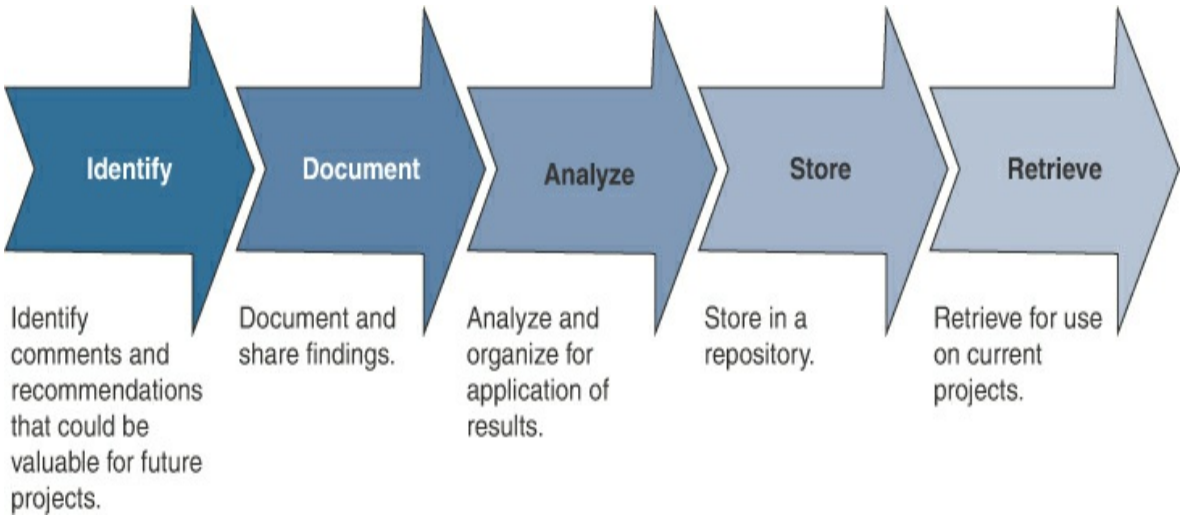

FIGURE 8-39 Post-Incident Response Process

In summary of these steps, the SOC leadership first must identify how the SOC performed during the incident. That feedback is documented, analyzed, and stored so it can be retrieved at any point to learn how the SOC responded to the previous incident. Let’s look closer at each of these steps, starting with the Identify step.

Step 1: Identify

The Identify step is all about gathering information regarding what occurred during the incident. Post incident research can be performed by SOC management or by an external party that specializes in post-incident response feedback. Whoever performs the post-incident research must keep a neutral viewpoint regarding internal politics and only document what happened with recommendations based on what would improve future responses. The following general questions about the SOC's response need to be addressed as part of the information gathering process:

- What happened and at what time?
- How well did the SOC handle the incident?
- Where can the incident response program improve?
- What information was needed sooner?
- Is there a possibility of the same incident occurring again?
- What was the estimated loss due to the incident?
- What changes need to be made to the security?
- How should employees be trained differently?
- What weakness did the breach exploit?
- How will you ensure a similar breach doesn't happen again?
- Was any legal action taken and what was the outcome?

Based on the answers to these questions, different changes can be recommended to adjust how the SOC responds to similar incidents that occurred, such as improving the SOC's playbook(s) and the organization's overall security practice. Maybe more training is needed. Maybe new security tools need to be acquired. Maybe more SOC staff are needed, or existing SOC staff daily tasks need to be adjusted. NIST SP 800-83 Rev. 1 offers the following examples of possible changes that can be made after a malware incident based on the outcome of the post-incident response activities:

- **Security policy changes:** Security policies might be modified to prevent similar incidents that were handled by the incident response program. For example, if connecting personally owned mobile devices to organization

laptops caused a serious infection, modifying the organization's policies to secure, restrict, or prohibit such device connections might be advisable.

- **Security awareness program changes:** Security awareness training for users might be changed to reduce the number of infections or to improve users' actions in reporting incidents and assisting with handling incidents on their own hosts.
- **Reconfiguration:** Network equipment, OS, or application settings might need to be changed to support security policy changes or to achieve compliance with existing policy.
- **Malware detection deployment:** If hosts were infected through a transmission mechanism that was unprotected by antivirus software or other malware detection tools, an incident might provide sufficient justification to purchase and deploy additional software or security tools.
- **Malware detection software reconfiguration:** Detection tools might need to be reconfigured in various ways, such as the following:
  - Increasing the frequency of software and signature updates
  - Improving the accuracy of detection
  - Increasing the scope of monitoring
  - Changing the action automatically performed in response to detected malware
  - Improving the efficiency of update distribution

Any of these changes would be discussed during the lessons learned meeting and then evaluated as part of analyzing and responding to what is documented in the Lessons Learned report. Essentially, conversations about changes will occur during the second and third steps of the post-incident activities.

Step 2: Document

The success of the post-incident response activities depends on the quality of the lessons learned meeting. This meeting is where all data that was collected about the incident is debated and documented. NIST SP 800-61 Rev. 2

recommends collecting feedback from all parties involved *before* the lessons learned meeting to ensure that all topics and participants' needs are identified prior to the meeting to avoid any surprises. Only parties involved with the response and who are authorized to have access to topics that will be discussed in the lessons learned meeting should be involved. Many topics will involve failures in people, process, and technology with corresponding sensitivity concerns that must be addressed. More than one lessons learned meeting may be needed if different levels of conversation need to occur regarding what is and what is not permitted to be talked about. For example, a lessons learned meeting that includes customers or general employees will be much different regarding the agenda than one that only involves the SOC.

Lessons Learned Meeting Speakers

Another NIST recommendation for the lessons learned meeting is to include a moderator who is skilled in group facilitation to help lead the conversation and ensure feedback from all parties is heard. Although NIST does not specify whether the moderator should be someone within the organization or an external consultant, I believe an external moderator is ideal to keep his or her viewpoint unbiased and focused on security without concern for internal politics. It is very possible that certain people's performance will be pointed out as not being good, which is easier delivered by an external resource. An external moderator can also help control a situation when conversations start to become toxic, such as when one team publicly shames another person or team during the meeting.

Note

It is critical to keep all conversations during the lessons learned meeting(s) private and secure from external sources. Many sensitive topics are likely to be covered, including if employees need to be reprimanded, weaknesses within the organization's security, and potential loss of sensitive data.

One best practice I have learned and recommend your SOC team lead

incorporate as part of the pre-lessons learned meeting activities is to develop an agenda for the meeting and validate it with the SOC sponsor. NIST SP 800-61 Rev. 2 also states that the success of such meetings depends on the agenda. Collecting input about expectations and needs from participants before the meeting increases the likelihood that participant needs will be met. In short, the agenda is critical to the success of the lessons learned meeting.

I also recommend the SOC sponsor open the meeting. There may be conflicts that can be predicted based on the data collected prior to the meeting, which incorporating the SOC sponsor can help address at the beginning of the meeting using a voice from leadership. I find that when a C-level member opens a meeting by addressing expected concerns, such as if somebody will be fired, it dramatically reduces tension and concerns. For example, if the SOC's sponsor starts the meeting by stating "I know the situation was bad, but it's over and I want to ensure you that nobody is going to be fired. Instead, we need to figure out how to avoid this situation in the future, which is why we are here," it will set the right tone for the rest of the meeting. The same message from a team lead can also work, but I find the message coming from leadership has much stronger impact.

Note

I recommend having your SOC sponsor introduce the moderator of the lessons learned meeting. This shows the moderator has the support of leadership, which helps when the moderator has to make strong statements or intervene when a conversation starts to become toxic.

Meeting Agenda

Validating the lessons learned meeting agenda with a SOC sponsor is important regarding confirming who should and should not be included in the meeting. Representation from human resources may be needed if employee rights are in question. Certain managers may be needed if their staff is expected to answer tough questions. Certain people may be excluded based on predicted conversations containing sensitive details that must be addressed but also must be kept protected from nonauthorized personnel. To avoid

having unauthorized people in the room or omitting people who should be there, I highly recommend validating the agenda and associated personnel with leadership, not only to plan for the correct personnel, but also to protect the meeting planner from being blamed that unauthorized people were invited.

Figure 8-40 is a generic meeting agenda template. Notice it includes key people such as the facilitator and who will be attending. It lays out the topics to be covered and provides a section called Other Information for diagrams, screenshots, or other artifacts that the attendees should be aware of prior to attending the meeting. Consider these to be the bare minimum aspects to include in your meeting agenda.

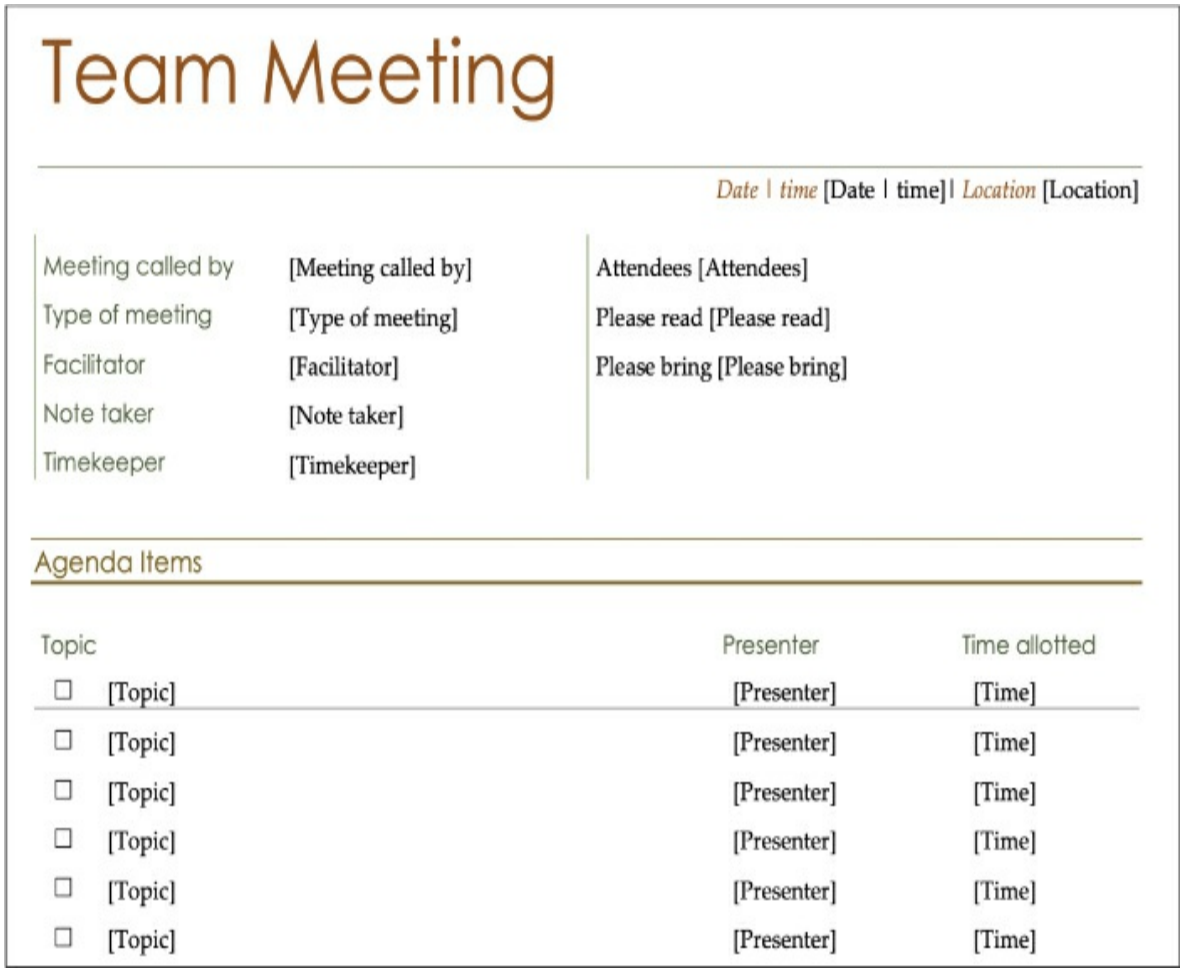

Team Meeting

Date | *time* [Date | time] | *Location* [Location]

Meeting called by	[Meeting called by]	Attendees [Attendees]
Type of meeting	[Type of meeting]	Please read [Please read]
Facilitator	[Facilitator]	Please bring [Please bring]
Note taker	[Note taker]	
Timekeeper	[Timekeeper]	

Agenda Items

Topic	Presenter	Time allotted
☐ [Topic]	[Presenter]	[Time]
☐ [Topic]	[Presenter]	[Time]
☐ [Topic]	[Presenter]	[Time]
☐ [Topic]	[Presenter]	[Time]
☐ [Topic]	[Presenter]	[Time]
☐ [Topic]	[Presenter]	[Time]

FIGURE 8-40 Lessons Learned Meeting Agenda Template Example

Step 3: Analysis of Incident Response

Analyzing how the SOC responded to the event is performed during the discussions in the lessons learned meeting as well as when reviewing all lessons learned documentation post meeting. The quality of what is documented will be based on what questions are asked and discussed during the lessons learned meeting, making the outline for the meeting critical.

Figure 8-41 is an example of a generic lessons learned meeting agenda that can be used during the meeting. Questions should be developed based on data collected prior to the meeting along with some generalized questions listed in Figure 8-41.

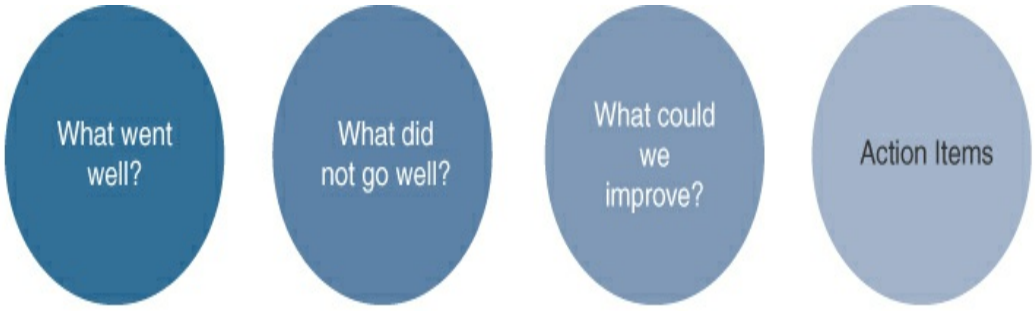

Questions	Discussed?	Observation
What happened and at what time?		
How well did the SOC handle the incident?		
Where can the incident response program improve?		
What information was needed sooner?		
Is there a possibility of the same incident occurring again?		
What was the estimated loss due to the incident?		
What changes need to be made to the security?		
How should employee be trained differently?		
What weakness did the breach exploit?		
How will you ensure a similar breach doesn't happen again?		
Was any legal action taken and what was the outcome?		

FIGURE 8-41 Lessons Learned Meeting Agenda Example

Key points raised when discussing each question should be recorded in an Observation (or similar) section. The template shown in Figure 8-41 doesn't have any lines in the Observation section because many answers will overlap, meaning the key points resulting from discussion of the questions will not correspond one-to-one with the questions. All the observations need to be validated as recorded properly, and any debated items should be noted. The observations will be cleaned up and documented in a final Lessons Learned report.

Lessons Learned Report

The details of the Lessons Learned report include the data captured during the lessons learned meeting as well as any additional input from participants who were not able to attend. It is important to note when feedback comes from somebody who did not attend the lessons learned meeting, since that feedback comes from somebody who did not hear the live discussion of the attendees and whose viewpoints were not necessarily considered during the meeting. I recommend asking those who were not involved in the live meeting to first look over the Lessons Learned report before submitting their feedback, and later allowing those who did attend the live meeting to review and respond to that feedback.

The facilitator of the lessons learned meeting should provide a draft of the Lessons Learned report to all attendees and allow time for feedback to be collected to ensure the accuracy of the report. After the report is finalized, the entire project team should receive a copy, even members who did not participate in the lessons learned session (unless the report contains sensitive details they are not permitted to see). The final report should be stored with the other project documentation and a copy should be provided to the SOC sponsor so that the SOC sponsor is aware of what action items are being planned as a result of the response and can support the resources required to make those recommendations happen.

The format of your Lessons Learned report will depend on the type of services your SOC offers and who you expect will read the report. In general, the following are good sections to include in your Lessons Learned report template:

- **Parties Involved:** Provide information about who was involved in the incident response. A template can look like Figure 8-42.

Incident Response Manager

Name	Email
Work Phone	Mobile Phone

Technical Contacts

Name	Email
Work Phone	Mobile Phone

Name	Email
Work Phone	Mobile Phone

Name	Email
Work Phone	Mobile Phone

Legal Counsel

Name	Email
Work Phone	Mobile Phone

Communications Specialist

Name	Email
Work Phone	Mobile Phone

Additional Members

In addition to those individuals listed above, additional experts may be included on the IRT, depending on the nature and scope of the incident. In particular, a software support expert from the team that supports the software in question will likely be necessary. These additional members will be chosen by the IRM.

FIGURE 8-42 Template for Documenting Parties Involved

- **Executive Investigation Summary:** This is a summary of the incident and how the SOC responded. This should be kept at a high level, targeting those who need the quick facts about what happened.
- **Affected Stakeholders:** All impacted parties need to be listed, including how they were impacted. This includes customers, employees, and other parties.
- **Report Management:** All data that was collected during the incident needs to be accounted for in case additional investigation needs to take place. This section should list where data that was used can be obtained.
- **Detailed Incident Overview:** A more detailed explanation is required for those that need a technical explanation of what occurred. The section that provides these details starts with an overview of what occurred and has the following subsections:
  - **Incident Types:** What incident categories were involved?
  - **Incident Symptoms:** What was detected to determine an incident occurred?
  - **Accessed:** What systems have been impacted? How widespread is it?
  - **Intelligence:** What data was collected about the incident? Are there hashes, signatures, vulnerabilities, or other characteristics that are associated with the incident?
- **Response:** The report should include a detailed overview of how the response was carried out. Subsection topics include the following:
  - **Initial Response:** What initial actions were taken? Which playbook was used?
  - **Temporary Security Controls:** What temporary security controls were launched during the containment phase?
  - **Eradication:** How were all of the threats eradiated?
  - **Remediation and Recover:** What recovery steps were performed to return impacted systems back to operational state?
- **Post Incident Recommendations:** Recommendations need to be listed regarding how the SOC can improve its response as well as how the

organization can reduce the risk of future compromise.

- **Appendix:** List any additional details that may need to be referenced, including definitions, screenshots, and evidence.

Final Steps: Store and Retrieve

The final steps of the post-incident response activities are to take the actions recommended in the Lessons Learned report. Again, it is important to deliver a copy of the Lessons Learned report to the SOC sponsor in order to receive executive support for the changes that need to be made. Changes that have an associated cost will require that the SOC sponsor obtain or allocate the appropriate budget or make the decision to choose an alternative path. Some recommendations will require more SOC staff to be recruited, which, again, will require the SOC sponsor to allocate budget or adjust personnel to support the new tasks. Some changes will require modification of policies and procedures, which need to be dictated from leadership to ensure they are followed immediately. All of these changes are best handled with the support of the SOC sponsor.

As changes are made, a project manager may be needed to monitor the status of each recommendation. During and after recommendations are enforced, the Lessons Learned report needs to be secured. Remember that a Lessons Learned report will contain sensitive information and it should never become a public document!

Phase 4: Post-Incident Response Summary

The list that follows summarizes the key points to keep in mind for the post-incident response phase of the incident response lifecycle:

- Post-incident response activities need to be well organized and mandatory after every incident.
- A lessons learned meeting is a gathering used to review how the SOC responded and to develop a Lessons Learned report. Only individuals involved with the incident or who would benefit from learning from the incident should attend the meeting.

- A Lessons Learned report contains details regarding what occurred during the incident and what recommendations should be implemented to improve the SOC's response to future similar incidents. Only authorized personnel, such as SOC management, should have access to the final report.
- All parties involved in the incident should have a voice within the Lessons Learned report unless security or other concerns apply.
- The SOC sponsor should receive a copy of the final Lessons Learned report and make decisions regarding how each recommendation will or will not be executed.
- The Lessons Learned report needs to be stored in a secure space so it can be referenced as recommendations are being followed.

Incident Response Guidelines

This chapter covered a lot of topics, many of which have entire books, standards, and certification programs dedicated to them and require specialized experience and skills to properly perform them. Example topics include malware analysis, digital forensics, threat hunting, and incident response. I have referenced a few industry guidelines throughout this chapter as resources to help your SOC build and deliver an incident response program. For playbook templates, I referenced the Incident Response Consortium. For guidelines, I referenced NIST SP 800-83 Rev. 1 and NIST SP 800-61 Rev. 2.

My final recommendation for solid resources you can use for incident response guidance is the FIRST.org Service Frameworks, introduced in previous chapters.

FIRST Services Frameworks

FIRST.org is a fantastic reference for guidelines for security best practices. FIRST provides two services frameworks, PSIRT and CSIRT. Both are

developed by experts within the security industry based on how they run their SOC. Experts come from companies such as Cisco, Microsoft, and other organizations responsible for billions of dollars, thousands of employees, and following industry certification standards such as ISO. All of these expert experiences and knowledge of how to best handle an incident are reflected in these guidelines.

PSIRT

The Product Security Incident Response Team (PSIRT) Services Framework focuses on how to respond to security events involving products an organization delivers. If the vendor of a firewall discovers a potential vulnerability in its product, the vendor must respond by validating whether the vulnerability is real and, if so, evaluating what is the potential risk, how to remediate the product for future delivery, how to assist impacted customers to reduce their risk due to the vulnerability, what to publish about the vulnerability, and how to avoid similar vulnerabilities in future versions of the product. If the SOC develops applications or tools used by the organization, the SOC will need a PSIRT to be responsible for the security of those products. Without a PSIRT, a SOC-built tool that is identified with a vulnerability will not have the proper support needed to handle all of the required tasks to deal with the vulnerability and return the system back to operational state post event.

Note

The PSIRT service within a SOC is a subset of the incident response service. It must have its own maturity grade levels to ensure it doesn't remain overlooked. If the developers of a SOC tool can't provide proper PSIRT support, I recommend migrating to a third-party vendor that can provide such support.

CSIRT

The Computer Security Incident Response Team (CSIRT) Services

Framework focuses on services and associated functions involving incident management for computers and networks. FIRST points out that the CSIRT Services Framework is developed through a collaboration of the Task Force CSIRT community and the International Telecommunications Union (ITU), giving the framework a broad and international perspective of providing incident response services.

Note

It is important to recognize that there is synergy between a PSIRT and a CSIRT. These teams do not operate independently of each other. As an example, when a vulnerability is announced by CSIRT based on the network it connects to, but because it impacts a product, the PSIRT is also engaged.

This CSIRT Services Framework breaks down incident response into four key areas:

- **Service areas:** Service areas group services related to a common aspect. For example, one service area focuses on Monitoring and Detecting, threats meaning anything within this topic would apply to this service area.
- **Services:** A service is a set of recognizable, coherent functions oriented towards a specific result. Services that fall under the service area Monitoring and Detecting would include log and sensor management, detection use case management, and contextual data management since all three of these are part of a Monitoring and Detecting service.
- **Functions:** A function is an activity or set of activities aimed at fulfilling the purpose of a particular service. Monitoring and Detecting has three services, which functions are what makes up each of those services. This means detection use case management would have one or more functions that relate to delivering the detection use case management service.
- **Sub-functions:** A sub-function is an activity or set of activities aimed at fulfilling the purpose of a particular function. The Monitoring and Detecting service will have services, and each service has its own

functions. Sometimes a subset of those functions will exist which are the sub-functions. For Monitoring and Detecting, there are not any subfunctions.

The CSIRT Services Framework has five service areas. Figure 8-43 is a high-level diagram of the different service areas with associated services you can find within the CSIRT Services Framework. I highly recommend downloading this guideline and using it as a template to develop your incident response playbooks.

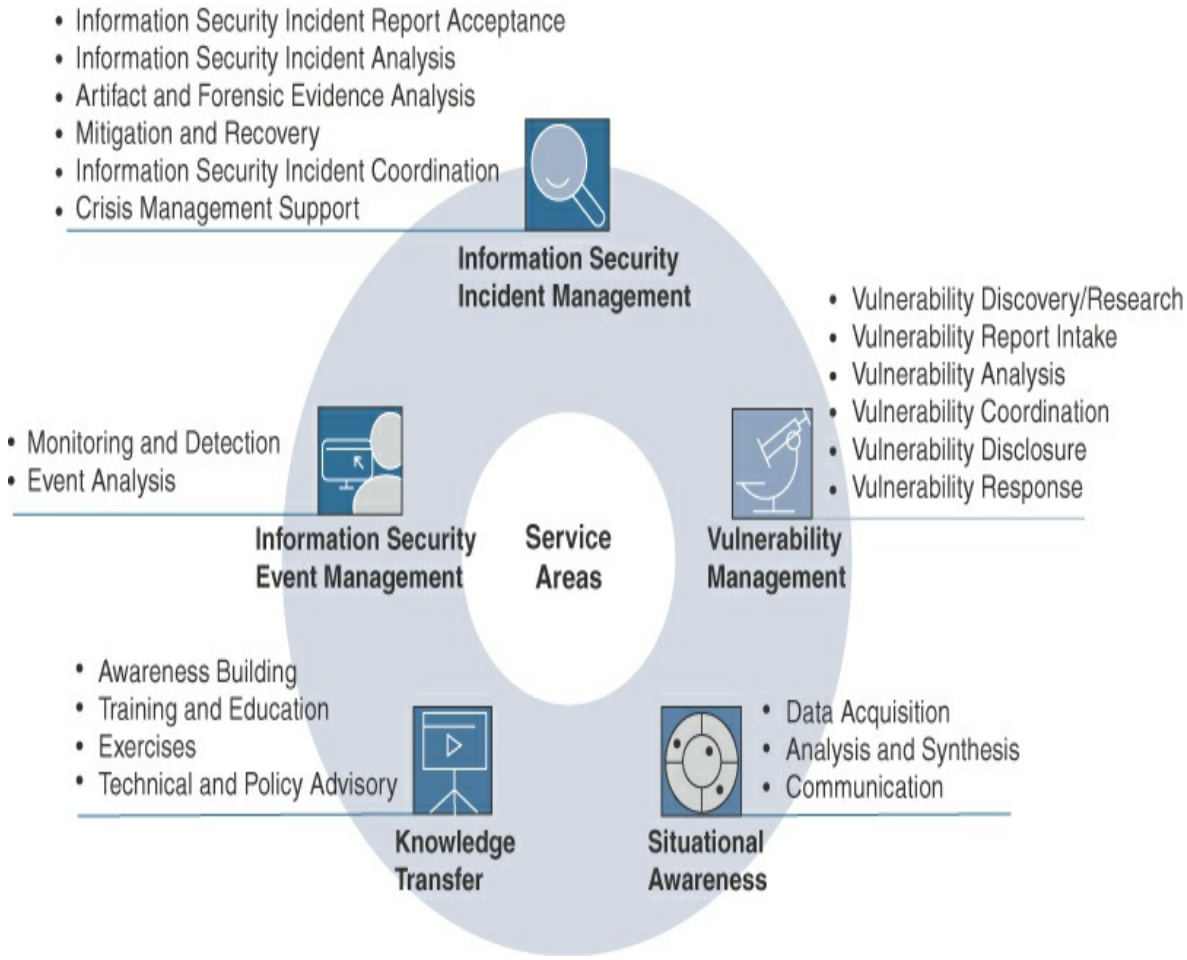

FIGURE 8-43 FIRST CSIRT Services and Service Areas

Note

The FIRST PSIRT and CSIRT Services Frameworks are available at

https://www.first.org/standards/frameworks/.

Summary

This chapter provided guidance for developing different aspects of an incident response service. It started by explaining the distinction between an incident and a breach, what that distinction means to a SOC, and why an incident response service needs to be developed in an organized manner. I next introduced the NIST SP 800-61 Rev. 2 four-phase lifecycle for performing incident response, which I used as the framework for the chapter. The section corresponding to the first phase focused on how to build an incident response program and prepare to go live. Next, in the section corresponding to the second phase you learned how to discover and analyze an incident to validate that it is a real threat as well as collect what is needed to perform a proper response.

I broke the third phase of the NIST incident response lifecycle (containment, eradication, and recovery) into two different paths. For events that don't involve potential legal action, I used the traditional NIST approach, which is to contain, eradicate, and recover from the threat. For incidents that might lead to legal action, I provided a dedicated section on delivering a digital forensic response, which is different from focusing on eradicating the threat because digital forensics is focused on preserving evidence and learning about the threat rather than just removing it. The chapter concluded with phase four post-incident response activities as well as additional references to help you build a strong incident response service.

Now that you should have a firm grasp on incident response, Chapter 9 will focus on proactive security measures to manage vulnerabilities and recovering from a security incident recovery.

References

24By7Security, Inc. (2020, February 25). How to Ensure Chain of Custody

After a Cybersecurity Incident. 24By7Security, Inc. https://blog.24by7security.com/chain-of-custody-and-digital-evidence-in-forensics

Atlantic Digital Forensics. (2017, December 18). Admissibility of Digital Evidence in Court. Atlantic Digital Forensics. https://www.atlanticdf.com/blog/2017/12/18/admissibility-of-digital-evidence-in-court/

Buchanan, W. (2020). Digital Forensics Magic Numbers. Asecuritysite. https://asecuritysite.com/forensics/magic

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012, August). Computer Security Incident Handling Guide. Computer Security Resource Center. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

Sqrrl. (n.d.). Hunt Evil: Your Practical Guide to Threat Hunting. Sqrrl. https://www.threathunting.net/files/hunt-evil-practical-guide-threat-hunting.pdf

Souppaya, M., & Scarfone, K. (2013, July). NIST Special Publication 800-83, Revision 1: Guide to Malware Incident Prevention and Handling for Desktops and Laptops. NIST. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-83r1.pdf

Teague, R., & Black, M. (2017). Evidence Verification Complications with Solid-State Drives. Journal of Digital Forensics, Security and Law: Vol. 12, Article 7. https://commons.erau.edu/jdfsl/vol12/iss4/7/

Tuma, S. (n.d.). Guide to Responding to Data Breaches and Reporting Cybersecurity Incidents to Law Enforcement and Governmental Agencies. Business Cyber Risk. https://shawnetuma.com/cyber-law-resources/guide-reporting-cybersecurity-incidents-law-enforcement-governmental-regulatory-agencies/#_Toc465982655

Chapter 9

Vulnerability Management

Cobra Kai is about strength. If you're not strong on the inside, you can't be strong on the outside.

—Sensei Johnny Lawrence

The focus of this chapter is managing vulnerabilities. As covered in Chapter 1, "Introducing Security Operations and the SOC," one of the key services found in successful SOCs around the world is *vulnerability management*, which is the service responsible for reducing the organization's risk by addressing identified vulnerabilities. The quality of this service impacts how the SOC functions. Cyberattacks target vulnerabilities with the goal of delivering malware through them. If there is not a vulnerability to exploit, the attack cannot happen. The more vulnerabilities that exist in the organization, the more attacks that can be expected, because more potential targets are available to exploit. A successful vulnerability management program reduces the number of vulnerabilities, which reduces the organization's attack surface, leading to fewer possible targets.

Note

Security professionals often say that you cannot eliminate all vulnerabilities. That might be true, but the closer your organization's environment gets to zero vulnerabilities, the more expensive it is for an adversary to attack it, which makes your organization a far less attractive target and greatly reduces the likelihood of attacks. Attackers are more likely to move on to more vulnerable targets.

Incident recovery is also a critical component of a vulnerability management program. If the results of an attack, such as malware infecting a host, are remediated by the SOC, then the immediate event might be handled, which is a good thing. The problem is that the vulnerability used to deliver the malware by the attacker can still exist if it is not addressed during the remediation steps. Ignoring the root of the problem allows for future exploitation and the potential delivery of more malware. By addressing only the results of an attack and not the vulnerability being exploited, the SOC is essentially leaving the door open for another attack. Therefore, a vulnerability management team needs to be involved with all security incidents.

With regard to incident response, the role of the vulnerability management team is to identify how the target was exploited and attempt to reduce the risk of a future attack by addressing the active vulnerabilities. The incident response team must also take into consideration residual risk from performing a remediation step to avoid causing more problems with the fix. Remediating a vulnerability is not always as simple as installing a patch or shutting down a service. There can be various dependencies, and if one service is interrupted, other services can be disrupted, causing more issues than prior to when the vulnerability was addressed.

Another challenge to responding to incidents is properly remediating an advanced persistent threat. Rarely will attackers or malware leave a system willingly, which means the cleanup process for sophisticated malware can be extremely tedious and difficult to properly address. Some versions of malware can detect when remediation is occurring and trigger unwanted responses, such as changing the focus of its operation to spread itself to other systems or to go into a stealth mode, reducing the chance of the SOC properly remediating all impacted systems. Therefore, simply cleaning up, replacing, or re-imaging a compromised system is not enough to protect an organization, because malware artifacts can be left behind or spread to other systems to avoid steps implemented by the SOC during the eradication process. It is critical for the incident response team not only to understand how the organization was compromised when an incident occurs, what systems were impacted, what vulnerabilities were exploited, and what was left behind post breach, but also to understand any risk associated with any planned remediation of vulnerabilities and threats. These topics are the focus of this chapter.

This chapter looks at best practices for vulnerability management. Topics include how to prepare for what should be scanned, how to detect and prioritize vulnerabilities, and how to enforce proper remediation if an option is available that doesn't cause more trouble than value. These best practices will be viewed through the lens of people, process, and technology found in SOCs with a successful vulnerability management methodology.

Before delving into how to run a vulnerability management service, the following section provides an overview of core vulnerability management concepts.

Vulnerability Management

The more effort your SOC invests in vulnerability management, the better prepared your SOC is for future events. At its core, the goal of vulnerability management is to identify and respond to potential weakness. According to ISO/IEC 27005:2018, *Information technology — Security techniques — Information security risk management*, vulnerabilities can be related to organization, process and procedures, management routines, personnel, physical security, IT systems, or a dependency on any of these items. This chapter focuses on identifying *technical vulnerabilities*—weaknesses in information technology. Vulnerabilities that are not related to IT, such as gaps in physical security or user training, fall under the purview of risk management, while vulnerability management is specific to IT-related vulnerabilities.

Organizations with limited to no vulnerability management capabilities (outside of having scanning tools and sporadically running vulnerability assessments) typically ask the question, "What is the best place to start regarding fixing vulnerabilities when there are too many to resolve?" Another common question I often hear is, "What should be involved with a vulnerability management practice to properly operationalize dealing with vulnerabilities?" The answer to both questions is to develop a repeatable process based on the vulnerability management lifecycle, shown in Figure 9-1.

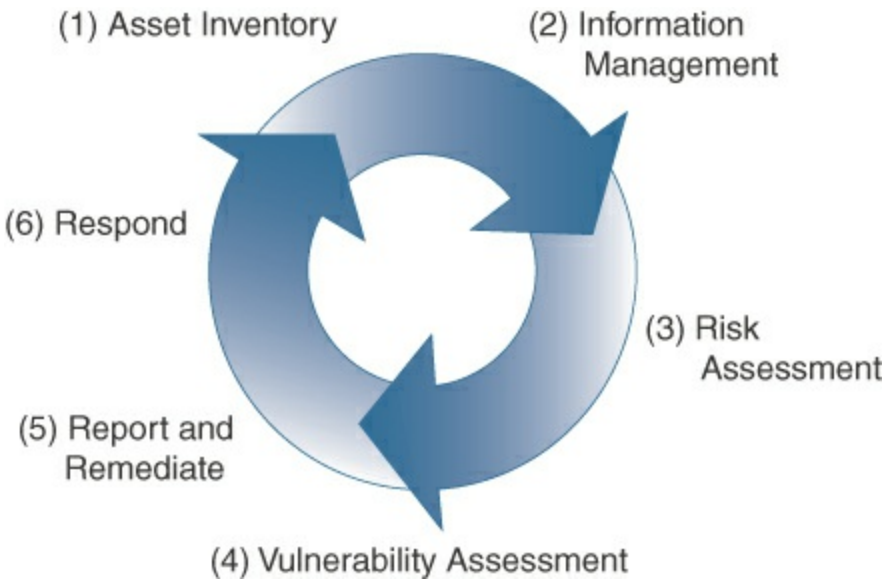

FIGURE 9-1 Best Practice for Vulnerability Management

Chapter 3, "SOC Services," briefly introduced this vulnerability management lifecycle in the context of introducing the SOC vulnerability management service. The following sections describe the concepts in more detail.

Phase 1: Asset Inventory

The model shown in Figure 9-1 begins with identifying all assets within scope of being evaluated by the vulnerability management service. In other words, what is on the network? If you don't know what's on your network, you can't properly assess everything that could be vulnerable. There is a whole segment of software products devoted to addressing the challenges of asset inventory for large organizations. Some of these offerings work well, while others have great marketing strategies but fail when it comes to finding assets on the network.

Note

One customer I was consulting was able to greatly reduce the licensing costs of several software packages after completing an asset inventory and discovery project, which identified that the organization

had overestimated the number of active computer workstations in its environment. This is just one example of the many benefits obtained by running a strong asset inventory program.

Good asset inventory means you are profiling not only endpoints but also every single network connection (wired and wireless), because many devices, such as IoT devices, do not log into centralized management or authentication servers.

Option 1: Automated Network Access Control

There are different ways you can enforce proper asset inventory practices. The industry standard for knowing who and what is on the network is to leverage network access control (NAC) based on popular guidelines such as NIST SP 800-53, *Security and Privacy Controls for Information Systems and Organizations*. By automating access control across your LAN, wireless, and VPN, you can have confidence that when a device connects to the network, you know about it. Chapter 2, "Developing a Security Operations Center," covered how NAC can be related to asset management; this chapter delves further into automated NAC concepts in the section, "Network Scanners and Network Access Control."

The biggest challenge I find in NAC solutions when speaking with customers is what is missed when NAC is not enforced across all network ports; for example, using NAC only for the wireless network or not enabling every port to be monitored by a NAC solution. Any port that is not enabled for NAC will not be able to track assets, causing potential blind spots to your list of assets that need to be validated for vulnerabilities. This issue can cause a false sense of security awareness regarding what is connected to the network, which leads to problems with the vulnerability management program.

Option 2: Manual Network Access Control

An alternative to using the automated benefits of NAC is to enforce security on a port-by-port basis, such as by using sticky MAC, and manually collect asset information from the network switch database. Sticky MAC allows a

switch to collect the MAC address of a trusted system and dedicate a port to that system. If another system attempts to connect to that port, the port will shut down.

One major flaw in this approach is how validation is done by the port checking for the MAC address. If an attacker were able to spoof the MAC address of a trusted system, the attacker could bypass port security, tricking the switch into believing that a system on the whitelist has connected. Also, port security has the same "omission" flaw in its approach as NAC. Ports that are not enabled with port security will allow devices to connect and not be tracked, again creating blind spots in your asset collection strategy.

Option 3: Network Scanners

Using network scanners is another option to collect information about devices on the network. A scanner can identify what is already on the network and develop an asset list. Challenges to this approach include how often the tool runs, if the tool has access across the entire network, and how the tool scans.

Combining network scanners with either a manual or automated NAC option is a recommended practice. This approach allows for a nice blend of collecting devices upon connection and validating that all devices have been accounted for using a periodic network scan. One of the most popular open-source tools used to identify what is on the network is Nmap (https://nmap.org/). Figure 9-2 shows an example of Nmap scanning a network segment for assets.

```
root@kali:~# nmap 198.19.20.0/24
Starting Nmap 7.70 ( https://nmap.org ) at 2020-04-19 17:09 EDT
Nmap scan report for 198.19.20.0
Host is up (0.00034s latency).
All 1000 scanned ports on 198.19.20.0 are filtered

Nmap scan report for 198.19.20.1
Host is up (0.00031s latency).
All 1000 scanned ports on 198.19.20.1 are filtered

Nmap scan report for 198.19.20.2
Host is up (0.00033s latency).
All 1000 scanned ports on 198.19.20.2 are filtered

Nmap scan report for 198.19.20.3
Host is up (0.00031s latency).
All 1000 scanned ports on 198.19.20.3 are filtered

Nmap scan report for 198.19.20.4
Host is up (0.00035s latency).
All 1000 scanned ports on 198.19.20.4 are filtered

Nmap scan report for www.hackmds.com (198.19.20.5)
Host is up (0.0024s latency).
```

FIGURE 9-2 Nmap Scanning the Network for Assets

Nmap and other network mapping tools can be helpful; however, most modern vulnerability scanners include the capability to scan for new devices. It is common to find that NAC, NAC profiling, and network scanners are normally part of an overall strategy composed of multiple technologies and solutions for asset inventory. The larger and more geographically dispersed the network, the more complicated it will be for a complete and accurate asset inventory. Best practice dictates leveraging a NAC technology as well as network scanner and integrating what is collected into a vulnerability management technology. The section "Network Scanners and Network Access Control," later in the chapter, covers this concept in more detail.

Phase 2: Information Management

The next phase of the vulnerability management model is information management. Information management consists of collecting information

about all devices on the asset list that you have created from port scans and access control data. Information such as what applications are installed and running on a system, level of patches deployed, and other factors all allow the SOC to identify potential vulnerabilities in a system. A server might not be vulnerable, but an application installed could be unsecure and vulnerable to an attacker who could use the application to gain access to the patched server. I often see this with vulnerable Internet browser plugins that allow a server that is patched and running the latest Internet browsing software to be exploited.

Looking back at NAC concepts, most modern NAC solutions offer the capability to gather information about systems accessing the network. Some examples of this feature are the Cisco ISE Profiling capability, Fortinet's FortiNAC Profiling feature, and Forescout's eyeSight feature. For most NAC vendors, including Cisco, Fortinet, and Forescout, various protocols such as DHCP, DNS, ARP, and Linkup traps are used to fingerprint a device as it connects and uses the network. Most modern NAC solutions can also use agents or Java to scan deeper behind the host's firewall, allowing full visibility of the system. Vulnerability scanners also have the capability of looking at ports and protocols to fingerprint devices. Many NAC systems incorporate their own scanners or provide integration with a major industry-recognized scanner. One popular example is, once again, using Nmap, but this time launching the **nmap -A** *target* command. This command will fingerprint a target as shown in Figure 9-3.

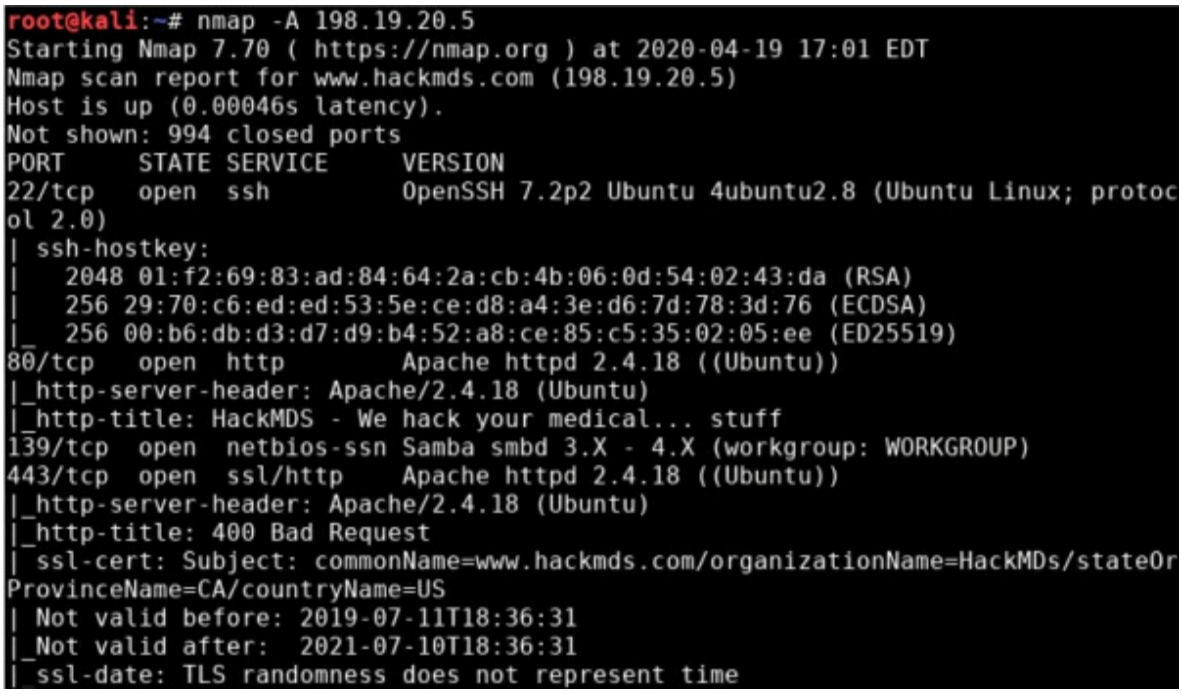

FIGURE 9-3 Nmap Fingerprinting a Targeted System

Part of the information management phase is preparing for the next phase of the vulnerability management program, which is risk assessment, in which you identify potential risk. It is critical to understand both the details of an asset and how the data was collected. The most powerful inventory tracking includes operating system information down to the patch level that is installed, applications (including installed extensions), installed libraries, configured data within the operating system, and services running on the asset. Missing these details will lead to less accurate risk assessment results.

How the collection of asset information occurs will make or break the accuracy of what you collect about an asset. The industry has an overwhelming assumption that vulnerability scanners are the most effective tool for collecting information about assets. There are many factors that impact the quality of the results from a vulnerability scanner, including if access to the system is part of the collection process versus just scanning from the network as well as what tools or other factors are limiting what data is collected by the vulnerability scanner. Using tools that can gain access to the system and understand its posture is the best way to collect true vulnerability data because those tools can get behind security tools such as a personal firewall. Security tools such as enterprise NAC offerings include

options for agents to be installed on endpoints to allow for quick and accurate profiling and posture assessments of endpoints. Device management platforms such as JAMF Pro can also access assets and capture very detailed information about an asset. Certain devices such as IoT will introduce challenges to collect asset details due to limitations on what can be installed and accessed.

Phase 3: Risk Assessment

Once you collect data about your assets, the next phase in the vulnerability management model is one that solves a challenge faced by many SOCs around the world. A SOC will scan for vulnerabilities and find there are too many vulnerable systems to be able to respond to in a reasonable amount of time. Imagine running a vulnerability assessment and the results show a report with more than 1000 potential vulnerabilities. The big question is, "Where should I start?" The answer to this question is to rank the value of each asset to the organization and target the most critical systems first. This can be accomplished using a risk assessment.

Note

The value of a risk assessment is based on the quality of the information being used. This is why it is critical to ensure the previous two phases of the vulnerability management program are performed properly.

A risk assessment allows the SOC to apply a value to each asset in question. This way, vulnerabilities found in critical resources are addressed before other vulnerabilities are handled. As a simple example, imagine that vulnerabilities are discovered in two systems, the organization's core datacenter and a demo lab. The vulnerability in the core datacenter is considered semi-risky while the demo lab's vulnerability is considered critical. Which one should you address first? If you were to look only at the vulnerability report, you would find that the demo center's vulnerability

listed higher on the report; however, the core datacenter is much more important to the organization and, therefore, associated vulnerabilities should be addressed prior to vulnerabilities in the demo lab. A risk assessment allows the SOC to prioritize the response to ensure critical systems are addressed first, rather than just working from the top of a vulnerability assessment report. Risk assessments can also be outsourced to contractors to get an objective point of view of your potential risks.

Part of the risk assessment process is to identify the best approach to address the risk, which leads back to why the vulnerability management team must also be part of the incident response program. Suppose you find a high-priority risk that is common across thousands of assets. The combination of the number of assets impacted by the risk can raise the risk of a vulnerability; hence, increasing its priority to be remediated. This also increases the chance of residual risk as well as the cost to perform remediation, which if too high can reduce the priority for remediation actions to be performed. If remediation requires hours of manual labor, performing remediation across 20,000 assets will have a huge cost and likely take an exceptionally long time, leading to a prolonged exposure of risk and an impossible task to accomplish. A risk assessment team can suggest different approaches to addressing risk so that the priority of remediation can be readjusted. Rather than addressing each asset, maybe enabling one or two IPS rules can address the risk across the 20,000 assets, dramatically reducing the time and cost for remediation and allowing tasks that are more likely achievable to be moved to the top of the SOC's to-do list.

Phase 4: Vulnerability Assessment

Once you have established the risk associated with all assets, you must figure out which vulnerabilities to address. Groups of high-risk systems will be your first target for remediation. You will work your way through those vulnerabilities before moving to systems of less value, and so on until you run out of resources to address the identified vulnerabilities. Most organizations will never address *all* vulnerabilities, as the known vulnerability list ends up being a moving target when proper asset collection and risk assessment are continuously performed. That is okay if all

vulnerabilities are eventually addressed based on their priority and time within the queue.

Vulnerability management best practice dictates using a case management system that tracks how vulnerabilities are being addressed to ensure that certain vulnerabilities are not overlooked. Overlooking a vulnerability even on a noncritical system could allow attackers to leverage that vulnerability to gain access and eventually pivot to where the critical resources live. You want to avoid having lower priority vulnerabilities from continuously being dropped from the to-do list or they will end up never being addressed. Time in the queue needs to be a factor that is included in evaluation of vulnerabilities. As vulnerabilities sit in the queue, they gain more priority based on the delay to address them. Many case management applications include a feature to monitor for prolonged delays addressing incidents.

Phase 5: Report and Remediate

Remediating a vulnerability might or might not be a simple fix. The work required to fix the vulnerability could break certain applications, require a system reboot taking the system offline, or affect the system in other ways that could lead to an impact that must be taken into consideration before a fix can be applied. That is why step 3 of the vulnerability management lifecycle includes assessing these types of risks and considering all options regarding how remediation can occur. For example, placing a security tool between the vulnerable system that prevents exploitation may be a much better approach than trying to fix the vulnerability. Or maybe the impacted system can be moved to a place on the network that is isolated from potential threats. All aspects of remediation must be considered before any action is taken.

At this point of the vulnerability management process, the SOC needs to consider all datapoints about the risk of the vulnerability as well as what is required for remediation and choose to accept the risk or take action. Before action is taken, the SOC needs to report what was identified, what actions are recommended, and why the recommendations are being made. Documentation will be used to track the vulnerability and referenced during any lessons learned review efforts. You learn more about understanding vulnerabilities, including how to properly remediate when applicable, in the

section, “Vulnerability Response.”

Phase 6: Respond and Repeat

The final phase of the vulnerability management model is to perform the planned response until it is time to repeat the cycle. Responses can range from deploying patches to implementing security measures that reduce the risk of the vulnerability being compromised. I cover best practices for responding to vulnerabilities later in this chapter.

A SOC could start the vulnerability management cycle on a Monday and by Thursday, certain systems could be addressed, while many others are still waiting to be assessed for remediation. When the following Monday rolls around, the SOC will repeat the cycle, which would validate vulnerabilities still exist that were not addressed as well as identify any new vulnerabilities. Vulnerabilities that are a repeat from the previous week will not be replaced with new data but instead be marked as carryover items, so they do not linger for weeks without being addressed. If all runs well, the number of new vulnerabilities combined with existing vulnerabilities should be reduced until a manageable number of events remain to be handled in each cycle. This will vary on the level of effort required to handle vulnerabilities, rules of making changes during maintenance windows, and other factors, which you will see in more detail later in this chapter in the section “Vulnerability Response.”

One additional aspect of the respond and repeat phase includes proactive measures to reduce the risk of repetitive vulnerabilities appearing in the network. Like with a lessons learned meeting, a lessons learned review needs to be included so that the SOC can adapt the vulnerability management program based on what it has seen and how effective the response was. After the SOC has identified known vulnerabilities, it can tune its scanning and posture-enforcement tools to automate vulnerability remediation, assuming the remediation has been proven to work and includes low residual risk. Steps in the vulnerability management lifecycle program can be adjusted based on what is and is not working. For example, manual efforts can be eventually automated, such as how asset information is collected, or how vulnerability and risk information is assessed. Eventually, the SOC can develop playbooks and apply orchestration, which is the focus of Chapter 10, “Data

Orchestration."

From a high level, the vulnerability management model is a great approach to managing vulnerabilities and a guideline you can validate your practice against. Before looking more deeply into how to manage vulnerabilities, you must first comprehend what vulnerabilities are and how the industry shares vulnerability data. This leads to the next topic, which is a closer look at how vulnerabilities are identified in the security industry.

Measuring Vulnerabilities

The opening of this chapter provided a general explanation of vulnerabilities. ISO/IEC 27000:2018 defines a *vulnerability* as "a weakness of an asset or control that can be exploited by one or more threats." But how can you know when something is vulnerable to an attack outside of having the device compromised by an unwanted party? Who should have the authority to determine the severity of a vulnerability so that all vendors, organizations, and researchers have a standardized language for vulnerabilities? How do you compare the potential threat of a compromised system that could gain access to the network against an attack that could take down a server when both are completely different weaknesses? Chapter 1 answered these questions in the content about measuring vulnerability metrics using the Common Vulnerabilities and Exposer as well as the Common Vulnerability Scoring System (CVSS) scoring system. Let's review each of these options for better understanding a vulnerability.

Common Vulnerabilities and Exposures

One popular resource for finding vulnerabilities is the Common Vulnerabilities and Exposures (CVE) representing a list of publicly disclosed vulnerabilities and exposures maintained by MITRE. The CVE includes a CVE ID, a description, dates, and comments. Figure 9-4 is an example of an Apache Struts 2 advisory including the (CVE) number CVE-2017-9793, which allows the industry to have one point of reference to this vulnerability. By searching CVE-2017-9793, you will be able to bring up details on this

particular vulnerability. The CVSS or CVE systems are far from perfect, and there are alternative methods to capture the characteristics of a vulnerability, but CVSS and CVE are the most widely used methods for tracking vulnerability information and is supported by multiple vendors and tools.

FIGURE 9-4 Struts Vulnerability Example

Common Vulnerability Scoring System

The CVSS provides a way to capture the principal characteristics of a vulnerability and produce a numerical score reflecting its severity. The CVSS is an open framework managed by FIRST (https://www.first.org) and supported by the National Vulnerability Database (NVD) at https://nvd.nist.gov/. The CVSS consists of three metric groups. Base, Temporal, and Environmental.

- **Base:** The Base group represents the intrinsic qualities of a vulnerability that are constant over time and across user environments.
- **Temporal:** The Temporal group reflects the characteristics of a vulnerability that change over time.
- **Environmental:** The Environmental group represents the characteristics of a vulnerability that are unique to a user's environment.

The Base metrics produce a score ranging from 0 to 10, which you can then modify by scoring the temporal and environmental metrics. The NVD only provides a computed base score for the vulnerabilities shown; however, it offers calculators for adding temporal and environmental conditions, which

you must manually enter.

The CVSS not only is used by most industry security tools that include any form of vulnerability management capabilities but also is the standard for industries, organizations, and governments that need accurate and consistent vulnerability severity scores.

Note

The difference between the CVSS and CVE system is that CVSS represents an overall score assigned to a vulnerability. CVE is simply a list of all publicly disclosed vulnerabilities. The CVSS score is not reported in the CVE listing. You must use the NVD to find assigned CVSS scores.

CVSS Standards

The NVD provides two standards for CVSS scores: CVSS version 2.0 and CVSS version 3.1. The provided base scores on the NVD can be translated to the baseline threat, meaning they represent the *innate characteristics of each vulnerability*. This is important to be aware of since the NVD does not adjust scores as the risk changes with a vulnerability and its value is just an estimated risk. The NVD does offer a CVSS calculator for both CVSS v2 and v3, allowing you to add temporal and environmental score data, which would take into consideration changes in the risk of a vulnerability. Even with that data, there are still factors in particular to your organization that you would need to account for to truly understand the associated risk; however, the CVSS is a great vendor- and organization-neutral starting place.

Note

Both CVSS v2 and v3 calculators can be found at https://nvd.nist.gov/vuln-metrics/cvss.

The severity of a vulnerability is weighed within different metrics that use three brackets for scoring CVSS v2. Any vulnerability 0–3.9 is considered low, any vulnerability 4.0–6.9 is considered medium, and any vulnerability 7.0–10 is high. CVSS v3.1 uses brackets adding more granularity to how a vulnerability can be weighed. Those brackets are None, Low, Medium, High, and Critical.

Table 9-1 represents the score ranges for each level of a CVSS 2.0 and 3.1 vulnerability severity. Metrics can be measurements, such as whether the attack vector is used or whether the confidentiality of data can be compromised. The NVD will look at a vulnerability and consider the potential impact against the metrics, which will result in a severity weight. If the attack vector could come from anywhere on the Internet, that factor will represent a very high severity. If the attack vector must be within a network local to the system, the severity of that specific metric goes down because it is difficult for an attacker to gain that level of access to deliver the attack. If the vulnerability allows for the full compromise of confidentiality of the impacted system, the severity for that specific metric will be high. If the vulnerability does not breach the confidentiality of the impacted system, the confidentiality metric will be very low. A denial-of-service attack can represent a high metric for an attack vector because it can be launched anywhere, but the confidentiality metric would be low because taking down a system doesn't breach its confidentiality. A CVSS final score represents the combination of all metrics computed together.

TABLE 9-1 CVSS v2 and v3 Ratings

CVSS v2.0 Ratings		**CVSS v3.1 Ratings**	
Severity	**Base Score Range**	**Severity**	**Base Score Range**
		None	0.0
Low	0.–3.9	Low	0.1–3.9
Medium	4.0–6.9	Medium	4.0–6.9
High	7.0–10.0	High	7.0–8.9
		Critical	9.0–10.0

Let’s look at each CVSS version to better understand the associated factors that are being evaluated for severity.

CVSS Version 2

To better understand CVSS v2, you can use the v2 available calculator found at https://nvd.nist.gov/vuln-metrics/cvss/v2-calculator. The base score calculator allows you to choose the three different severity options for six different metrics, meaning low, medium, or high for each of six metrics. Three metrics are grouped under Exploitability, meaning the attack vector, access complexity, and authentication all have to do with how easy it is for an attacker to exploit the vulnerability. The other three metrics are grouped under Impact, meaning confidentiality, integrity, and availability all can impact the targeted system if the vulnerability is exploited. Figure 9-5 shows the base score calculator, with which you select one of three severity options for six different metrics and the calculator combines the results to calculate the overall potential severity of a vulnerability.

FIGURE 9-5 CVSS v2 Base Score Calculator

The following explains each of the CVSS v2 metrics found within the base score calculator in more detail:

- **Attack Vector (AV):** This is how the vulnerability could be accessed and exploited. There are three possible attack vectors:
  - **Local (L):** The attacker must have physical access to a vulnerable system or to a local account that gives such access. Local represents a

low score severity because most attackers won't obtain local level of access to a system before exploiting a vulnerability.

- **Adjacent Network (A):** The attacker must have access to the broadcast or collision domain of the vulnerable system, such as using ARP spoofing or Bluetooth attacks.
- **Network (N):** The attacker can access the vulnerable system remotely. This option has the highest value because the attacker can access the vulnerability from anywhere.

- **Access Complexity (AC):** Complexity means how easy or difficult it is to exploit the vulnerability. There are three levels, High (H), Medium (M), and Low (L). Low is the highest value, meaning it doesn't take a complex attack to exploit the vulnerability.
- **Authentication (Au):** This measurement covers the number of times that an attacker must authenticate to a target to exploit it. An example is requiring multifactor authentication to access the vulnerability. Accessing other elements such as the network is not considered—the focus is just on authentication. The three values are Multiple (M), meaning the attacker must authenticate two or more times to exploit the vulnerability, Single (S), and None. None has the highest value because that means authentication is not required to access the vulnerability.

  The next three metrics are the impact metrics, which cover the core security defenses concepts identified in pretty much every security certification program, which are confidentiality, integrity, and availability (commonly referred to as the CIA triad). As a security professional, your job is to protect these three aspects of your data. The three metrics corresponding to these aspects all have three rankings, which are None (N), Partial (P), and Complete (C). Complete has the highest value representing a complete loss of confidentiality, integrity, or availability to the data.
- **Confidentiality Impact (C):** This factor considers the impact on the confidentiality of the data processed by the system if the vulnerability is exploited.
- **Integrity Impact (I):** This factor relates to the integrity of the exploited system and data if the vulnerability is compromised.

- **Availability Impact (A):** Availability relates to the impact to what is available if the vulnerability is exploited. Examples include if the network bandwidth is consumed, the processes exhausted, or other resources are negatively impacted, no longer allowing access to the data.

Temporal and Environmental Metrics

The CVSS is great to understand the potential risk of vulnerabilities, but what if things change? Once the NVD releases information about a vulnerability, that data remains static. For example, the NVD can assign an extremely high value to a new vulnerability, but that value needs to be reduced as the industry responds to the vulnerability with remediation steps and better understands how the vulnerability can be exploited. The NVD does not continuously adjust for change, but it does offer temporal and environmental metrics that can accommodate for change.

Temporal metrics accommodate characteristics that have evolved over the lifetime of the vulnerability. This additional calculation allows the SOC to accommodate the various changes that could occur and impact the true severity of a vulnerability. Examples of change include whether the exploit of the vulnerability has or has not been proven, the availability of a fix at the current moment, and the overall confidence that the vulnerability is a true weakness versus a false positive. All these aspects change as a vulnerability becomes known and is addressed by the industry.

Environmental metrics take into consideration a specific implementation or environment, which also will impact the severity. For some environments, the collateral damage that would occur if the vulnerability were exploited would be extremely severe, while other environments that have impacted systems in a demo environment would provide little to no impact to the organization if those systems with the vulnerability were exploited. Figure 9-6 shows the CVSS v2 Temporal and Environmental Scoring metrics with additional factors that are taken into consideration to adjust the severity score of a vulnerability. As with the base scoring system, you choose the factor and the CVSS calculator will apply your input to the calculation of the severity of the vulnerability.

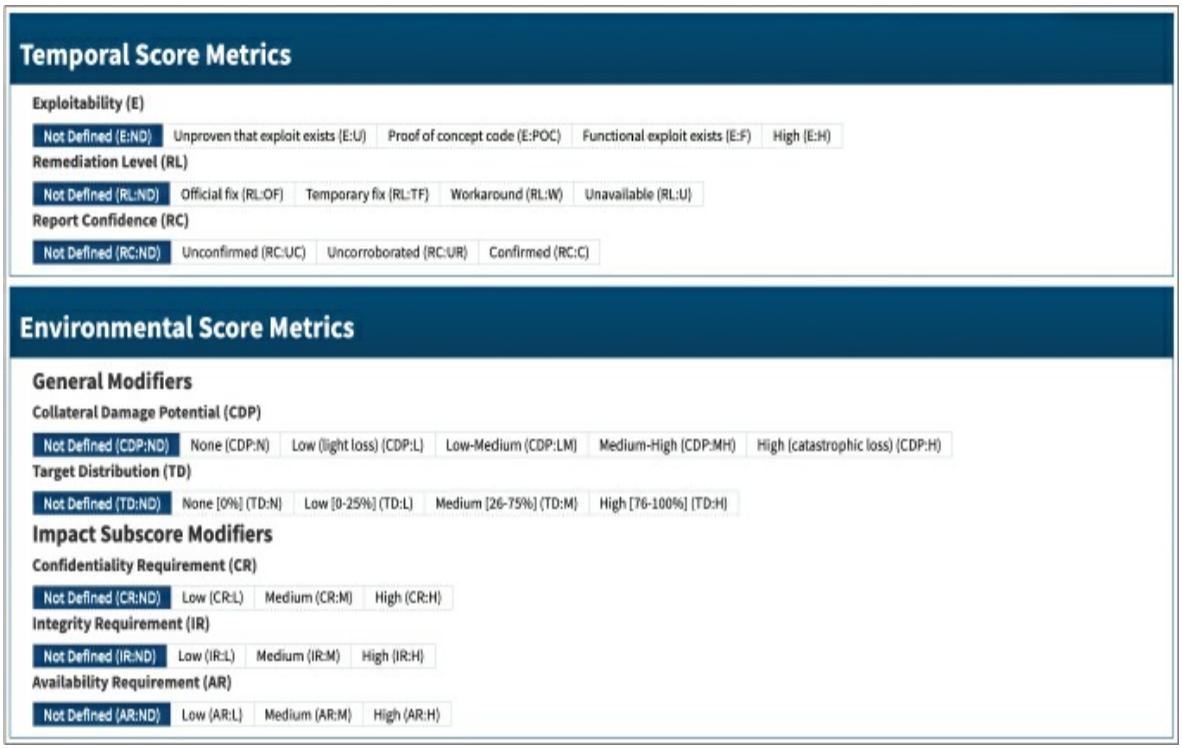

FIGURE 9-6 CVSS Temporal and Environmental Calculators

Understanding CVSS Shorthand

FIRST doesn’t expect you to manually compute a value for each CVSS value, as that would be extremely tedious. The severity of a vulnerability is already computed as a risk score based on all the metric values, and metrics are summarized using a form of shorthand. The reason the scoring and shorthand of the metrics is provided within the NVD is to show how the risk value was calculated, so you can better understand the threat and why the score is what is shown rather than using the calculators to generate your own score.

To simplify understanding what metric weights were considered for a risk score, CVSS v2 uses a shortened summary of each score to represent how a CVSS score was computed. Notice that each metric and weight showed included brackets ( ) with a shortened phrase within the brackets. The shortened phrase within the brackets is used followed by a colon : to represent which of the options was selected for the particular metrics. For example, AV is used to represent attack vector and L is used to represent

local. This would be represented as AV:L in the CVSS shorthand. A backslash (/) is used to separate the next factor, as all size factors are displayed as one shorthand explanation of the vulnerability's risk factors. Another example is continuing with using AC for access complexity, which is the factor to follow attack vector. A full CVSS 2.0 state can look like AV:N/AC:L/Au:N/C:P/I:P/A:C. If you look back at Figure 9-5, you will notice AV:N/AC:L/Au:N/C:P/I:P/A:C matches what was selected within the CVSS calculator as its shorthand value.

To put this all together, Figure 9-7 shows the Severity section of the NVD page for the Apache Struts vulnerability CVE-2017-9793 (mentioned earlier) based on CVSS Version 2.0. This shows a base score of 5.0 and an explanation of how that was computed using the CVSS v2 shorthand.

FIGURE 9-7 Struts CVSS v2 Example

CVSS Version 3

CVSS v2 is still an industry standard; however, it has received some criticism from vendors and customers. To address such concerns, FIRST developed CVSS v3 to provide vendors with a method to better analyze security vulnerability impact as well as permit customers to provide more detail so they can determine the urgency needed to respond to vulnerabilities. Industry expert Omar Santos (who happens to be a good buddy of mine) published an article on Cisco Blogs (https://blogs.cisco.com/security/cvssv3-study) explaining his research on the changes between CVSS v2 and CVSS v3, including how he found CVSS v3 produces an increase in base score for many vulnerabilities when calculated using the new system. A slight increase in concern can cause major changes in how a SOC functions, including triggering different work orchestrations and causing changes in the order in which customers address vulnerabilities.

In June 2019, FIRST released a minor update, CVSS v3.1, designed to clarify and further improve the 3.0 standard without introducing new metrics or metric values. Figure 9-8 illustrates the changes in the CVSS v3 base scoring system from the CVSS v2 system, shown previously in Figure 9-5. An example of a change is that the attack vector now includes a Physical (AV:P) component.

FIGURE 9-8 CVSS v3 Base Score Metrics

NVD has a single calculator for both CVSS 3 versions, 3.0 and 3.1, called the 3.x calculator and found at https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator. Usage is like the v2 calculator, including the shorthand summary to understand how a severity score was calculated. Figure 9-9 shows the same Struts vulnerability with the CVE number of CVE-2017-9793 represented in the v3 scoring system. Notice that the vector shorthand is a longer result to accommodate the new factors added to version 3.x. The results using the version 3.x calculator have a base score of 7.5 versus the 5.0 shown using the version 2.0 calculator.

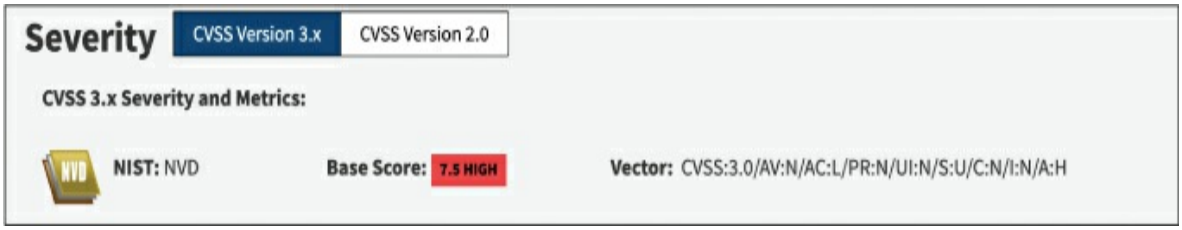

FIGURE 9-9 Struts CVSS v3 Example

When you look up a vulnerability in the NVD, you see not only the base score but also some other important information. Below where the severity score is listed are references to advisories, solutions, and tools. Advisories are formal warnings about the vulnerability from manufactures of the impacted technology as well as other security advisory groups. Figure 9-10 shows the references to advisories, solutions, and tools for the CVE-2017-9793 Struts vulnerability, the sources for which span from security trackers to Apache's own advisory and solution to the issue in the form of a patch. Each resource is labeled and includes a hyperlink.

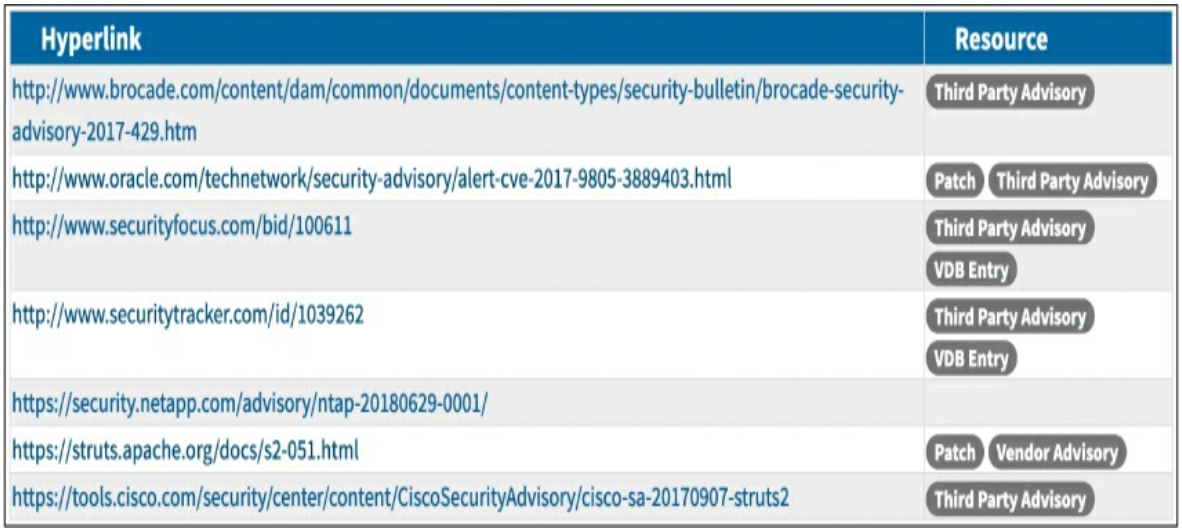

Hyperlink	Resource
http://www.brocade.com/content/dam/common/documents/content-types/security-bulletin/brocade-security-advisory-2017-429.htm	Third Party Advisory
http://www.oracle.com/technetwork/security-advisory/alert-cve-2017-9805-3889403.html	Patch, Third Party Advisory
http://www.securityfocus.com/bid/100611	Third Party Advisory, VDB Entry
http://www.securitytracker.com/id/1039262	Third Party Advisory, VDB Entry
https://security.netapp.com/advisory/ntap-20180629-0001/	
https://struts.apache.org/docs/s2-051.html	Patch, Vendor Advisory
https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20170907-struts2	Third Party Advisory

FIGURE 9-10 Struts CVE-2017-9793 Resource Example

Note

You should highly consider validating any vulnerability with the vendor to learn the latest information, including the vendor's advised steps to remediate the associated risk.

Another key piece of information included in the NVD is a list of all impacted systems. For the Struts example, the Known Affected Software Configurations section lists many versions of Apache Struts, including cpe:2.3:a:apache:struts:2.3.7:*:*:*:*:*:*:*:*,

cpe:2.3:a:apache:struts:2.3.8:*:*:*:*:*:*:*, and many others. Once again, I highly recommend validating directly with the vendor whether your specific system and software would be impacted, to ensure the NVD hasn't inaccurately listed or not listed your system or software as potentially vulnerable.

Most security tool vendors use either or both CVSS scoring systems in their products when refencing vulnerabilities. Figure 9-11 is an example of a Struts vulnerability shown in Rapid7's Nexpose. Notice that both CVSS v2 and v3 shorthand is shown as well as the difference in CVSS scores for both versions. In this example, the CVSS v2 score is much higher (9.3) than the v3 score (8.1).

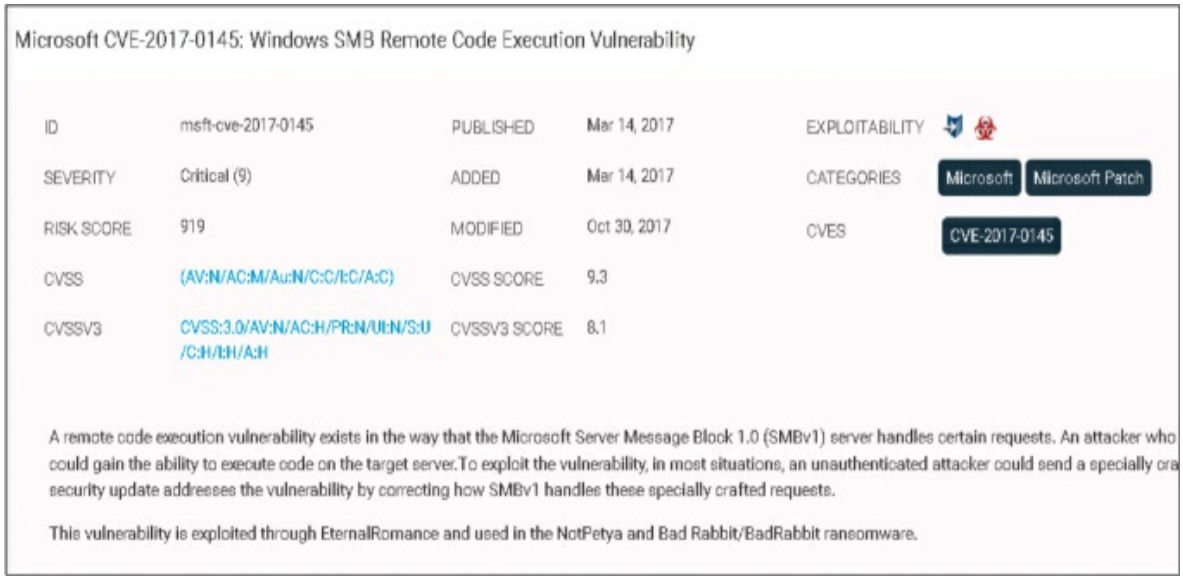

FIGURE 9-11 Struts Vulnerability Shown in Rapid7's Nexpose

Vulnerability Technology

Now that you have an understanding of what goes into calculating the severity of a vulnerability, it's time to operationalize that knowledge into your practice. To do so, your SOC needs some key tools that include capabilities to identify what is on the network, to scan the network for vulnerabilities, and to conduct some form of case management to ensure all

vulnerabilities are tracked and addressed. The categories covered in the sections that follow can be used to address these needs.

Vulnerability Scanners

The most important tool and the first to use is a vulnerability scanner. A good vulnerability scanner can evaluate a target for potential vulnerabilities and assign corresponding CVSS scores to any vulnerabilities that it discovers. Research firms such as Gartner name Tenable, Rapid7, Qualys, and Beyond Security are some of the top vendors in the vulnerability management space. There are also open-source options, such as Open Vulnerability Assessment Scanner (OpenVAS), that you can use, but they lack some of the capabilities included with enterprise offerings. There are lots of fans of OpenVAS, and I think it enhances the capabilities of a SOC as a secondary tool; however, I find that most organizations need an enterprise-level, commercial tool. Figure 9-12 is an example of results displayed in Rapid7's Nexpose showing a summary of vulnerabilities based on the CVSS score (left) and based on the skill level required to exploit the vulnerabilities (right).

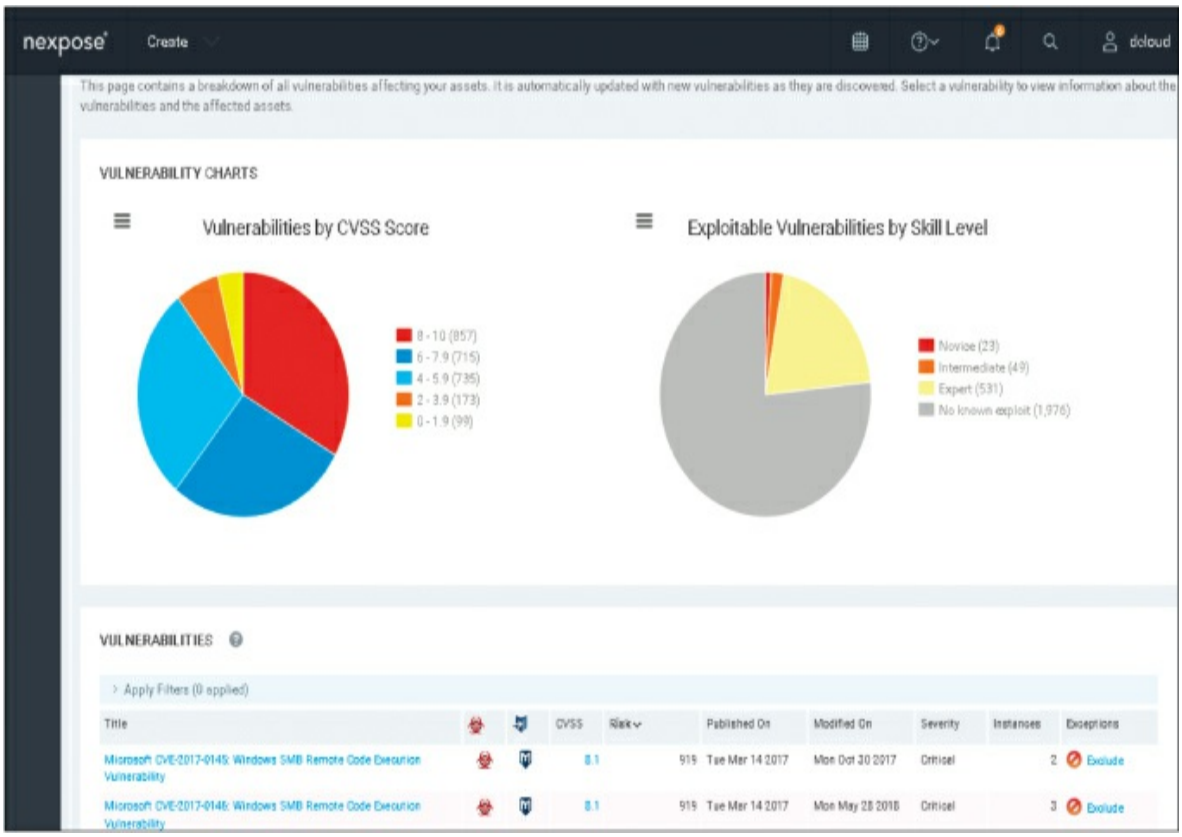

FIGURE 9-12 Rapid7 Nexpose Dashboard

Active Scanning

Scanning for vulnerabilities can be broken into two active scanning practices: network scanning and host scanning. *Network scanning* means a vulnerability scanner can evaluate a target over the network to identify potential vulnerabilities looking from the outside in. The benefits of this approach include not requiring permission to scan any target and gathering what the industry would deem as the “attacker’s point of view.” This means the scanner would see the same vulnerabilities as any outside entity, including attackers looking to compromise the vulnerability. The cons of network scanning include not having a true evaluation of vulnerabilities because host-level access is not permitted, meaning host security tools will block visibility into some vulnerabilities. Another con is how network resources are required to deliver the scan, which can lead to stress on the network. Network scanning can also trigger security tools and possibly harm systems being evaluated. I have seen situations where vulnerability scanners have taken

older printers offline, causing upset customers.

The other type of active scanning is *host scanning*, which means allowing a vulnerability scanner to access host-level details, including potential vulnerabilities. The biggest benefit of this approach is that it provides very accurate data because the data is collected behind security tools, such as the host firewall, that can prevent external parties from identifying vulnerabilities. The capability of host security tools to prevent a scanner from seeing their vulnerabilities should make sense; it would be bad if outsiders could easily identify how a host is vulnerable to attack.

Using a blend of both network scanning and host scanning is a recommended practice as both offer unique value. Host-level details will be your baseline for all vulnerabilities because the data comes directly from the system level. Network vulnerability findings, however, can influence the SOC's response because a real adversary is more likely to target vulnerabilities found in the results from a network scan than to target vulnerabilities that are being protected by host security capabilities.

Note

Keep in mind that just because a vulnerability scanner does not detect a vulnerability doesn't mean the host is protected! The quality of a scan depends on how the scanner is tuned. I will cover tuning vulnerability scanners as the next topic.

Passive Scanning

A third type of scanning is *passive scanning*, meaning the vulnerability scanner collects data about systems as it monitors traffic and compares that data against a database of vulnerabilities. The benefit of this approach is that no additional scanning is needed to gain vulnerability information about a system. The biggest gap in this approach is that a passive scanner can only see the traffic that crosses its path. If the system never generates traffic that crosses the tool, the system will never be scanned. Security tools such as application-layer firewalls have this capability.

To find vulnerabilities on systems that are not validated by a passive scanning tool, active scanning such as provided by a traditional vulnerability scanning tool will be required to scan against a specific part of the network for any vulnerabilities that exist. Figure 9-13 shows an example of a Cisco next-generation firewall listing vulnerabilities identified using a passive vulnerability evaluation, meaning none of these vulnerabilities would be seen if the traffic didn't cross the security appliance.

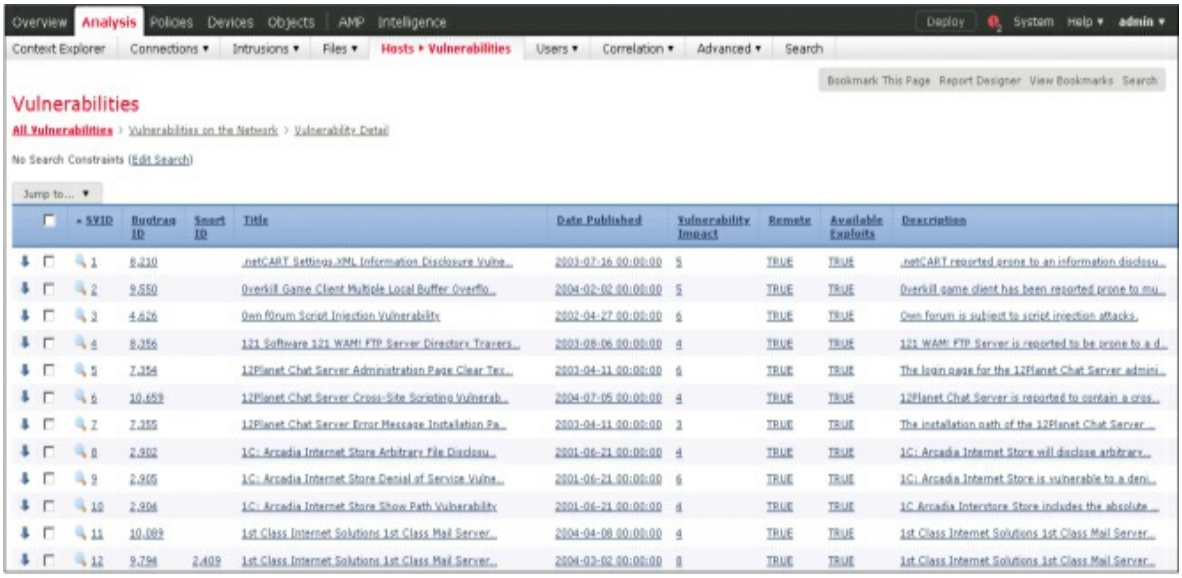

FIGURE 9-13 Passive Vulnerability Scanning Example

Currency and Coverage

Regardless of the approach you use to scan for vulnerabilities, you need to consider how current the vulnerability data being used is as well as what technologies are part of the coverage model. If your vulnerability scanner is kept current through regular updates, part of your vulnerability management program must be to ensure updates are enforced. Not doing so will render your results much less valuable because vulnerability data is continuously changing. All device types and vulnerability types need to be included if your SOC is to provide a thorough vulnerability management service. When evaluating a vulnerability assessment technology, consider all system types, including mobile devices, applications, cloud-based technology, and security

tools, that may prevent vulnerability data from being collected.

If you want to test vulnerability scanning using free options, two good candidates are Nmap and OpenVAS. OpenVAS can be found at https://www.openvas.org/. Figure 9-14 is a screenshot of OpenVAS running from Kali Linux. Tenable also offers a free version of its Nessus scanner, which at the time of writing allows for up to 16 IP addresses to be scanned.

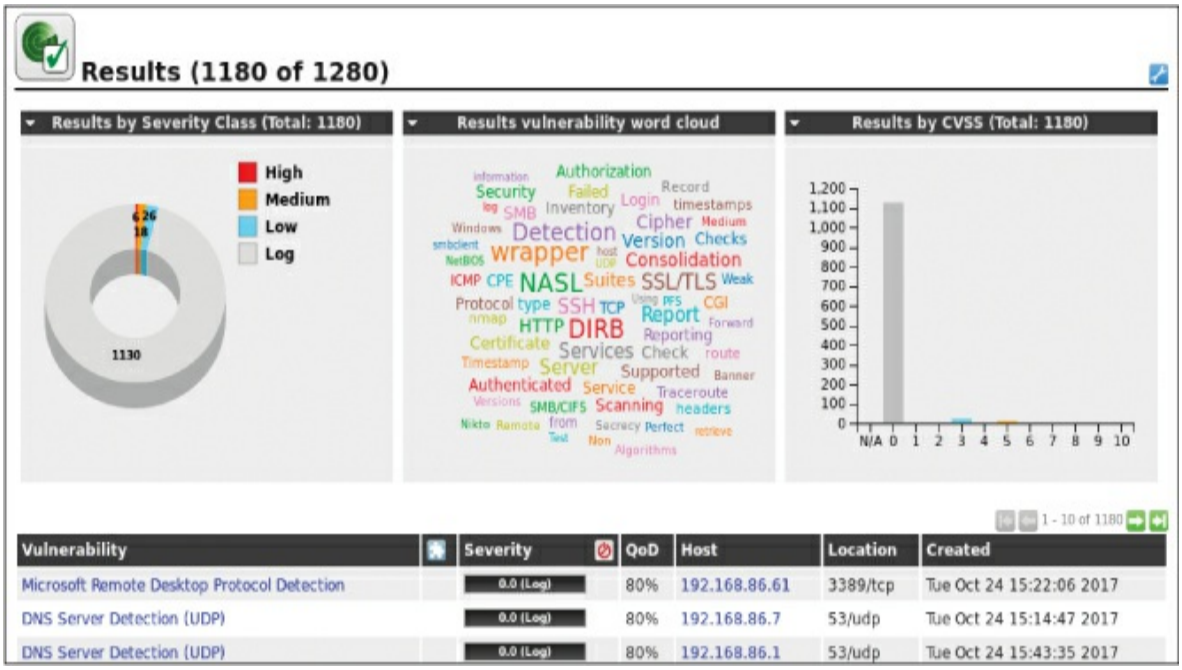

FIGURE 9-14 OpenVAS Example

Tuning Vulnerability Scanners

Another critical concept to cover is tuning vulnerability technology. Vulnerability evaluation solutions are just like any other security product regarding expected results. The quality of the tool, how much time is allowed for scanning, and how the vulnerability scanner is tuned (what it is configured to look for) will determine how effective the tool is at identifying vulnerabilities. Security tools miss vulnerabilities for many reasons. For example, the vulnerability database might not contain details about a vulnerability, if the expected traffic and data are available to be scanned that are needed to identify vulnerabilities, when the scan occurred, ensuring all

systems are available to be scanned, and if the vulnerability within the system is capable of being identified using a vulnerability scanner.

Rapid7, the creator of Nexpose, explains that there are three factors to consider when tuning a vulnerability scanner, as illustrated in Figure 9-15 and described in the list that follows.

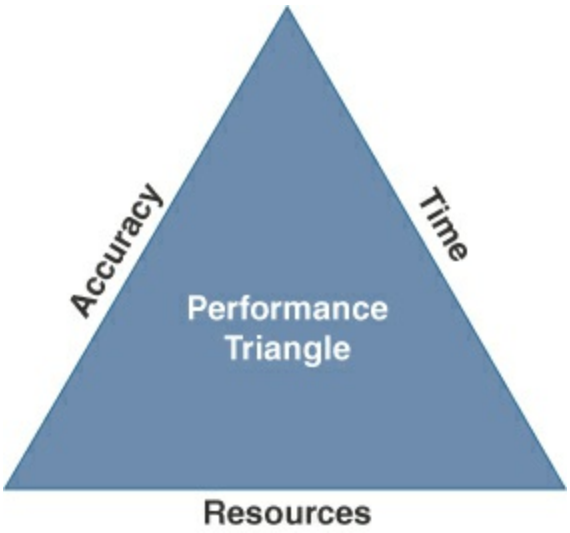

FIGURE 9-15 Triangle of Tuning Vulnerability Scanners

- **Accuracy:** How much data is collected to be evaluated (the quality is important, not the quantity)
- **Time:** How much time is allowed for the scanner to collect data
- **Resources:** How many resources are available to allow the scan

As Figure 9-15 illustrates, Rapid7 uses a triangle to represent each side as a category you can favor as you tune with that as the focus. In this example, Rapid7 explains that one or two sides of the triangle can be lengthened, representing investment in focus, but all three sides can't be the focus. One area would need to be sacrificed while the others are improved.

As an example use case to better understand the triangle tuning concept, consider an organization with locations spread around the world that wants its SOC to scan its datacenters. There are the three tuning factors to consider:

- **Accuracy:** If the SOC wants accurate results, it needs to perform a full scan.

- **Resources:** Performing a full scan will take away resources on the datacenter servers, which could have a negative impact to either the network if a network scan is performed or the servers themselves if a client/agent is performing the scan.
- **Time:** How often the SOC wants to push this scan and for how long it can scan. If the scan window is only after business hours, that means the scan will have to be paused and reenabled, increasing the time needed to collect the required data.

For this example and any situation that involves tuning a vulnerability scanner, decisions will have to be made regarding which of these three factors will be improved based on two key goals, which are the following:

- What are your goals and priorities for tuning the vulnerability scanner?
- What factors of the vulnerability tuning triangle are you willing to compromise on?

There are different ways to improve the three factors of the vulnerability tuning triangle. For this example, the following list offers a few ways to improve each factor for this use case.

- Improving accuracy:
  - Increase the time dedicated to a scan
  - Allocate more scan threads to the vulnerability scanner
  - Increase the ports and protocols that are being scanned
- Improving timing:
  - Increase the number of assets that are scanned simultaneously
  - Add scan engines
  - Use a less exhaustive scan template
- Improving resources:
  - Scan during nonoperational hours
  - Use a less exhaustive scan template
  - Increase scanning resources and deploy local scanners if network

resources are a concern

If the organization wants to improve accuracy by increasing what is scanned, then more time will be needed to conduct the scan, which will hurt the timing factor. Also, resources would be impacted if scanning occurs during normal business hours. Improvement to accuracy would negatively impact timing and resources. If timing and resources are critical, less exhaustive templates can be used and scanning could occur during nonwork hours; however, accuracy and timing would be negatively impacted. Timing would not be hit as hard as the accuracy factor because a less exhaustive scan template is being used; however, the accuracy would take a much larger hit. As you can see, there is no way to improve all three factors, so decisions need to be made regarding what is most important and attempt to balance between these three factors. The options for adjusting for these three factors will depend on the vendor of choice, available resources, and business goals.

Most vulnerability management solutions offer templates and the ability to run different templates against different parts of the network. This will allow you to balance the tuning for different jobs that apply to different parts of your network. An example of this is using a less exhaustive template for critical systems during business hours, but using a more accurate scanning template over the weekend or during nonbusiness hours. Combining different templates allows one scan that has a specific focus to fill the gaps in another scan type, giving a clear view of all the vulnerabilities within an organization.

Note

If you are having trouble deciding whether accuracy, timing, or resources is the most critical focus for your tuning, evaluate your tuning strategy against your mission statement. If the SOC's mission statement is deeply focused on sustaining services, resources would be your focus. If the mission statement is focused on protection, accuracy and timing will be more important.

Exploitation Tools

Vulnerability assessments identify potential vulnerabilities, while penetration testing further evaluates a vulnerability by exploiting it in the same manner as a malicious party would. The list of results from a penetration test will be much shorter than what is found with a vulnerability assessment because a vulnerability scanner is a generalized assumption that something is vulnerable without any validation. The results of a penetration test are validated vulnerabilities and should be prioritized by the SOC.

Penetration testers use tools to simplify the exploitation evaluation process during a penetration test. Exploitation tools can be the same tools used by attackers or commercial tools designed to replicate real-world exploitation tactics. One of the most popular open-source frameworks for identifying and executing exploits is Rapid7's Metasploit project (https://www.metasploit.com/). One basic feature of Metasploit is its search engine for weaponized exploits against known vulnerabilities. If your SOC needs to evaluate a vulnerability, such as the Struts example used earlier in this chapter, you could search for key terms such as Struts, Apache, or even the CVE number, CVE-2017-9793. Searching these terms in Metasploit will result in a list of potential exploits that could be executed against a system with the Struts weakness.

Note

Most vulnerability reporting sources used by penetration testing tools base their research disclosure and coverage on application use, meaning applications that are sold and commonly used within organizations. If an organization uses an in-house app-development process, vulnerability data likely will not be included in an enterprise exploitation tool since the vendor of the tool will not have data on custom-built applications.

Figure 9-16 shows an example of Metasploit listing the possible exploits against CVE-2017-9793.

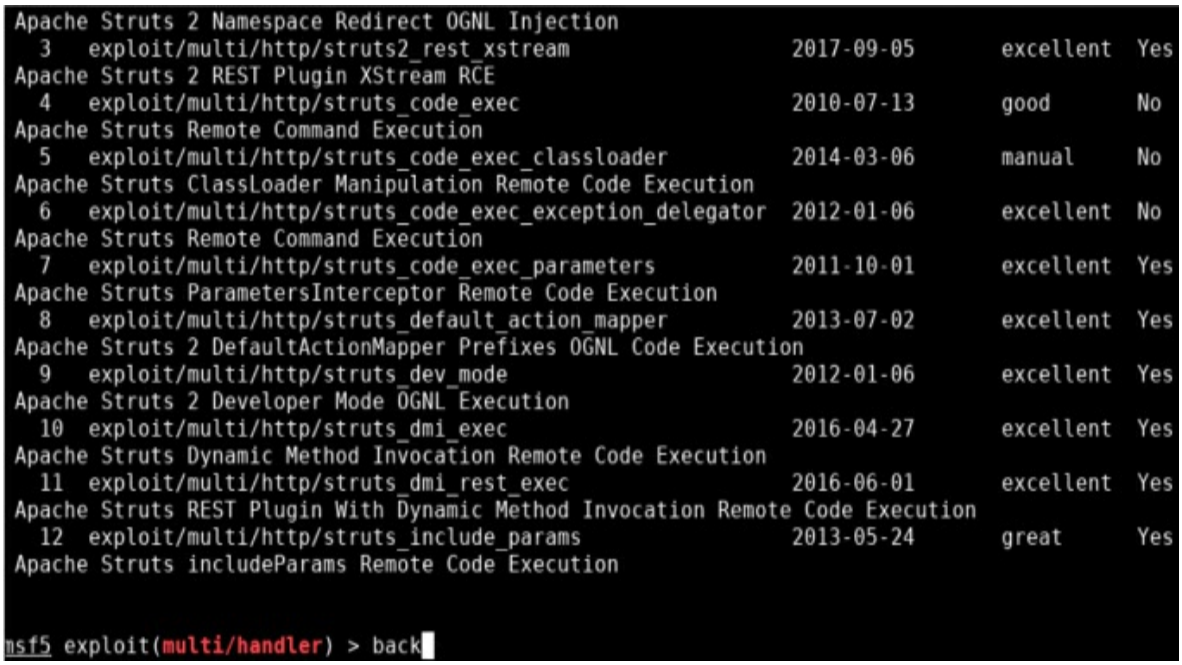

FIGURE 9-16 Metasploit Exploits Against CVE-2017-9793

> **Note**
>
> To be crystal clear, obtaining an exploitation tool does *not* make you a penetration tester. Professional penetration testers use custom scripts and many other tactics outside of penetration testing tools; however, a good baseline for evaluating a vulnerability is leveraging an exploitation tool to simplify the process.

Using an exploitation tool is not a required step for a vulnerability management practice; however, the following are a few of the big benefits that make an exploitation tool a good investment:

- ou can search for any potential vulnerability found during your vulnerability scanning process within the exploitation tool to see what exploits the penetration industry has found and provided to the open-source community.
- You have a way to test systems that are flagged as vulnerable to properly determine whether the vulnerability is a real threat or a false positive.

- You have the ability to execute an exploit (when authorized by the asset owner) to prove the possible outcome if the vulnerability is not addressed and to prove the value of the vulnerability management program.

> **Note**
>
> I had one customer tell me that although his SOC team flagged vulnerabilities in systems, the desktop team continued to ignore the vulnerabilities. To prove his case that these vulnerabilities were real, the customer cloned the impacted system, created a short video of the cloned system being exploited, and provided that video to leadership to show that he was able to gain root access. In response to this demonstration, leadership required the vulnerabilities that were previously ignored to be patched.

Asset Management and Compliance Tools

Industry best practices such as the CIS Controls (introduced in Chapter 6, "Reducing Risk and Exceeding Compliance") for securing host systems point out the importance of standardizing host system configurations. By doing so, a desktop support team can develop policies and standards for what should be installed and running and then establish a baseline for what is considered acceptable. The desktop support team can deploy controls when systems fall out of compliance with the standardized build designated in the baseline. Without standardization, host systems could have any form of operating system, updates, and security tools installed, making it extremely difficult to understand the current risk associated with host systems. Lack of standardization also makes it tough to accommodate all the different types of systems along with their unique requirements for keeping them secure. Standardization helps with vulnerability scanning because it clarifies expectations around what should be installed and evaluated. I recommend including a vulnerability scanning agent along with the standard build to simplify the vulnerability scanning process.

Asset management platforms not only help desktop teams gain visibility into

what host systems are running but also give them the capability to enforce standardization and to push remediation steps when violations in compliance are identified. Asset management tools can also include asset vulnerability assessment tools and vulnerability assessment capabilities. The results of these assessments provide a list of assets that need to be scanned along with their associated vulnerabilities once scanned. Examples of asset management tools include BMC Helix ITSM, Certero, and Snow Software Asset Management. Figure 9-17 shows an example of the Certero dashboard.

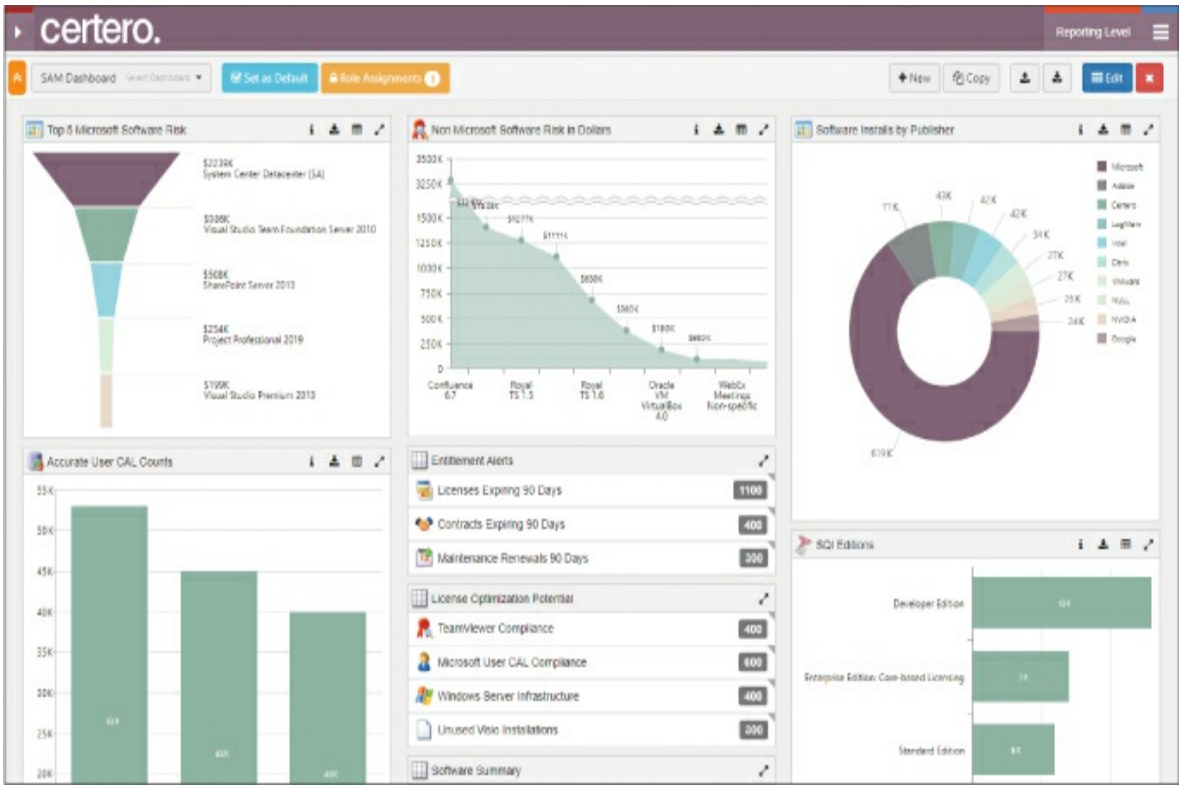

FIGURE 9-17 Certero Dashboard Example

Network Scanners and Network Access Control

Another option to collect an inventory list of assets on the network is to use a network access control (NAC) solution (introduced earlier in this chapter). A NAC solution can collect many details, including hardware addresses, IP addresses, user information, and what applications are installed, and export

those details to a vulnerability scanner. Some NAC solutions can also be configured to launch a vulnerability scan against devices that connect to the network. Including this NAC setup within your vulnerability management program guarantees that new devices are scanned prior to provisioning network access, which is a requirement that I have seen organizations include in their policies but have trouble enforcing without automation offered by NAC tools. That guarantee would apply only to NAC-enabled networks. Figure 9-18 shows an example of Cisco's NAC solution listing devices on the network.

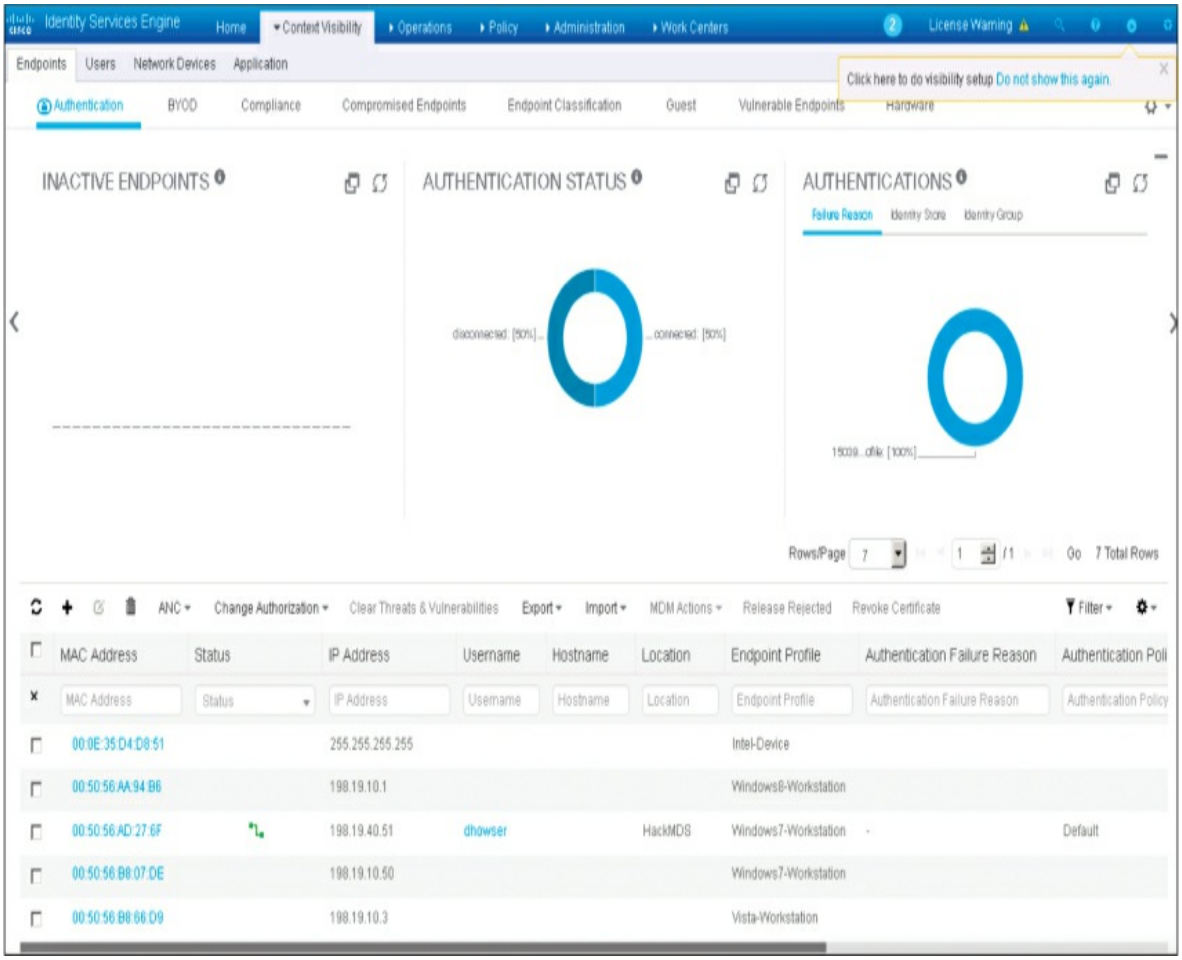

FIGURE 9-18 Network Access Control Asset List Example

Note

NAC technology and inventory management tools help augment each other. A NAC tool can function as an inventory management tool but won't provide as much detail. NAC tools are ideal for validating what is collected by an inventory management tool.

Network devices can also be identified using simple scanning technology such as Nmap, Angry IP, or Zenmap. The weaknesses in this approach include missing devices between scanning, the impact on the network while running scanners, and the tuning requirements of scanning depending on similar factors that impact a vulnerability scanner (time, accuracy, and resources). Figure 9-19 is an example of Zenmap doing a port scan.

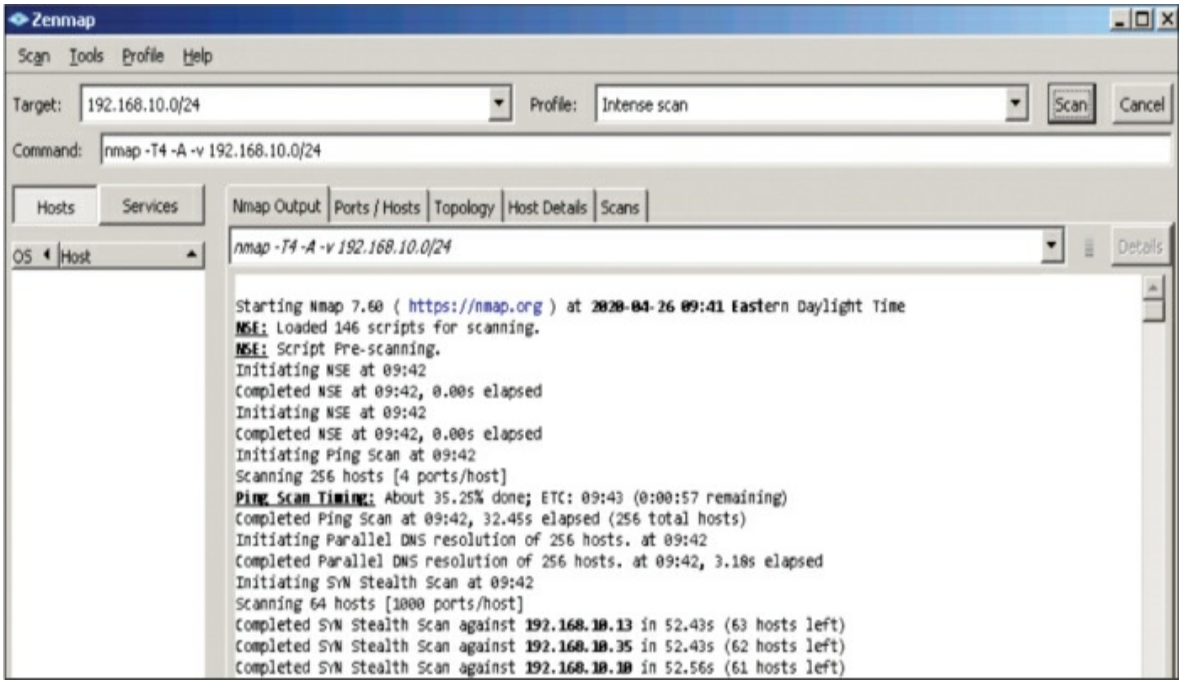

FIGURE 9-19 Zenmap

NAC can enforce security only where it is enabled on the network, opening the potential for threats to access the network if NAC is not enabled on a LAN, VPN, or wireless network. Combining NAC and network scanners allows for the network scanner to validate that the NAC technology has identified all devices that have connected to the network.

Most modern vulnerability scanners include a network scanner, reducing the need of a dedicated tool or additional scanning. Asset management solutions, NAC solutions, and dedicated scanners are additional options that can provide asset lists to a vulnerability scanner to assist with developing vulnerability scanning schedules and expectations for the type of assets that will be assessed by the vulnerability management solution. Recommended practice dictates integrating tools that have asset lists and vulnerability data to maximize investments and ensure all systems have the latest data regarding assets.

Threat Detection Tools

Some threat detection tools include vulnerability assessment capabilities as part of their toolkit to defend network weaknesses in an organization. For example, an intrusion prevention system (IPS) is designed to prevent exploitation against known vulnerabilities. The best way to tune an IPS is to match what it is configured to protect against and what is vulnerable in a network. This should make sense and explain why IPS vendors are adding vulnerability assessment capabilities as a way to learn what their solutions should protect.

Figure 9-20 shows an example of Cisco's next-generation firewall (Firepower) using passive vulnerability data collected from its application firewall capability as well as additional vulnerability data pulled from Rapid7's Nexpose vulnerability scanner to tune the Firepower IPS. All of the vulnerability data is used to identify vulnerabilities in the network so that the IPS can shut off signatures matching assets that don't exist in the network and enable signatures to defend things that are currently not being protected.

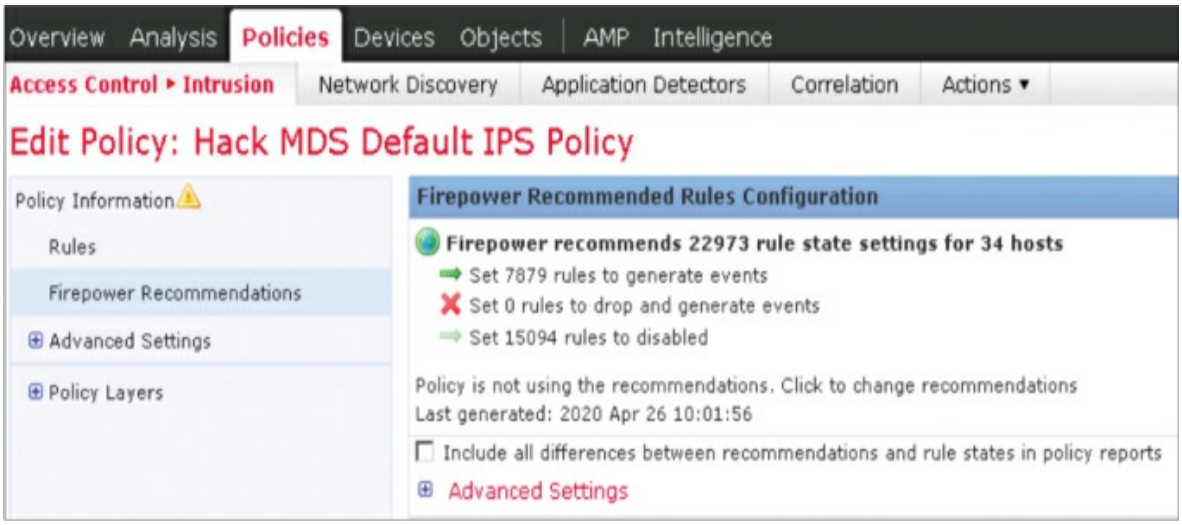

FIGURE 9-20 Cisco Firepower Tuning with Vulnerability Data

Vulnerability scanners, asset management tools, security appliances, and network access control technology are all examples of tools that can provide value to your vulnerability management practice. There are variations of these tools that are becoming popular as technology shifts to the cloud. Scanners with a focus on web applications are becoming popular features requested by customers. Also, penetration testing tools are adding breach simulation testing to speed up the penetration testing process, meaning the test can be performed without having to attack any live systems. Breach simulations are based on continuously updated cloud feeds that represent current attack campaigns, which allows organizations to better understand their risk exposure based on real-world threat data. Chapter 11, "Future of the SOC," reviews these and other cloud trends as well as provides a future look at other security concepts.

With coverage of vulnerability concepts and tools behind you, next is a look at how to deliver a vulnerability management practice.

Vulnerability Management Service

A vulnerability management practice does not boil down to just performing vulnerability scanning. As the name indicates, vulnerability management includes all aspects of managing vulnerabilities, including identification,

tracking, remediation, and recovery. This section walks you through what you need to know to develop and run a mature vulnerability management service, one of the eight foundational services all SOCs need to provide.

You can think of vulnerability management as the process that surrounds your vulnerability scanning program with a focus on IT vulnerabilities. With this in mind, the first point of focus for a vulnerability management service is scanning.

Scanning Services

Scanning can occur at different points in time and at different points in the network. One option is to scan upon connection, meaning that anytime a device connects to the network, it is immediately evaluated for vulnerabilities. The benefit of this approach is that systems with existing vulnerabilities are identified and, if a NAC solution is used, an action such as remediation can be immediately enforced. The downside of this approach is that a device can become vulnerable while on the network, which would occur after it has been evaluated by the NAC solution. One workaround to this downside is to have the NAC solution periodically reauthenticate and reassess systems; however, even this approach leaves gaps of time that a system could become vulnerable and become a liability to the organization.

A second approach to identifying vulnerabilities is to periodically scan devices in the network. Most vulnerability scanning technology offers the capability to set up a schedule for frequent scanning. Schedules can be tuned to use different types of templates to accommodate business operations, such as using less aggressive scanning techniques during business hours. One benefit of scanning systems while they are on the network is the ability to catch systems that develop new vulnerabilities after connecting to the network. Another benefit is that it provides a method to check the status of a system with a vulnerability that is being remediated both before and after the work is done. Ideally, the vulnerability should not exist after remediation steps are applied; however, there could be residual vulnerabilities as a result of remediation, which is a topic I will cover shortly.

The potential downside to periodically scanning the network and hosts

depends on how frequent your scanning program is enforced. Regardless of how often you scan, changes could occur after a scan is complete. For smaller networks, it is easy to continuously scan devices; however, larger networks will have larger breaks of time before a scan can be repeated due to the number of devices that have to be accounted for. Figure 9-21 shows a diagram representing the gap in a scheduled scanning that is a weakness of periodic scanning. This gap ranges in size. Some organizations perform annual scans, which means the gap between scans is at least 365 days.

FIGURE 9-21 Scanning Gap Vulnerability

Note

NAC technology and asset management tools can help address challenges with gaps in vulnerability scanning since they are designed to track the state of systems as they connect to and use the network.

Ideally, you don’t want a large gap between scanning periods. If you have a lot of devices to account for, one option is to perform light scanning often and deeper scanning less frequently. Other factors, such as the value of systems, could determine what level of scanning should occur. For example, maybe devices connected to the office ports are not scanned as often as servers within the datacenter that contain critical data. Adjusting scanning through tuning (covered previously) can help balance your business goals

with your risk appetite regarding the time between scanning.

The data captured through scanning is of little value until someone operationalizes it. This bring us to the people within a vulnerability management program.

Vulnerability Management Service Roles

There are certain roles that are key to the success of a vulnerability management program. Sometimes these roles are covered by the same person; however, it is more common for each role to be part of a different group within the organization. Roles include people that are responsible for the systems impacted by vulnerabilities, people that identify and validate vulnerabilities, people that run the vulnerability program, and people that maintain the network between systems. The roles and responsibilities found in vulnerability management programs at successful SOCs are as follows:

- **Security officer:** The owner of the vulnerability management program. This person does not necessarily work in the SOC. Typically the security officer is an executive sponsor that supports the SOC's mission and develops policy (which are critical for a SOC to have), but the security officer could be someone who works in the SOC and manages daily activities. Either way, the security officer is responsible for the vulnerability management program.
- **Security engineer:** Handles the vulnerabilities. Tasks include running the scans, selecting the vulnerabilities to address, working with other teams to remediate a vulnerability, and documenting the entire lifecycle of the vulnerability.
- **IT team:** Responsible for the systems that make up the network that hosts the assets being evaluated for vulnerabilities. IT will play various different roles in the security program, including estimating and reporting impact of the vulnerability management program to the network, assisting with proper functioning of vulnerability detection tools, and helping to address asset owners.
- **Asset owners:** Play a role in determining what vulnerabilities found in systems (that are not owned by security) can or can't be addressed, when

vulnerabilities can be addressed, and how they are addressed. The level of decision power an asset owner has is based on a few factors, including their role in the organization, the organization's policy on addressing vulnerabilities, the impact the resource has on the organization, the level of effort required to address the vulnerability, along with many other case-by-case situations.

Figure 9-22 shows an example of a generic vulnerability management program that includes these four roles. The scanning process starts with the security engineer launching the scan. IT monitors the stability and performance of systems under scan and monitors the scan for any problems that are detected. If a problem is detected, the issue is reported to the security engineer. When the scan completes, the results are sent to the security engineer. The security engineer sends a copy of the results to any impacted asset owner to address mitigation and sends a summary to the security officer.

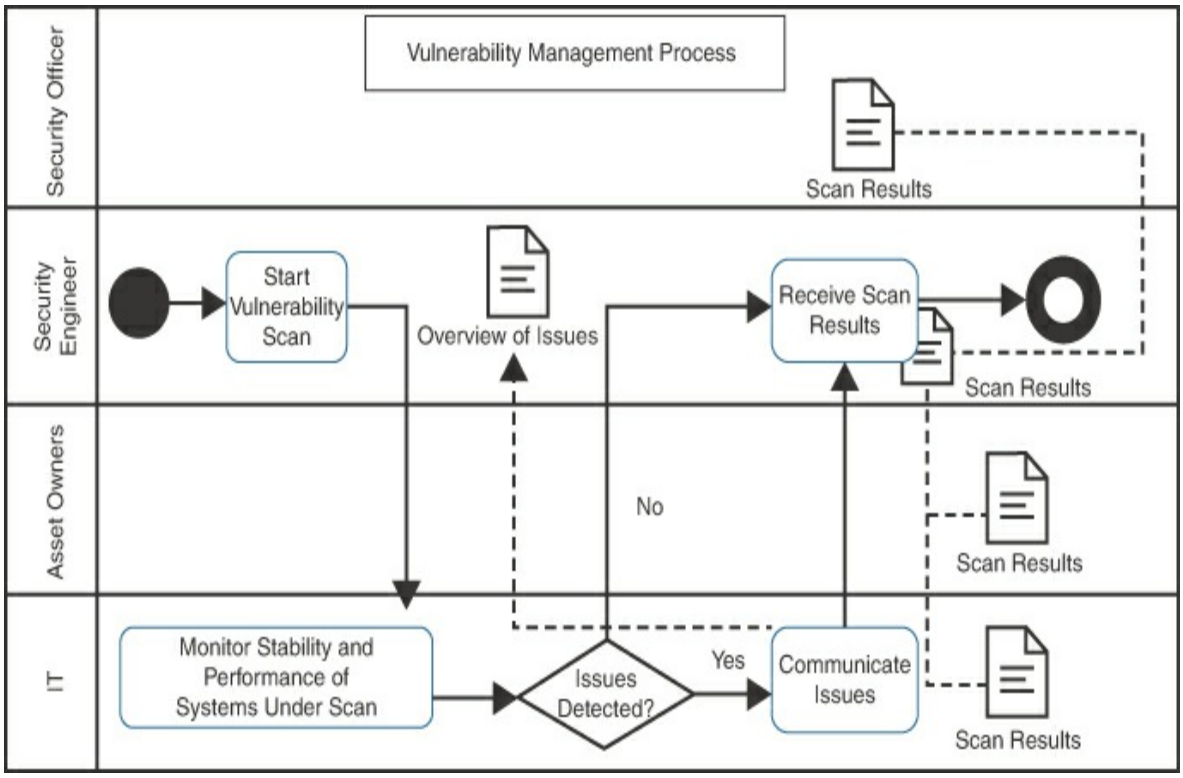

FIGURE 9-22 Vulnerability Management Program Diagram

To better understand the details involved with Figure 9-22, the next topic is a deeper look at the process of evaluating vulnerabilities.

Vulnerability Evaluation Procedures

A vulnerability management practice is made up of people, process, and technology as demonstrated at a high level earlier in this chapter and depicted in Figure 9-1. A more detailed explanation of vulnerability evaluation best practices can be found in industry guidelines such as ISO 27002 A.12.6.1. The following summarizes those recommendations.

- Define responsibilities, including vulnerability monitoring, risk assessment of vulnerability, asset patching, asset tracking, and any necessary coordination responsibilities
- Define resources to identify and raise awareness about the relevant technical vulnerabilities
- Define a timeline to respond to potentially relevant technical vulnerabilities.
- Define the organization that will recognize the associated risks and acts when a potential technological weakness has been identified; these acts can include patching compromised systems or enforcing other controls
- Deal with vulnerabilities through defined procedures. High-risk systems should be dealt with first.
- Measure the risk of the installation of a patch or other remediation
- Define reference sources to measure against for vulnerabilities
- Make an asset inventory
- Make records for post-event analysis

To ensure its efficiency and effectiveness, the technical vulnerability management process should be monitored and assessed regularly.

ISO 27002 A.12.6.1 is a solid guideline but includes steps that you may or may not need, depending on the focus of the vulnerability management service you are providing. Regardless of the type of vulnerability

management service you plan to provide, you need to create policies and procedures. Policies ensure the program's goals are met while procedures explain what needs to be done to meet all expectations set by the policies. ISO 27002 A.12.6.1 includes many actions that need to be included in vulnerability management policies and procedures, including defining who is responsible for what, how vulnerabilities and the risk associated with fixing vulnerabilities is handled, how to identify assets, in what order of priority vulnerabilities should be addressed, and a firm recommendation on documenting everything during the vulnerability management lifecycle.

Figure 9-1 was a general look at the steps you can take to evaluate vulnerabilities, but next I'll provide a deeper look at my recommended procedures for delivering a mature vulnerability management service. For example, I pointed out you need to collect assets but how should you do so? What tools and procedures should you use? The next part provides such details.

I have broken my recommendations into four steps that incorporate guideline recommendations such as ISO 27002 A.12.6.1:

Step 1. Asset collection

Step 2. Planning and asset evaluation

Step 3. Launch scanning

Step 4. Corrective actions or accept risk

A short explanation of my recommended procedures for delivering a vulnerability management service is to first collect information about what you must evaluate. Next, use that data to develop a plan regarding which systems, networks, or applications to address first and how to address those targets. Once you have a plan, start your scanning by working through your list of targets to evaluate. Based on your results, either apply corrective actions or accept the risk of not performing a corrective action.

Let's look closer at what is involved with each of these four steps.

Step 1: Asset Collection

Looking back at Figure 9-1 representing best practices for vulnerability

management as well as my recommended procedures for delivering vulnerability management services, the focus for the first step is collecting information about assets. You need to know what needs to be evaluated before you can start performing any form of evaluation services. Having a complete and reliable asset list will simplify the entire vulnerability management program. *The most common failure of a vulnerability management program is a gap in what is within scope to be evaluated versus what exists on the network.*

Popular tools covered earlier in this chapter that can be used to develop an asset list include the following:

- **Network scanners:** Run scans against network subnets and collect a list of identified devices
- **Network access control:** Collect a list of devices as they connect to the network
- **Asset management tools:** Pull asset lists from tools that have collected details about assets

Many organizations use segmentation as a method to keep track of assets and deliver different scanning processes based on the value of the assets within a subnet of the organization. For example, the critical servers in a datacenter can be placed on a separate network segment, which would be scanned more frequently than the general network. Engineers are often able to see an IP address and know what the system is and where it is located based on how the network is segmented and how systems are provided with IP addresses.

Throughout this book, you have seen examples of best practices that include segmentation strategies based on device type. Devices such as IoT devices are associated with high risk and should be placed on a different network than systems containing sensitive data. Scanning the “IoT network segment” makes expectations for asset identification and collection simpler and allows for additional detection because rules can be developed to respond to devices that are out of scope for the specific network. A common example found in organizations is to dedicate a network segment to IP phones. This is done to ensure that resources in the segment are dedicated to IP phones, simplifying management of phone traffic, ensuring proper bandwidth is available in the network segment, and identifying when a device that is not a phone has connected to the network segment. Investing in a robust and well-planned

segmentation strategy provides many benefits, including asset management simplification and improved incident response to unauthorized devices.

Asset Collection Best Practices

To summarize best practices for asset identification, consider the following:

- Use a NAC solution to monitor who and what connects to the network. This allows the SOC to collect information about new devices as they connect to the network.
- If you use a NAC solution, ensure that the NAC technology is enabled across all parts of the network, including the LAN, VPN, and wireless networks.
- If you use a NAC solution, configure it to be able to account for unauthorized hubs, switches, and routers, meaning when such devices are detected, an action such as blocking these devices or moving them to another network is taken. Not doing so will allow devices to connect and not be identified by the NAC technology.
- If you use a NAC solution, configure it to reevaluate devices that have been connected to the network for a certain length of time. The reevaluation process can occur upon a forced network reconnection of the device or by using a network scanner that doesn't require a network reconnection.
- Segment devices by device type and value to the organization. Examples include providing separate networks for mobile devices, IP phones, printers, servers, and IoT devices.
- If you use a NAC solution, configure it to monitor for devices that don't meet an expected device type attempting to connect to a network segment.
- Use network scanning periodically to identify devices that have already connected to the network, even if you are using a NAC technology.
- Ensure that network scanning accommodates all parts of the network and that it is scheduled to occur on a regular basis so that there are no long periods of time between scanning a network segment.
- Use a centralized master list to consolidate information about all collected

assets.

- Ensure that any deduplication steps performed against the master list of devices do not remove assets that look similar but are different.
- Periodically test asset collection tools and manually evaluate the results to ensure that all asset collection is performed correctly.
- Back up all results. Ideally, the backup should be sent to a different location than the primary database, such as a cloud storage option or offsite location that wouldn't experience similar impact as the primary database if a natural disaster occurred.

Asset Collection Challenges

The preceding suggested best practices are designed to ensure that you are capturing any new device that connects to the network as well as evaluating authorized devices that already have been approved network access. The following are some of the many reasons that a device might exist on the network yet not be part of the NAC asset list:

- The NAC technology is not enabled on the port used to provide access.
- Additional devices such as hubs or switches conceal the device.
- An authorized device was modified while connected.
- The device was connected prior to the NAC technology being enabled, or the device connected while the NAC technology was not functioning properly.

You also learned previously that most NAC technologies do not evaluate devices once they are connected to the network. This concept exposes a potential blind spot for collecting devices, which a network scanner will address as long as scanning is performed often and across the entire network.

Asset Collection Value

As covered in the next section, step 2 of the vulnerability management service best practices includes planning for how to evaluate identified assets for vulnerabilities. Part of that process is to develop a list that shows the order

of how assets you identified during step 1 are evaluated. That order is based on the value of the asset to the organization. Knowing the value of an asset allows critical systems to be prioritized regarding evaluation and the approach to mitigation. Proper planning cannot occur if asset value is not identified during the asset collection step.

The value of most assets can be generalized within a range, allowing the SOC to quickly determine where an asset should fit according to a risk assessment. A popular formula for calculating the value of an event is using the single loss expectancy (SLE). SLE is measured as the asset value (AV, or $) times the exposure factor (EF). EF is measured as the negative effect or impact that would be realized if the asset was lost. So essentially, you multiple the value of an asset by its exposure to get your expected value, shown as SLE = $ × EF. The top-valued systems will have the highest importance regarding remediation of identified vulnerabilities. Chapter 3 covered risk management as a SOC service and the steps to follow to identify risk and to estimate the value of an asset.

Asset Collection Master List

One final point to highlight is how you manage the master list of assets. Using different tools will produce a separate list of assets, which will have overlapping data. To address this, using an asset management platform that can accommodate for data deduplication, protect your asset lists, track when assets were last evaluated, and manage details such as when the device was identified and who owns it not only will be critical to having an up-to-date and accurate list of assets but also will assist with other steps in the vulnerability management process. Steps include understanding associating risk and value of an asset, meeting compliance requirements, and monitoring the status of the mitigation process, if applicable. Redundancy and data backup of the master list are equally important to prevent data being lost, which would prevent the vulnerability management program from performing properly.

Step 2: Planning and Asset Evaluation

Having a reliable list of assets is important, but knowing how to use that list

is equally important. For smaller networks, a SOC can quickly set up periodic network and vulnerability scanning to obtain acceptable results. A small business such as a gym or a home network typically has fewer than 20 devices on the network, which means very little planning needs to be done regarding what should be scanned. For larger networks, there are many different approaches that can be taken, but regardless of which approach chosen, you must develop a repeatable process to ensure that all devices are accounted for and that potential speedbumps to performing scanning are identified.

You need to consider the following factors that will impact the development of your vulnerability scanning program:

- Number of network segments
- Available bandwidth
- Risk associated with scanning
- Network access (remote, cloud, required VPN/routing, etc.)
- Hours of operation
- Peak workload hours
- Required compliance for scanning
- Asset value to organization
- Asset evaluation method

Accessing Assets

When you are evaluating the previous list of factors, the first datapoint to focus on is availability to assets. Smaller networks consist of a few network subnets, making your vulnerability scanning design simple. Larger networks, however, can have thousands of subnets spread across the globe, which poses multiple challenges that must be addressed for a complete scan of the network to be accomplished. Therefore, having an accurate asset list is essential. An asset list can be used to map out asset locations. By knowing the asset location, the SOC can work with the IT team that supports the associated network regarding how access to devices can be obtained so

devices can be properly scanned.

As covered previously, three possible options for scanning devices are host scanning, network scanning, and passive scanning. Each approach has deployment considerations as well as limitations based on the vendor being used. Starting with host scanning, this approach either uses software installed on the endpoints being evaluated or provides access rights (which should be read-only) to the vulnerability scanner so that it can log in to the devices and collect data. For either option, the vulnerability management system must be able to collect results from the endpoints. If you choose the agent-based option, you also need to consider how the agents are installed on endpoints. If administration rights are required to install the software, those rights must be granted. Also, any asset owner pushback must be considered.

Note

Many asset management tools include the ability to evaluate assets for vulnerabilities. This feature is typically a form of passive vulnerability scanning since collected data is compared against a database of known vulnerabilities (rather than specific assets being actively evaluated). An exception to this would be if the asset management tool has a built-in active host vulnerability scanner.

Network scanning must also have access to systems in order to perform its analysis. This means regardless of your scanning approach (host or network), you may encounter potential hurdles in the network that you must overcome. Some common network hurdles to plan for are as follows:

- **Application-layer firewall:** An application-layer firewall will most likely inspect every packet before allowing it to be routed, unless special exceptions are configured. A typical vulnerability scan can generate thousands of connections in a short period, which can overload the application-layer firewall state table to state tracking mechanism.

  Workaround: Place scanners behind the firewall and use host scanning or allow list for scanning traffic from an authorized scanner if possible.
- **Intrusion detection/prevention system (IDS/IPS):** Performing

vulnerability scanning through an IDS/IPS can overload the device or generate an excessive number of unwanted alerts. An IDS/IPS can drop packets, resulting in lost vulnerability scanning results, or mask the true characteristics of a target, such as showing blocked ports as open, which is part of deception defense techniques.

Workaround: Deploy scanners behind the IDS/IPS, use host scanning, or disable the security tools during a scanning window.

- **Network address transition:** NAT can cause a vulnerability scan to slow down because the security tool enforcing NAT can become a chokepoint for the traffic going between the scanner and target, providing a limited number of packets per second. NAT can also obfuscate network scan data, which means vulnerability data will not align to the scanned hosts.

  Workaround: Place scanners within the NAT environment or use host scanning.
- **Virtual private network (VPN):** Scanning across a VPN can slow down the process even if a high-bandwidth VPN connection is used. Delays are caused by the number of connection attempts made during a vulnerability scan, during which a VPN can become a bottleneck as it performs the encryption and decryption process. Remember that VPN technology is designed to protect traffic through encryption, which means this same effort will be applied to every packet sent and received by the vulnerability scanner.

  Workaround: Place scanners within a network rather than scanning over VPN, use cloud scanning, or use host scanning.
- **Segmentation:** Many networks use firewalls, portals, containers, VLANs, and access control lists (ACLs) to segment networks. Scanning across a subnet can violate corporate security policies and must be controlled and monitored. Security between network divisions can be severely restricted, resulting in slower scans. Scanning across segments can also restrict ports and protocols, limiting the quality of results from the vulnerability scan.

  Workaround: Place scanners within the network segments or use host scanning.
- **Distance:** Scanning across thousands of miles can slow down scanning. Scanning across different countries and service providers can introduce different legal challenges as well as policies that automatically block

certain ports and protocols.

Workaround: Place scanners closer to endpoints, use cloud scanning, or use host scanning.

- **Perimeter networks (DMZ):** Systems within the DMZ are considered out in the open and top targets for attack.

  Workaround: Use additional scanners, including network and host scanning, to accommodate additional focus on these systems.

Overcoming Asset Access Hurdles

Vulnerability scanning vendors offer different options to overcome network challenges. One option is to use a distributed deployment, such as shown in Figure 9-23, which depicts a blend of network and host scanning options placed at locations behind security tools and other systems that would otherwise cause scanning interference. For this example, the vulnerability scanning vendor offers appliances—one dedicated to the DMZ and one dedicated to the inside network. For the datacenter, the example uses network scanning that occurs over SSL/TLS from a scanner located on another network. This example assumes an accommodation has been made to allow a network scan to perform at an optimal level over SSL/TLS. An alternative could be placing a network scanner within the datacenter network segment or using host scanning. One other scanner type shown in this example is a cloud scanner used to perform network scanning of remote devices. Not all vendors support all these features; however, most enterprise vendors can support a similar architecture.

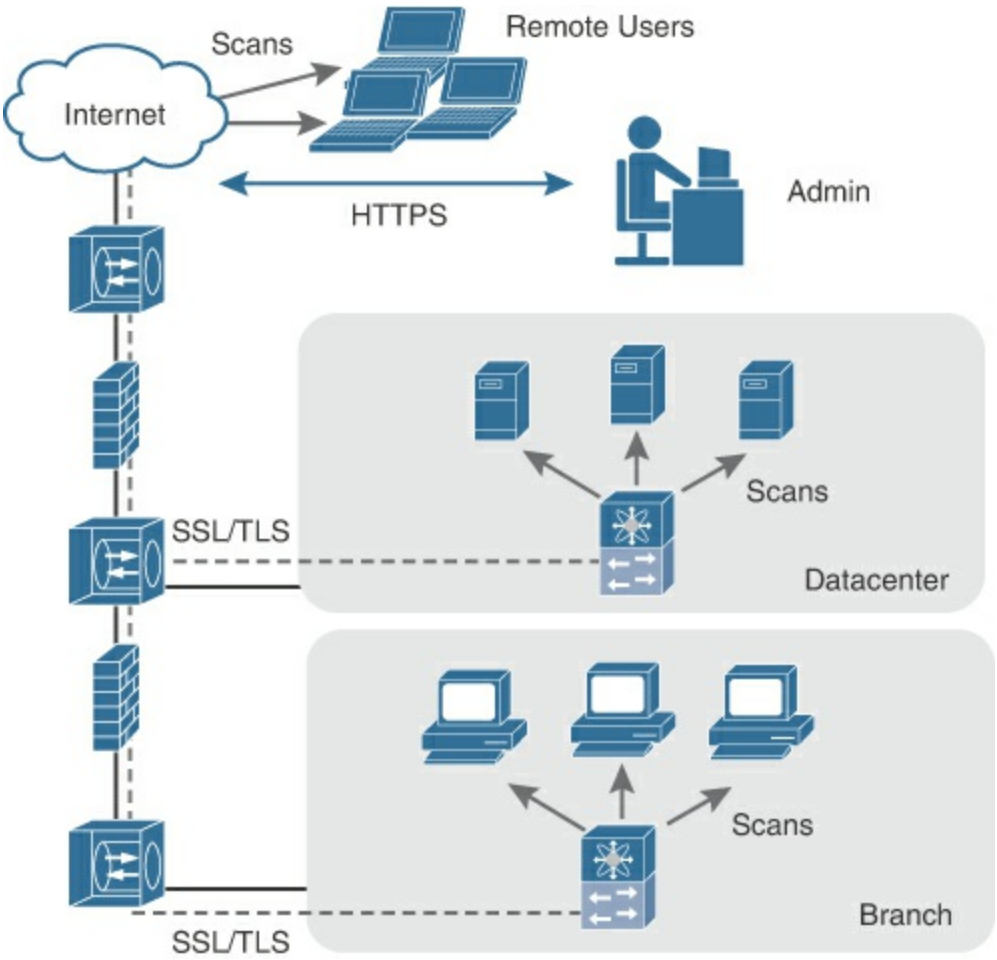

FIGURE 9-23 Example of Vulnerability Deployment

Scanning Templates

The final planning steps involve the type of scanning being used. As previously covered, scanning tuning concepts include leveraging different scanning templates based on the type of risk and devices being evaluated. Less aggressive templates can be used during business hours or against systems that would be at risk of experiencing negative impact if an aggressive scan were performed. The same concept applies for host and network scanning. In both situations, either network or host resources can be impacted and must be considered as you choose the scanning template to use. You want to avoid consuming too many host resources from critical systems, such as datacenter servers, during normal business hours. If both network and host scanning will negatively impact the availability and performance of a

system, limited scanning or scanning during nonbusiness hours might be the only possible options. Remember to consider the three key vulnerability management tuning points: accuracy, time, and resources.

Prioritizing Assets

Part of your planning is deciding the order in which to evaluate the systems, networks, and applications. As previously stated, this list is created based on the value of each asset to the organization. However, an additional underlining factor to consider is what type of risk the vulnerability poses to other assets. Suppose you have a lower priority system with a critical vulnerability and a critical system with a mild vulnerability. If you base your priority on asset value, the critical system with the mild vulnerability would be your starting point; however, what if the lower priority system can be easily exploited as a beachhead to other systems, including critical systems? This would dramatically raise the importance of responding to the less critical system, to avoid that system being used as a starting point for a much larger attack.

Figure 9-24 represents the concept of rating a system based on each asset's value to the organization and its potential to expose the organization to a critical breach.

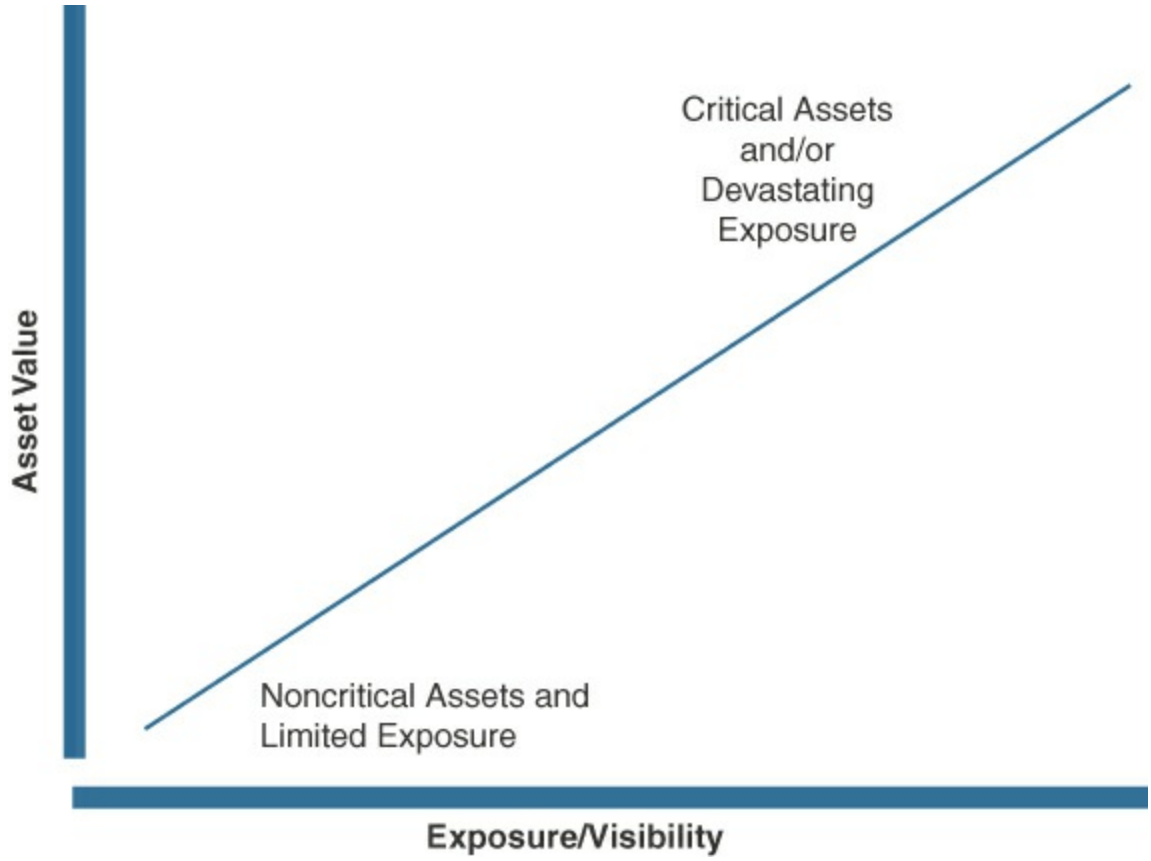

FIGURE 9-24 Evaluating the Risk of an Asset to the Organization

Planning and Asset Evaluation Summarized

To summarize the planning and asset evaluation phase, you need to do the following:

- Confirm how many scan engines you need, where to place the scanning engines, and the type of scanning to perform.
- Consider potential obstacles on the network, the expected impact on a host when scanned, and available resources during and after business hours.
- Speak with relevant asset owners if there are any internal politics or compliance concerns that need to be addressed before an asset can be evaluated.
- Develop a list of targets to evaluate, including the order of evaluation based on their value and potential exposure to breach.

- Perform final housekeeping items, which include obtaining the latest vulnerability threat intelligence for your tools as well as any deployment or updates needed to the vulnerability infrastructure.
- Consider how to aggregate all scanning results. An asset management tool is ideal for accomplishing this purpose when multiple types of scanning are being performed.

Once this is all accomplished, it is time to launch your vulnerability scanning.

Step 3: Launch Scanning

Once you properly plan your vulnerability management infrastructure and scanning strategy, it is time to start evaluating assets for vulnerabilities. Best practice is to automate the scanning process using a vulnerability management tool or orchestration tool that can manage the scanning schedule as well as which type of scanning templates are being used. You learned earlier how different scanning templates can be used and tuned to business requirements. An example is using a less aggressive scanning template during business hours and using a more aggressive template when the network is not heavily utilized. Using a centralized system allows estimation for completion because all tasks can be accounted for along with any interruptions.

When launching vulnerability scanning, you want to schedule scanning against the most critical systems or systems that have a high exposure rate first and work your way down the priority list. Using your vulnerability management tool, you can build an automated scanning schedule based on multiple tasks that make up your current network-wide evaluation.

The following are factors that would increase the priority for scanning as you develop the order of what is scanned:

- The criticality of the asset to the organization in terms of confidentiality, integrity, and availability.
- Whether the asset is in scope for a required standard such as PCI DSS.
- The number of other systems that are dependent upon that system, such as Active Directory, which is used to manage multiple other systems.
- Whether if by exploiting the asset, other high value assets become visible

and put at risk for attack.

- Whether certain compensating controls like encryption are in place and need to be accounted for.
- Whether there is a need to log in to the system and, if so, whether you are authorized to do so.
- How difficult is it to deal with potential disruptions that are associated with scanning the targeted system(s) or network(s). For example, scanning a gateway router could pose the risk of taking it out, leading to a network outage.

Scanning Best Practice

When scanning, you should first target automation of scanning tasks that disrupt as few people as possible. This means scanning systems that are less likely to be impacted by being scanned, systems that don't require special permissions from the asset owners, and systems that are easy to access. Figure 9-25 shows an example of Rapid7's Nexpose vulnerability scanner Automated Actions configuration screen. For this example, I can choose the scanner to act when a new asset is identified. A common action would be to launch a specific scanning template, such as a less aggressive one during business hours when scanning a network subnet containing user desktops. Also notice there is a calendar allowing scheduling of actions to occur only during specific times. One possible configuration for scanning during business hours is to scan only new assets that haven't been previously identified and scanned. Automated scanning actions allow for the accommodation of these types of scenarios.

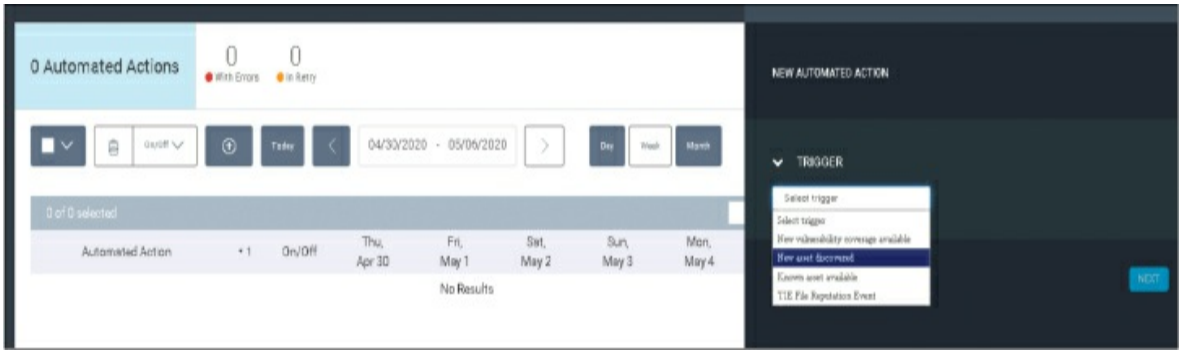

FIGURE 9-25 Rapid7 Nexpose Automated Actions Configuration Example

Another scanning best practice is to group scans together based on required logins, if user login is needed. Network scanning that does not require host scanning should be separated from host scanning that requires a login process using a read-only account. If different read-only accounts are needed, it is recommended to separate those scans to simplify troubleshooting situations in which a login fails. You should also include asset owner contact information with asset lists and pull such information within any ticket created when a scan fails to access the system. This will allow the SOC to track down the asset owner and troubleshoot why the scanner was unable to log into the asset. The most common reasons are either the asset wasn't available or the password is incorrect.

Scanning will likely identify systems that are not part of your asset list. Most enterprise vulnerability scanners include the ability to identify and discover services on new systems based on scanning a network subnet for anything connected. The discovery phase uses a high-speed ping system to identify systems; however, the discovery of a new system will require more time as the scanner attempts to connect to various ports. Some attributes will require configuration, which the scanner can be tuned to accommodate. The goal is to accurately identify as much as possible about the new asset so a proper vulnerability analysis can be performed. Figure 9-26 shows an example of Rapid7's Nexpose asset dashboard. Notice systems are broken down by discovery, operating system, and type of potential risk. Using a dashboard such as this allows a security engineer to quickly identify any device that is not part of the original asset list.

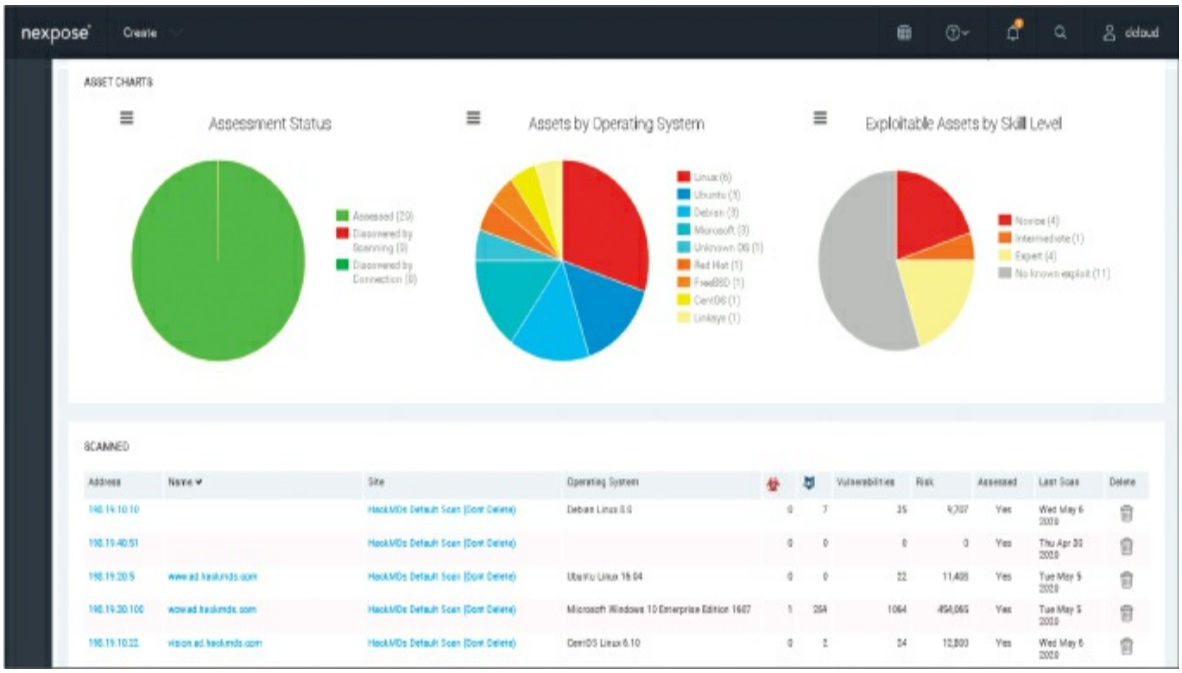

FIGURE 9-26 Nexpose Asset Dashboard

Step 4: Choose Corrective Actions or Accept Risk

Once scanning results show up in your vulnerability management solution, you will have a list of results as well as tickets for scans that did not complete or execute properly. Tickets regarding scans that did not complete must be addressed. Failure to complete could be caused by any of several factors, including unpredicted network interference, not having the correct host login credentials, or changes in the environment. The remaining tickets will be related to identified vulnerabilities.

A dedicated ticketing system is critical for ensuring that all scanning is launched and completed properly as well as for tracking how any found vulnerability is mitigated. Certain systems have predefined timeframes for remediation based on the risk of the vulnerability to the organization. A ticketing system can be set up to prioritize tickets associated with high risk systems. Best practice is to use a ticketing system that does not simply create a ticket per vulnerability instance, or you will quickly be overloaded with too many tickets. Ticketing systems need to create tickets based on vulnerability grouping, which means identifying the same vulnerability within multiple systems and approximating the patch or mitigation for all associated system

types. This will allow a more realistic understanding of how many systems are impacted by a vulnerability and the required steps for system-wide mitigation.

As you view tickets, you will need to decide which action to take. Should you attempt to mitigate the risk or does the mitigation process require too much effort to deem it worth doing? This decision merits much further discussion, which leads us to the next section, which is dedicated to responding to vulnerabilities.

Vulnerability Response

As covered in the previous section, the fourth step for delivering a vulnerability management practice is focused on the SOC's response to an identified vulnerability. All vulnerabilities should not and will not be remediated. If remediation outweighs the risk of leaving the asset vulnerable, then no action should be taken. If the risk of the vulnerability is extremely high, a different approach to remediation needs to be considered that does not put the organization at risk. I will cover these concepts in this section, but first I will cover how to confirm the true risk associated with an identified vulnerability.

Vulnerability Accuracy

Before any decision can be made about how to address a vulnerability, the SOC needs to properly qualify the vulnerability. The good news is that most enterprise vulnerability detection tools will do this work for you. Keep in mind that the results are a best guess unless you have performed a penetration test against the vulnerability. For example, a vulnerability scanner will show if a system has a vulnerability, but that identification is based on the version of software the scanner "believes" it is seeing as well as the ports and protocols the scanner "believes" it identified. If any one of these datapoints is not collected correctly, the result will be different. Simply put, the quality of the results from a vulnerability scanner is based on the accuracy of data being processed compared against the quality of the intelligence of current

vulnerabilities. If either the accuracy of the data or the quality of the vulnerability database is not good, the identified vulnerabilities will not be accurate.

There are many factors that could reduce the accuracy of a vulnerability scan. Security tools are designed to hide or display the wrong information when being scanned by external tools, since their purpose is to prevent exploitation of what they are protecting. Also, the general use of ports can be changed to not follow industry standards for associated services. Email services are expected to be seen unencrypted from port 25 or 110, but that does not mean email can't run over other ports. Some protocols may not be used within networks, reducing the available content a vulnerability scanner can collect. These types of changes and use of tools will result in inaccurate matching to vulnerabilities, equating to false positives. Keep these concepts in mind as you evaluate the results of a vulnerability scan.

Note

Don't assume what was reported by a vulnerability scanner is the true risk!

When evaluating vulnerabilities, answer the following questions before you plan any action:

- Is this vulnerability a true positive or a false positive?

  How confident are you in this decision? Remember not to assume the scanner is accurate!
- How accessible is the vulnerability?

  More importantly, could someone directly exploit this vulnerability from the Internet?
- How difficult is it to exploit this vulnerability?

  More importantly, can a scripted exploit be run against it and, if so, is the exploit publicly known/published?
- What would be the impact to the business if this vulnerability were exploited?

- Are there any other security controls in place that reduce the likelihood and/or impact of this vulnerability being exploited?

  Examples include firewalls, IPS, or antivirus/antimalware defending the impacted system.
- How old is the vulnerability/how long has it been on the network?
- What is the impact to the business/system if the vulnerability on the device is patched?
- Does this vulnerability violate any mandatory compliance or regulation?

  Could this make you liable for fines or cause your cyber insurance to be out of compliance?

Many of these questions are answered in the CVSS. With the CVSS shorthand, the example AV:N/AC:L/Au:N/C:P/I:P/A:C (introduced earlier in this chapter) indicates that the access vector is network (AV:N), the access complexity is low (AC:L), and no authentication is required (Au:N), all of which is very important information. Make sure to include these results in your ticketing system so the vulnerability's risk score has relevance to its potential impact. In particular, highlight when AV:N (access vector is network), AC:L (access complexity is low), and Au:N (no authentication required).

Understanding if an exploit has been published against a vulnerability might require additional research. Using an exploitation tool such as Metasploit as your search engine for published exploits is extremely helpful. Penetration testing tools will collect weaponized exploits against the latest vulnerabilities as long as they are kept up to date. Also, searching the Internet for key terms associated with a vulnerability can return possible exploits. Key terms can include the CVE number representing the vulnerability or specifics of the vulnerability. For the Apache Struts vulnerability, searching the Internet for "Apache Struts Vulnerability" will find many documents that include details on published exploits.

Responding to Vulnerabilities

Once you have qualified a vulnerability, you need to decide your response.

According to ISO/IEC 27005:2018, you have four options when responding to the risk associated with a vulnerability:

- **Risk retention:** This option means you accept that there is a risk but take no action. This decision typically is based on the determination that the risk is at an acceptable level in terms of your organization's risk tolerance. Imagine the risk is that a demo lab could be exploited. If you don't care if that happens, you can accept the risk and leave it.
- **Risk avoidance:** Avoiding a risk is similar to risk retention in that no action is taken. The difference regarding avoiding risk is a risk is found that is unacceptable to your organization's tolerance level. However, actions to remediate the risk are not suitable due to factors such as cost or residual risk. This means you do not accept the risk, which is different from risk retention. In this situation, you find it is not worth reducing the risk but document it as a real risk.
- **Risk modification:** Modifying a risk means applying safeguards and patches, altering controls, or making other efforts to address the risk until the residual risk is acceptable.
- **Risk transfer/sharing:** Transferring the risk means pushing the risk to another party. In most cases, this occurs with insurance. Cyber insurance is an insurance product used to protect businesses or users from risks related to information technology. Assets could also have their own insurance policy that could be used if an event occurs.

Figure 9-27 depicts the risk response process. First, you get your risk assessment results, which are the general warnings identified by the vulnerability scanner. Next is risk decision point 1, which is when a security engineer reviews the ticket and qualifies the vulnerability by answering the questions covered in the previous section, "Vulnerability Accuracy." After the engineer qualifies the vulnerability, you start the second risk decision process, risk treatment, which is deciding which action to take. Treatment options are reviewed to understand what choices can be taken, including any patches, security tools, or other modifications. Once remediation options are identified, one of the four action options in the preceding list are taken, which will result in some form of residual risk. Finally, the vulnerability is reevaluated to see if the treatment reduced the risk to satisfactory levels or if residual risk introduced new factors to consider before closing the ticket

shown as residual risks. Once a satisfactory treatment is identified, it is applied.

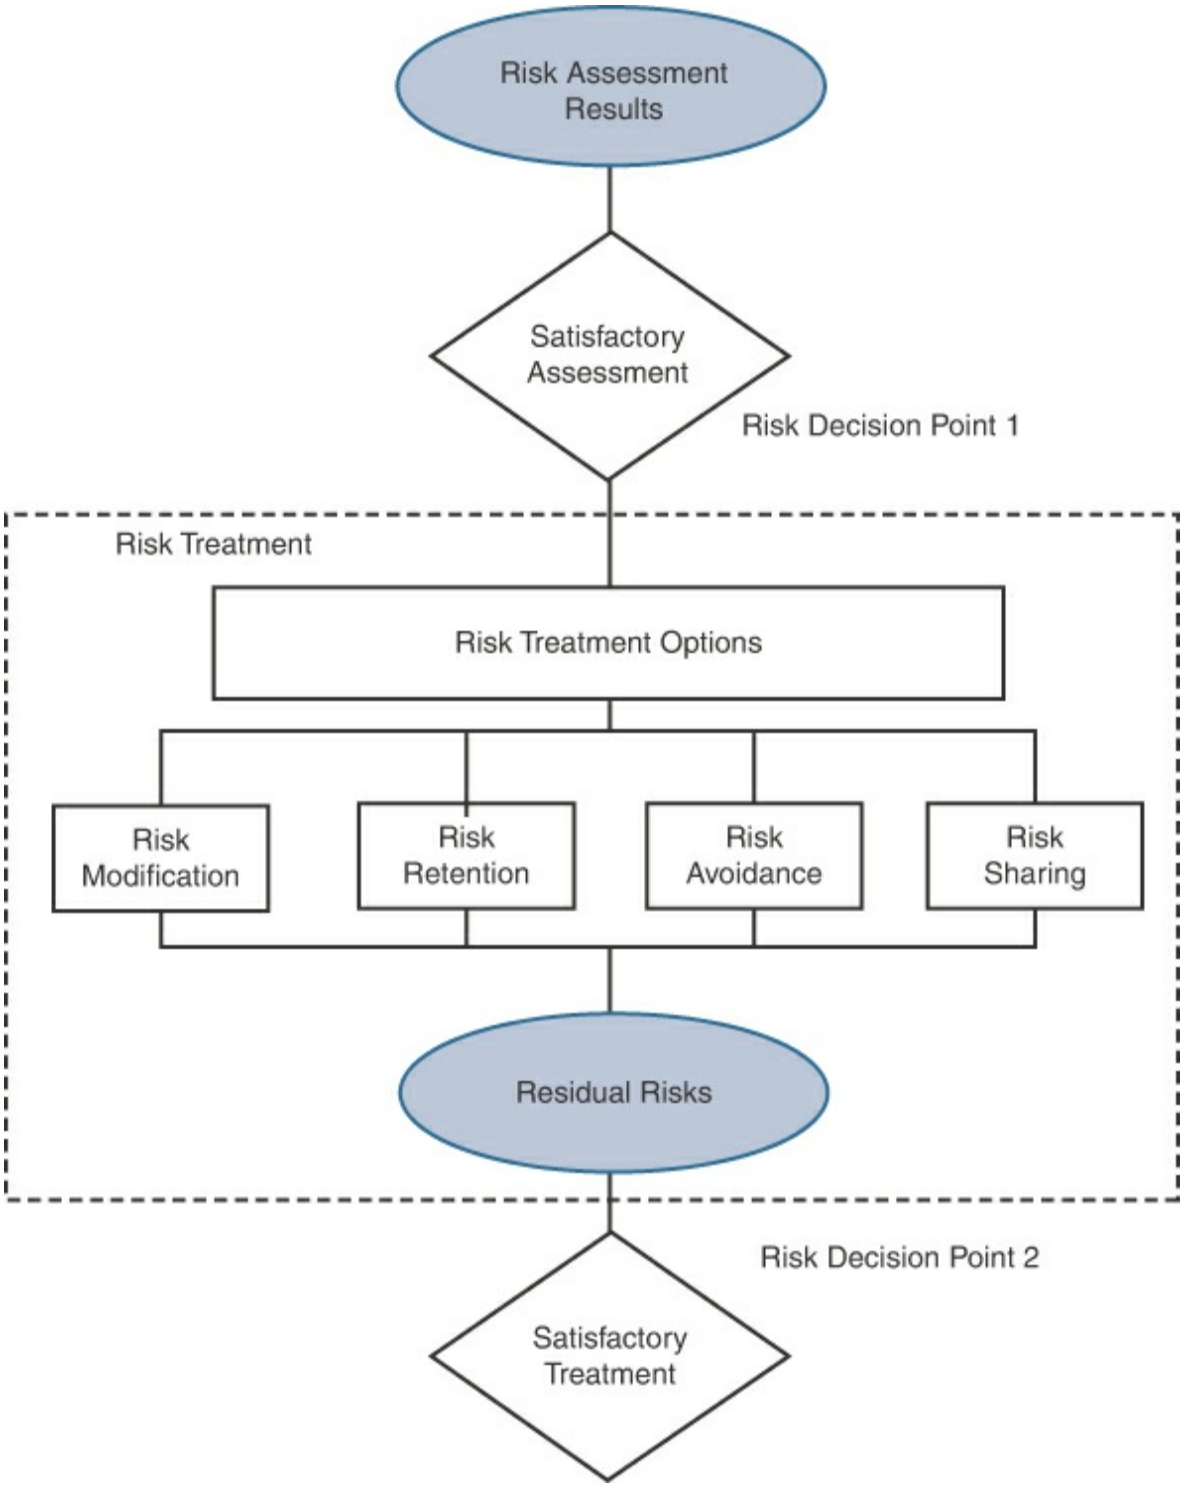

FIGURE 9-27 Responding to Risk Flowchart

There will be times that accepting the risk is the right decision. Doing so means there won't be any new residual risks because no changes are applied. Accepting risk means the existing risk is within tolerance, allowing for the ticket to be closed. For example, imagine if the potential residual risk of fixing a vulnerability within a system could take down that system based on what is required to fix a system. If nothing is done, the system will continue to introduce a current risk of a small amount of noncritical data being available if an attacker abused the vulnerability that system holds. The risk of losing noncritical data versus losing the entire system will result in a decision to avoid the fix. Another example could be a required cost that outweighs the loss. Imagine if the only fix to a vulnerability is to acquire a tool that costs more than the impacted system. At that point, it would not be cost effective, assuming there are other, more pressing needs that could use the available budget.

Note

There is logical difference between accepting risk and avoiding risk. Accepting risk means you have identified a vulnerability and determined that if it is exploited, you are not worried about the impact because the risk falls within acceptable levels. Avoiding the risk means you *do not* accept the impact caused by the risk of a vulnerability; however, you determine that the mitigation outweighs addressing the risk. It is absolutely critical to identify when either of these choices is selected and document why the decision was made as well as the potential impact post assessment. This will allow a better response during future assessments when the vulnerability is brought back into evaluation.

Ideally, you will attempt to fix an identified vulnerability so that it can no longer be exploited. An example of preventing access to a vulnerability is patching a vulnerable application so the vulnerability no longer exits. However, another option to reduce the risk is to prevent access to or use of the vulnerable application. An example of preventing access to a vulnerable

application is placing a security tool in front of the vulnerable system that prevents exploitation of the vulnerable system. The vulnerability can still be exploited, but the security tool prevents exploitation behavior. Both options prevent exploitation of the vulnerability, which results in reduction of risk.

A final option for dealing with risk is the transfer or sharing of risk. This can be tricky and commonly is accomplished by acquiring cyber insurance.

Cyber Insurance

According to Nationwide, “Cyber insurance generally covers your business’ liability for a data breach involving sensitive customer information, such as Social Security numbers, credit card numbers, account numbers, driver’s license numbers and health records.” Do your standard business insurance policies cover this? In most cases, the answer is no. Cyber insurance specializes in loss due to cyber-related events. Besides legal fees and expenses, cyber insurance typically helps with

- Notifying customers about a data breach
- Restoring personal identities of affected customers
- Recovering compromised data
- Repairing damaged computer systems

The specifics of what losses are covered will vary depending on the insurance provider; however, there are some general categories of cyber insurance:

- **Data protection:** This type of insurance includes credit monitoring and other services that monitor for data loss, such as loss of company trade secrets. This category is commonly labeled *data compromise protection* or *media liability*.
- **Identity protection:** The focus of this type of insurance is identifying fraud and restoring credit history. Also called *privacy liability*, this category is popular for both organizations and individuals. Many banks and credit card companies offer this protection for their customers.
- **Cyber concepts:** This type of insurance protects the cost to restore IT systems and re-create data post compromise. This may be called *network*

business interruption or something more generic such as *cyber security insurance*.

It is common that a cyber security insurance plan will include coverage for all three of these categories because a cyberbreach can result in loss of data, stolen user identities, and damage to IT systems. Some insurance providers might break down these categories into further subcategories, such as protection specific to the technology labeled “network security.” Specific coverage will depend on how your insurance provider writes up the policy.

Many organizations question if cyber insurance is worth the investment. When an organization doesn’t have insurance and continues to become more dependent on technology, it eventually will need to consider cyber insurance. From a leadership viewpoint, if a cyber incident occurs, the leader of the organization has two responses to the question of whether cyber insurance is worth the investment. One is “Yes, we have a partner and are protected” and the other is “Sorry, but I didn’t think it was important.” Most leaders do not want to state the latter, which is why I continue to see leaders interested in cyber insurance. Other justifications for cyber insurance include the rising concern about being breached and the growth in the cyber insurance market, meaning the average organization owns more technology that needs to be protected, which is leading to an increase in cyber insurance investments.

One absolutely critical takeaway is that *cyber insurance is not a substitute for a security program*! Cyber insurance is just a protection plan if security fails to perform. Some things to keep in mind are as follows:

- Cyber insurance only helps recover certain aspects of your loss. Losses of reputation/brand, time, and money invested in your incident response, and customer trust are just a few aspects that insurance can’t help with.
- Insurance companies will not back up their protection if they find your organization is considered unsecure by industry standards or has violated the policy requirements. You will not find a policy that doesn’t have minimal security requirements to meet.
- You can reduce the cost of cyber insurance with a healthy security program. Sometimes investments in security pay for themselves due to the savings in cyber insurance.
- Many mandatory compliance policies and regulations do not allow

insurance as an option to reduce the risk of a violation. If insurance is required, it’s in combination with other security controls being enabled.

Cyber Insurance Pitfalls

You might be wondering how an insurance provider would be able to not pay on a claim. Imagine you purchase car insurance and subsequently crash your car. What occurs is an investigation of the incident. That investigation determines liability and results in a response from the insurance provider. If you are liable, different insurance policies are invoked and, in some cases, insurance is not paid. If the associated actions are not covered, the insurance provider will reject your claim. If you are in violation of a requirement for your protection, such as driving while intoxicated, the insurance provider will reject your claim. Similar concepts apply to cyber insurance. If your organization suffers a data breach and the insurance company’s investigation concludes that the breach occurred because your organization lacked minimum security measures or was otherwise out of compliance with the insurance policy, the insurance company will deny your claim. Liability could result in a reduction or complete loss of coverage.

The best way to avoid having an insurance provider deny coverage to your organization for a cyberbreach is to ensure that you understand and have met all requirements stated in the policy. Insurance companies include very specific requirements that are designed to reduce the likelihood that the policy will provide coverage if the insured didn’t take reasonable measures to protect itself. For example, if the cyber insurance policy requires your organization to have a packet-capture technology monitoring your network, if a breach occurs, you better have evidence that you were using that technology or the insurance company might deny coverage. Do not overlook any requirements in a cyber insurance policy; develop a method to routinely validate that you are currently in compliance. If your organization’s cyber insurance policy requires use of a specific type of technology, your organization likely is not compliant if it owns the technology but doesn’t keep it powered on. For example, if the policy requires a packet-capture tool and a breach occurs while that tool is offline, the insurance company is likely to deny coverage.

Shopping for Cyber Insurance

Definitely include your security team in the early stages of evaluating and acquiring cyber insurance. The conversation about obtaining cyber insurance commonly starts with nontechnical leadership, even if leadership doesn't completely understand the technical functionality of the systems being insured. Do not let an insurance provider fill out technical requirements or coverage concepts without the security team's review, since the objective of all insurance providers is to limit the scope of coverage to their advantage. Best practice is to align leadership's business goals with what needs to be protected and consult with both the asset owner and team responsible to protect the asset to ensure all aspects of the asset are covered in the insurance. Aspects include the cost of the system, loss to the business if the system is breached and data becomes unavailable, and recovery of lost data.

When shopping for a cyber insurance policy, make sure leaders and technical staff evaluate the cost to *meet and maintain* all requirements to be in compliance with the policy. It is common that requirements will involve additional hardware and people, which sometimes can outweigh the value of the insurance, especially if additional managed service contracts are required. The insurance provider might require an audit or assessment of your security before providing a draft policy. By having a mature security program, the results of an assessment will reduce your proposed cost for cyber insurance. It is common for insurance company auditors to look for specific security technologies and practices as indicators of a mature security program, thus leading to a reduction of the cost of the insurance policy. Be sure to request a list of those indicators along with potential savings from your insurance provider. You might find the associated savings pay for the level of effort to meet the missing technology or practice.

Note

I had a customer that was able to reduce the cost of its cyber insurance by 25% simply by enabling IPS licenses for existing hardware. The cost savings obtained by enabling the IPS licenses was 20 times the cost to obtain the IPS licenses.

As recommended throughout the book, you should implement periodic security assessments and penetration testing. The results from these efforts can serve a few purposes regarding cyber insurance. First, positive results can help justify a reduction in the cost of cyber insurance. This especially holds true when an unbiased third-party contractor performs the services rather than doing an in-house assessment. Second, the results can be used as evidence that your organization has performed due diligence by investing in industry-recommended countermeasures to cyberattacks. Having evidence on hand will be critical if an insurance company denies an insurance claim on the basis that your organization was not secure at the time of the breach and should be liable. Having both evidence that your security tools were enabled at the time of breach and results of an outside party stating your organization has an industry-acceptable security posture will be ideal if a claim goes to court.

Patching Systems

As previously stated, the most ideal approach to mitigating a vulnerability is to perform an action that reduces the risk rather than deploying security tools around the vulnerable system while leaving it vulnerable. A common method to fix a vulnerable system is to apply a *patch*, a piece of software designed to update a computer program or data to fix or improve it. Patching doesn't necessary mean fixing vulnerabilities, but patches often include fixes for security vulnerabilities. Sometimes patches are focused on adding or enhancing features. Those details are included in the patch notes. *Patch management* is the process of identifying, acquiring, installing, and verifying patches for products and systems.

Patching Challenges

It is critical that the engineer deploying a patch first research the potential impact from deploying the patch. Most patches include release notes stating what the patch does, how to deploy it, and potential risks. Many vendors also include links to open cases resolving the patch, which the engineer must review to identify potential problems applicable to the specific system being

patched.

According to NIST SP 800-40 Rev. 3, *Guide to Enterprise Patch Management Technologies*, there are challenges with patching that must be overcome before a patch is to be deployed. Among those challenges are timing, prioritization, and testing. Timing challenges include when the patch is available, when the system can be patched, and how much time is required to deploy the patch. Vulnerability management practices have many tasks to handle, which leaves limited time for deploying patches. This leads to the second challenge, which is the prioritization of patches. The priority of other systems might cause some patches to not be addressed. Finally, best practice is to test a patch to validate that it won't introduce residual risk; however, that once again requires time and other resources such as a test lab. You should expect your patch management practice to find that the timing, prioritization, and testing requirements of patches will often conflict.

Some additional challenges come from the vendors responsible for the patches. It is common for a vendor to bundle patches together, saving your team time with the patching process; however, bundling also makes it difficult to break apart and test specific patches. Some vendors that bundle patches do so in a monthly or quarterly fashion, which means there will be a period of time that a vulnerability patch is not available. If an attacker finds a vulnerability before a patch becomes available, your only defense will be the security surrounding the system. Also, the vendor might offer multiple ways to obtain the patch, which can cause conflicts. Patches might be automatically applied, circumventing the patch management program.

Yet more challenges can come when assets are unmanaged, out of the office, or nonstandard IT components. All of these situations need to be accounted for, and as challenges come up, the timing, prioritization, and testing of patching need to be considered.

No Patch

In some cases, a patch won't be available. If you encounter this scenario, be sure to reach out to the vendor and inquire about any identified vulnerabilities. The vendor might have patches available that have not been published or might be able to advise on the best course of action. If a patch is

not available for a vulnerability, your mitigation option is to secure around the vulnerability. Segmentation, firewall, IPS, and antivirus are all examples of security solutions that can protect a vulnerable system from being compromised. If you are using a signature-based security tool, validate that you have the proper signature(s) downloaded and enabled that are associated with the vulnerability. Figure 9-28 is an example of Cisco Firepower showing the available signatures for the Apache Struts vulnerability mentioned earlier in this chapter. If signatures are not available, you should speak with the vendor of the security tool to request protection against the vulnerability of concern.

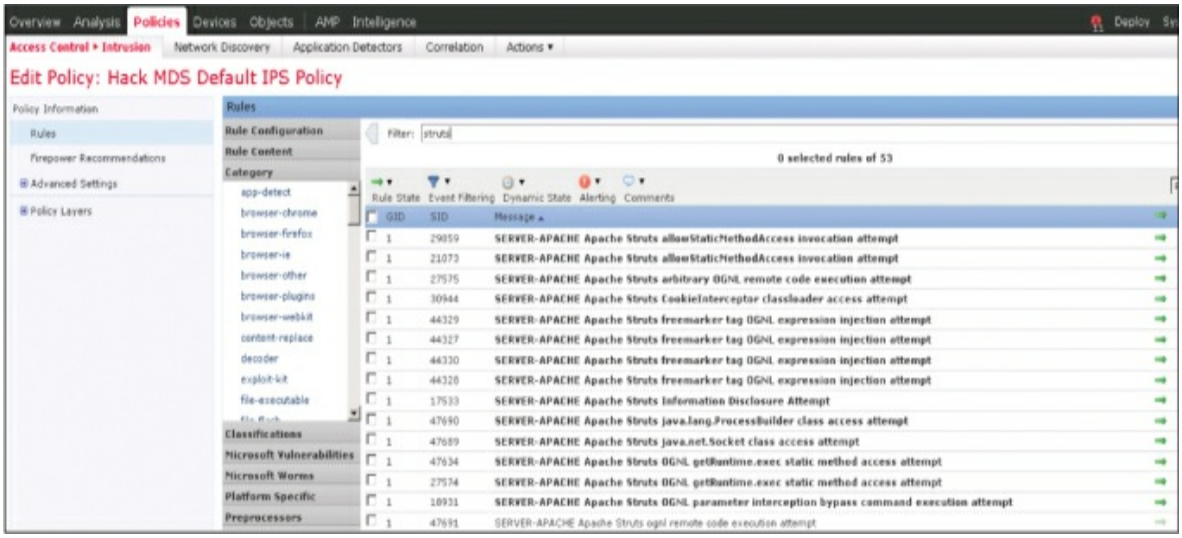

FIGURE 9-28 Cisco Firepower Apache Struts Rules

It is common for a security tool to not display a signature for a new vulnerability or exploit when news of a recent vulnerability makes major headlines. The reason is that the security vendor does not want to inform the attackers how the defense signature works until the vendor has had sufficient time to release a patch.

Note

You can learn more about hidden signatures by speaking directly with the vendor. Further delays in showcasing signatures can occur for

devices that never release a patch. I find this to occur often with IoT devices.

Three general reasons why your SOC might not find details on the Internet or in a security tool about how to remediate a vulnerability are as follows:

- The industry or vendor researchers are not aware of the vulnerability. This is common when custom in-house applications are built, which external vendors do not have access to and hence cannot release vulnerability data about. This could be considered a day zero; however, the industry typically refers to a vulnerability within commonly used technology as a day zero rather than a vulnerability within a customized tool that is used by only a specific organization and therefore never tested by industry researchers. In both cases, a patch will not be available because a commercial vendor is not involved.
- The industry or vendor researchers know about the vulnerability but have decided not to disclose details about it (yet). Among the many reasons for withholding this information are to avoid exposing the defense measures to the attackers until a patch is available and to comply with contractual nondisclosure guidelines.
- The vulnerability is a true zero-day vulnerability, meaning threat actors know about the vulnerability before the industry and vendors do. These vulnerabilities tend to be released on the black market and based on widely used technology rather than a specific customized tool.

All hope is not lost when data isn't available about a vulnerability or associated threats. Chapter 1 focused on how important it is for every organization to use a defense-in-depth approach to defending against cyberthreats. Part of that approach is to use behavior and anomaly detection, both of which are designed to detect threats that bypass signature-based security tools. When new threats are found that have bypassed your security tools, it is critical to understand why your tools failed. Maybe tuning is needed or maybe there is not a signature available based on the reasons previously covered. Recommend practice dictates addressing successful attacks with the vendors of products that failed so that you better understand why the failure occurred and what you can do to prevent future failure.

Residual Risk

Making changes such as patching systems or implementing new security controls will have some level of residual risk, meaning existing or new problems. Patches can require a reboot of a system, which could lead to residual risk. Patches could upgrade some systems, causing other programs, such as Java, to no longer function properly. Some common examples of issues that are a result of a reboot are lost configuration that wasn't saved, breaking integration with other systems, or nontechnical problems such as customers or employees complaining the system is not available. Deploying a new security tool could prevent certain tools from functioning, once again causing unplanned downtime until a workout can be deployed. Keep in mind that proper patch management is a major responsibility of a SOC. Attackers know organizations are slow to patch; therefore, they wait until patches are released and attempt to exploit the associated vulnerabilities, knowing organizations will take weeks or months to patch critical systems.

Vulnerability mitigation documentation must include details associated with expected residual risk before a final decision is made regarding whether to implement the mitigation. If residual risk is unknown, testing must be performed, or the impacted vendor must be contacted to validate what residual risk should be expected. Residual risk evaluation must consider not only vendor-specific risk but also operational risk. Operational risk includes backing up the data, saving the current configuration, and all dependencies to the impacted system. Residual risk is one major component that is evaluated before remediation steps are approved or denied.

Remediation Approval

Approving remediation requires having a procedure that is repeatable to ensure all vulnerabilities are addressed in a controlled fashion. The first part of accomplishing a repeatable practice is determining how often approved mitigation will take place. The period for implementing mitigation is often called the *maintenance window*. Maintenance windows tend to occur during nonbusiness hours and weekends when systems are at low utilization. There is not a set rule for how often a maintenance window should exist. Some

organizations schedule windows every weekend, while other organizations schedule larger windows once a quarter or even only once a year. Best practice dictates scheduling recurring windows at least once a month to avoid the need for multiple exceptions due to critical patches and to accommodate the frequency with which the industry is currently releasing updates.

The following are best practices regarding obtaining and delivering patches or other mitigation steps:

- Verify the integrity of patches before installing them. This is commonly done by using a checksum.
- Make sure patches are secured from download to delivery and delivered from a verified source.
- Encrypt network communication as patching takes place.
- Test patches before deploying.
- Validate all residual risk, including operational impact due to a reboot or other changes to the system post patch.
- Determine expected downtime and alert all impacted parties.
- Ensure permission from the asset owner is obtained.
- Deliver patches during approved maintenance windows.
- Ensure there is a backup plan if the patch fails.

Your goal is to reduce the approval time through process efficiency because you will have a limited maintenance window for applying patches. Using batch software updates that have been properly validated will allow you to avoid multiple updates and minimize downtime. Some patches will require what is sometimes called “warm patching,” meaning only a service goes down rather than the entire system. Another ideal situation is when redundant systems are being patched, because one system can remain up while its backup system is patched. This scenario works by allowing the backup to act as the primary while the primary system is patched. Then the process is switched, limiting the downtime and reducing the risk of residual negative impact. If redundant systems exist, the approval time can be reduced. One final patching option is *rolling patches* where patches are put in a patch queue and automatically rolled out over multiple days. All risk assessment of

patches can be done before placing the patches in the queue so that the time to deliver patches can be shortened. Any of these approaches can help improve the time to approve and deliver patches.

Once patching is complete, the security engineer responsible for the ticket must rescan the system to validate the vulnerability was mitigated and that the system is operating properly. If a problem is found, it will be up to the asset owner regarding the next step. Options can include implementing a rollback plan using backed-up data, retrying the mitigation, or consulting with the vendor to obtain further instructions. If the vulnerability still exists but operations are returned to acceptable levels, the ticket will move back to the decision process regarding how the mitigation should proceed. It is important to include all steps taken within the vulnerability ticket to avoid repeating the same steps that led to the unsatisfactory outcome.

Note

Be sure to consult with the vendor of a system that is negatively impacted by a patch. Vendors are motivated to develop custom patches as well as help troubleshoot patches when they are liable for the negative impact.

Reporting

NIST SP 800-40 Rev. 3 recommends that organizations establish repeatable metrics for patch management approval. This enables the SOC to demonstrate progress in implementing security programs, meaning a SOC can show improvement based on reducing the risk to the organization through patch management. Showing progress allows for obtaining more buy-in from leadership, establishing a baseline for a reward system, and providing the capability to target maturity metrics, which leads to obtaining more funding and support. Having metrics allows the SOC to demonstrate current investments are meeting a desired outcome, which is critical for demonstrating a healthy SOC program.

There are options regarding how your SOC can develop metrics regarding the patch management program. The SOC can track the percentage of desktops and laptops covered by the patch management program. Metrics can include trends on identified and remediated vulnerabilities to demonstrate the SOC's impact on host system risk reduction. A similar strategy can be applied to other systems, such as those in the datacenter. Reports for these types of metrics should include how often hosts/systems are checked for missing updates, how often updates are applied, what is the average time to mitigate a host system, and what is the average time between when a patch is released and it is delivered to host systems. Reports must also include impact metrics to show the result of the SOC's patch management efforts. Impact can include cost savings, estimated risk reduction, and cost trends as the SOC improves its vulnerability management program. Trends should show a reduction in costs and growth as the SOC's vulnerability management program matures using automation and playbooks and improves in mitigating identified vulnerabilities.

Note

It is critical to not overlook the value of reporting metrics. Nontechnical team members including leadership will need this information to justify the investment in the vulnerability management program.

Exceptions

There will be situations that put the organization at enough risk that a policy needs to be violated. An example is a critical vulnerability found within a system that needs to be patched before an approved maintenance window. For these situations, an exception policy needs to be invoked. Not having a policy will lead to sporadic and unmonitored exceptions occurring as well as exceptions existing far longer than needed.

An exception program starts with a method to request the exception. Anybody will be able to request an exception, but all exceptions will not be

granted. Requests can come through email, over the phone, or via other communication mechanisms, any of which results in an exception ticket being created. It is recommended to use a form that limits the request to specific datapoints and categories. Doing so will allow for the development of playbooks leading to automation and efficient response.

The team receiving exception requests will have a few tasks:

- Assess the risks created by the exception
- Evaluate potential alternatives
- Provide recommendations
- Determine the appropriate departmental and, if applicable, asset custodian(s) approval

A response will be delivered by the SOC member responsible for the incident ticket and an accept or deny, both of which must contain certain details. A response from the one responsible for the ticket of accepting an exception needs to include a time limit to ensure that the exception terminates as soon as the task requiring the exception is completed. A common example this solves is an exception to open a firewall port to allow traffic that is typically denied, such as opening a port to traffic so that a patch can be downloaded to a vulnerable system. If no timeline is set, the port could remain open, exposing the organization to unwanted risk. An acceptance will include details regarding what is being accepted, whether assistance is needed, what permissions or other requirements must be met prior to the exception being used, and when the exception expires. Exceptions must come through the security team. General managers are not permitted to create exceptions. They must also put requests in to the security team to ensure a unified and control exception practice.

All exceptions must be documented. Documentation must contain the following elements:

- Individuals or systems involved
- The scope of the exception
- Limitation of the exception
- Mitigating controls required

- Dates/times
- Reasons for exception
- Approval

If an exception is not granted, a message will be returned explaining why it was not granted as well as providing alternatives to the exception. Common reasons for an exception not being granted are that a feasible alternative exists or that the risks outweigh the projected benefit of allowing the exception. To avoid repetitive requests, some organizations specify that another request regarding the impacted system can't be submitted for a certain period of time; this typically isn't required unless repetitive requests become a problem.

Vulnerability Management Process Summarized

Figure 9-29 shows a summary of the vulnerability management process based on the job roles covered earlier in this chapter. The asset owner will be notified of the vulnerability since he/she will be closest to the system and be the best person to qualify its relevance. The asset owner can determine whether to accept the risk or specify that another action, such as a risk reduction, should be implemented. The security officer will be notified to track the status of the vulnerability throughout the mitigation process. The security officer can also provide input regarding potential risk reduction steps. The IT team will also be notified to determine the vulnerability's impact to the network. Both the IT team and security officer could offer network or security tool options that reduce the risk without making changes to the device. All input will be combined and turned into a corrective action or a waiver to accept the risk.

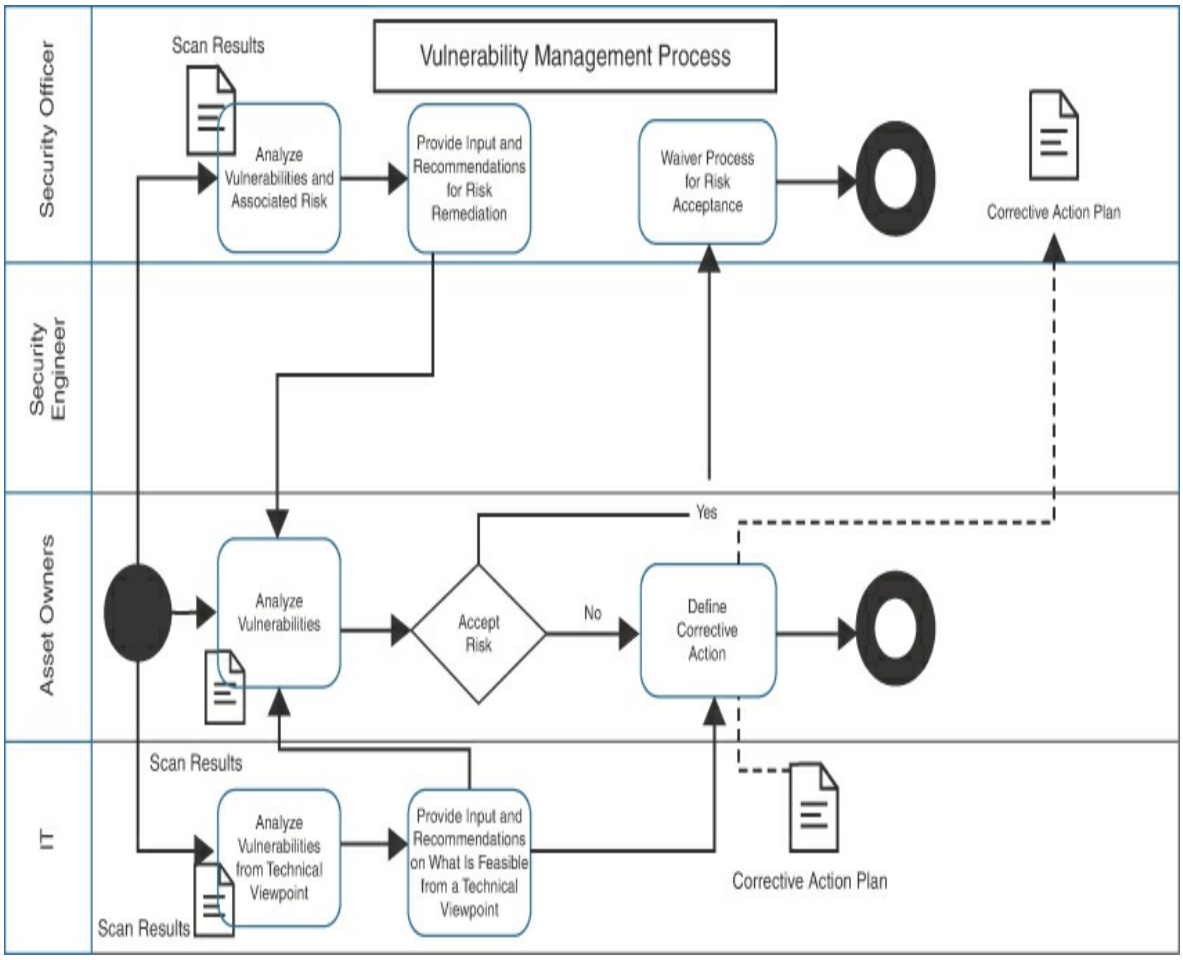

FIGURE 9-29 Summarizing Vulnerability Management Process

Summary

This chapter covered everything involved with the lifecycle of a vulnerability management program. You learned how to identify assets and assess them for value and potential risk to the rest of the organization. You learned about different tools that can be used within a vulnerability management program. I covered all the ingredients needed to develop and run a vulnerability management service. You learned about the different ways to tune existing scanning practices and address when special exceptions need to be made. The chapter concluded with a look at how and when to deploy remediation to a vulnerability as well as the potential risk.

Chapter 10 will focus on data orchestration, which can help improve many

SOC services, including vulnerability management. In my experience, every organization that deals with lots of events is investing heavily in automation and orchestration, making Chapter 10 a must-read for SOC professionals who are responsible for mid- to large-sized organizations.

References

Joint Task Force. (2020, September). NIST Special Publication 800-53, Revision 5: Security and Privacy Controls for Information Systems and Organizations. NIST. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Nationwide. (n.d.). What Is Cyber Insurance? Nationwide. https://www.nationwide.com/lc/resources/small-business/articles/what-is-cyber-insurance

NIST Information Technology Laboratory. (n.d.). Vulnerability Metrics. NIST. https://nvd.nist.gov/vuln-metrics/cvss

Rapid7. (n.d.). Working with Scan Templates and Tuning Scan Performance. Rapid7. https://docs.rapid7.com/nexpose/working-with-scan-templates-and-tuning-scan-performance#section-keep-the-triangle-in-mind-when-you-tune

Risk Based Security. (2017, May 2). CVSSv3: When Every Vulnerability Appears to Be High Priority. Risk Based Security. https://www.riskbasedsecurity.com/2017/05/02/cvssv3-when-every-vulnerability-appears-to-be-high-priority/

Santos, O. (2016, October 31). The Evolution of Scoring Security Vulnerabilities: The Sequel. Cisco Blogs. https://blogs.cisco.com/security/cvssv3-study

Shakeel, I. (2019, July 5). Risk Management Concepts and the CISSP (Part 1). Infosec Resources. https://resources.infosecinstitute.com/category/certifications-training/cissp/domains/security-and-risk-management/cissp-risk-management-concepts/#gref

Souppaya, M., & Scarfone, K. (2013, July). NIST Special Publication 800-

40, Revision 3: Guide to Enterprise Patch Management Technologies. NIST. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-40r3.pdf

Chapter 10

Data Orchestration

I always try not to overload my music with orchestration and to use only those instruments that are absolutely necessary.

—Abel Korzeniowski

The term *data orchestration* has gained popularity among many groups, including cloud and datacenter specialists who deal with lots of data. This increasing popularity of data orchestration is due to challenges that have occurred as data increases in volume and velocity. One challenge is expecting an analyst to be able to view and understand large amounts of data. Even tools such as security information and event management (SIEM) can quickly overwhelm an analyst with event data, the alarms for which represent consolidated events that take time to investigate. Another challenge is choosing the best response to an event when there are many factors to consider, mixed with too many events occurring at once that need prompt attention. As public cloud and threat intelligence resources are introduced into the SOC, more effort is required to tune data, leading to more work based on technology upkeep. These types of challenges are causing SOCs to experience a reduction in productivity, the response to which is to invest heavily in data orchestration.

Chapter 5, "Centralizing Data," covered various types of data that SOCs can leverage as well as how to centralize that data to simplify visibility across multiple tools. What Chapter 5 did not cover is how to take action on that data. This chapter focuses on how to put data to use through automation and orchestration. The security industry calls this *security orchestration, automation, and response (SOAR)*, which integrates the data collection concepts covered in Chapter 5 and the incident response concepts covered in Chapter 8, "Threat Hunting and Incident Response," and focuses on

automating through orchestration. This chapter also covers required concepts and skills necessary to leverage SOAR technology such as playbooks and DevOps. There is a lot involved with SOAR technology, which this chapter breaks down into more digestible modules.

> **Note**
>
> In this chapter, you will also see coverage of cross-layered detection and response (XDR), which has many of the same capabilities offered by SOAR technology.

Before defining what SOAR is, a logical place to start is to look at what is involved with data orchestration, a fundamental component of SOAR.

Introduction to Data Orchestration

A key takeaway from Chapter 5 and Chapter 7, "Threat Intelligence," is that the value of data is judged based on how it is used and understood. Data is used and understood based on its situational context. A basic example of the concept can be seen in John Boyd's OODA loop, which is a strategy for responding to events. The OODA loop states that one observes, orients, decides, and acts based on feedback, meaning (in this book's example) how a SOC responds to a situation. Feedback, as referenced by the OODA loop, means applying the context of new data received to adjust one's response, which can be interpreted as viewing the available data from security tools, threat intelligence, and input from SOC analysts. Simply put, a SOC needs to obtain the right data to understand and respond to a situation properly. Figure 10-1 shows a simple OODA loop diagram.

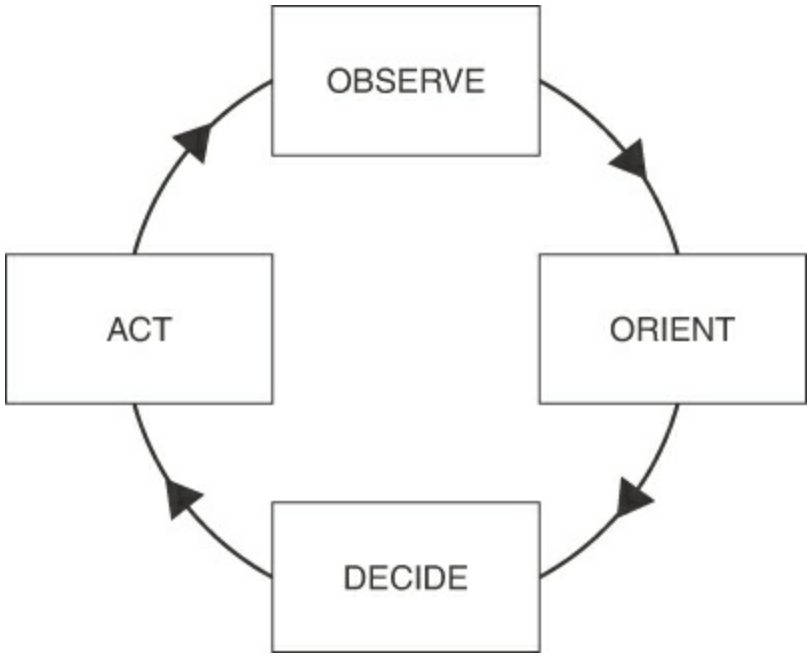

FIGURE 10-1 OODA Loop Diagram

Chapter 5 covered how to consume various types of data, including leveraging a SIEM solution to be a centralized view of data. What Chapter 5 did not cover is how to take action based on the feedback—the act part of the OODA loop. Many SOCs are finding that their processes for collecting and interpreting data are ineffective for keeping up with the number of cases that they are responsible for, leading to very slow responses to requests and delayed reactions to events. As a result of this problem, SOCs are seeking solutions that enable their SOC services to improve their efficiency by providing quicker and more effective responses. This demand for improving responsiveness to events is what led to the creation of the SOAR market.

To further explain where the SOAR market came from and demystify the different industry tools that can fall under an orchestration-based solution set, you need to understand the difference between SIEM, SOAR, and XDR. Many vendors mix and match the terms SIEM, SOAR, and XDR in their marketing campaigns, causing a lot of confusion regarding which capabilities to expect from their solutions. To demystify which of the different industry tools a true SOAR solution should include, and to further explain what gave rise to the SOAR market, the sections that follow address these industry terms using vendor-agnostic language.

Comparing SIEM and SOAR

The SOAR market actually is a subset of a market focused on the value derived from a combination of a SIEM and SOAR solution. There are a few reasons for this. First, SIEM technology predates SOAR technology—SOAR was created in response to the deficiencies in SIEM solutions. A SIEM solution allows a SOC to centralize data, but adding a SOAR solution extends what a SIEM solution can do. Capabilities offered by adding a SOAR solution include case management, orchestration, automation, and response. Most SOAR products do not store data, meaning they require a SIEM solution for a place to store data and retrieve data commonly referred to as a *data lake*.

To summarize the difference between a SIEM and SOAR solution, the primary focus of a SIEM solution is collecting and analyzing data, mostly in the form of logs, which are stored in a data lake for access by SOC analysts. SIEM solutions also interpret data-based patterns, establish behavioral baselines, and have machine learning training sets that help to simplify searching data. SOAR solutions focus on grouping existing data from the SIEM solution and converting alarms into a single case, which is handled using the playbook concept and improved with automation and orchestration. Table 10-1 highlights the common features and focus areas for SIEM and SOAR solution.

TABLE 10-1 SIEM and SOAR Solution Common Features

SIEM	SOAR
Collects logs	Consolidates indicators of compromise (IOCs) in cases
Syslog based	Alert and IOC focused
Automates commands	Automates playbooks (such as Ansible and Puppet)
Optimized for true and raw event search	Suited for hunting suspicious activities
Integrates with vulnerability scans	Integrates with third-party threat feeds

Incorporates user behavior analytics	Automates processes for testing for false positives

Many organizations are investing in a SIEM/SOAR solution combination to receive the value from combining these technologies. As a response to the market demand, many SIEM solution vendors have acquired or developed SOAR solution capabilities. Splunk acquired Phantom for its SOAR solution. IBM added the SOAR solution platform named Security Resilient to complement its SIEM solution known as QRadar. In both examples, the SIEM solution is required for the SOAR technology to function. Exabeam is an example of an offering that does not have this requirement—one solution does both. The specifics of how SIEM/SOAR solution capabilities are offered will vary by vendor.

The Rise of XDR

As the market demand for the SIEM/SOAR solution combination increased, another security market came into the mix. Some endpoint-focused security tools vendors found that antivirus capabilities such as pattern matching using threat signatures weren't capable of detecting complex malware. Also, the same endpoint-focused security vendors found that the capabilities for responding to antivirus needed improvement. Over time, the endpoint detection and response (EDR) market developed, representing what the SIEM/SOAR solution market can do but only from an endpoint viewpoint. With an EDR offering, a SOC can collect endpoint logs, monitor the endpoint for the three key security detection capabilities (signature, behavior, and anomaly), and respond using orchestration and automation through the use of the EDR endpoint client.

In response to the increase in demand for the SIEM/SOAR solution combination, some EDR vendors started to add capabilities outside of the endpoint—performing EDR but on email, applications, networks, and anywhere else they could expand their capabilities. Eventually, some EDR tools were able to offer similar value as the SIEM/SOAR solution combination by extending EDR to other systems. Those vendors rebranded their offering from EDR to "extended" detection and response, more

commonly referred to as cross-layered detection and response (XDR).

Each vendor's specific capabilities will be slightly different, however. When a vendor claims to offer an XDR, it is referring to offering similar capabilities to what a SIEM and SOAR solution can offer, which are the following:

- Collecting logs and other data for centralized display of data
- Deduplication and correlation of events
- Consolidation of IOCs in cases using case management capabilities
- Support for multiple system types
- Automation and orchestration
- Various methods of response

SIEM vendors have adapted their product marketing to emphasize either a SIEM/SOAR value or an XDR value. For example, LogRhythm has a substantial share of the SIEM market but today markets several of its products as XDR solutions. Other SIEM solution market share leaders, such as Splunk, focus on the SIEM/SOAR solution combination value, which is very similar to how an XDR solution is marketed from a capabilities viewpoint.

That is a high-level explanation of where SOAR solutions fit into the data collection and event response market. The next section delves into SOAR, which is followed by a section covering EDR/XDR. You'll be introduced to examples of popular enterprise offerings that relate back to common capabilities found across the SOAR and XDR solution market. You can find additional coverage of SIEM in Chapter 5.

Security Orchestration, Automation, and Response

SOAR solutions are built on four engines as defined by Gartner:

- Workflow and collaboration
- Ticket and case management
- Orchestration and automation

- Threat intelligence management

The fusion of these capabilities improves SOC productivity and incident response times by bringing together people, process, and technology. As such, these engines also provide an ideal basis for a robust security stack.

Data orchestration is the foundation of a SOAR solution and is defined as the automation of data-driven processes from end to end, including preparing data, making decisions based on that data, and taking actions based on those decisions. Automation does bring up some interesting questions such as:

- Can every aspect regarding your response to events be automated?
- Will automation lead a reduction of jobs within the SOC?
- Are the robots taking over the world?

The answer to all of these questions is emphatically no. People are needed to address complex decisions.

Automation is ideal for simple repetitive tasks—it can address freeing up time for an analyst to deal other work. Automation is also only one part of a successful response to an event. Orchestration combines well-structured processes (often referred to as playbooks) with automation to deliver the result of a faster and more effective response. Some tasks in the overall orchestration can be automated while others that require human manual efforts will not be automated.

There is a ton of value from leveraging orchestration technology, which is why many SOCs are investing heavily in this technology space. A summary of specific technical value points offered by SOAR-type tools is as follows (these topics are the focus for this chapter):

- **Playbook automation:** Automate a series of tasks with one or more tools
- **Context enrichment:** Bring data from different tools into useful context
- **Enhanced investigation:** See how security sources relate to an event
- **Key performance indicators (KPIs) business intelligence:** Capture trends to prove return on investment (ROI) in security tools, processes, and personnel
- **Case management:** Ensure the lifecycle of an event is properly managed

- **Collaboration:** Allow for collaboration across different teams within and outside the SOC

All these value points are dependent on an existing solid foundation for SOC processes. If a SOC's services are not mature and function in a sporadic manner, orchestration will not provide any value. The reality of orchestration is that it will leverage whatever is available, including existing problems, meaning orchestration can make things less effective by automating your problems unless those problems are first addressed.

Note

If your dataset has problems, orchestration will just automate those problems! You must first ensure the process works before you attempt to speed it up.

SOAR Example: Phantom

The best way to understand SOAR concepts is by exploring a SOAR product as an example. This section uses Phantom by Splunk to better explain each concept. Know that many industry leaders in the SOAR solution market offer similar features, but the specifics vary based on the vendor. Also know that many of the commercial SOAR features can be developed by using open-source options; however, commercial SOAR solutions are more user friendly, have simpler plug-and-play configuration, and offer various levels of support from the vendor that you wouldn't get with an open-source option.

Splunk is a market leader for SIEM solutions with a heavy focus on its self-developed application community. Splunk saw the need to add SOAR features and acquired Phantom back in 2018. Like many SOAR offerings, Phantom as a SOAR requires a SIEM solution in order for it to have data, meaning Splunk customers looking to add SOAR capabilities must stand up a separate Phantom system.

Note

All SOAR offerings do not require a SIEM; however, this is a very common design in the industry. Make sure to consider how many systems need to be stood up to perform SIEM and SOAR functionality as you evaluate a vendor's offering.

The best place to start is by exploring the main Phantom dashboard. Figure 10-2 shows the Splunk Phantom dashboard.

FIGURE 10-2 Splunk Phantom Main Dashboard Example

The top widgets in the Phantom dashboard focus on core SOAR solution features. Resolved events represents any events that have been addressed by a playbook or analyst. Following that are stats regarding how fast a case is resolved, because the focus of a SOAR solution is speeding up and improving how events are resolved. Stats include everything from time to dollars saved by the SOAR solution. As you can see, the focus of Phantom is not identifying events, like you would expect with a SIEM solution; instead, this SOAR solution heavily focuses on case management and saving money

through orchestration and automation.

Case Management Example Using Phantom

Case management means dealing with the entire lifecycle of an event, from recognizing an event is occurring, to tracking how the event is handled, and eventually closing out the case once all steps of a playbook have been executed successfully. If multiple incidents are found related to an event, the SOAR case management feature needs to add those incidents to ensure the entire event is handled during the case management process. Case management can also include other forms of data enrichment designed to help the analysts assigned to the case be more effective at responding to an event.

Figure 10-3 shows the Splunk Phantom case management dashboard. The top widgets focus on the number of events, severity of events, what cases are new, open, or closed, and who the case owners are. To better navigate the cases handled by an organization, search features enable a SOC manager to pull up previous cases, track how cases are being managed, and search for other key terms in relation to what has been or is being handled in Splunk Phantom.

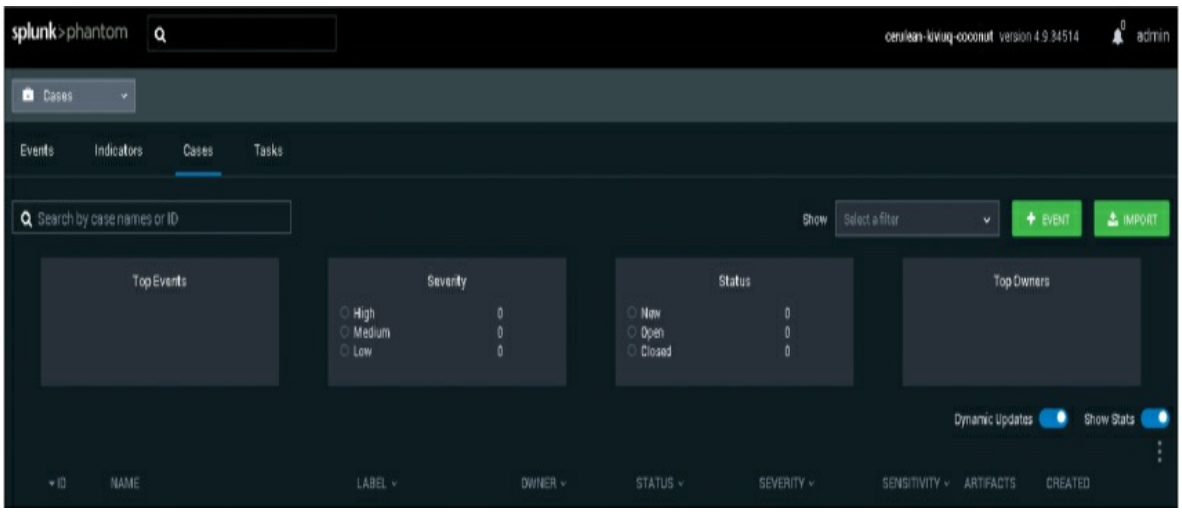

FIGURE 10-3 Splunk Phantom Case Management Dashboard Example

Playbook Example Using Phantom

One important concept to understand is *playbooks* (covered in more detail throughout the chapter). It is common for enterprise SOAR solutions to include a wealth of playbook templates to help organizations match common tasks that can be automated in a prebuilt template. Phantom has many templates that can and should be customized for your SOC use. A new playbook can also be either uploaded or built from scratch. Figure 10-4 shows the search area in Phantom for viewing the playbook templates that are available.

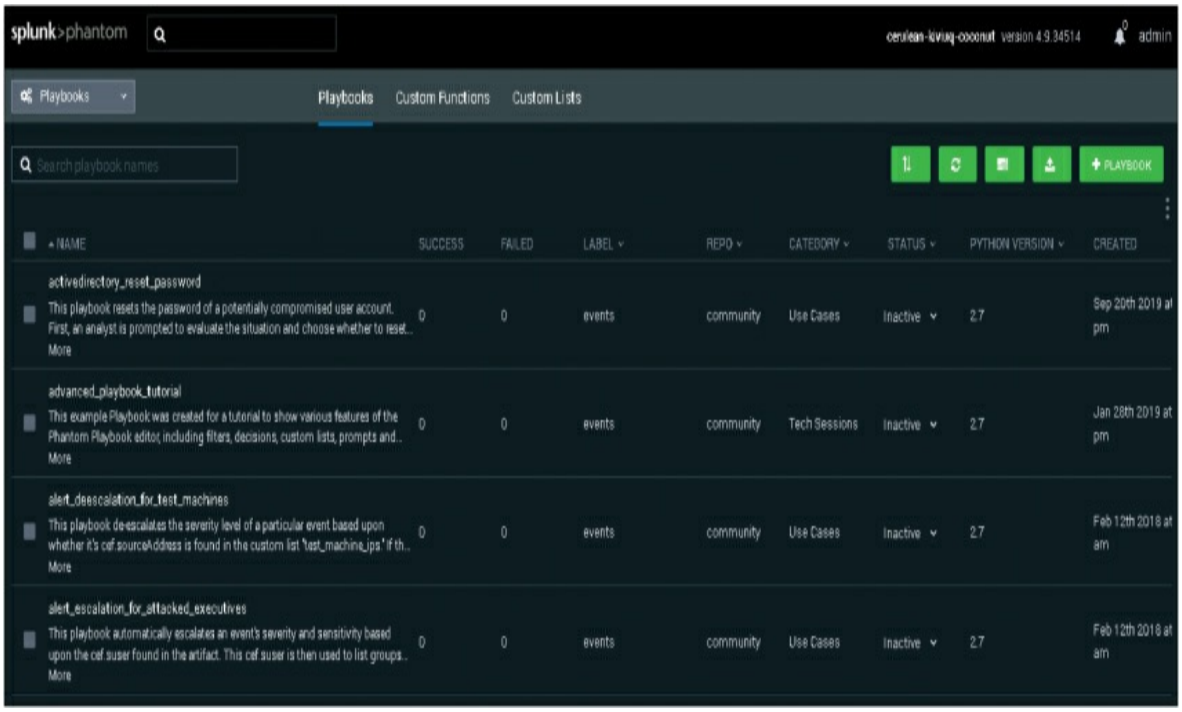

FIGURE 10-4 Splunk Phantom Playbook Template List Example

Figure 10-5 shows an example of what Phantom playbook looks like when you open it up. For this example, I have clicked into a playbook template focused on responding to a compromised email. Phantom uses the industry-standard workflow symbols, including a start icon, decision diamond, and rectangles representing actions that are taken. This format makes it easy for a SOC analyst of any skill level to view and understand what steps to take in response to an email compromise. The section "Playbook Components" later in this chapter covers these symbols in more depth.

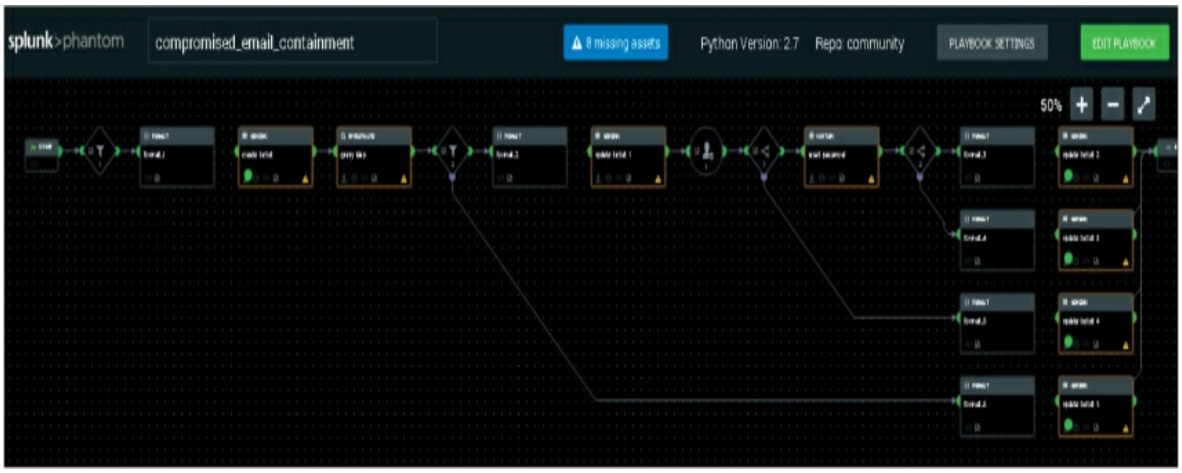

FIGURE 10-5 High-Level Splunk Phantom Playbook Example

Figure 10-6 shows a zoomed-in look at the first part of the Phantom playbook template for an email compromise. Boxes that have an orange border represent steps that have not been configured properly, which means this playbook cannot be run until these items are resolved. This makes sense based on how this workflow is designed to deal with a specific organization's response to email compromise. Your SOC must identify some specific datapoints that relate only to your organization, such as who should be alerted, what should be done when a potential email compromise is seen, and what represents an email compromise.

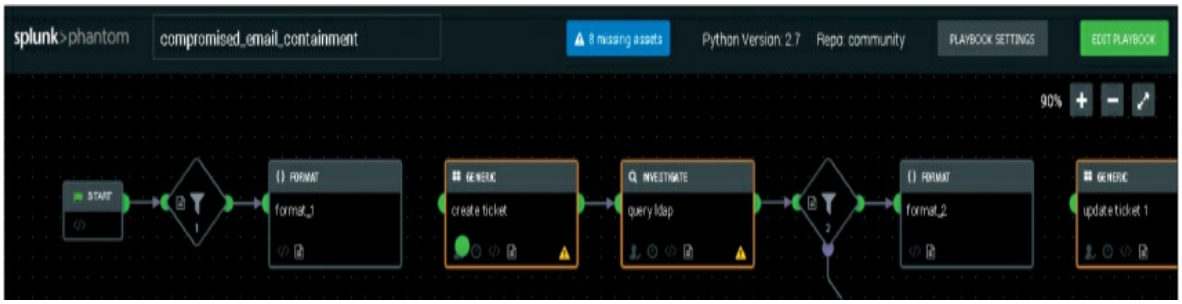

FIGURE 10-6 Zoomed-in Phantom Playbook Example

In this example, one area that needs to be completed is identification of who is responsible to investigate an email compromise. If this is left blank, the logic of the workflow can't identify who should receive an alarm when an event occurs, which is why this is a mandatory step to establish before this

playbook can take action. Phantom offers workflow wizards that walk analysts through the steps of setting up playbook templates, explaining what is needed for required steps as well as why these steps are required, dramatically reducing the chance of creating a broken playbook.

DevOps Example Using Phantom

Commercial SOARs simplify the playbook creation process; however, that does not mean programming is completely removed. It is common for enterprise SOAR solutions to help simplify playbook creation by using templates, wizards, and prebuilt functions that can be called upon; however, a certain type of programming called DevOps must be done in order for a workflow to be created. Figure 10-7 shows an example of clicking into the "create ticket" function in the Phantom email compromise playbook template.

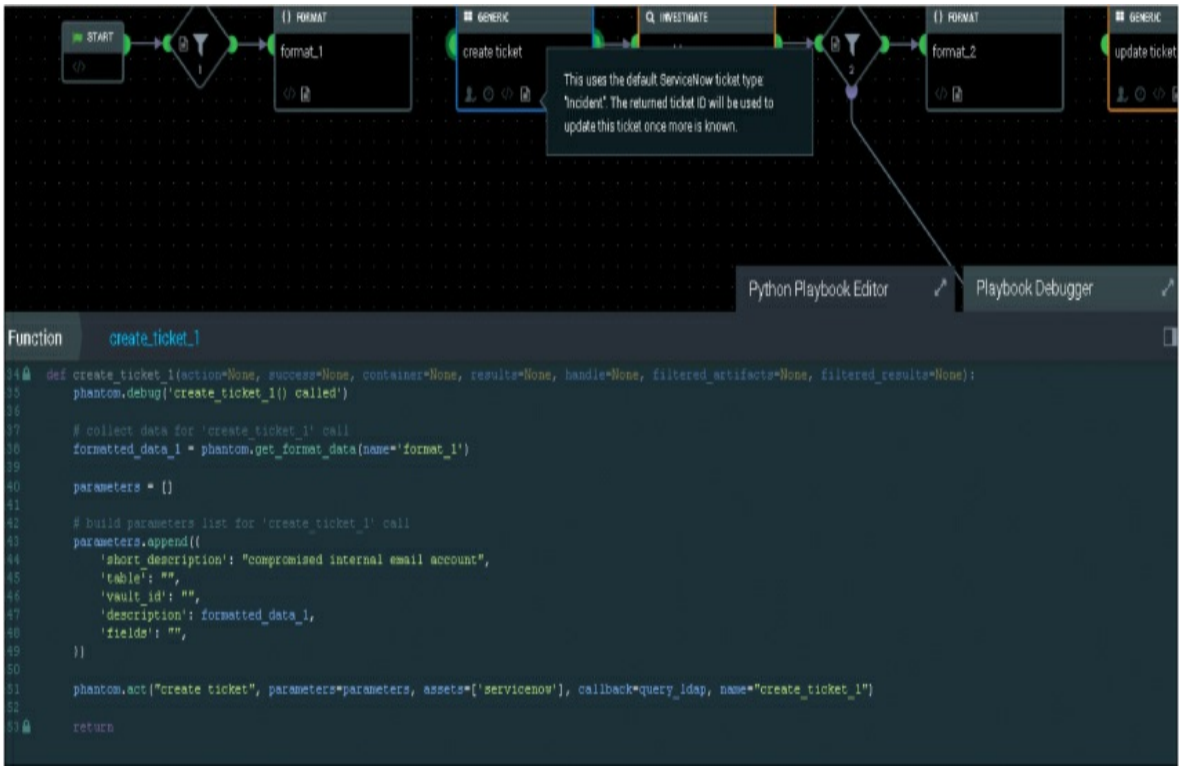

FIGURE 10-7 Phantom Example of DevOps Coding

Notice Phantom brings up the code associated with generating a ticket; wizards and other built-in tools can help an analyst quickly adjust the code to

meet the organization's needs. As with most programming languages, Phantom includes a debugger (Playbook Debugger), which will flag errors found in a playbook's code and attempt to highlight which line of code is causing the problem. This example emphasizes why DevOps skillsets are increasing in demand as well as why DevOps skills are viewed as a different skillset than general programming experience with programming languages such as C++ or Pearl. The focus of this type of programming skillset is to understand workflow objectives and write code that can accomplish those objectives using the most efficient and accurate method possible. Another way to say this is that this type of programming combines the goal of IT operations and software development. The section, "DevOps Programming," later in the chapter, takes a closer look at this type of programming.

The previous examples just scratch the surface of what a SOC can do with an enterprise SOAR solution such as Splunk Phantom. Key points to take away from this use case are as follow:

- SOAR solutions focus on concepts such as managing cases and time saved.
- Playbooks can be manual, but the value of a SOAR is automating steps when it's appropriate.
- Automation can be achieved using templates and wizards; however, some programming is needed.
- Programming in a SOAR solution is a different type of programming focused on making technology work with other technology rather than creating something new.
- SOAR solutions do not focus on data mining or event management like a SIEM solution offers but are great for event consolidation, which can be part of a playbook. For example, when a specific event occurs, the SOC can launch a playbook action.

The next topic to examine closer is EDR and XDR solutions as they offer similar value as what a SIEM/SOAR solution combination can offer. You'll have a close look at a market leader in the EDR market to give you a feel of where XDR came from.

Endpoint Detection and Response

So far, the discussion of SIEM and SOAR solution concepts has used examples of automation at the network layer with a focus on the workflow of playbooks. Automation can also relate to endpoints and software. The National Institute of Standards (NIST) published NIST Interagency Report (NISTIR) 8011, *Automation Support for Security Control Assessments*, which goes into detail of a specific methodology on how to test and automate cybersecurity concerns around software automation. In other words, NISTIR 8011 addresses questions regarding how to audit automation software. Table 10-2 lists some common attack methods documented in NISTIR 8011 and the corresponding mitigation processes.

TABLE 10-2 NISTIR 8011 Attack Methodologies

Attack Method	Mitigation Process
Internal Entry	Block local access: restrict admin access and software installation privileges
Insider Attacks	Reduce and disallow software to run unless previously authorized
Gain a Foothold	Restrict software installation
Land and Expand	Reduce existing vulnerabilities on systems and enforce segmentation between networks.

Endpoint security vendors have taken on the challenge to address these concerns in a similar manner as the SOAR market, including offering case management, playbooks, automation, and orchestration, commonly categorized as endpoint detection and response (EDR). Like a SOAR playbook, an EDR playbook considers the entire lifecycle of an event, including who is responsible for performing what actions and what should be accomplished before a case is closed.

EDR Example: CrowdStrike

One example of a popular EDR enterprise solution offering comes from CrowdStrike, known as Falcon. Falcon is a cloud-delivered EDR solution that leverages an agent installed on an endpoint. CrowdStrike Falcon offers endpoint antivirus capabilities mixed with automation and orchestration of event management; advanced threat prevention using various forms of data enrichment and threat intelligence; and some endpoint device control capabilities, all of which are based on what type of licensing is purchased and how the system is configured.

Figure 10-8 shows an example of the CrowdStrike Falcon dashboard. The first key datapoint to notice is the heavy focus on detecting events. There are widgets covering the number of endpoints that have been infected by all accumulated events, the secure hash algorithms (SHAs) associated with identified malware, what type of tactics are used by the malware, and the most recent events. EDR operates similarly to SIEM solution with regard to events; the difference with EDR solution is that the scope is limited to endpoints.

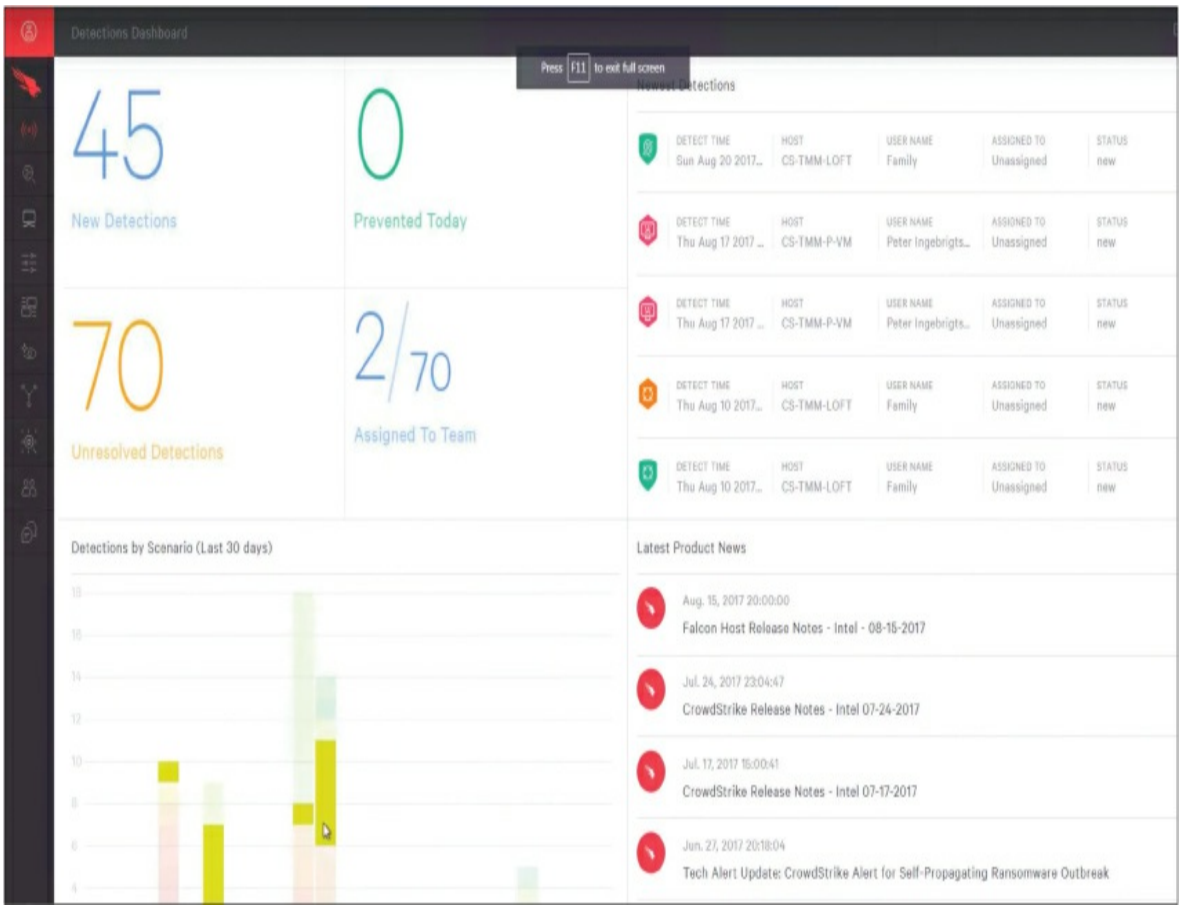

FIGURE 10-8 CrowdStrike Falcon Dashboard Example

Events are displayed in a graphic format in Falcon so that the analyst can track the lifecycle of the event from start to remediation. Figure 10-9 shows an example of an event in Falcon that started through Microsoft Outlook and launched through notepad.exe. Each step of the attack can be analyzed to help the SOC analyst understand what occurred, why the action is malicious, and what should be done to prevent additional actions taken by the threat—in other words, prevent an outbreak. This approach to developing a graphic view of each step taken by malware during an event is very common in EDR solution platforms.

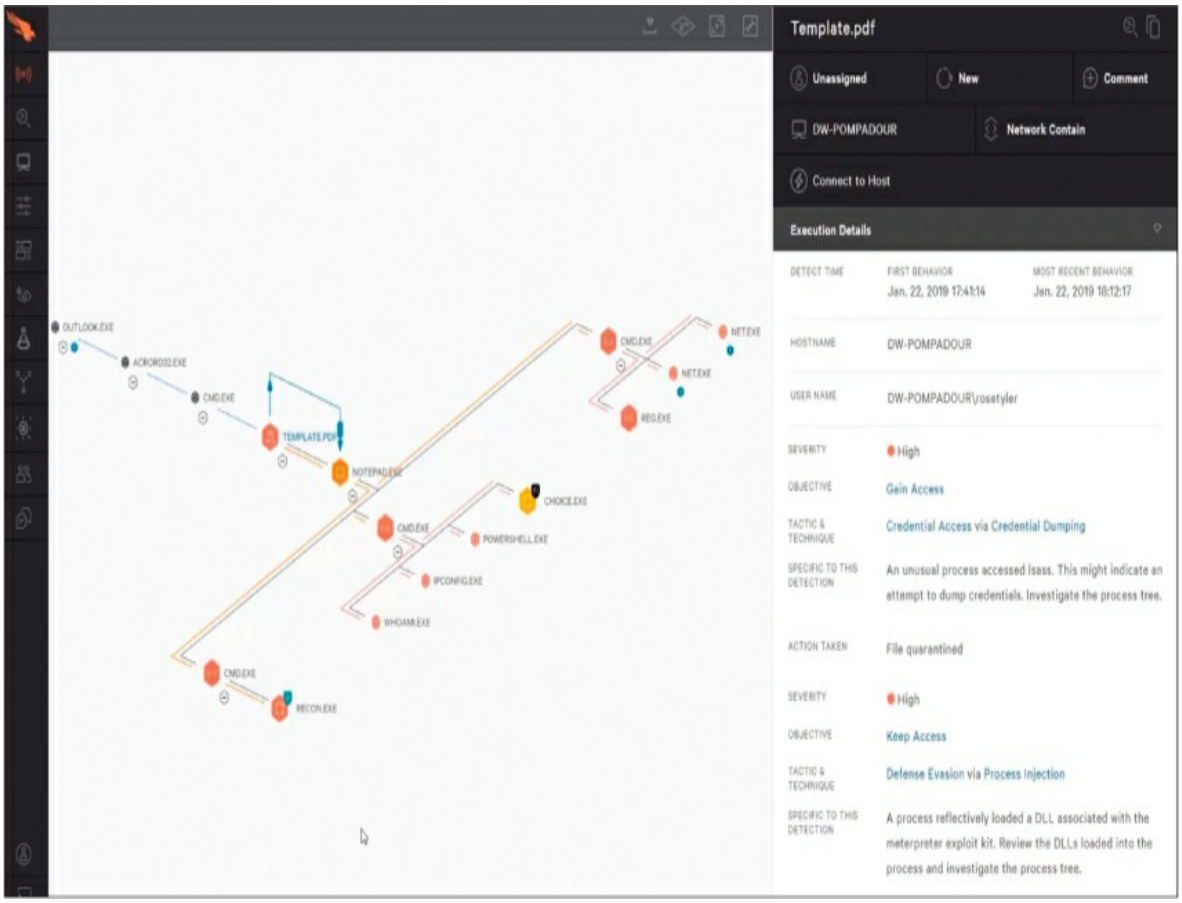

FIGURE 10-9 Falcon Event Graph Example

Actions can be taken leveraging the Falcon agent installed on the endpoint, including pushing scripts, manually pushing commands, or executing programs. From the interface shown in Figure 10-9, you can view details on why a step is potentially malicious by clicking into a notation in the attack timeline. Figure 10-10 shows an example of clicking into the details explaining why a file is identified as malicious, including relating it to a known threat actor. It is common for enterprise EDR solutions to provide much more detail about threats than host-based antivirus offerings because EDR solution vendors include in the solutions far more capabilities than signature-based pattern matching.

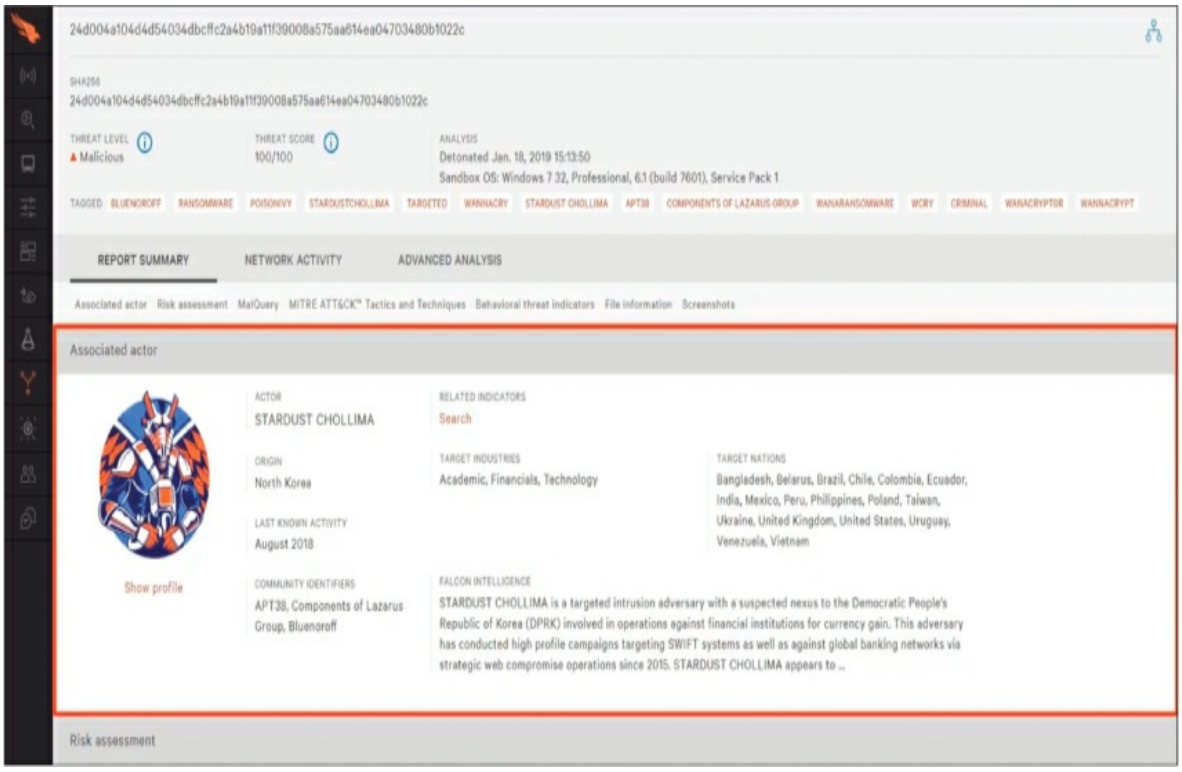

FIGURE 10-10 CrowdStrike Falcon Example of Event Details

Responding to endpoint threats is important, but EDR offerings are limited to endpoints. This is why many EDR vendors have looked beyond the endpoint to meet the market demand for capabilities offered in a SIEM/SOAR solution combination. As previously introduced, this new breed of EDR solution that can support more than endpoints is the extended detection and response (XDR) solution marketspace. From a capability standpoint, an XDR solution has very similar capabilities to the previously covered SIEM and SOAR solution capabilities, so I am not including an XDR solution demo here. In other words, XDR solutions offer the same general features as a SIEM/SOAR solution combination.

At this point, you should understand what SIEM, SOAR, EDR, and XDR solution are and what they can do based on the provided technology examples. The rest of the chapter covers concepts that apply to what both a SOAR and an XDR solution can offer. One important value these technologies can provide to your SOC is the capability to develop and enforce repeatable responses to specific events. To build repeatable responses, you need to create playbooks.

Playbooks

Every SOC with an incident response service has some form of planned response when an incident occurs that could impact the organization the SOC is protecting. That plan might be very basic and reactive, such as sending an expert to investigate the situation. More mature plans cover everything from how potential events are identified to who is notified, how new cases are tracked, what actions should be taken based on the event type, and what is needed to recover. As with any SOC service, higher-maturity services have repeatable processes that provide all members of the organization with a clear understanding of their responsibilities before, during, and after a security incident. In short, *a playbook is a representation of a standardized SOC process*.

A playbook is a linear-style checklist of required steps and actions that are followed to successfully respond to a specific incident type and threat. Playbooks not only support implementation of automation but also help to keep the incident response program consistent, predictable, and measurable, which enables organizations to meet and comply with regulatory frameworks such as NIST SP 800-61 R2 or the EU General Data Protection Regulation (GDPR). The contents of playbooks are a combination of what can be automated with manual tasks performed by humans and systems. Basically, a playbook is the entire end-to-end response to an incident. This is why I pointed out that automation is a subset of orchestration: orchestration involves everything in a playbook.

Playbook Components

The security industry uses common components or symbols when developing and representing playbooks. Components of a playbook are represented as parts of a workflow diagram so that any person can understand how the workflow should occur from start to finish. The list that follows provides an explanation of each playbook component. Be aware that components can vary slightly among playbook creators.

- **Initiating condition:** The event that triggers the start of the playbook. An initiating condition is required to know when to launch a playbook.

- **Process steps:** The actions performed during each step of the playbook. Steps can include generating responses, authorizing responses, and quarantining something, to name a few common examples. There must be at least one process step or a playbook doesn't provide any value.
- **End state:** The end goal and final step of the playbook. The end state represents the desired outcome based on predefined conditions that have been met.

Workflow diagrams use some variation of the symbols shown in Figure 10-11. Start and end states are represented by ovals, while rectangles and diamonds make up process steps. Steps are connected by arrows showing how the workflow should occur. The diamond represents decisions, such as the answer to a yes or no question, while rectangles are specific steps that are taken and must be completed before moving to the next step.

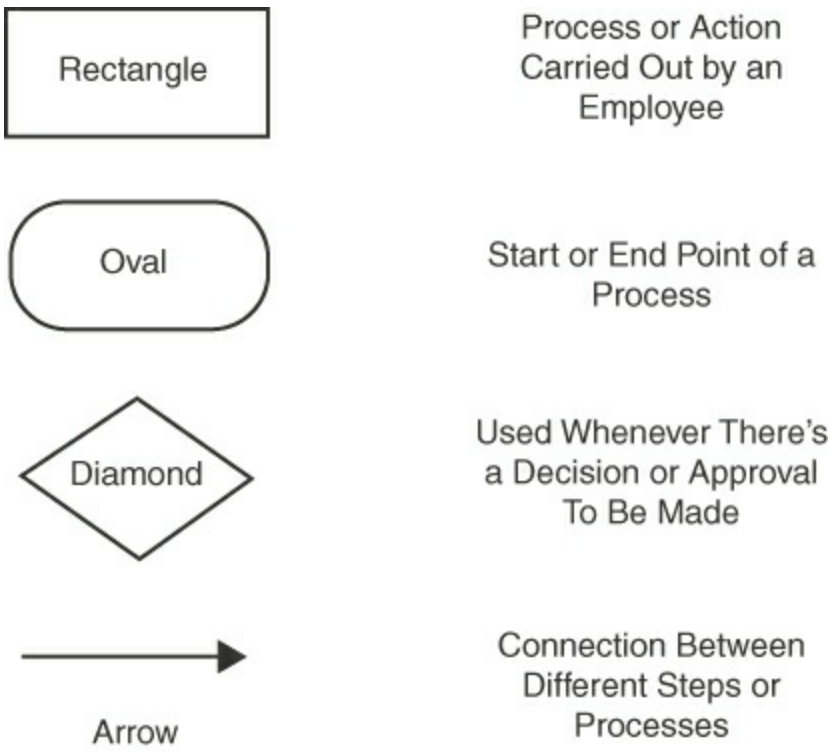

FIGURE 10-11 Workflow Symbols Defined

Constructing Playbooks

Well-built playbooks accommodate all situations, including when a step fails by sending the workflow back to a previous step until all steps have been accomplished to move forward. Figure 10-12 is an example of a playbook.

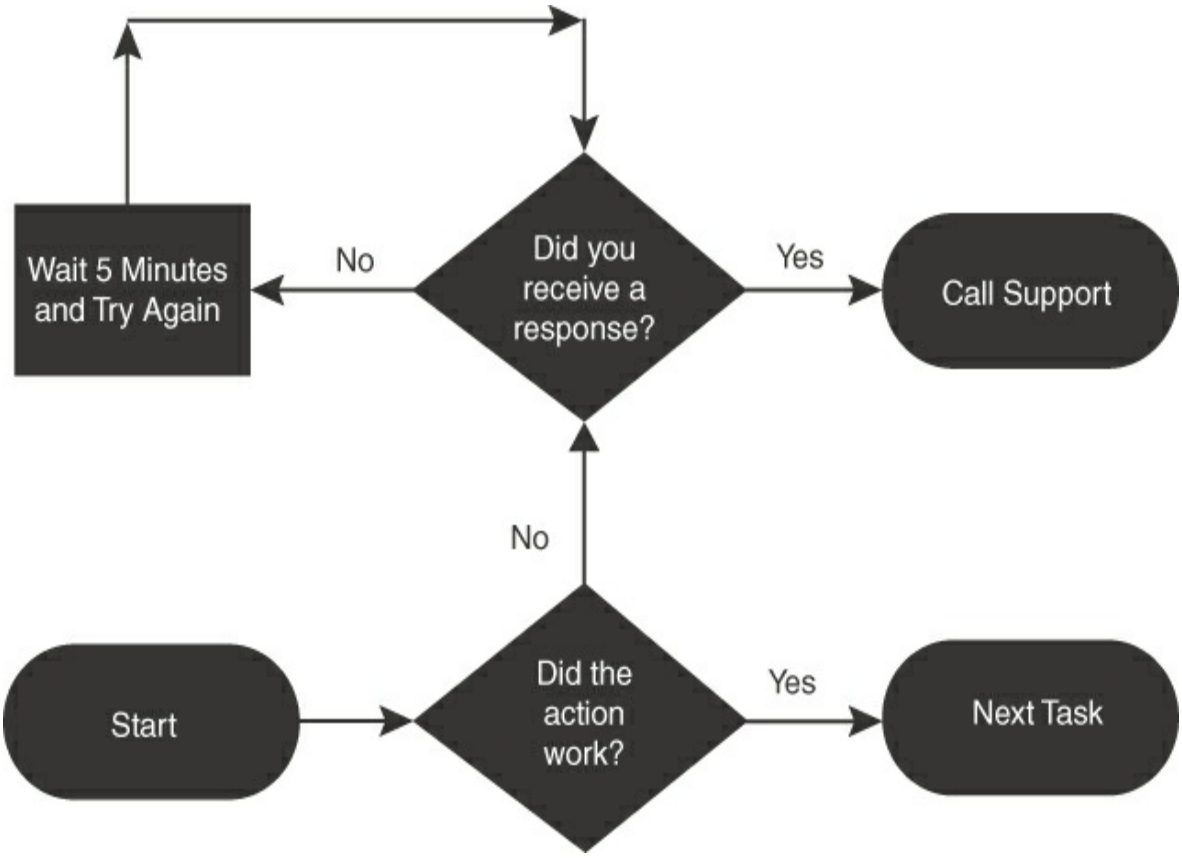

FIGURE 10-12 Sample Playbook Workflow

As shown in Figure 10-12, the question in the bottom diamond asks "Did the action work?" A response of yes sends the workflow forward to the next task, while a response of no sends the workflow to another question, "Did you receive a response?", which has another set of yes and no options. Workflows can have more than one start state and end state as well as infinite loops. As an example of more than one end state, Figure 10-12 shows that the workflow can lead to an end state of either calling IT support or proceeding to the next task. An example of an infinite loop is "Wait 5 minutes and try again"; until a response is received, the workflow continues looping in the wait period. The actions in the diagram in Figure 10-12 can be summarized generically as you do some action, and if it works, you do the next task; if the

action fails and you get a response, you call support; if the action fails and you don't get a response, you wait until you get an error message, and once you have that message, you call support.

Incident Response Consortium

The Incident Response Consortium (IRC) at https://www.incidentresponse.com is one of many resources you can use to find playbook templates and is a great resource I will use to teach you playbook best practices. I have already referenced the IRC in Chapters 3, 6 and 8. The IRC approach to developing playbooks is based on different parts of the lifecycle of responding to a specific type of incident. The following is a summary of the steps of IRC's playbook lifecycle. These are common steps within the lifecycle of developing and delivering SOC services.

1. **Prepare:** Actions taken to prepare the SOC for a certain type of incident
2. **Detect:** Actions and tools set up to detect a certain type of incident
3. **Analyze:** Actions used to analyze a potential type of incident
4. **Contain:** Actions implemented to prevent the spread or outbreak of an incident
5. **Eradicate:** Actions involved with removing the risk/negative impact associated with a specific incident
6. **Recover:** Actions deployed to return the organization back to normal operations post incident
7. **Post-Incident Handling:** Steps regarding lessons learned to improve the playbook

Playbook Examples: Malware Outbreak

Let's review a few playbooks found under the IRC's recommendations for handling a malware outbreak. Figure 10-13 shows the Prepare playbook for a malware incident taken from IRC's templates. This is a great starting point not only to learn about developing playbooks, but also to understand what is expected for a playbook regarding malware outbreaks. According to IRC's

format, there are different playbooks for other parts of the malware lifecycle such as detecting and analyzing malware; however, IRC recommends that you first prepare the SOC for dealing with malware following the “Prepare” focused playbook. This means that at this point of the malware response lifecycle, you do not have a service established within your SOC and will follow this prepared playbook to develop a malware response service based on IRC’s recommended process.

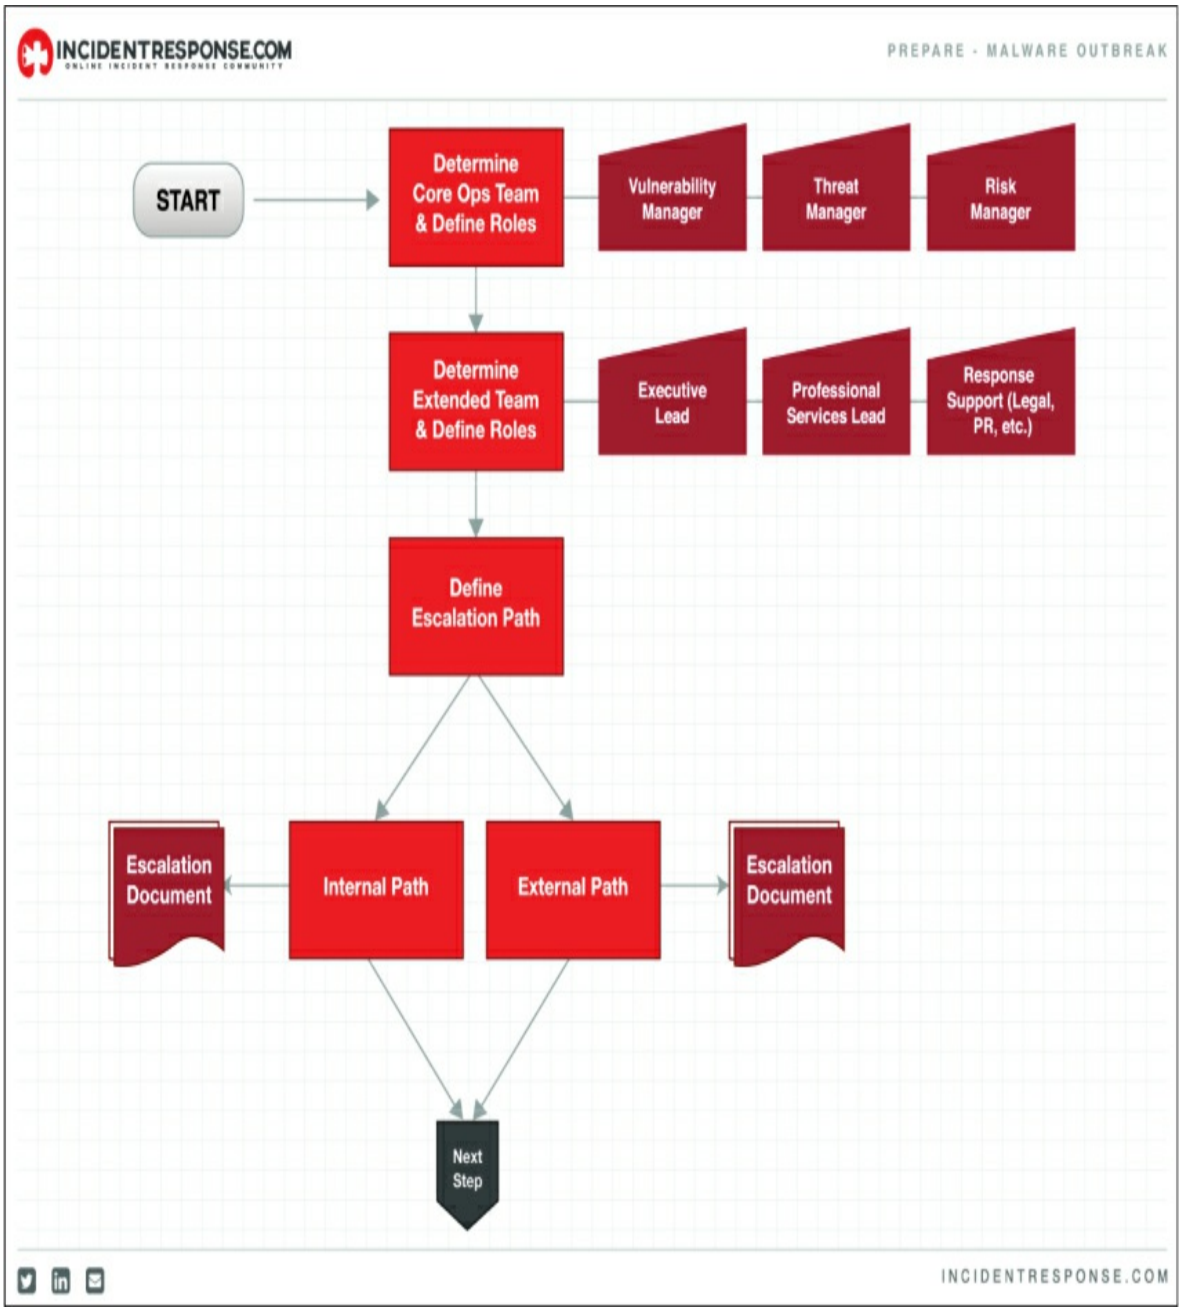

FIGURE 10-13 IRC's Prepare Playbook for Malware Outbreak

Looking at the malware prepare playbook, it starts with the task Determine Core Ops Team & Define Roles, in which the SOC assigns a vulnerability manager, a threat manager, and a risk manager. You must complete this step before moving forward. Taking this step first makes sense because a SOC service needs a leader before it can take any further actions. This playbook enforces this concept. Next, the playbook has the SOC Determine Extended Team & Define Roles, which again must be done before the SOC can move forward. This task recognizes that the manager of each service identified in the first task needs the support and representation of other departments in the organization to help develop a coordinated response to a malware outbreak. Next, this playbook tasks the SOC with defining escalation paths, which will leverage the team that was previously constructed in the previous tasks. In summary, this playbook states that preparing escalation documentation is the focus of the Prepare step of the malware outbreak playbook. These documents will identify who is involved with a malware outbreak program and what their escalation path would look like—in other words, determining when each party is engaged and their assigned roles as the malware outbreak playbook is further developed. The IRC is basically saying that roles must be assigned and documented before any other actions can be taken.

Next, let's move forward in time regarding the malware lifecycle playbook list and dig into the Analyze malware outbreak playbook, shown in Figure 10-14.

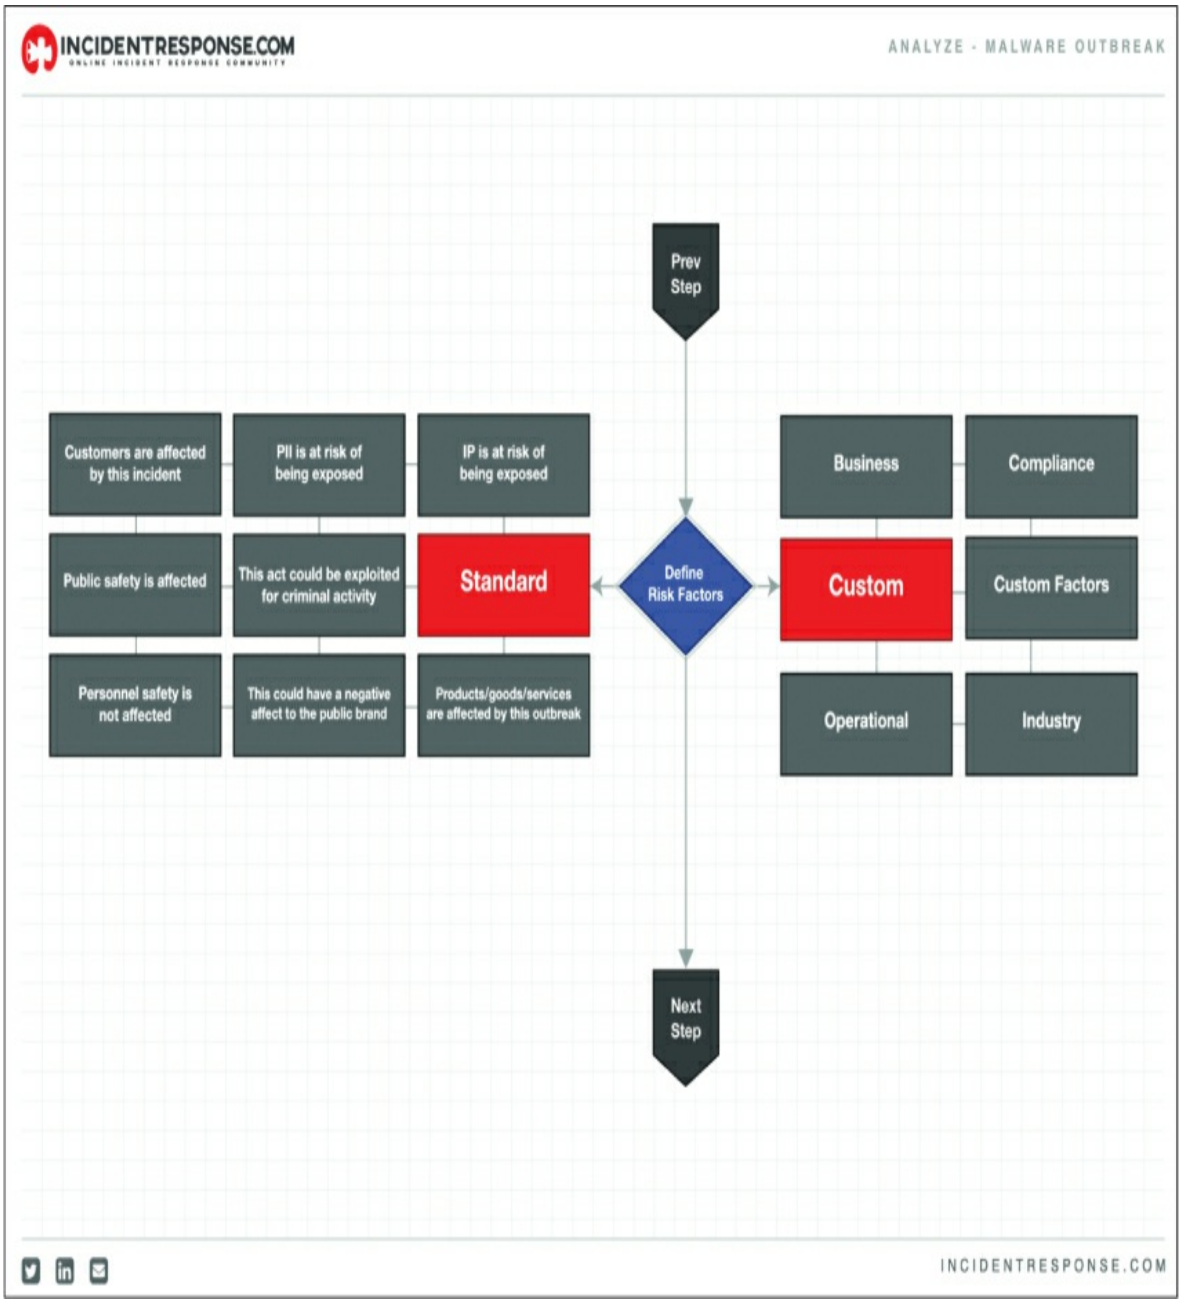

FIGURE 10-14 IRC's Analyze Playbook for Malware Outbreak

The first step of the analyze playbook is executing the previous playbook covering how to detect any potential detected threats. This means the assumption is that you have detected a possible malware outbreak and you are now launching this playbook to start the analysis process to determine if it is a real threat or false positive. This playbook starts with a diamond that has the analyst define the risk factors associated with the detected threat. On the

left are standard industry factors to consider, while on the right are what this playbook calls “custom” checks, which can be based on specific factors such as operation, industry, or compliance factors. Keep in mind that the risk evaluation process can be manually computed or automation can be used to speed up the process. Regardless of approach, this analyze playbook defines what factors must be considered and calculated before the analyst is permitted to move to the next playbook, which would be how the SOC would respond through containing the threat.

Key point

Check out https://www.incidentresponse.com/playbooks/ to see more playbook templates from IRC.

In addition to playbook resources such as the IRC, playbook templates are available in some security tools. Regardless of the source, no template will represent everything required for your specific needs. You are expected to customize a template based on your unique technology, architecture, available skillsets, policies, and procedures.

Playbooks are a critical prerequisite for automation. At first, all playbook steps are manual and should be well vetted to ensure the process continuously delivers a desired result. The IRC includes a final playbook within the malware outbreak playbook category called Post-Incident Handling, which has the focus on improving previous playbooks based on lesson-learned actions targeting discovering what works and doesn’t work with playbooks that are put into action. Maybe adjustments within the steps used within the Detect playbook could be made. Maybe new tools are needed to improve the Eradication playbook. No playbook will be perfect, but as repeated success is seen, that playbook will be ready to consider applying automation.

Note

Mature SOCs can have dozens of playbooks as procedures are formalized, converted into a playbook, and implemented as part of a

SOC service. I recommend to start creating playbooks for fundamental services, such as responding to common events, because your SOC will have experience with the event and be able to outline what the playbook should look like. Examples include responding to malware outbreaks, responding to stolen passwords, and dealing with unauthorized devices accessing the network.

Automation

Playbooks can be manually implemented; however, the real value from developing playbooks is derived by applying automation when applicable. Security-based automation can be defined as the execution of a series of actions with the purpose of detecting, investigating, and remediating cyberthreats with or without human intervention. The result of successfully applying automation is that it frees up your SOC staff from doing tedious and mundane tasks so that they can prioritize their time on more complex tasks. Automation can also improve response time and guarantee action is taken on specific events. For example, if a system connects to the network, automated responses can include initiating a vulnerability scan of the system and, if a critical vulnerability is found, denying or limiting that system's access to the network and sending an alert to the SOC. Without automation, manual execution of the playbook would require that a SOC member identify that a new system is connecting to the network, provision it limited access, manually initiate the vulnerability scanning process, and, if a vulnerability is found, pass the alert about the vulnerability to the vulnerability management service and wait for them to respond, all of which delay the new system's full access to the network.

I recall earlier in my career working for a U.S. government agency that stated "All new devices must be scanned by IT"; however, IT would only scan systems when either they identified them or when a new system owner called IT regarding having their device scanned. Due to the lack of automated enforcement, many new systems would connect and not be scanned. Automation removes the SOC from having to monitor for newly connected devices, launching vulnerability scanning against those systems, and applying

network access based on the results of the vulnerability scan. Figure 10-15 represents this example workflow to accommodate new systems connecting to the network as a playbook. The icons in red represent what is being automated in this example. Notice in this example that most steps can be automated, leading to a lot of saved time!

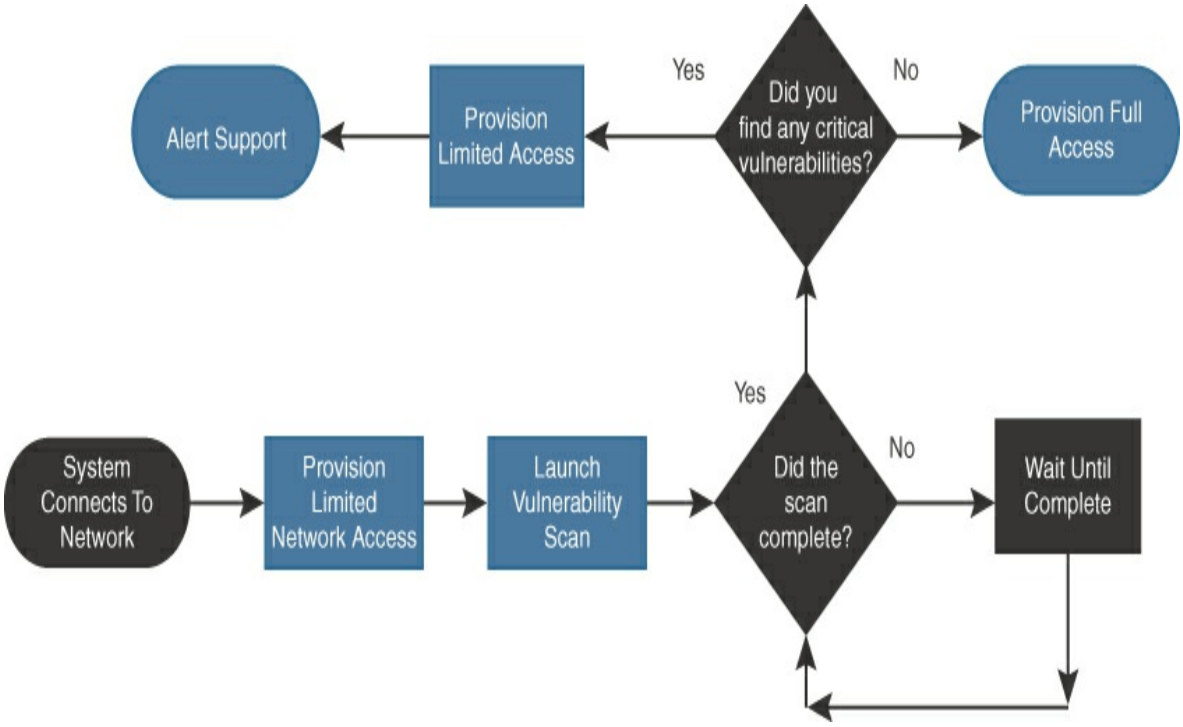

FIGURE 10-15 Example of Automated Playbook for Vulnerability Compliance Enforcement

Automating Playbooks

There are different approaches to automate each step represented in the Figure 10-15 playbook. Provisioning access is commonly automated using a network access control (NAC) technology. Vulnerability scanning can be triggered based on the event of a device connecting to the network, identified either through monitoring switch connection states, such as leveraging link-up traps, or through direct integration with a NAC technology. A SIEM, vulnerability scanner, and NAC technology can be configured to generate alerts based on any of these actions, and those alerts can include notifying

parties responsible for delivering the services described by this playbook.

Figure 10-16 shows an example of this configuration with Cisco Identity Services Engine (ISE), a NAC technology, working with Nexpose, a vulnerability scanner by Rapid7. Figure 10-16 shows a Cisco ISE policy stating that if a system connects to the network and Rapid7 Nexpose identifies that it has a threat ranked as a CVSS score of greater than 7, then enforce the Quarantine policy. The Quarantine policy could mean moving the system to a different network, applying access control lists to limit what the system can connect to, or even removing the system from the network. This setup automates the response to when a system has a vulnerability, which saves the SOC time and ensures the same response is issued anytime this situation occurs.

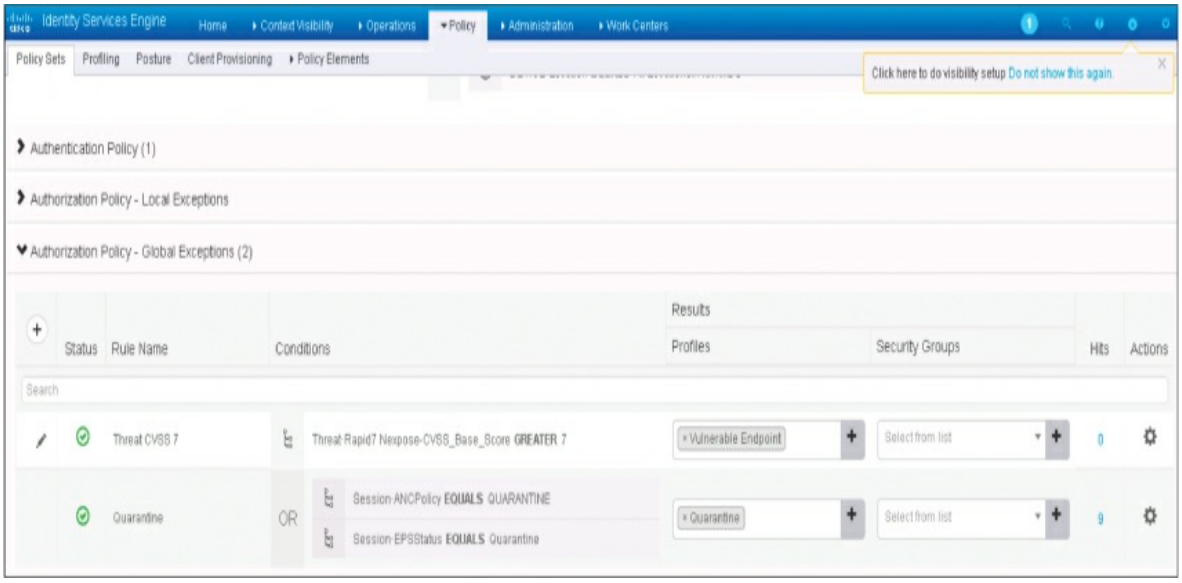

FIGURE 10-16 Cisco ISE Configured with Rapid7 Nexpose Example

Another example is using the same Cisco ISE network access control technology to automate the response of threats seen by an internal IDS/IPS. Figure 10-17 shows rules in Cisco Firepower, a next-generation firewall with IPS capabilities, configured to be automatically enforced by Cisco ISE. Figure 10-17 shows that Cisco Firepower is set up to look for malware and, if a malware event occurs, automatically instruct Cisco ISE to move the associated endpoint to the Quarantine policy. Once again, the SOC can save time by automating this response and can be assured that the response is

repeatable, regardless of the time of day. This ensures trust that the same response is consistently delivered.

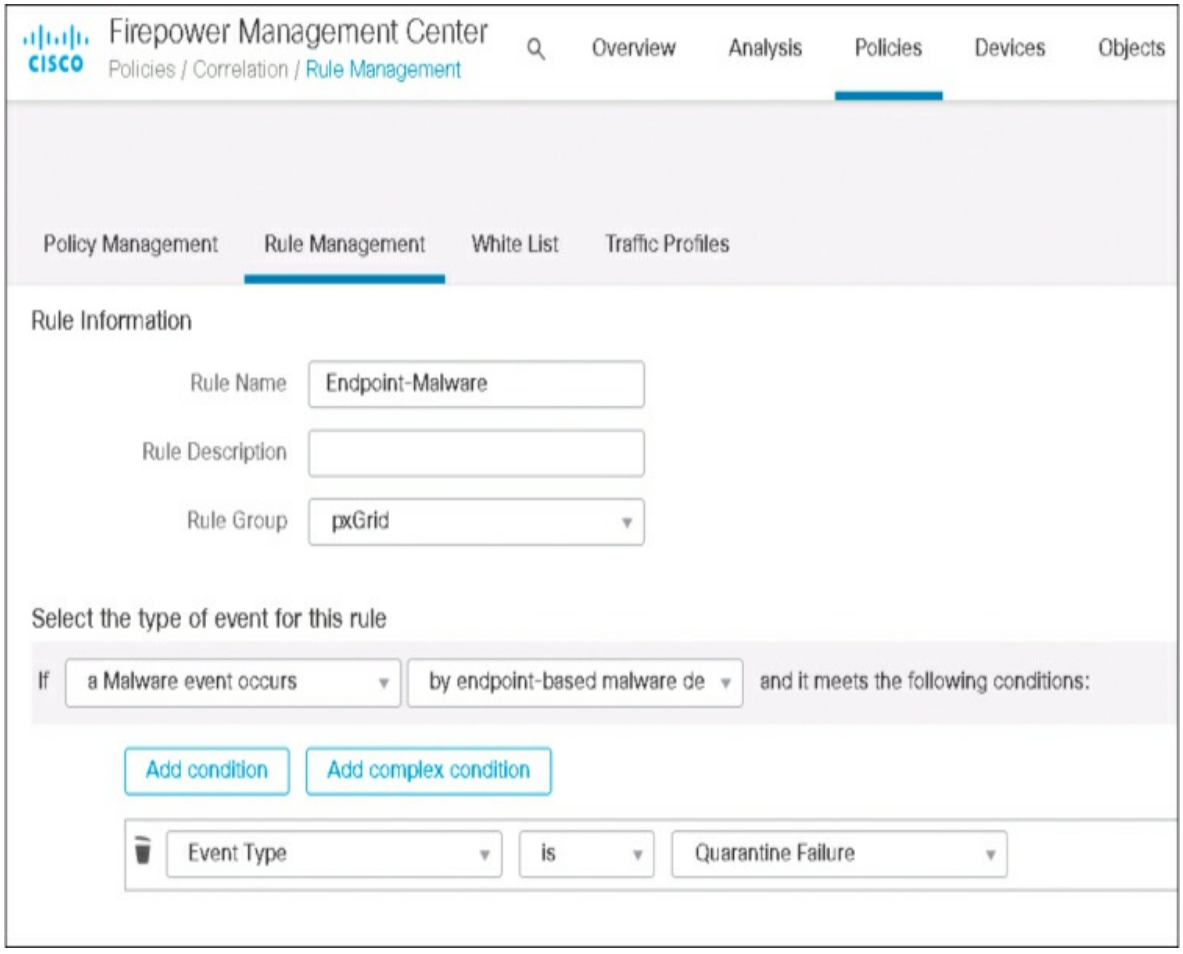

FIGURE 10-17 Cisco Firepower Configuration Rule Example

Common Targets for Automation

Almost every task can take on some form of automation based on how repeatable and complex the task can be. In fact, instead of restricting automation to specific features that the vendor must build and maintain, many vendors are including application programming interfaces (APIs) and other open platforms that support almost any type of automation to occur through specialized programming known as DevOps. (DevOps was introduced earlier in the context of SOAR and will be covered in more detail

later in the "DevOps Programming" section.) However, there are several common tasks that are ideal candidates for automation.

The following tasks are prime targets for automation:

- Identity and access management
- Patch management
- Unsophisticated malware detection
- Data protection
- Reputation lookups
- Risk scoring
- Blocking users
- Reporting thresholds
- Notification and task assignments
- Launching remediation tools
- Automating specific responses such as blocking or shutting down services

Automation Pitfalls

In contrast to the previous list, there are certain tasks you will not want to automate. More complicated tasks that require decisions based on human behavior or generally complex tasks will have too many variables to map out into a playbook. An example is responding to social engineering attacks. There are certain defense mechanisms that are possible to automate; however, social engineering is based on deception, which can be difficult to predict. Social engineering attacks can target email, social media, or even occur over a phone call. Penetration testing is another example of a task that continuously adjusts based on data obtained from the attacker's reconnaissance and must consider both technology and human factors of the target, leading to hundreds of potential variables that are impossible to predict.

My recommendation is to attempt to create a playbook diagram of a process

before considering automation. You will quickly find if a workflow is too complex for automation based on the number of outcomes required to completely capture what needs to be covered within the playbook. Looking back at the IRC playbooks, you should take the same approach they prescribe regarding breaking complex playbooks into smaller, more focused playbooks. By taking that approach, even the most complex tasks can be isolated into smaller, more digestible playbooks.

Playbook Workflow

Building automation is based on workflow. A *workflow* is a combination of multiple tasks, and *automation* is executing those tasks in a more efficient manner. Relating the workflow concept to playbooks, think of a workflow as how all the tasks (regardless of whether they are automated) are executed, ensuring that the playbook is properly followed so the playbook generates repeatable results. Automation of each step in a workflow can be accomplished using raw scripting; however, some SOAR tools offer a drag-and-drop, template-based approach. You can also leverage existing integrations, such as the Cisco ISE, Rapid7 Nexpose, and Cisco Firepower examples previously covered. Drag-and-drop templates found in many SOAR and XDR offerings provide a simple visualization of workflow, which allows the analyst to choose which steps of the workflow can be automated.

Workflow Example 1: Cisco SecureX

This example looks at building a workflow in Cisco SecureX. Cisco SecureX offers SOAR capabilities that include developing and automating playbooks. The left side of Figure 10-18 shows a workflow of a playbook with the purpose of sending a URL to be evaluated by a cloud security sandbox service called Cisco Secure Malware Analytics (formally Cisco Threat Grid).

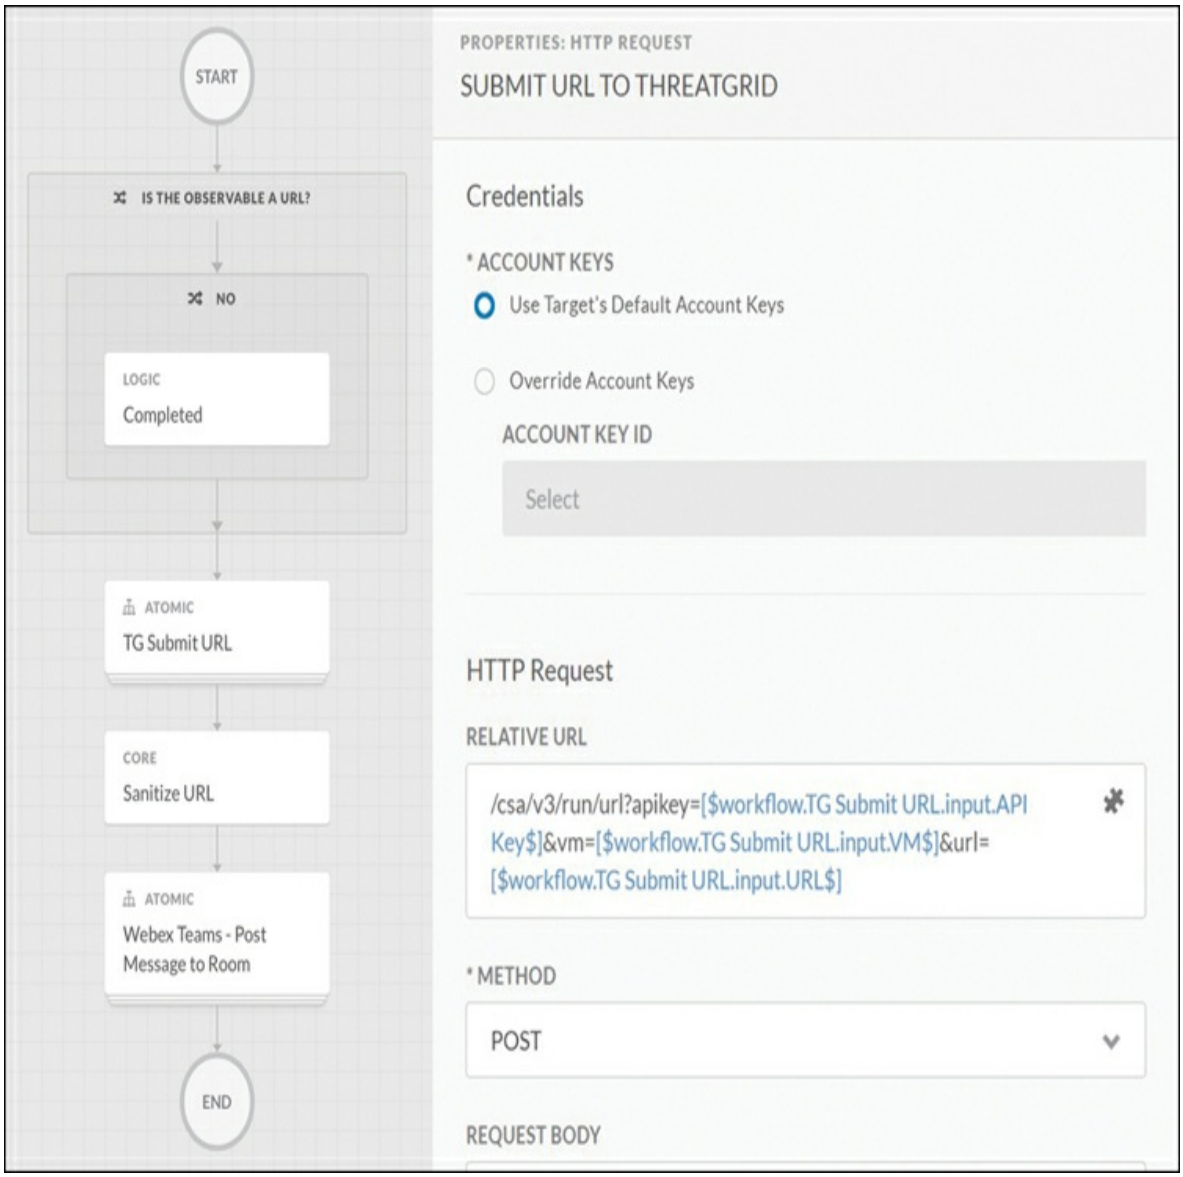

FIGURE 10-18 Cisco SecureX Orchestration Example

This workflow starts with validating if the data being submitted is a URL. If it is, the workflow continues to the next step. If the URL provided is not a URL, a logic completed step is launched, which kills the workflow. The analyst configured this as a response to when the wrong data type is inputted to avoid the workflow crashing or causing an unwanted response. Keep this concept in mind as you review developing workflows. It might seem tedious to add checks for the wrong input during each step, but the end result will be a much more robust and effective workflow.

If the right data is provided, a URL in this case, that URL is submitted to Cisco Secure Malware Analytics using the authorized account information. Think of this step as automating the task of a SOC analyst copying a URL, logging into Cisco Secure Malware Analytics, navigating to where a URL is inputted to be evaluated, pasting that URL, and submitting it to be evaluated. This entire process can be done manually, but why do that if it can be automated, ensuring the same response occurs every time? Automation not only speeds up the process but removes the possibility of the analyst mistyping the URL being submitted, mistyping the login information to access Cisco Secure Malware Analytics, or clicking the wrong thing while trying to find where to submit the URL. Automation makes this entire process easier!

Finally, once the URL is submitted, a result is pulled and sent to a SOC social space, such as a message board or dashboard dedicated to monitoring alerts. Once again, this automation saves the analyst time regarding waiting for the response, copying the results from Cisco Secure Malware Analytics, opening the SOC social space, and pasting the results into that space. Each of these smaller tasks might not seem like a huge effort to perform manually, but think about having to do them 500 times in one day. If you consider all the steps and possible issues that could occur, regardless of how minor they can be, automating this entire process can add up to hours of time savings, resulting in a more efficient SOC service. A more important result is removing these tedious tasks from an analyst's workload, not only making the analyst's job much more enjoyable but also enabling the analyst to focus on more complex tasks that require human intervention. I've found over time that the reason many unhappy SOC analysts don't like their job is that the majority of their tasks are repetitive and mundane!

On the right side of Figure 10-18 is the configuration of the workflow step "TG Submit URL" representing how to configure and make the process of submitting the URL to Cisco Secure Malware Analytics function properly. In order for this step to work, a few actions need to be taken correctly, such as logging into the Cisco Threat Grid system to prove you are authorized to submit URLs. This can be accomplished by using an API that automates the process of proving your SOC has a subscription to a commercial lookup service. Another action that must occur is going to the proper web source and delivering the URL in question. Configuration options for workflow widgets

in enterprise SOAR products are built this way to simplify creating automation of steps of a playbook with properly functioning logic. Similar to programming code, vendors develop workflow templates that include validation of what is entered to reduce the chance that a workflow will not provide the desire results. Without these checks, inputting something other than a URL for this example workflow would cause unwanted results.

Figure 10-19 shows an example of the options to validate that a workflow configuration is correct in Cisco SecureX.

FIGURE 10-19 Example Workflow Validation and Run Options

Workflow Example 2: Splunk Phantom

Playbook workflows function in a specific order. It is important to have each step of the playbook launch in the correct order or you will not get the expected result. Many times, variables in a workflow are dependent on items being executed before the variables can function properly. Looking back at the last workflow example, if the login credentials were sent to Cisco Secure Malware Analytics prior to inputting the URL that needed to be researched, the remote source (Cisco Secure Malware Analytics) would not see the login and would view the request as one coming from an unauthorized party, leading to an error message.

Figure 10-20 shows Splunk Phantom's use of workflow when running a playbook. Notice the timeline at the bottom of this playbook, which shows when steps of the workflow are executed in the playbook. When a workflow step is complete, a checkmark is added, indicating that step's data can be accessed to see the results as well as be used by future steps in the playbook.

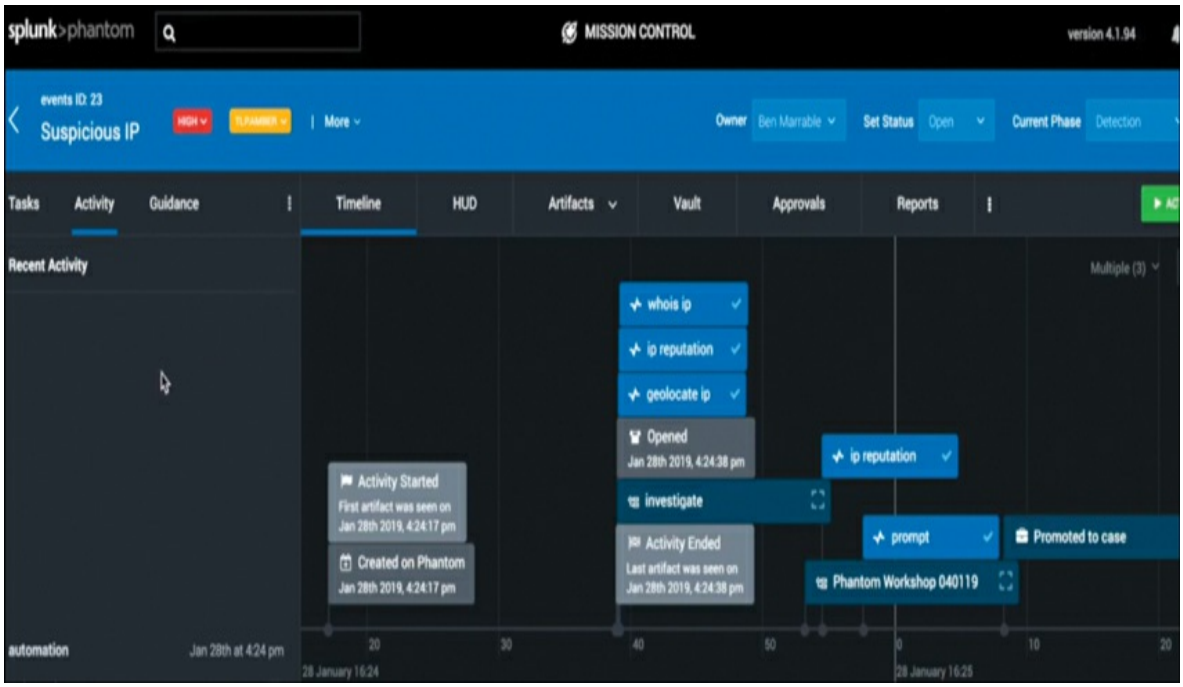

FIGURE 10-20 Splunk Phantom Workflow Execution Example

Many enterprise tools that offer template-based workflow programming are designed with the goal of simplifying open-source, industry-recognized programming languages. This brings us to another approach to accomplishing workflows, which is building the automation from scratch or modifying existing workflow templates using the code they make available. To use this approach, you need to understand workflow programming (the ability to configure servers using a programming language).

If you were hoping orchestration technologies would eliminate the need for coding skills, you're going to be disappointed, because the orchestration market is encouraging *more* programming requirements for the future IT professional. This leads us to the conversation about DevOps.

DevOps Programming

There are many programming language options for developing an application. According to Applitools (https://applitools.com), the five most common programming languages used to build simple scripts are Java,

Python, JavaScript, C/C++, and Perl. Any of these languages can be used to develop scripts that can automate existing steps or add new functionality to improve steps within a playbook. By searching the Internet, you can find hundreds of tutorials, program snippets, and even drag-and-drop tools that simplify the programming process, all aimed at helping you develop a program.

> **Note**
>
> My personal opinion is that Python is the best popular programming language to start learning if your goal is to use your skills in a SOC to automate security tools. Whatever language you choose to learn, the underlying concepts are applicable to all languages.

The focus of traditional programming languages is to build programs. This is not the same focus as what IT operations hope to achieve with programming. IT operations care about configuration management—using programming to automate workflows by enabling tools to work together in a more streamlined fashion. This is why the industry has created a separate practice that focuses on orchestration representing the combination of software development and IT operations, commonly called *DevOps*.

DevOps bridges the gap between development and operations, which is key to improving playbooks. Think of DevOps as using programming skills to apply automation between steps of a playbook to improve the workflow of that playbook. This means you want to use APIs and other capabilities that already exist in tools rather than writing a program from scratch, which leads to a whole different set of coding skills and software development lifecycle.

In the DevOps world, tools like Puppet and Ansible are better suited for creating automation compared to writing a program using a traditional programming language. There are also platforms such as wikis that are built from the ground up with workflow in mind and are much better options than developing a similar platform from scratch. DevOps aims to customize what has already been created for your needs (rather than creating from scratch) as well as improve the data management process. This focus addresses some of

the biggest challenges with the traditional data management process that tends to be slow and siloed. Let’s look more closely at those challenges.

Data Management

Data management is the process of ingesting, storing, organizing, and maintaining the data created and collected by your organization. Chapter 5 covered different data formats, how to collect data, and how to optimize it for your SOC. Effective data management is a crucial piece of running a successful SOC. The goal of proper data management is to make the data accurate, available, and accessible to all required parties.

Figure 10-21 represents a strong data management practice. First, the SOC must identify what data is required for its services. The combination of this data that the SOC can use can be viewed as the SOC’s data architecture (all of the desired data that is useful to the SOC). The SOC uses data models to map workflows and relationships in datasets so that information can be organized to meet the SOC’s needs. As data is generated or collected by the SOC, that data is processed and stored in a database. The database holding the SOC’s data can be installed within the SOC or provisioned through cloud storage. As more data is collected and stored by the SOC, data is integrated by the SOC tools, which allows for the SOC to perform analysis across multiple datasets. To ensure data accuracy, the SOC must check for errors or inconsistencies so they resolve any problems using data cleaning tasks. Finally, any data governance policies developed by the SOC are applied to ensure data is stored and used in compliance with any requirements as well as kept consistent across systems.

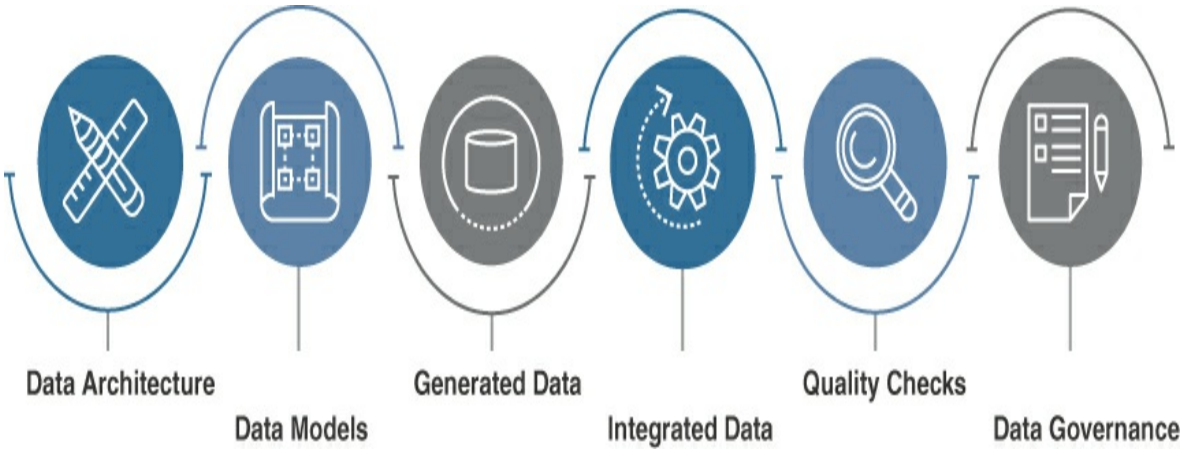

FIGURE 10-21 Data Management Model Example

Traditional methods to store and share data tend to be slow due to protocols used for sharing data and limited capabilities regarding how different tools share data. Today's data tends to not be homogeneous, meaning it doesn't fit easily into a relational schema based on rows in tables. Modern applications must process data that includes records, documents, video, text, and semantic triples in a variety of formats including XML, JSON, txt, and binary. The SOC needs a way to overcome data format challenges and quickly abstract what is required for the mission at hand. Data also needs to be shared more quickly and more often to allow multiple systems to share their insight on an event.

DevOps brings agile operations to the SOC, which can greatly improve the data retrieval stage in the data analysis process. Data models can be enforced using a data modeling language that allows for any data type to be converted into an understandable text-file format. Data integration can be dramatically improved by enabling tools to share with other tools rather than forcing all data to be centrally stored before it can be retrieved. Automation can be used to remove steps in the collection, processing, storage, and sharing steps, which has the cumulative effect of reducing the required data custodian tasks involved with managing large datasets. As a result, the SOC can work with larger datasets more quickly, leading to more productive SOC services.

There are a few fundamental DevOps concepts that you need to be familiar with to properly grasp the value of DevOps. First, you need to know how DevOps can work with different datasets. This includes understanding file formats and data modeling languages. You also need to know how

automation can be used to enable tools to share data with other tools as well as how to reduce manual steps when collecting, processing, storing, and sharing data. In the sections that follow, you will learn how to build your own DevOps lab as well as consult many DevOps resources you can use to gain a better understanding of how DevOps can improve your SOC.

The first fundamental DevOps topic you need to understand is text-file formats. DevOps is all about working with existing tools, which means that you will be sharing data and commands between tools. You must understand how to communicate with another tool in order to have the receiving tool understand what is being said. This makes text-file formats an essential ingredient of DevOps programming.

Text-File Formats

DevOps requires tools that have an open platform, meaning the tools are based on open standards and thus can be integrated with other tools that are based on those standards. In the past, many security vendors didn't support open platforms, meaning their tools would not work with any third-party tools. Closed platforms are becoming extremely unpopular in the security community because every organization needs to leverage multiple security vendors to accomplish all of their security needs. SOCs expect security tools to support an open platform and will specify open platform support as one of the requirements for purchase.

In order for tools to be open, they must understand how to communicate with each other. The technology industry has developed standards for this purpose, enabling vendors to know what to support in order to be accepted as an open platform that can function within the DevOps community. One type of standard specifies how data must be formatted for tools to be able to communicate. A data format is what it sounds like, which is the arrangement of data fields for a specific shape. Let's say I want to present a file that gives a quick summary of this book to my daughter. That file format could look like the following, which is easy for my daughter to read and understand:

    The book ID is 1

The author is Joey Muniz

The title is The Modern Security Operations Center

It is a technology book

A summary of the book is a technical journey to understanding a SOC

This file format is easy for any human to read, but what about computers? Computers consider aspects such as whitespace, the case of characters, and many other factors that humans don't care about. This data can cause a computer to see this file different than intended. This is why it is important to understand the data format being used. By understanding the format, the system will know how to handle everything from text to whitespace as well as have expectations for certain key words and characters.

Next, let's look at the most common data formats used in DevOps environments.

Common Data Formats

In the DevOps world, three common formats that computers understand for data are XML, JSON, and YAML. It doesn't matter which option you use as long as the systems in play support it. If a system supports all three formats, then the decision of which format you should use comes down to which you are most comfortable using. There isn't a best option—the choice between using XML, JSON, and YAML is based on if the systems being used support the format and which you decide to use.

Regardless of the data format you choose, expect to see the following concepts:

- **Objects:** Objects represent characteristics of something you are referencing. For example, an object can be this book and characteristics are that it's a technology book and its author is Joey Muniz.
- **Keys/values:** Keys label a value type, while the value itself is the data. An example is the term *time*, the value for which can be 4:20.

- **Arrays (or list notation):** An array is a list of options. For example, the recipe for an apple pie can be a list of ingredients that include sugar, apples, flour, butter, cinnamon, and a pie crust.

The sections that follow review each of these common data formats.

XML

Extensible Markup Language (XML) is a very popular method to use to format data. As its name implies, XML was designed to be extensible, meaning generic enough for a wide variety of applications. The look and feel of XML should make you think of a web page design, as that was XML's original intent when created. XML is somewhat human readable and simple to create. A "< >" tag tells the operating system to open a section in XML and a "</>" tag represents to close a section. XML is a common way to encode data for APIs, which makes being able to read XML useful to IT pros looking to master automation.

The following code converts the plaintext file example I created for my daughter into the XML data format. Notice the list concept is represented in XML as <tag> to start the list and </> to end the list.

Click here to view code image

```
<? xml version="1.0" encoding="UTF-8" ?>
   <books>
      <book bookID="1">
         <lead>Joey Muniz</lead>
         <title>The Modern Security Operations Center</title>
         <genre>Technology</genre>
         <desc>A technical journey to understanding a SOC</desc>
      </book>
   </books>
```

JSON

JavaScript Object Notation (JSON) is a popular option for encoding API data. If you query a RESTful API using HTTP(S), the response you see will likely be in JSON format. Even though "JavaScript" is in the JSON name, JSON works with a wide variety of languages and applications. The format of

JSON is different from XML but also is coded in a human-readable format. JSON heavily uses curly brackets ({ }) rather than the angle brackets (< >) used in XML, but the overall data structure is very similar to XML. JSON is also somewhat human readable and easy to understand and format.

The next bit of code is the same file example from the previous XML section but now formatted in JSON. Notice how the key value uses the ":" tag and lists are indicated in JSON by square brackets ([ ]).

Click here to view code image

```
{
   "books":{
      "book":[
      "author":"Joey Muniz",
      "title":"The Modern Security Operations Center",
      "genre":"Technology",
      "desc":"A technical journey to understanding a SOC",
      "_bookID":"1"
      ]
   }
}
```

YAML

YAML (representing the recursive acronym YAML Ain't Markup Language) is a much more human-friendly data serialization standard for all programming languages based on the makeup of bits and pieces of other languages. This means YAML can be used with other programming languages, typically with the purpose to write configuration files. YAML is a superset of JSON, meaning it can do everything JSON can do. To break down how YAML relates to other popular programming languages, YAML uses scalars, lists, and associative arrays based on how they are used in Perl. YAML uses the document separator ---, which is also used in the Multipurpose Internet Mail Extensions (MIME) standard. YAML escape sequences are similar to C and whitespace wrapping is done exactly like you would do with HTML. Due to these overlapping characteristics with popular programming languages, YAML can be used with nearly any application that needs to transmit and store data.

The next big of code is converting the previous XML and JSON example

code into YAML.

Click here to view code image

```
---
- books:
      book:
- author:Joey Muniz
- title:The Modern Security Operations Center
- genre:Technology
- desc:A technical journey to understanding a SOC
- bookID:1
```

I'll further explain YAML based on how it is used by Ansible, which is a tool covered later in the "Ansible" section. The first key concept to know about YAML is its core components. The contents of a YAML file define a single data structure that is composed of nested nodes, which means the whitespaces matter. The most common special characters in YAML you need to know are the following:

- **:** Used between key/value pairs
- **-** Denotes a sequence entry, meaning a list item
- # starts a comment

Note

There are other special characters in YAML that you need to be aware of so that you don't enter one incorrectly and cause YAML to fail to parse your data. See https://www.tutorialspoint.com/yaml/yaml_syntax_characters.htm for more information on YAML special characters.

The basic structure of a YAML file is a hash map, which is a data structure that implements an associative array (a structure that can map keys to values). An example of this hash map/key value system is mapping Jenny to the value 8675309. When I reference Jenny, I can abstract the value 8675309, making 8675309 a characteristic of Jenny. This concept applies to most data formats, including XML and JSON.

Much like other programming languages such as Python, YAML uses indentations to denote a change in scope level, which makes whitespaces significant. You can see a similar characteristic in the previous XML and JSON examples with a subset concept, such as the title of the book is indented with whitespace while the main topic (books) is not indented, allowing the understanding that the title is part of the book concept.

The following is a second YAML code snippet example that demonstrates these basic YAML programming concepts. Notice the - depicts a sequence entry, : is used between key values as well as to start the **do** statement, and # is used for comments to explain what each part of the code does.

Click here to view code image

```
#This program will divide one number by another
- divider:
    do:
      divide:
        - dividend: ${input1}
        - divisor: ${input2}
    publish:
      - answer: ${quotient} #Answer represents the result of the
division
    navigate: [{ILLEGAL: FAILURE}, {SUCCESS: printer}]
```

Let’s look at a third example YAML file that can be run in Ansible to better understand YAML programming logic and see a more complicated YAML program. This YAML program represents a playbook that has two parts. The first part focuses on updating web servers, while the second part updates database servers.

Click here to view code image

```
---
- name: Update web servers
  hosts: webservers
  remote_user: root

  tasks:
  - name: Ensure apache is at the latest version
    ansible.builtin.yum:
      name: httpd
      state: latest
  - name: Write the apache config file
```

```
    ansible.builtin.template:
      src: /srv/httpd.j2
      dest: /etc/httpd.conf

- name: Update db servers
  hosts: databases
  remote_user: root

  tasks:
  - name: Ensure postgresql is at the latest version
    ansible.builtin.yum:
      name: postgresql
      state: latest
  - name: Ensure that postgresql is started
    ansible.builtin.service:
      name: postgresql
      state: started
```

Notice that --- starts the program. The next part calls out variables, also known as hash maps. For example, when the remote_user info is referenced, the value of “root” is pulled. The tasks section tells the program what to do. Feel free to load this program into Ansible after I walk you through how to set up your Ansible environment later in this chapter.

If you want to run this YAML example or any playbook in Ansible, you need to use the command **ansible-playbook playbook.yml -f 10**. I recommend checking out YAML tutorials and trying out very simple program examples before attempting to develop your own code. Later in this chapter, you will learn how to build your own SOAR lab using Ansible. Ansible is based on YAML, meaning it’s an application that builds YAML files as you use it. This also means you are expected to be able to program in YAML if you plan to use Ansible.

Note

Ansible provides a tutorial for understanding YAML at https://docs.ansible.com/ansible/latest/reference_appendices/YAMLSyntax.html.

Understanding the details of this YAML program or the previous examples is not important at this point. The three YAML examples were intended to show you that YAML is easy to read and follow regardless of your programming background, to give you a general idea of how YAML is structured, and to enable you to recognize YAML code, which you are likely to run into as you learn more about DevOps.

Note

There are many good free YAML programming tutorials available on the Internet. One example is the Tutorials Point tutorial at https://www.tutorialspoint.com/yaml/index.htm.

Data Modeling

Another DevOps concept that is important is data modeling, which is different from what YAML, JSON, and XML offer as data formats. YANG is viewed as a data modeling language used to model configuration and state data. This means YANG defines data schemas that JSON or XML must follow. Think of YANG as a guardrail as well as a translator setting the rules for what is allowed and translating what is allowed into an understood format. Anything that falls outside of the guardrail will be rejected, forcing only approved formatted data to be translated, sent, or received. This is critical for accommodating the various data formats the SOC will encounter as it leverages data from different internal and external databases.

When working with a device that supports YANG data stores, a specific XML or JSON encoded message would be deemed valid or invalid based on YANG models. YANG can be used with other protocols; however, it is most commonly used with the Network Configuration Protocol (NETCONF). The reason for this is that YANG provides a way to show how configurations are modeled, while NETCONF is a protocol that can modify them.

Figure 10-22 provides a dataflow example of the YANG serializer's capability of encoding and decoding arbitrary structures into something like

XML or JSON. For decoding, YANG pulls data into memory and decodes it to a data node. Once data is ready to be encoded, it is encoded from a data node back to a data stream such as XML or JSON. As you can see from this example, data modeling allows for data to be converted into a desired data format such as JSON or XML.

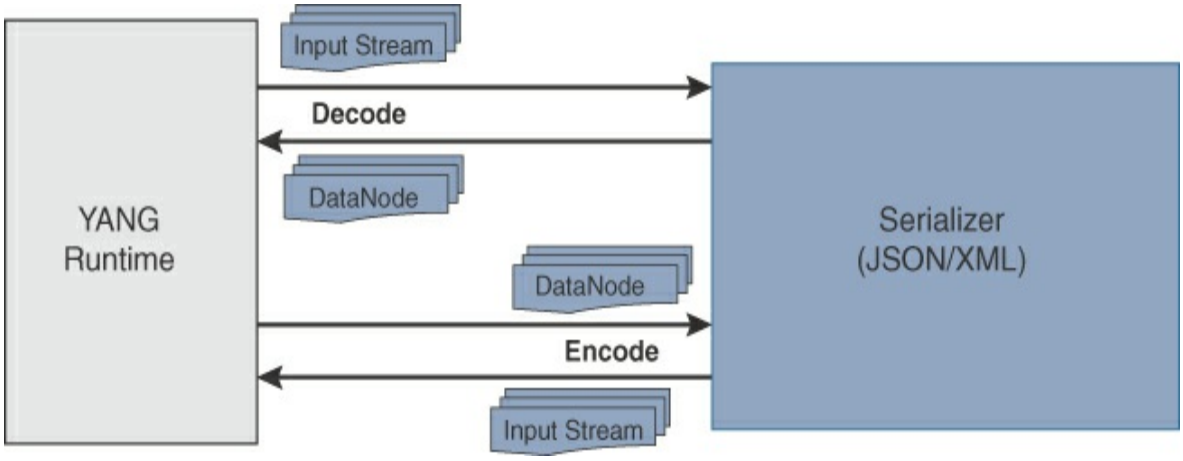

FIGURE 10-22 YANG Serializer Flowchart

NETCONF

To best explain NETCONF, first consider the Simple Network Management Protocol (SNMP), which has been the traditional method used by network administrators for monitoring and modifying configurations. An example of using SNMP is pulling status from a running device or pushing a change such as changing an interface VLAN. Certain versions of older NAC technology would use SNMP as a way to push changes to switch ports. SNMP is not going away anytime soon; however, SNMP is not optimized for automation, leading to the need of a more programmable version of SNMP as network administrators want to explore DevOps concepts. NETCONF provides programmability based on a simple, standards-based, and robust API to read and modify configurations ideal for orchestration.

NETCONF follows a traditional client/server model using remote procedure calls (RPCs). Server configurations are stored in a NETCONF configuration datastore that follows a YANG-defined data format. To query or modify data, a client sends an XML-based RPC over one of the supported secure transfer

methods, and the server replies with XML-encoded data. You can reference RFC 6241 to get a better understanding of NETCONF operations, which are defined by YANG. Figure 10-23 represents the YANG NETCONF process.

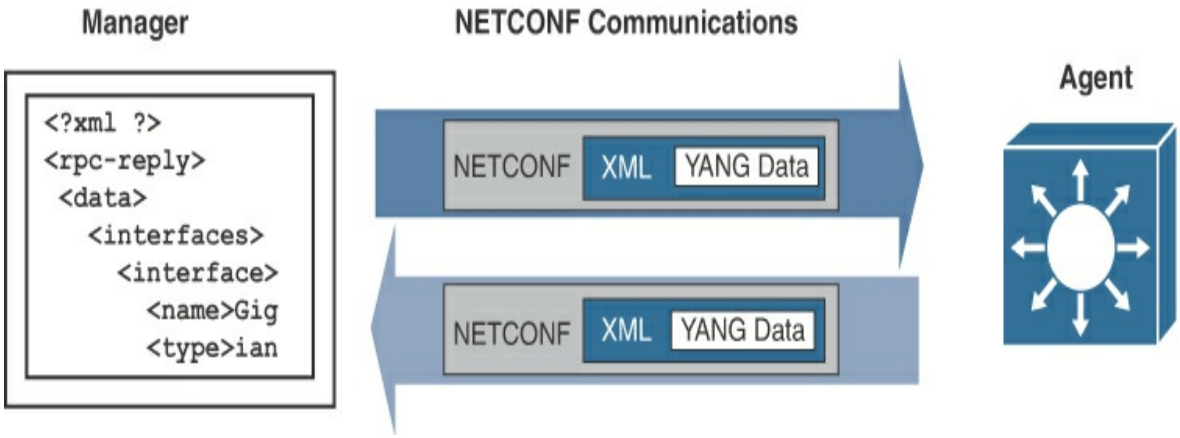

FIGURE 10-23 Example DevOps Communication with YANG and NETCONF

RESTCONF

A protocol similar to NETCONF (not a replacement) is RESTCONF. RESTCONF uses the RESTful HTTP interface that can be used to query and configure devices with NETCONF configuration datastores. RESTCONF is an HTTP-based protocol that supports both JSON and XML, whereas NETCONF uses just XML. When you leverage RESTCONF with a scripting language like Python, you can automate a wide variety of network administration tasks.

To see examples of NETCON and RESTCONF used with YANG, check out the Cisco DevNet GitHub at https://github.com/CiscoDevNet. One example is a script that can provision or delete a VLAN interface on an existing interface that is connected to an upstream switch using a 802.1Q trunk. (You can find this example script at https://github.com/CiscoDevNet/restconf-examples/tree/master/restconf-samplecode/vlans.)

DevOps Tools

So far, you've learned a lot of general DevOps fundamental concepts, including file formats and data modeling. Now let's look at different types of DevOps tools that help you to create orchestration and automation in your SOC services. DevOps is a fundamental component of enabling SOAR capabilities.

It is common to start your DevOps journey by creating playbooks and applying various DevOps techniques to different steps of a workflow based on a desired outcome. Maybe you want to push changes to switches when an event occurs. Maybe you want to send alerts to specific team members when security tools identify a potential malware outbreak. How you leverage DevOps to develop SOAR capabilities, ranging from case management to automated response, depends on what tools you have available and what is the best option to accommodate your workflow.

DevOps Targets

The following list highlights aspects of playbooks that are ideal for applying the DevOps concepts covered in Chapter 9. There are many other good use cases; however, these use cases are very common areas to which DevOps programming is applied to simplify elements of common workflows.

- **Analysis of code and configurations:** Analyze code or configurations so that the team can quickly identify potential vulnerabilities or problems
- **Submitting changes:** Allow for anybody to submit changes to increase productivity yet ensure all changes meet compliance
- **Monitor compliance:** Continuously validate that the SOC is following all policies and procedures and can collect data to prove compliance when requested
- **Investigate threats:** Collect what is needed for an analyst or an automated system to provide the quickest yet most effective response to an incident
- **Vulnerability management:** Identify and manage IT-related vulnerabilities in the organization
- **Training and enforcing security:** Ensure security practices are enforced

and offer notifications to encourage more secure behavior

Delivering value from DevOps does not always require creating a new tool or coding a script. Suppose you need to change the switch port of a switch when an event occurs. A SOC could first manually use a port security approach and later either leverage a vendor tool such as NAC or rely on a programmer to create a program that is triggered when a certain event occurs. That program could be part of a script that is executed in a SOAR platform, representing the automation of steps in a playbook. That program could also be a dedicated tool that is called upon by the SOAR to perform the changes to the switch. It all depends on what tools your SOC has available and what is the best approach to accomplish each SOC service goal.

DevOps tools can essentially be broken down into two different categories of functionality—manual tools and automation tools. Keep in mind that both manual and automated tools are designed to provide configuration management and improve the workflow of playbooks but use different approaches to accomplish this goal. Both tool options can also work with a SOAR or be part of a SOAR's feature set.

Manual DevOps

Manual DevOps tools can provide three main areas of value to a SOC. The first is knowledge management—the ability to create an accessible knowledge base and make it searchable. Doing so allows the accumulated wisdom to be replicated to other teams, which supports more effective responses to similar situations. As fantastic as this concept sounds, you need the right tool to benefit from this approach of knowledge sharing. I have seen knowledge management tools not properly utilized, leading to huge forums of data that take users hours to crawl through to find what they are looking for. I highly recommend evaluating how created content is organized and referenced. ChatOps, covered later in this chapter, make this concept even more interesting by enabling you to quickly reference material by voice. A SOC analyst can simply pull information without having to touch a keyboard.

The second key area of value that manual DevOps tools can provide is streamlining a SOC's incident management service. The first value point,

knowledge management, impacts this value point because a well-built knowledge management system will improve how a SOC executes playbooks according to its incident management established processes. I highly recommend any SOC with incident response management services convert playbooks into a knowledge management system, which is a very common practice found in mature SOCs around the world.

The third value from manual DevOps tools is their ability to leverage automation. I use the word “leverage” because manual DevOps tools do not offer automation capabilities such as what I’ve referenced in this chapter when speaking about building automation. Manual DevOps tools launch existing automation scripts such as asking a ChatOps tool to run something rather than spending time within a ChatOps tool writing an automation script.

There is a lot of value in having manual DevOps tools launch scripts. Automation leads to repeatable results and control over how automation is launched, which keeps the workflow functioning properly. For example, if a common situation is reported in a wiki, a DevOps programmer could write a script that automates the suggested response and post a link in the wiki for others to benefit from when they encounter the same situation. The wiki tool would not be used to create the script but rather to host a link to the script as part of a suggested response.

The DevOps manual tool categories of wikis and ChatOps are both based on manual processes, but they enhance the workflow of playbooks in slightly different ways as described in the sections that follow.

Wikis

Wikis are hypertext publications composed of collaborative edits and are managed by their own audiences via web browsers. When a SOC develops a playbook and applies a wiki tool, steps not only can be documented with responses to predicted questions, but can also be adjusted as the playbook is put into action, allowing for continuous improvement. Wikis are focused on process management and information flow across the course of an action designed for a SOC analyst to reference as a playbook is executed similar to a manual covering how to use a tool.

Regardless of the platform you choose, wikis must be programmed and

maintained to avoid responses becoming obsolete as various aspects of the playbook change over time. Figure 10-24 is a screenshot of a fresh install of MediaWiki, which shows what the basic platform looks like. To get an idea of what is involved with installing MediaWiki on Linux, for example, check out the configuration guide at https://www.linuxtechi.com/install-mediawiki-on-linux/.

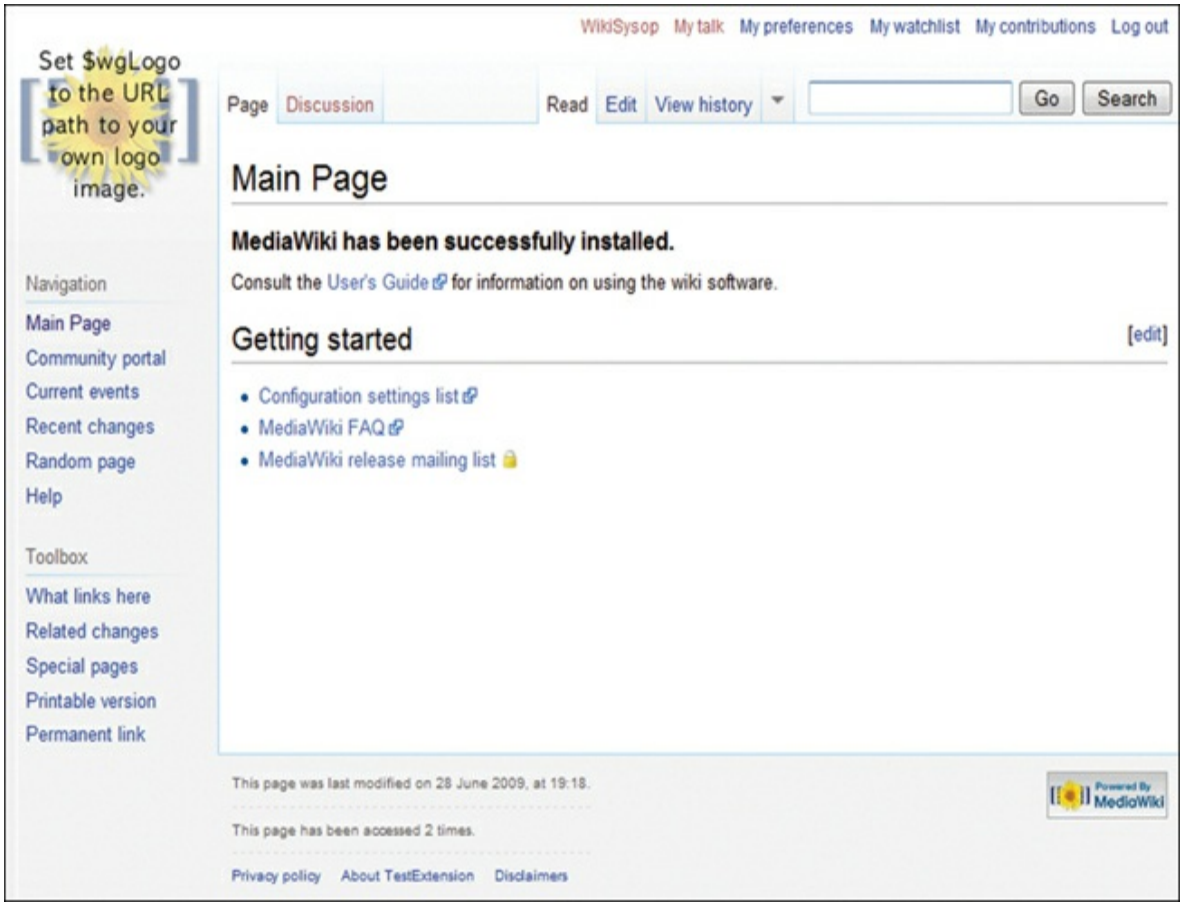

FIGURE 10-24 New Installation of MediaWiki

The following lists are popular wiki platform options to consider for your SOC:

- **MediaWiki:** Free tool that Wikipedia was built on. MediaWiki facilitates effective management of a wide range of information types with easy to insert and edit options using any database such as Oracle, PostgreSQL, SQLite, or MySQL.

- **Confluence:** A commercial wiki option providing management of tasks through managed ticketing functionality. SOC analysts can create, edit, assign, and track tasks with the Confluence tool.
- **TikiWiki:** A business-focused platform for developing blogs, forums, wikis, and calendars.
- **DokuWiki:** A free option for document management, where you can save your documents as files rather than saved within a database. This means documents are plaintext format, allowing for various read and create options.

ChatOps

ChatOps tools are designed to increase collaboration through simplified communication similar to wikis but with a focus on communication rather than creating dynamic workflows based on documenting steps on a wiki forum. ChatOps systems provide instant message options as well as ways to document and reference conversations, allowing for ongoing improvement of processes based on daily activity. Unlike wikis, where workflows are configured, ChatOps workflows are built based on communication that occurs. SOCs around the world have found ChatOps to be an essential tool to improve communication and increase documentation of the workflow of a playbook.

ChatOps automates tasks using a quasi-intelligent, conversational interface to query information and initiate scripts that perform repetitive tasks. A great example with similar functions is Amazon Alexa. Using a voice-activated system, you can launch tasks to connected devices, including turning on/off the alarm system for your house, turning on/off lights, and playing music. As you use Amazon Alexa, it learns your behaviors and is able to recommend things based on what you do, representing a form of machine learning.

There are a handful of industry-recognized free and commercial ChatOps options to consider. The following are some of the more popular options:

- **Cisco Webex Teams:** Cisco’s version of ChatOps bringing together messaging, file sharing, video meetings, whiteboarding, calling, and other tools targeting improving teamwork.

- **Microsoft Teams:** Microsoft's business-oriented communication and collaboration platform. Tools include messaging, file sharing, video meeting, and other collaboration tools.
- **HipChat:** Ticket-based chat system that allows for creation of groups for easy collaboration. HipChat offers public or private rooms, exchanging files, and URLs to streamline team collaboration. Both free and paid options are available.
- **Slack:** A popular platform for facilitating team conversations while also allowing the creation of private or public channels. There are both paid and free options of Slack.

Automated DevOps

Building automation leads us to the other major category of DevOps tools, those that are designed for creating automation rather than just referencing existing automation. There are many variations of automation tools, but in the world of DevOps, the most popular ones provide configuration management capabilities, which means they are designed to deploy, configure, and manage servers. All applications have some form of a settings menu for configuration management, so the concept of configuring servers shouldn't be new to you. What makes DevOps special is that it offers automation across any vendor to improve the workflow of a playbook.

When it comes to market share for DevOps automation tools, Ansible, Puppet, Chef, Jenkins, SaltStack, and Microsoft Endpoint Configuration Manager are all popular options. Many DevOps training programs leverage Ansible, which tends to fall in the top five DevOps automation options used in organizations and is open source, which makes it easy to acquire for testing purposes.

The next section shows you how to build a basic Ansible environment to practice DevOps configurations. Enterprise offerings can be easier to use and typically provide more features, including wizards and plug-and-play programming; however, Ansible can be more flexible because it is open source, and it's a great tool for learning how a SOAR functions.

The following provides a brief summary of Ansible and other popular

DevOps automation tools:

- **Ansible:** Ansible is an open-source software provisioning, configuration management, and application deployment tool enabling infrastructure as code. Ansible runs on most Unix systems but can configure both Unix and Microsoft systems. You will learn how to configure Ansible on a Linux system in the following section.
- **Chef:** Chef is another popular-open source automation software that uses build blocks to configure systems similar to recipes in a cookbook, which is how it got its name. Chef is used for infrastructure automation and can help reduce manual and repetitive tasks for infrastructure management.
- **Jenkins:** Jenkins is an open-source automation tool built for continuous-integration purposes. Jenkins offers hundreds of plugins to support building, deploying, and automating most projects.
- **SaltStack:** SaltStack is a configuration management and orchestration tool that helps in the deployment of dynamic applications and general-purpose infrastructure they run on.
- **Puppet:** Puppet is another open-source configuration management and deployment tool. The value of Puppet is its capability to spread automation across an organization.

The best way to learn a concept is to see it in action through hands-on experience. Next, I'll walk you through how to build a basic Ansible setup. I'll also show you automation examples and provide resources for further learning. Let's get started!

DevOps Lab Using Ansible

One good way to wrap together all of the topics I covered around playbooks, automation of workflows, and programming workflows is through hands-on experience. A very popular automation tool you can use for this purpose is Ansible. Ansible calls itself a universal language used to develop repeatable playbooks. Ansible functions in an agentless format using Secure Shell (SSH) and network traffic; however, some specific software and hardware manufactures offer plugin enhancements, but these are optional.

Ansible allows SOC teams to create Ansible playbooks. An Ansible playbook is essentially a configuration file for Ansible, written in YAML, that is a set of tasks, device settings, and configurations that happen in a sequence specified by the creator of the playbook, hence a workflow. Playbooks can be simple tasks, such as backing up configurations, or as complicated as setting up multiple network devices, servers, VLANs, and firewall rules. Because the steps in the Ansible playbook are automated, the workflow is improved, giving time back to the SOC analyst as well as ensuring a repeatable delivery is executed and tracked by Ansible.

> **Note**
>
> Check out https://www.ansible.com/ for more details on what Ansible can do for your SOC.

Installing Ansible

The first thing you need to do for your SOAR lab is install Ansible. You can run Ansible on a variety of operating systems, but most SOCs I've encountered run it on a Linux system. For purposes of this example, I will explain how to install Ansible on Ubuntu Linux. You should not have any issues running Ansible on any recent version of Ubuntu, but I recommend validating the minimal required specs at https://docs.ansible.com.

The first Ansible installation step is to perform basic housekeeping of your system. For Linux, you do this by running **update** and **upgrade** commands on your Ubuntu system. The fastest way to do this is to run the command **sudo app-get update -y && sudo apt-get upgrade -y**. Once the update is complete, you will want to check that Ansible is not already installed on the system. You can do this by issuing the command **ansible --version** from the terminal. You should receive a message stating the program is not installed.

Next, run the command **sudo apt-get install software-properties-common**, which installs some standard libraries and shares software Ansible uses. It's possible the update you ran to prepare your system will have updated and installed this software, so don't worry if there is nothing to update or run after

running the **software-properties-common** command.

Next, add the Ansible repository to your operating system by issuing the command **sudo apt-add-repository ppa:ansible/ansible**. You may be prompted to accept user agreements, which you must accept to proceed. The final step is to install the Ansible software by issuing the command **sudo apt-get install ansible**. You can verify the installation is complete by using the command **ansible --version**.

A summary of the installation steps is as follows:

Click here to view code image

```
sudo app-get update -y && sudo apt-get upgrade -y
ansible –version
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get install ansible
ansible –version
```

Setting Up Ansible

Ansible, by default, is installed under /etc/ansible. The first file you will examine is the Ansible hosts file. The Ansible hosts file works exactly like the system OS file, which means it resolves hostnames that Ansible interacts with to their IP addresses. Even if you have DNS or predefined hosts files on your system, you still need to set up the Ansible hosts file. Think of the Ansible hosts file as the inventory of useable systems that Ansible has access to configure. Figure 10-25 shows the my host file.

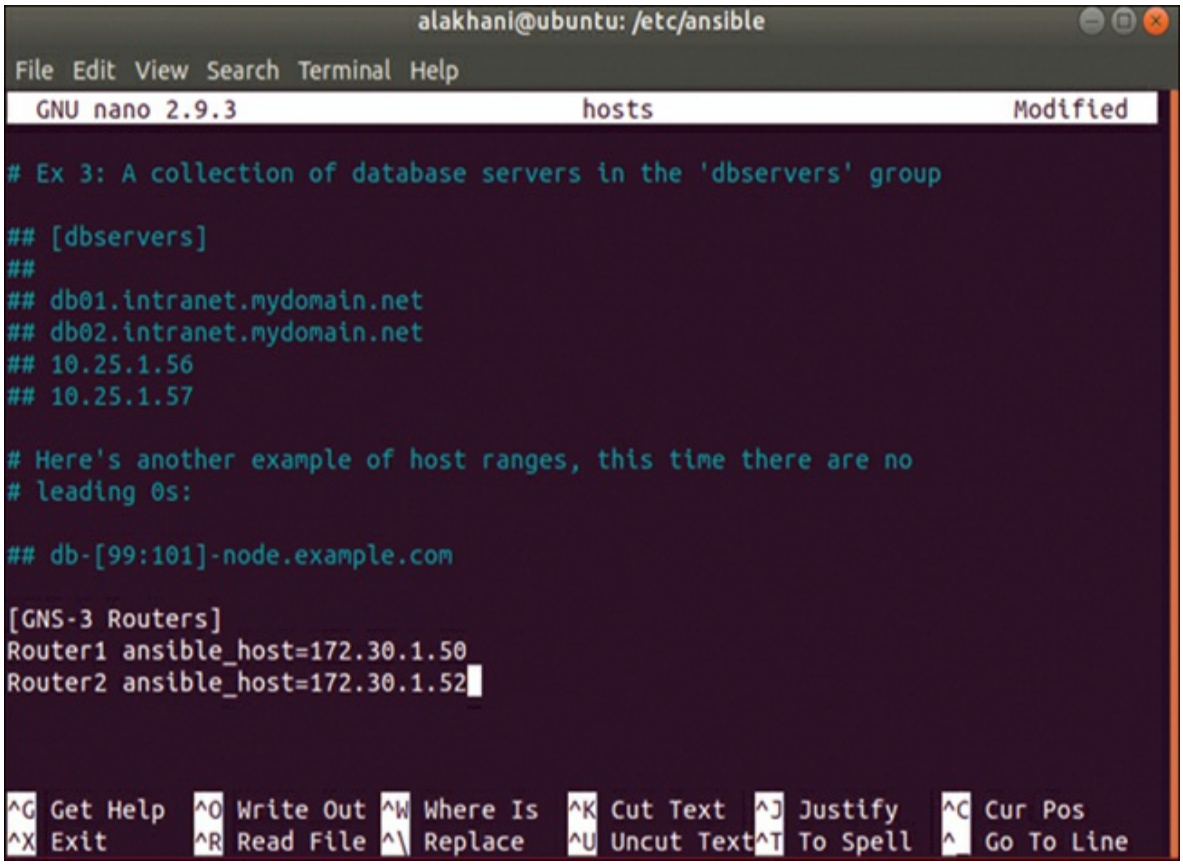

FIGURE 10-25 Ansible Hosts File

For this example, I have created a group for my routers and have named that group GNS-3_Routers (you want to avoid spaces, if possible, in your names). I also have two different entries under the group I named GNS-3_Routers. Ansible is known for its scaling features, and this is one of the ways it achieves that. Any playbook that is called by Ansible that uses these specific group names as defined in my hosts file will know which systems to reference. In my example, I only have two hosts, but it is not uncommon to find Ansible groups containing thousands of hosts. At this point, your basic Ansible installation is complete and you are ready to create your first playbook.

Ansible Playbooks

After you populate the Ansible hosts file, you are ready to create your first

Ansible playbook. Remember that an Ansible playbook is a configuration file in YAML format representing a set of tasks you want to perform. The first and most common task needed for automating a workflow is accessing a system. It is common for a SOC to set up an SSH digital certificate for Ansible to authenticate to the device it is accessing during an automated workflow step. You can also use usernames and passwords, but certificates are more common. This can be accomplished leveraging the YAML program language.

To see how this works, let's examine a basic Ansible playbook that allows me to back up Cisco routers I have already defined in my example Ansible hosts file. The following code is an example of my playbook that accomplishes this goal. If the YAML format seems confusing to follow or understand, look back earlier in this chapter for my brief overview of YAML, which explains the core components of what to expect in a YAML file.

Click here to view code image

```
---
- hosts: Router1
  gather_facts: true
  connection: local
  tasks:
   - name: show run
   ios_command:
   commands:
     - show run
     host: "{{ ansible_host }}"
     username: admin
     password: cisco
     register: config

  - name: save output to /etc/ansible/router_backups/
  copy:
  content: "{{ config.stdout[0] }}"
  dest: " /etc/ansible/router_backups/ show_run_{{ inventory_hostname
}}.txt"
```

Looking at this example, the first thing to notice is how the file starts with three dashes. Remember that every YAML file starts with three dashes to inform the computer that this is a YAML file. Next, there is a single dash followed by a space and the word hosts. This is where you enter the name of

a host that is referenced in the Ansible inventory. The “connection: local” line indicates that an SSH connection is sourced from this server, meaning the server running Ansible. When you are running larger Ansible deployments, you might want connections to be sourced from different geographic regions to improve efficiency. This is a key point. Recall from earlier in this chapter that one key value from a SOAR is reducing task overhead by leveraging more effect searches. This is an example of how time can be saved through tactical device access built into the automation.

The tasks and commands part of the previous configuration file is the most critical part of the configuration file because it includes the actual commands I want to run. When you see things in brackets or double brackets, they represent variables. You can replace those variables with the name of the system you are connecting to. For example, you might decide to run this playbook on a hundred hosts but want Ansible to automatically name the backup configuration file as the same name as the host being evaluated. Figure 10-26 shows the example playbook file in the Nano text editor.

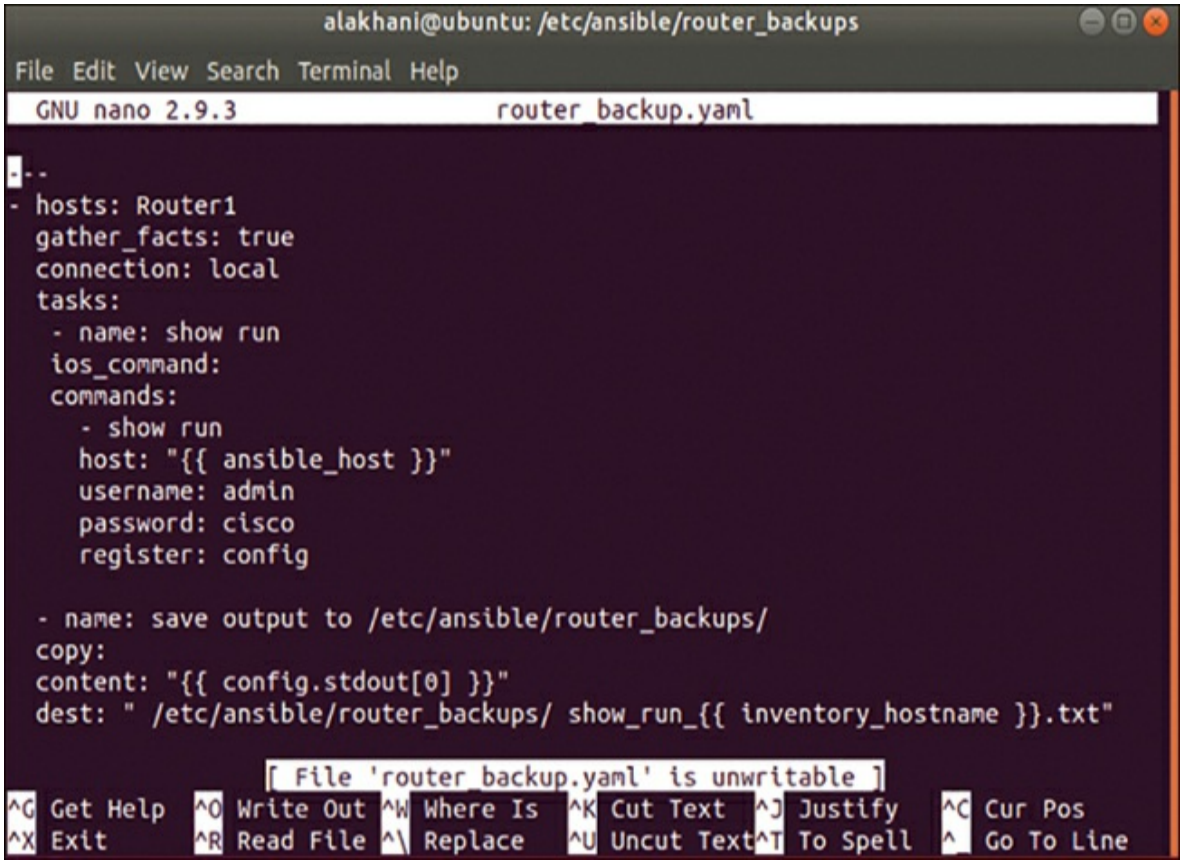

FIGURE 10-26 Executed Ansible Playbook Example

There are many resources you can access to obtain more Ansible playbook examples. The purpose of walking through Ansible is for you to get a feel of setting up and developing a basic playbook using an open-source SOAR that can automate a simple task such as backing up a router configuration. If you need to brush up on your YAML programming skills, look back to the YAML section in this chapter for references to YAML tutorials provided by the creators of Ansible.

> **Note**
>
> Colin McCarthy of many useful Ansible playbook examples on his GitHub repository, located at https://github.com/colin-mccarthy/ansible-playbooks-for-cisco-ios.

Now that you have a SOAR lab that you can use to practice DevOps concepts, you need a task to automate. One extremely popular task to automate is blueprinting an endpoint, meaning pulling details about an endpoint. This task is extremely useful for various SOC services, including vulnerability management so that you know what to scan for vulnerabilities, as well as incident response so that you know what is occurring on a system impacted by an incident. Next, I'll work through how to use Osquery to perform endpoint blueprinting. You can use your Ansible lab to automate this process.

Blueprinting with Osquery

Guidelines such as NIST 800-53 Rev4 CM-11 outline a need for identifying what is installed and running on systems within an organization. Identifying what is installed and running on host systems is sometimes referred to as establishing a blueprint of that system. Blueprinting queries what hardware and software are in your environment and attempts to understand the

capabilities related to what is collected. This is slightly different from the *baselining* concept referenced earlier in this book, which develops an understanding of what is considered normal and looks to identify outliers. Compiling a blueprint of the organization's managed endpoints enables the SOC to have visibility into endpoints so compliance to security standards can be monitored and enforced.

The goal of blueprinting is to identify systems that do not meet corporate standards for security, proactively address vulnerabilities, understand what is operating in the environment, including applications installed on systems, and map out an overview of risk associated with what is connected within an organization. It is common for security tools such as vulnerability scanners, SIEM/SOARs, and NAC technology to include variations of blueprinting capabilities.

The use case in this section presents Osquery, an open-source tool for investigating metadata from endpoints. If you want to query an operation, such as a specific PowerShell command, Osquery allows you to build predefined rules to search for such functions. Many EDR offerings have similar functionality. Among the values offered by Osquery are its ability to be customized, scalability options, use of common languages, and industry support.

In order for Osquery to function, an organization must have control of the systems on the network, including having administrative privileges and having the ability to install software. It is common practice for organizations to use orchestration tools such as Ansible, Puppet, and Chef to push out and deploy Osquery, which makes Osquery a great use case to understand and test with your Ansible lab. Osquery can provide event information such as new accounts being created on a system, failed login events, the creation of files, and what processes are running on a system. The real power of Osquery's visibility, however, is grouping what is seen together into a potential risk, such as flagging when PowerShell is launched after an Excel file is executed. Next, I will walk through a simple Osquery installation so that you can add this tool to your lab.

Note

The website https://osquery.io/schema/ has a list of the schemas supported for each version of Osquery. This gives you an idea of all the data being collected from endpoints.

Running Osquery

There are two different approaches to installing and running Osquery. One option is to deploy a centralized managed and reporting method. The other option is a standalone method. For simplicity, I will walk through the standalone install option on Ubuntu. Step one is to go to the Osquery website, https://osquery.io/, and click the Downloads menu. On the Downloads page, you have the option to choose which version you want to download. Always make sure to use the latest version available.

Note

The Release Type option offers the choice of an official release or a debug version. The debug version is targeted for software developers.

Next, choose one of the download options: macOS, Linux, RPM, Debian, or Windows. For purposes of this example, I downloaded the Debian version of Osquery. To install what I downloaded, I opened a terminal, navigated to the folder where the Osquery package was downloaded, and used the command **dpkg -i osquery**, as shown in Figure 10-27.

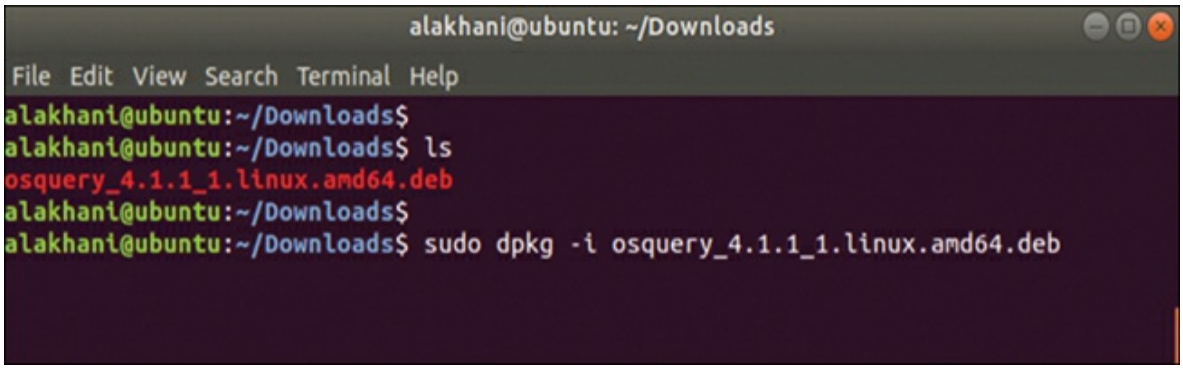

FIGURE 10-27 Installing Osquery on Ubuntu

To start using Osquery from the command line, type **osqueryi**. At this point, you can issue any SQL command Osquery supports. Figure 10-28 shows an example of running **select * from uptime;** to see the uptime and running **select * from logged_in_users** to see logged-in users.

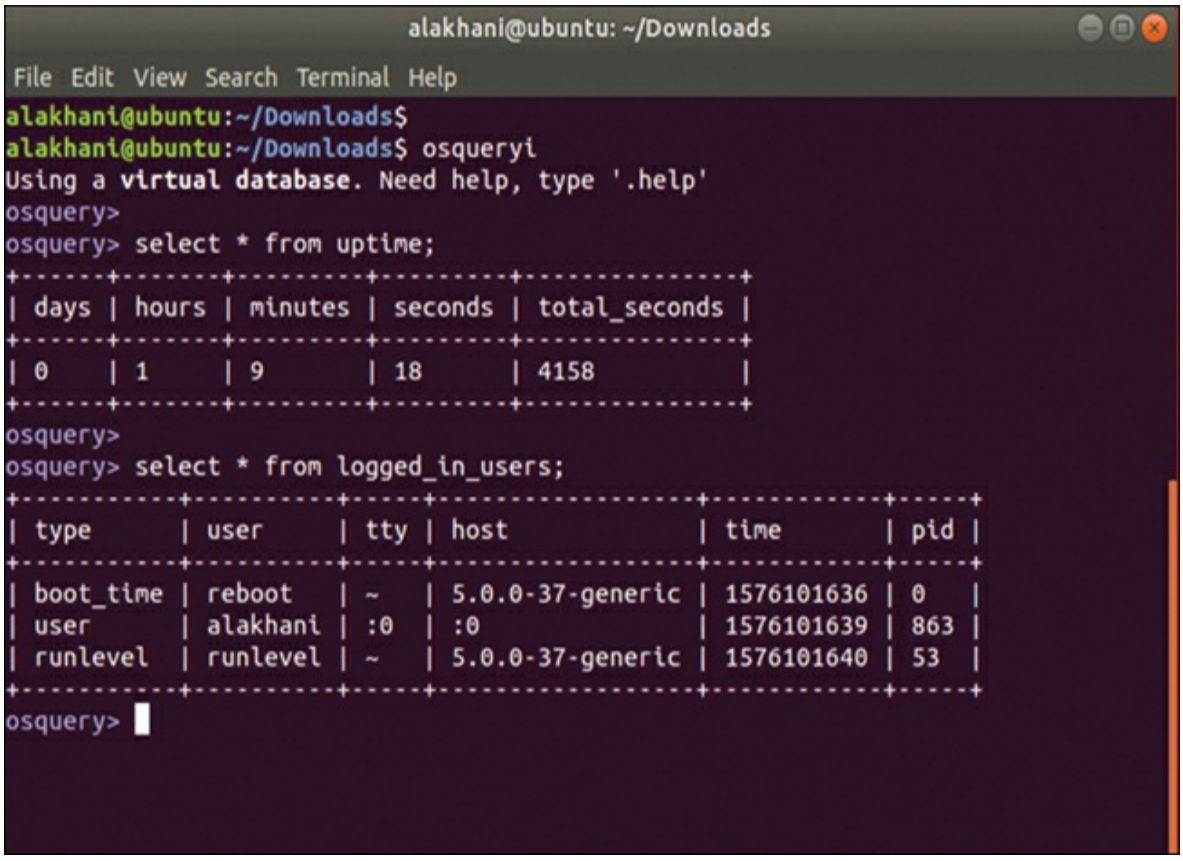

FIGURE 10-28 Osquery Uptime and Logged-in Users

I highly recommend reviewing https://osquery.io/schema to better understand what queries you can run using Osquery. Two common Osquery use cases used by SOCs is to search for configuration settings or be alerted when there are changes in settings for disk encryption, software, firewall settings, or querying usernames. Figure 10-29 shows Osquery running a simple SQL command on macOS to show all the usernames on the system.

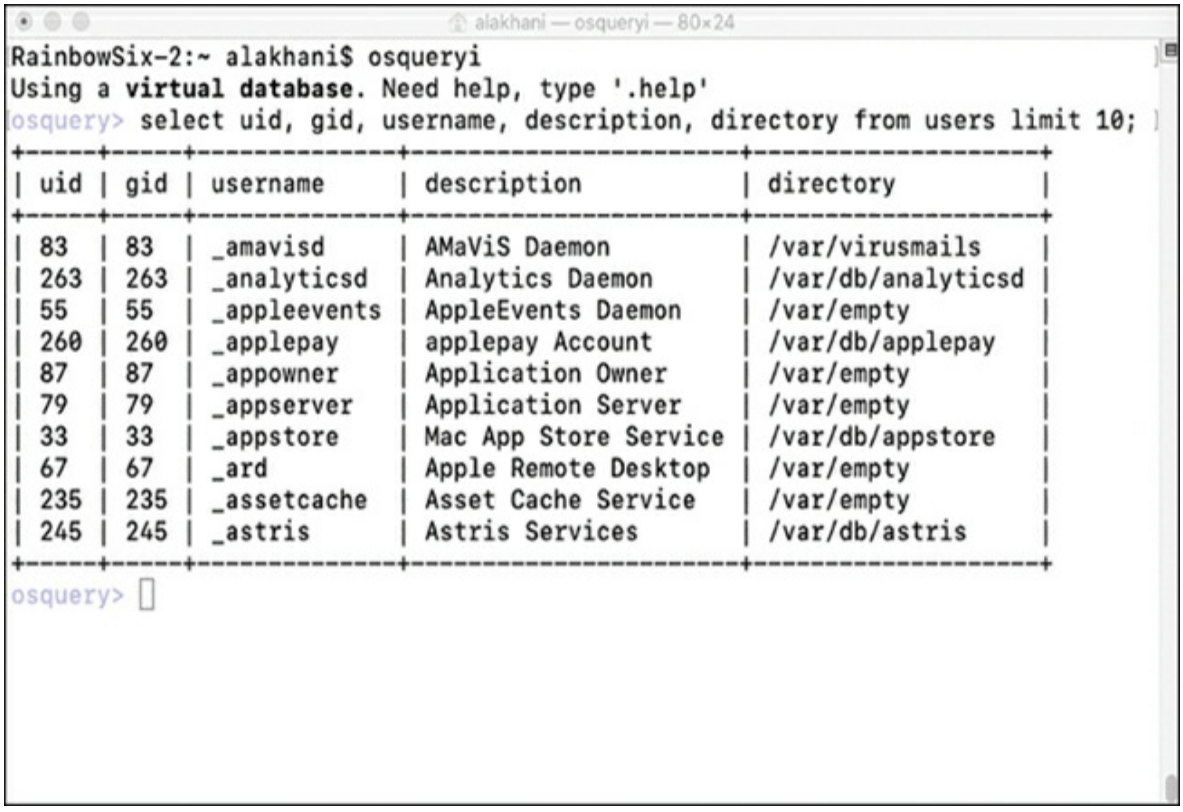

FIGURE 10-29 macOS User Query

Apple device owners, such as those using a Mac Book Pro, might be surprised to know that your Apple device has dozens of default usernames that are included with the standard operating system install. To check the details, perform a SQL query to find out exactly how many usernames your device has. As shown in Figure 10-30, according to Osquery, the example system has 101 different usernames! Surprisingly, this is expected with a standard macOS install with one person using it.

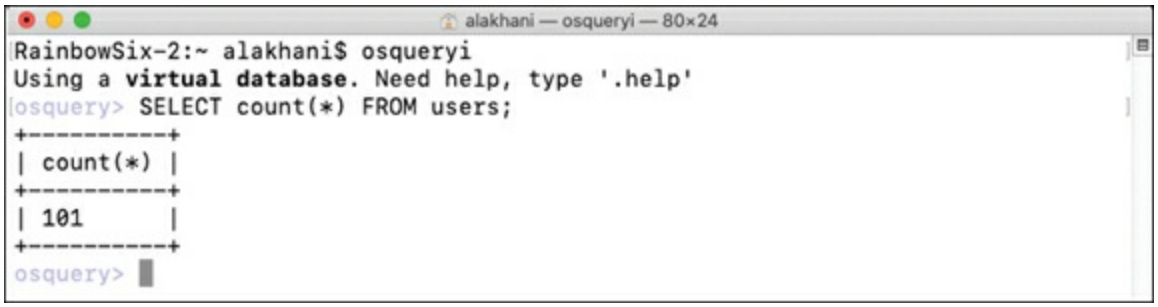

FIGURE 10-30 Osquery Username Count

The visibility capabilities that Osquery offers are very useful, but the real value comes from the potential of its automation capabilities. For example, when conducting remediation for a network-wide event, Osquery can push a fix to all systems, versus requiring the SOC incident response team to address each system individually. A similar concept can apply to developing widgets in a SIEM solution that monitor the risk associated with endpoints through continuous evaluation based on what data Osquery collects from endpoint systems. I recommend trying to automate some simple Osquery tasks with Ansible to get a better understanding of how a SOAR can leverage tools like Osquery.

There are other forms of DevOps, including those focused on the network and on the cloud. This leads us to the next topic, DevOps focused on the network.

Network Programmability

Network programmability, which many in the industry call NetDevOps, applies the DevOps principles of software development to network programming. Older tools for monitoring and configuring networks, such as SNMP, syslog, and ping, work, but there are much more effective ways to accomplish the same outcomes. Technology concepts such as NetFlow used for identifying security anomalies, machine learning, and streaming telemetry open up new network management concepts leading to automated and much more effective management of the network.

Think about the concept of “the network” and what that means to you and

your organization. Today, the network no longer just means a routing and switching network. The network can include cloud technologies that host your applications, and your local users will need a way to access those cloud-based applications. Your local datacenter will also have applications that users offsite will need to access. A network will always be between your applications, and the SOC needs a way to configure and monitor that network. Traditional network management heavily relies on SNMP and syslog; however, a better method is needed to keep up with current trends in automation and collaboration requirements between network architects and IT operators, allowing for a more streamlined flow of running and securing the network, wherever it may be.

These collaboration challenges among different teams in an organization have been painful for many organizations. Consider as an example performing the steps to add a server to the network. The datacenter team understands how to provision a server, but they need IT operations to provide network access and need the security team to open a firewall and other security tools to allow traffic to and from the server. The value DevOps brings to the table is huge in regard to reducing the number of manual siloed steps that need to occur to get things done. By allowing tools to work together, cross-team tasks can be automated.

Learning NetDevOps

Manufacturers of networking equipment, such as Cisco, have adjusted their education and certification programs to accommodate a change in behavior toward managing networks following DevOps trends. For example, Figure 10-31 shows an overview of the Cisco engineering and software development certification programs, indicating the software program is targeting DevOps skills with a focus on NetDevOps, meaning how to make the network work with other tools. Cisco expects certification candidates to demonstrate skills such as Python, NETCONF, and YAML used to manage and modify existing networks rather than how to build new Cisco networks.

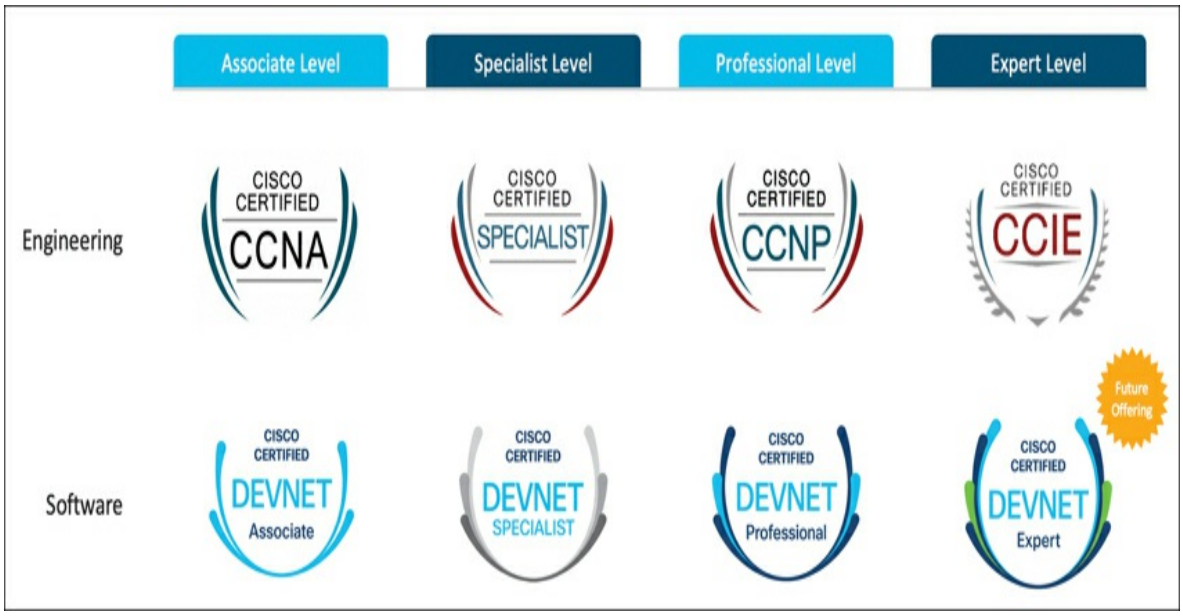

FIGURE 10-31 Cisco Engineering and Software Certification Programs

Another network market leader, Juniper Networks, offers its own flavor of DevOps certifications based on an associate and specialist Dev Ops program (see https://www.juniper.net/us/en/training/certification/certification-tracks/devops?tab=jnciadevops). Juniper's programs include mastering Juniper as well as open-source DevOps tools such as Python and Ansible. Similar programs are starting to be offered across the industry based on the drive for managing networks using the latest DevOps concepts. In short, if you are looking to manage IT networks, DevOps is a topic you are either now or eventually going to need to master.

DevOps certification programs require learning programming concepts but also have a heavy focus on application programming interfaces, which enable machines to talk to machines. In the DevOps world, this means allowing a tool such as Ansible the access rights to make changes to a network switch in the same manner a network administrator would work. APIs come with a set of requirements that govern how one application can talk to another, which is why tools like YANG exist that are designed to ensure API specifications are met. Many vendors follow industry specifications that allow API support to work across multiple vendor products. A good analogy is thinking about a power plug. In theory, your home's power plugs could shape the holes in the plug or adjust the amount of power delivered; however, the manufacture of the plug decides that won't be useful for its purpose. Instead, there are

general standards that are expected for how the holes look in the plug and the amount of power that is generated. Vendor API rules are very similar insofar as they could be designed however the vendor sees fit but instead typically follow the industry expectations of what is supported. This is great for developers because a universal API structure means scripts and programming can work across many vendor platforms.

APIs

In a non-DevOps world, a network manager communicates with network devices using a graphical user interface (GUI) or connects to a system using its command-line interface (CLI). To simplify tasks, a network operations manager can use scripts that pass multiple lines of configuration through the CLI. This concept is essentially using the CLI like an API, but it's a user pushing code rather than one system pushing the configuration to another. Regarding GUIs, SOAP is a way to perform API-like commands for web services; however, the representational state transfer (REST) standard was later developed to offer a much simpler approach to building web services based on API-like concepts. SOAP and REST are HTTP focused, whereas XML and JSON have a remote procedure call (RCP) API function, meaning one system requesting another system to do something and return a value. For network-focused tasks, SNMP is available, but the newer and more DevOps-friendly option is NETCONF and RESTCONF, which provide REST API-like interfacing with the network. These are essentially different examples of API-like protocols you can use to allow your tools to communicate with each other.

NetDevOps Example

To tie together concepts of using APIs to automate tasks, I'll use a free lab offered by Cisco, DevNet Sandbox (https://developer.cisco.com/site/sandbox), to show you how to test how APIs work. I encourage you to perform the same exercise if API concepts are new to you. When you access the Cisco DevNet Sandbox lab resource center, you will find that it offers many always-on and startup labs that allow you to test

many of the concepts from this chapter. Figure 10-32 is a screenshot of the Cisco DevNet Sandbox lab catalog. Any lab with a Reserve button requires you to schedule the lab, while the labs with the Always On button are available immediately. Know that “always on” means the lab is read-only, because anybody can access it at any time, while reserved labs allow you to configure and essentially break things if you so desire.

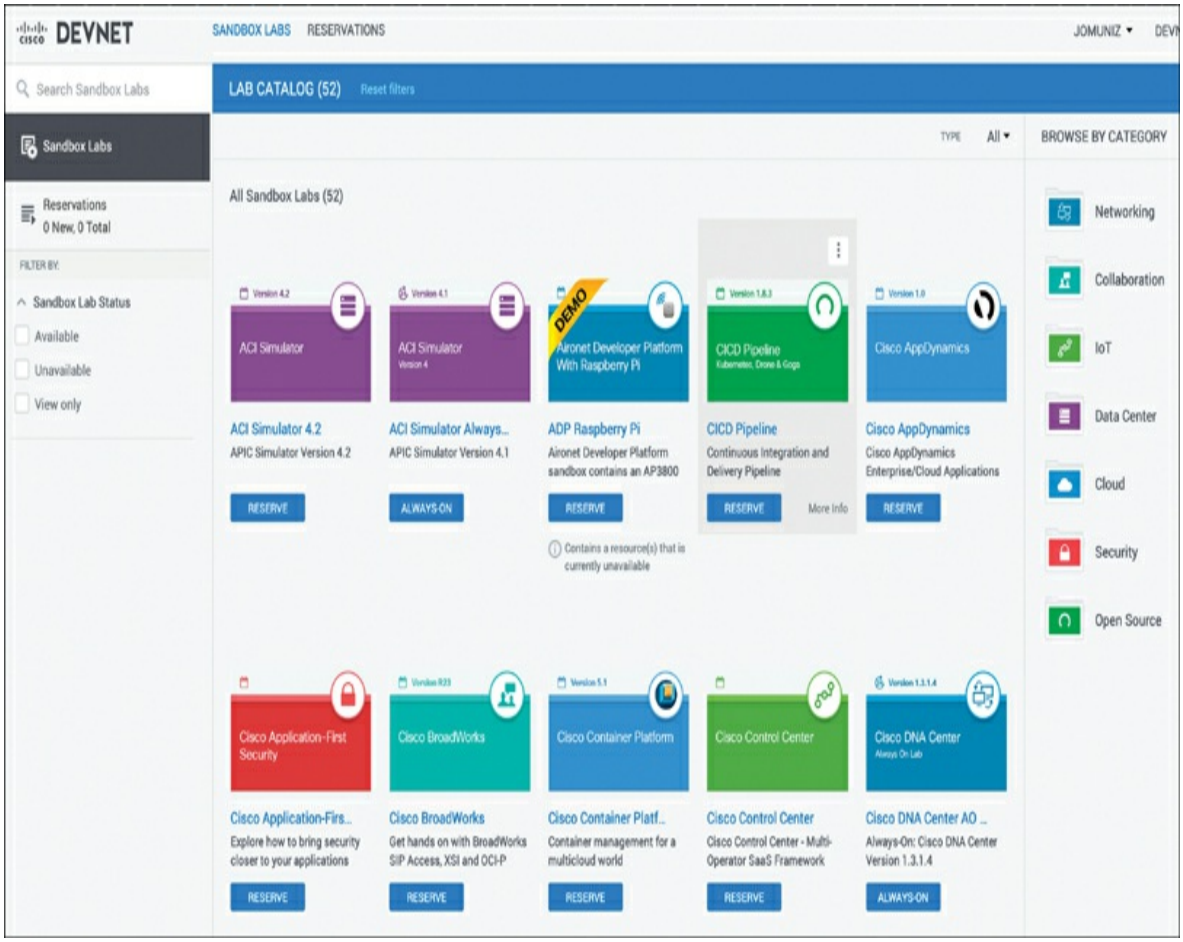

FIGURE 10-32 Cisco DevNet Sandbox Lab Catalog of Free Labs

This next exercise demonstrates the always-on lab called IOS XE on CSR Latest Code AlwaysOn. This lab provides access to a Cisco Cloud Services Router (CSR) running IOS XE code, meaning I can try to push configuration changes to this system using API calls. To simplify sending APIs commands, I’ll use Postman.

Note

You can download Postman from https://www.postman.com/downloads/. Both free and commercial versions are available. You will be required to create an account, which will be part of the installation process.

NetDevOps Example Part 1: Using Postman

Postman is a tool that enables you to send REST API calls to any type of endpoint. Postman also allows you to save your calls, known as Postman collections, which represent example API calls you frequently use and want to bookmark. API requests use an HTTP method based on GET, POST, PATCH, PUT, and DELETE. Figure 10-33 shows the Postman interface running on a Mac. You can use Postman to push commands from your local device to the external DevNet Lab environment, representing how a SOC administrator can push commands from anywhere in the world to his/her organization's network equipment.

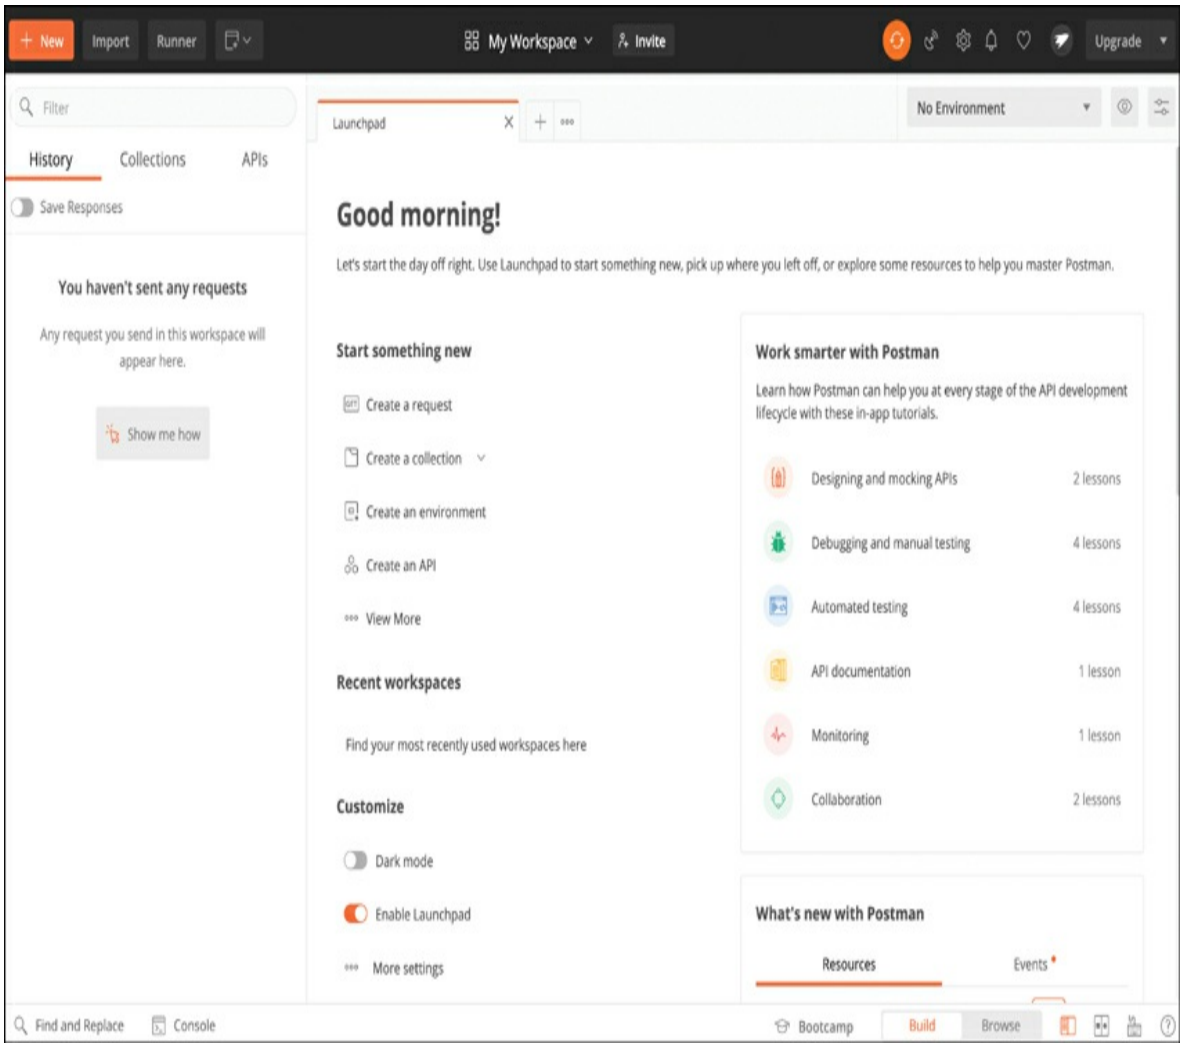

FIGURE 10-33 Postman Dashboard

Within the DevOps environment are the details of the Cisco CSR router that is waiting for me to send API commands to it. The following information comes from the always-on lab showing how to access the Cisco CSR router.

Access Details:

Developers and network engineers access the IOS XE on CSR Latest Code Always On Sandbox directly using the following information:

- CSR1000V Host: ios-xe-mgmt-latest.cisco.com
  - SSH Port: 8181

- NETCONF Port: 10000
- RESTCONF Ports: 9443 (HTTPS)

- Username: developer
- Password: C1sco12345

NetDevOps Example Part 2: Capturing Configuration

This exercise uses Postman to capture the configuration of the Cisco CSR router using RESTCON. First go to Postman and click the New button at the top left, as indicated in Figure 10-34, to open a new workspace. Within the workspace, the GET field is where you enter the address and interface of the device you want Postman to communicate with. Because the Cisco router is located at ios-xe-mgmt-latest.cisco.com, you will want to use HTTPS, the address of the Cisco router, :9443 since that is the port opened for RESTCON, and restconf/data/ietf-interfaces:interfaces. You also need to click the Authorization tab and enter the username and password. Then, click Send to collect the configuration. Figure 10-34 shows how this looks in Postman.

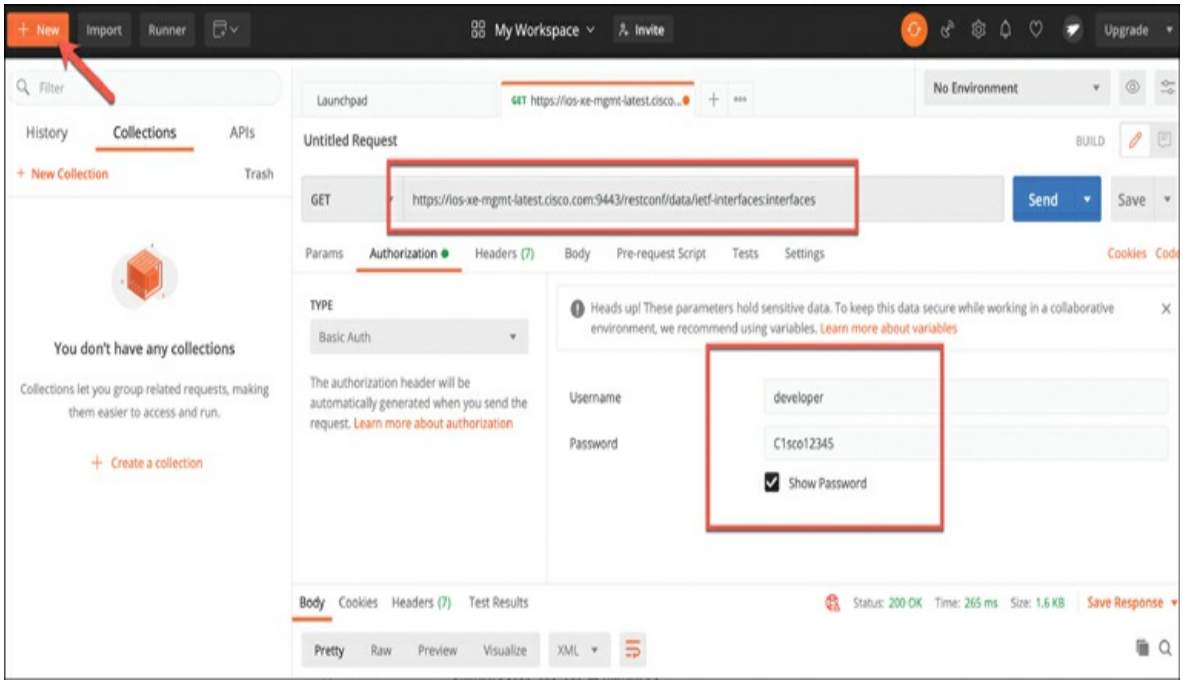

FIGURE 10-34 Configuring Postman to Communicate with a Cisco Router

The results of running this successfully will be the configuration shown within the window below where you ran the commands. Notice that you can change the format by simply clicking the format drop-down menu; Figure 10-35 shows the configuration I pulled down in JSON format. Because this is a shared lab, you will find various configuration changes that have been pushed and labeled by others that have used this lab. The only configuration that can't be changed is that of the interface (as indicated by DON'T TOUCH ME in Figure 10-35), as that is required for the Cisco router to be publicly available. Modifying that can take down the lab for everybody.

FIGURE 10-35 Configuration Pulled into Postman Example

This free DevOps lab is great for practicing how to leverage topics I covered including NETCONF, RESTCONF, and YANG to interact with various types of network, server, and security products. You can download YANG models, practice developing scripts using Python, and have systems online and available for you to interact with. If you have interest in learning more about NetDevOps, Cisco offers free classes at https://developer.cisco.com/ to complement the free DevOps labs that are available. I highly recommend leveraging this and other free online resources to gain hands-on experience with DevOps concepts so you can eventually translate those skills into

building value within a SOAR platform such as Ansible or Phantom.

> **Note**
>
> You can find Cisco YANG models at https://github.com/YangModels/yang/tree/master/vendor/cisco.

Cloud Programmability

Another area of focus for the SOC is automating the configuration and management of cloud services. What is available to access in the cloud depends on what type of cloud service is being considered. Figure 10-36 shows a summary of comparing a traditional on-premises datacenter to different versions of cloud services. The easiest way to understand this diagram is to view what you, “the customer,” are responsible for and what the cloud service provider (CSP) is responsible for. In the traditional datacenter, the CSP is not involved, hence you are responsible for everything.

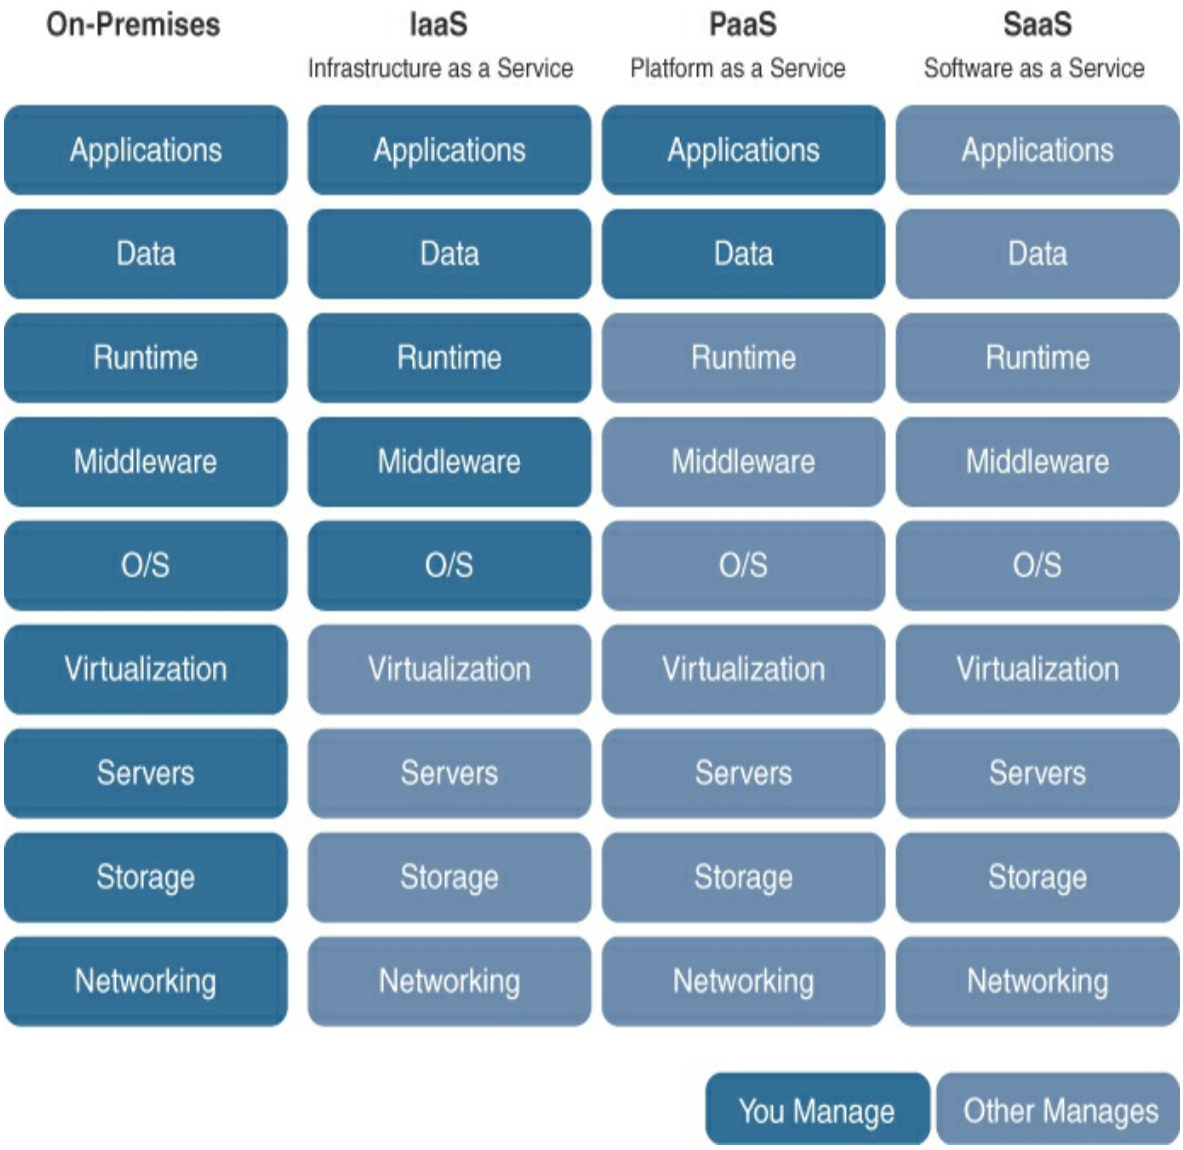

FIGURE 10-36 Comparing Cloud Services to On-Premises Services

The term *Infrastructure as a Service (IaaS)* represents handing off physical responsibilities to a CSP while maintaining control of what is installed and how that technology is configured. Physical responsibilities that are being outsourced to a CSP include powering the servers, racking and stacking equipment, and physically protecting the hardware. Everything else is still your responsibility as the customer.

The next level of outsourcing responsibility is *Platform as a Service (PaaS)*, in which you not only outsource the infrastructure but also outsource what is installed, including the operating system and middleware. For example, if

you don’t want to deal with installing or upgrading a generic Linux system, you can use PaaS to deliver a prebuilt Linux server, allowing you to deal only with the applications that are installed. PaaS is great for anyone who wants a standard build of an operating system without the worry about software updates; however, if your SOC needs customized systems, you will want to use IaaS.

The final option for outsourcing services to the cloud is *Software as a Service (SaaS)*, in which everything is outsourced to the CSP. In a SaaS environment, you the customer have limited responsibilities that include the data and how the service is being used, but the CSP handles everything from what is installed to where it exits. With SaaS, you can’t install anything because you don’t have any backend access to the application. The benefits of this approach are that you don’t have to build the solution and the CSP handles all aspects of maintenance.

An analogy I like to use when explaining different cloud services and why customers choose each option is comparing choosing a cloud service to choosing what to eat for dinner. Your on-premises option is to plan your own meal, grow your own ingredients, and cook and eat your meal at a time of your choosing. The downside of that approach is the amount of time and work needed to do everything yourself. You will have to spend months growing all of the ingredients, preparing the meal, and cleaning everything up. This example is similar to running your own datacenter.

Another option for your dinner is to go to a grocery store and buy the ingredients but cook the food yourself. You save a lot of time by buying the ingredients rather than growing them, but the quality of the ingredients is probably not as good as if you grew them on your own, unless you pay for very high-quality ingredients. You can still customize the meal with the ingredients you purchased, meaning you maintain a lot of the responsibility for the meal and you also must clean everything up afterward. This second example represents IaaS, because you outsourced growing the ingredients to the grocery store (and its supply chain).

One dinner option growing in popularity is meal kits such as what Hello Fresh or Blue Apron offer. You are provided a box full of ingredients and instructions for how to create the meal; however, you can use those ingredients however you see fit. The cost for this option tends to be lower

than going to a restaurant but higher than if you were to go to the grocery store and buy the ingredients on your own. You save time by having what you need to cook based on your goals for the meal you select for each day. This approach is similar to PaaS, in which you are provisioned an OS by the CSP you can use for any purposes, but that OS is typically sized for a specific need.

Finally, you could just go to a restaurant and order your dinner. This approach is the most efficient way to get a meal; however, it is also the most expensive and the option that you can't modify. You also don't have to clean up after you eat, which is really nice. This final example represents SaaS, where you outsource everything to a restaurant.

In each example, the more you outsource, the higher the cost for the meal and the less control you have over the meal, which can also mean lower or higher quality depending on what you are willing to pay. This is very similar to the logic with selecting cloud services. Organizations choose a service based on the level of control they want, what responsibilities they do not want to handle, what they are allowed to outsource, and what they can afford.

Orchestration in the Cloud

Your SOC is or soon will be using different cloud services to accomplish goals in your SOC services. Applications will be moved from the intranet to an IaaS environment or outright replaced by a comparable SaaS offering. Your SOC will need to be able to continue to improve services using orchestration and automation regardless of where the technology or data sits. You cannot let a cloud migration break your orchestration capabilities.

Looking closer at Figure 10-36, the different levels of outsourcing also represent what you as a customer can configure. Limitation in configuration options can also mean limitations in automation and orchestration depending on the CSP and what you are looking to accomplish. Chapter 1 already covered addressing some of these limitations in the context of securing different cloud services such as Amazon IaaS versus a SaaS service such as Dropbox. For IaaS, I covered various tools that can be installed, while SaaS environments are limited to cloud access security brokers (CASBs) and

access control technologies.

Configuring security is different than configuring DevOps. The good news regarding DevOps concepts in the cloud is that the leading CSPs are heavily investing in pushing DevOps capabilities with the intent to encourage customers to use them for orchestration. The CSPs are doing this based on the cost savings organizations are seeing when leveraging orchestration and automation for services being outsourced to the cloud. As cost savings and capabilities offered by CSPs improve and show value to customers, more customers will be willing to hand off responsibilities to the CSPs, leading to more demand for cloud DevOps services; hence, more profits for CSPs, making cloud-focused DevOps a booming marketplace.

Amazon DevOps

One example of cloud DevOps options from a leading provider is what is available from Amazon at https://aws.amazon.com/devops/. Amazon defines DevOps as "the combination of cultural philosophies, practices, and tools that increases an organization's ability to deliver applications and services at high velocity: evolving and improving products at a faster pace than organizations using traditional software development and infrastructure management processes." A shorter way to summarize this explanation is Amazon sees DevOps as a way to improve how you can offer services from the cloud.

As a SOC, you are offering various security services, and leveraging cloud services can add value such as converting capital expenditures to operating costs using a pay-as-you-go model for your applications, scalability and elasticity of services, and outsourcing of certain responsibilities to allow for much more flexible SOC services. Elasticity is a huge value add, meaning you can always scale up if a major event occurs and throttle down when services are not being utilized. Considering Amazon's size, your SOC could offer unlimited capacity and only be charged for what you need for your mission.

Amazon has categorized its DevOps-focused tools into the following groups:

- **Continuous Integration and Continuous Delivery:** These tools allow you to securely store source code as well as automatically build, test, and

deploy applications. Tools include workflows for releasing software, code testing, deployment automation, and even a user-friendly interface Amazon calls AWS CodeStar, making the entire process streamlined.

- **Microservices:** Container and serverless computing is a massively growing field that enables more specific workflows and applications to be developed. Amazon offers support for Docker and AWS Lambda to meet such needs.
- **Infrastructure as Code:** For customers who don't want to reinvent the wheel, Amazon offers code templates that can be leveraged to quickly enable automation of monitoring and enforcing infrastructure configuration.
- **Monitoring and Logging:** Amazon offers specific services targeting cloud-based network and application monitoring. Services include Amazon CloudWatch and AWS X-Ray.
- **Platform as a Service:** Amazon offers the option to deploy web applications without needing to provision and manage the infrastructure and application stack.
- **Version Control:** Customers can use Amazon to host secure, highly scalable Git repositories in the cloud.

Many of these services target IaaS; however, services such as AWS Elastic Beanstalk are focused on customers seeking DevOps capabilities within a PaaS offering. Elastic Beanstalk's automation can provision the deployment of the desired platform, including considering the required capacity, load balancing, and enabling application health monitoring, all based on the code you upload to Amazon. Learn more at https://aws.amazon.com/elasticbeanstalk/.

SaaS DevOps

SaaS is becoming an extremely popular option for delivering SOC services. SaaS, however, can have some limitations to accomplishing DevOps goals.

SaaS-based DevOps will be much more limited depending on what APIs are available and the purpose of the service. Typical CSP service-level

agreements (SLAs) not only explain what the responsibility of the service provider is, but also include responsibilities for allowing or not allowing different DevOps functionality. Know that you as a cloud customer can negotiate anything with a CSP, but the more you ask of the CSP, the more its service will cost you.

There are generally three types of CSP service-level agreements:

- **Single-customer SLA:** This type of SLA addresses a single service for a specific customer. If you provide service to three departments, such as HR, administration, and marketing, a single-customer SLA would apply the same guidelines to all three departments.
- **Customer-level SLA:** This type of SLA is more aligned with contracting with CSPs that market services to the general public. Because these CSPs offer similar services to thousands of different customers, they rely on generic SLAs that don't have to be tailored to each customer. Each service being provided will have a separate section in the SLA.
- **Multi-level SLA:** This type of SLA is a combination of a single-customer and customer-level agreement. This type of SLA is more common when a large organization outsources IT services, such as a large manufacture with hundreds of employees needing a place to create and store data. This agreement type covers services for each customer in the organization.

When evaluating an SLA, I stress the same advice I provided in Chapter 9 regarding cyber insurance, which is to make sure to include technical staff during the negotiation of what is being contracted. The most successful SLAs are simple, measurable, achievable, relevant, and time bound. *Simple* means the language is clear regarding what is expected from both the CSP and you, the customer. *Measurable* allows you to keep track of service use so that billing is predictable. *Achievable* means you have developed an SLA based on goals you can accomplish with the service. Measurable is also linked to achievable in that you must be able to measure your success rate toward achieving your goals. *Relevant* validates that the services are providing enough value to deem them worth the investment. *Time bound* ensures that you reevaluate services at some point in time.

Keep these points in mind as you evaluate potential SLAs with CSPs. In addition to including technical staff to help with defining your requirements, include legal experts to validate the language accurately depicts what you are

hoping to achieve. Always keep in mind when evaluating cloud service providers that almost anything is possible, but at a price!

Summary

Security orchestration, automation, and response (SOAR) allows SOC engineers to concentrate on meaningful tasks, problems, and investigations in their environments and rely on automated scripts for more repetitive tasks. However, a SOAR is not just about automatically repetitive tasks. The power of a SOAR is that it dramatically reduces an organization's attack surface by creating standard-based deployments of technology. Human error is eliminated, and deployments are based on tested playbooks that make the SOC more secure. Standard-based deployments also make it easier for the SOC to find potentially vulnerable software and services, in addition to threat actors that might be targeting their organization. Implementing SOAR concepts is not only an essential part of any security operations center but also one of the pillars a SOC must be built on.

This chapter covered all aspects of SOAR. First, you learned about the SOAR marketspace and the differences between SIEM, SOAR, EDR, and XDR platforms. You also saw examples of enterprise SOARs and learned what capabilities are common across SOARs and XDR offerings. You also learned how SOARs perform orchestration and automation by diving into DevOps programming. You learned about file formats, how to enforce data modeling, and how DevOps applies to your SOC's services regardless if they use tools in the datacenter, across the network, or in the cloud. You also learned how to set up an open-source SOAR known as Ansible and how to use that lab to automate blueprinting using Osquery. The chapter also cited many resources for SOAR and DevOps concepts, including free labs and training guides that can help your SOC maximize a SOAR's capabilities.

At this point of the book, you should have a firm grasp of the key concepts to planning, building, funding, and maintaining a mature SOC environment. The only topic left is what you should expect the future to hold for your SOC and other SOCs. Chapter 11 wraps up the book with a look at the future of the SOC.

References

Ansible Project. (2019). Installing Ansible. In *Ansible Installation Guide*. Red Hat, Inc. https://docs.ansible.com/ansible/latest/installation_guide/intro_installation.html

Amazon Web Services (AWS). (n.d.). What Is DevOps? AWS. https://aws.amazon.com/devops/what-is-devops/

Creately. (2020, December 1). Ultimate Flowchart Tutorial (Complete Flowchart Guide with Examples). Creatley Blog. https://creately.com/blog/diagrams/flowchart-guide-flowchart-tutorial/

Incident Response Consortium (IRC). Playbook – Virus Outbreak. IRC. https://www.incidentresponse.com/playbooks/virus-outbreak

ISACA. (n.d.). COBIT 2019 Framework (various publications). ISACA. https://www.isaca.org/resources/cobit

Oren Ben-Kiki, O., Evans, C., & döt Net, I. (2009, October1). YAML Ain't Markup Language (YAML) Version 1.2. YAML.org. https://yaml.org/spec/1.2/spec.html

Vertis. (n.d.). DevOps Automation Tools: Configuration Tools for Continuous Deployment. Vertis. https://www.veritis.com/solutions/devops/made-easier-with-devops-tools/#:~:text=DevOps%20tools%20facilitate%20ways%20for,for%20an%20effective%20product%20output

W3Schools. (n.d.). Introduction to XML. Refsnes Data. https://www.w3schools.com/xml/xml_whatis.asp

W3Schools. (n.d.). JSON – Introduction. Refsnes Data. https://www.w3schools.com/js/js_json_intro.asp

Yuan, S. (2019, May 22). The Anatomy of NetDevOps. Cisco Blogs. https://blogs.cisco.com/developer/anatomy-of-netdevops

Eric Ahlm, Augusto Barros, Michael Clark. SOAR: Assessing Readiness Through Use-Case Analysis. https://www.gartner.com/en/documents/3981938/soar-assessing-readiness-

through-use-case-analysis

Chapter 11

Future of the SOC

The best way to predict the future is to create it.

—Abraham Lincoln

Welcome to the last chapter of this journey into the SOC. This book has covered topics ranging from how to build a SOC to what types of practices are common in mature SOCs around the world. You have explored what works today in a mature SOC, but this leaves one final topic open, which is what you should expect to see in future SOCs. Technology is rapidly changing and so are the tactics used by cybercriminals. Responding to these changes requires adapting your SOC practice or eventually you will encounter a compromise. Security is a journey, not a destination. You don't become secure and move on to another project. Instead, you continuously observe and adapt as described in the OODA loop (covered in Chapter 10, "Data Orchestrtion").

This final chapter provides an overview of security industry trends and what I believe will be key elements for the success of the future SOC. Topics are based on a combination of what industry analysts have predicted for the security industry, expectations industry analysts have for how technology will improve, and what I'm hearing in the industry based on industry conferences, meetings with customers, and expectations industry analysts have for future compute power. Think about these concepts and push yourself and SOC to be forward thinking as you plan your next SOC project. It is critical to be forward thinking as you mature your SOC services.

All Eyes on SD-WAN and SASE

The first concept is one that every organization's C levels have been talking about since early 2020, which is how to address the need for a new network architecture approach designated by Gartner as *secure access service edge*, also known as SASE (pronounced "sassy"). SASE represents merging WAN and security technologies into a single cloud-delivered service. The crazy thing is there wasn't a single vendor that offered SASE as it was first defined by Gartner in 2019, but since late 2019 just about every security vendor has claimed that it can provide SASE. The reality is that many vendors are offering only some of the capabilities that are part of what SASE promises, causing an overload of inaccurate marketing with no specific reference to how Gartner explains SASE should be provided by vendors and service providers. This has caused a lot of confusion since Gartner introduced SASE, which is why it is the first future-looking topic that needs to be clearly explained in this chapter, as it will impact your SOC today and in the future.

If you want to be forward thinking, you need to develop a SASE strategy. Next, I'll explain why doing so is critical for the future of your SOC.

VoIP Adoption As Prologue to SD-WAN Adoption

To understand SASE, let's first step back and review a technology trend that was similar in its transformational impact on the IT industry but focused on how organizations run their phone systems. Prior to the invention of Voice over IP (VoIP), organizations would have telephone service provided by a telephone company and Internet service provided by an ISP. Phone system technology was very basic, providing dial tone and limited calling features. Sometime in the 1990s, VoIP started to be introduced as a way to combine phone service and Internet service into a single plan, saving organizations tons of money. On top of the savings, organizations also got to use bleeding-edge VoIP phones, which included centralized management systems with simplified user interfaces and other cool features that dramatically improved end-user satisfaction.

All the benefits of VoIP caused a major transformation in how organizations viewed what they needed for voice communication. By 2000, VoIP was considered a common technology in organizations. Retail stores to large enterprises all had computer-based phone systems. VoIP was an industry-

transforming technology, causing many engineers skilled in deploying and managing analog phone systems to have to switch to digital phone systems or face job obsolescence. Today, it is rare to find an enterprise using analog phone systems.

VoIP Benefits Summarized

Some of the reasons why VoIP changed the IT marketspace can be summarized as follows:

- Combined Internet bill and phone bill, leading to savings
- Offered bleeding-edge technology:
  - New IP phones
  - Centralized management
  - Bleeding-edge features such as customized ring tones and Music on Hold
  - Better performance monitoring to understand return on investment
  - More options for security
- Became a trend everybody else was investing in
- Perceived as providing a competitive advantage

Introduction of SD-WAN

Around 2018, a new transformational technology was introduced that had similar characteristics as VoIP in regard to the value being provided and the reasons why organizations became interested in adopting the technology: *software-defined wide-area networks*, more commonly referred to as *SD-WANs*. The traditional way to host multiple offices is to connect each branch office over a WAN link connection using a network protocol. The early days used point-to-point or Frame Relay lines; however, those technologies were replaced with Multiprotocol Label Switching (MPLS).

With MPLS, the organization is required to pay one service provider for each

branch location and traffic needs to be managed from a router connected to the service-provided MPLS connection in order for communication to take place between offices. The service provider controls throughput and is essentially a blind spot between branch offices that are managed by the organization's network administrators because organizations can send and receive traffic only through the service provider connection, but network administrators of the organization don't see what occurs while the traffic crosses into the service provider's network. Managing multiple branch offices means the organization must manage multiple routers that provide the MPLS connections to the service provider's network. If the router at either side of the MPLS connection experiences problems, the entire site will lose connectivity unless a high-availability link is in place, costing the organization even more for the additional MPLS service. Figure 11-1 represents a traditional MPLS network design.

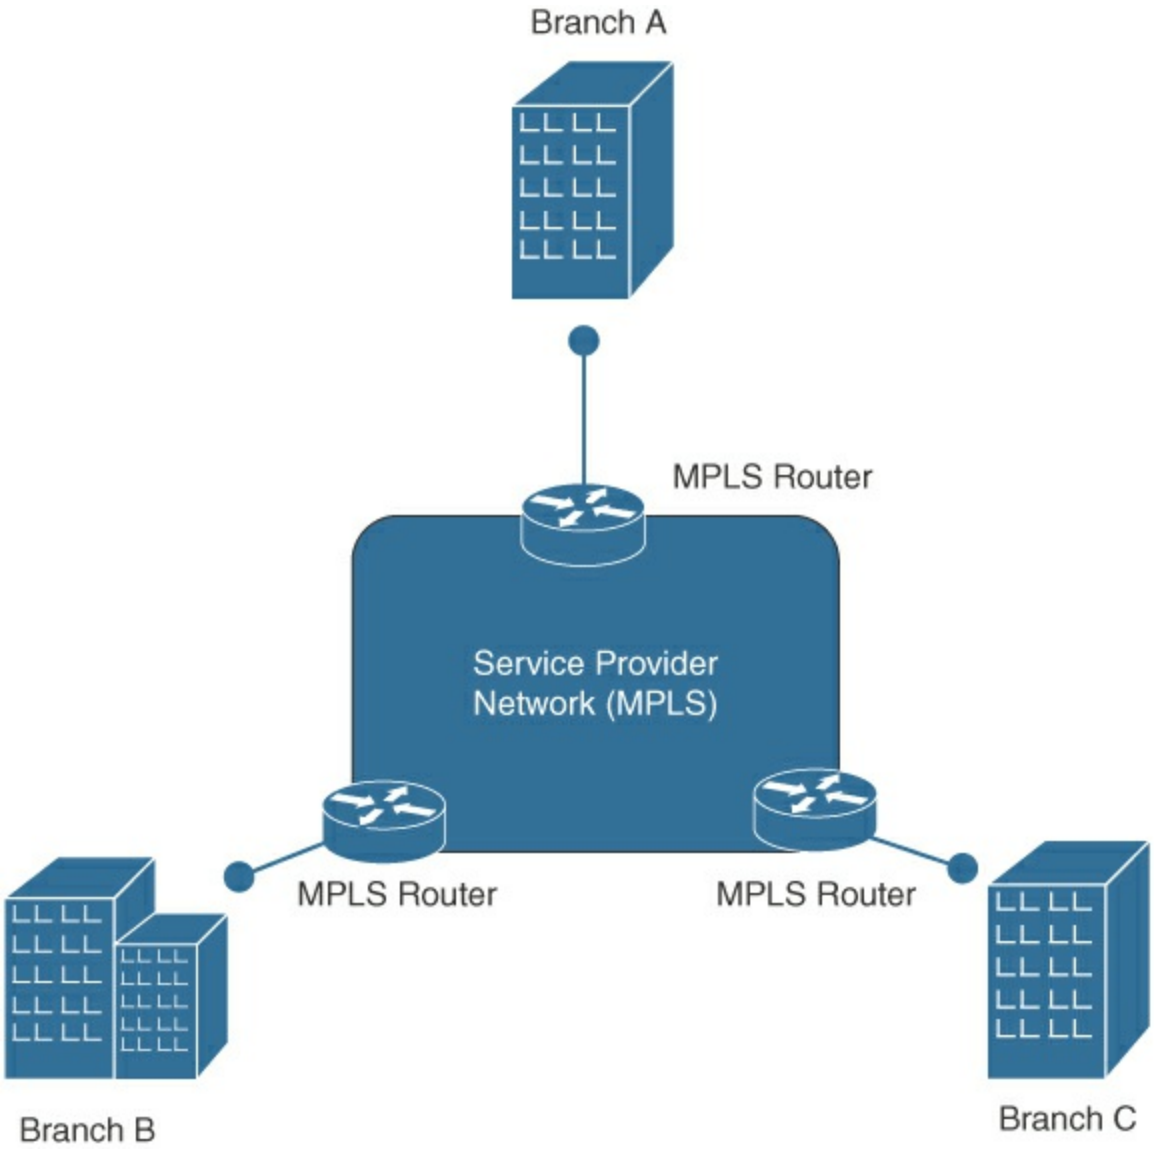

FIGURE 11-1 Traditional WAN Connection Between Different Branch Locations

Challenges with the Traditional WAN

In the past, the traditional network traffic was predictable and security needs were simple to understand regarding securing access to critical resources in the datacenter. The majority of datacenter-based traffic would be internal and the remaining traffic that required being sent outside the organization would be forced through a centralized point known as the network gateway, which

included a stack of security capabilities to protect the traffic from cyberattack. The gateway security stack would typically have a firewall, IPS, and possibly additional file analysis technologies that could run files in a sandbox to better understand their intent or send them to a cloud team to analyze. Any remote users would be required to use a VPN to gain access to the datacenter, and all of their remote traffic would be forced through the security stack if the traffic was going in or outbound based on what was sent through the VPN tunnel back to the internal network. Figure 11-2 represents a high-level view of traditional network and security for the datacenter.

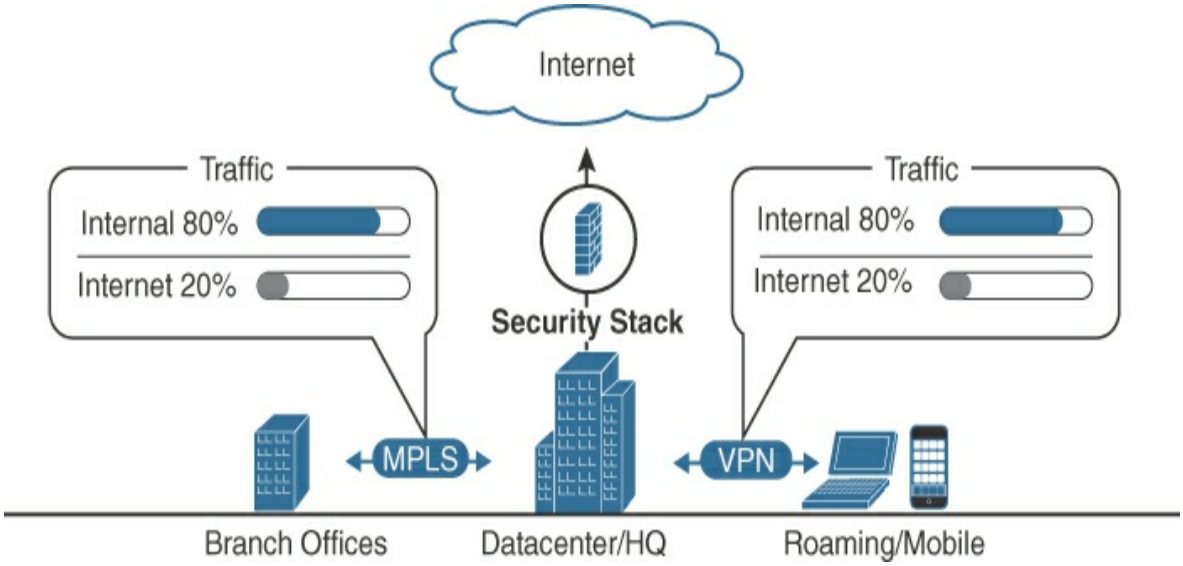

FIGURE 11-2 Traditional Network and Datacenter Architecture

Over time, new challenges required multiple exceptions to be made in order to maintain operations. First, the datacenter's traffic needed more external access to the Internet, swinging from an 80/20 ratio of internal versus external daily traffic to a 20/80 ratio. The increase in external-bound traffic put stress on the gateway, causing a performance bottleneck to occur. The bottleneck problem increased as remote-user traffic was forced through the gateway, leading to the entire organization experiencing performance issues. The VPN overload created the need for *split tunneling*, which separates normal user traffic from traffic that needs to go over the VPN. Split tunneling reduced the stress on the gateway; however, it also exposed user traffic to the risk of attack because it was no longer protected by the security tools in the

gateway. On top of all of these challenges, cloud technologies such as Software as a Service (SaaS) and Infrastructure as a Service (IaaS) started to be used without controls. This created an unknown data-sharing and operations space that the organization's security team could not see (this space is sometimes referred to as *dark operations* for this reason). Figure 11-3 represents a high-level view of these challenges applied against the traditional datacenter and network architecture.

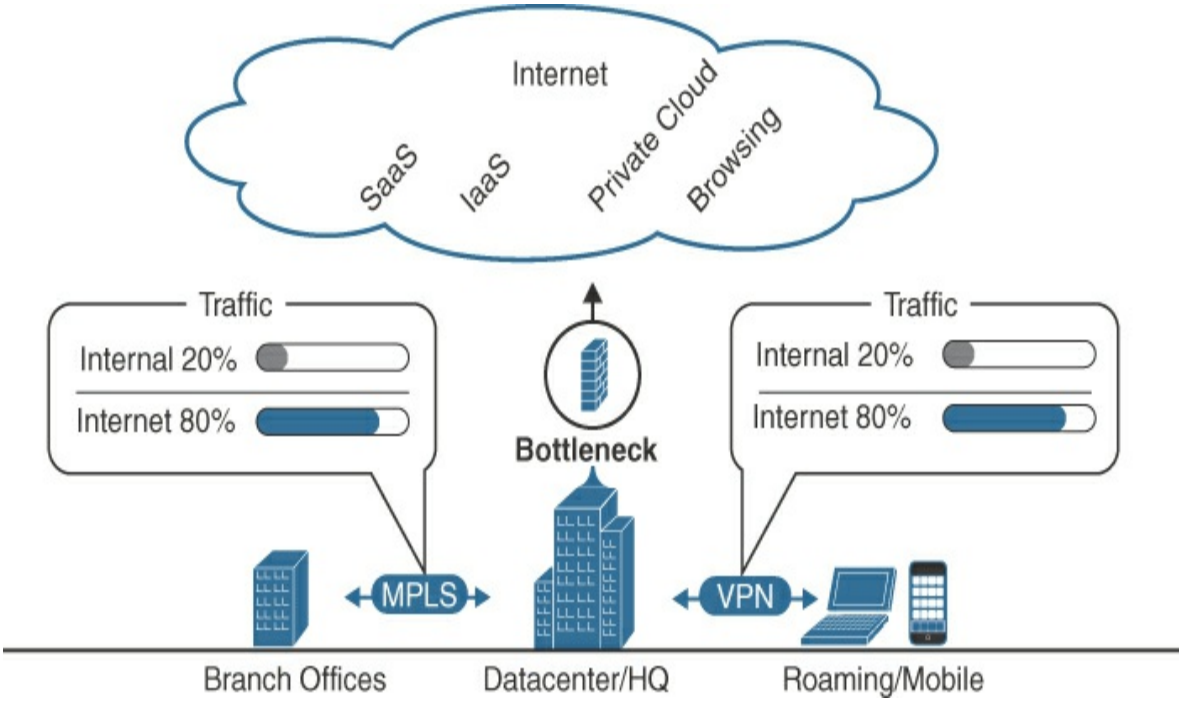

FIGURE 11-3 Challenges with Traditional Networking, Datacenters, and Roaming Users

---

MPLS Failure!

On a side note, I have a true story regarding what can go wrong when an MPLS line goes down. Early in my career when I was a network administrator, I was required to be on call one week a month, which entailed carrying an emergency cell phone for use during network outages and answering it regardless of the time of day. One night around 12 a.m., as I was starting to DJ at a local Washington DC

nightclub, the cleaning crew at my employer's datacenter bumped a router's MPLS link, taking a critical branch office offline. Half an hour later I received a call from a tier one monitoring service informing me that a site had gone down and I needed to go in to fix the problem.

A friend filled in for my DJ gig and I summoned a cab to the office. Meanwhile, the network was down, and IT was waiting for me to arrive onsite to troubleshoot the problem. Upon arriving 45 minutes later, sure enough, within a few seconds I was able to see the link lights were off on the router port connecting one branch office and the organization's headquarters. I saw that the cable connecting the branch office was slightly plugged into the router and simply pushed the cable back in. Within a minute, I had confirmed the branch office was back up based on the MPLS link reestablishing connectivity and things were back to normal, minus my night being ruined.

To reduce the risk of the same situation occurring, my team racked the routers differently by locking the racks holding the equipment and rerouting the cables under the raised floor, allowing the cleaners to work without the risk of bumping into a cable and taking down critical services. We also added additional MPLS lines for each branch office, dramatically increasing our service provider bill. The network going out caused sufficient panic in leadership to make them realize the need for more high availability, turning a bad night into a hard lesson learned regarding the need for proper redundancy in a core network infrastructure.

The key point to this story is that the service provider could only tell tier one support that they were not seeing traffic, tier one support could only see that a line was down, and it was up to one person (me) to figure things out by physically looking at the router. This approach to managing critical networks is not effective; hence, the need for a better way to monitor, troubleshoot, and manage networks!

SD-WAN to the Rescue

SD-WAN changes the old model of connecting branch offices and managing

north–south traffic from the datacenter by allowing an organization to use other protocols such as Long-Term Evolution (LTE) and Digital Subscriber Line (DSL) along with MPLS when a more reliable connection is needed between branch offices. Allowing the use of other protocols translates to saving organizations a ton of money as well as allowing for the most optimal path based on the application being accessed. The key to the value of SD-WAN is that the SD-WAN service, managed by the organization (not the service provider), controls the allocation of bandwidth. In addition, SD-WAN supports much improved high availability because it allows multiple services to run in an active/active type of architecture—if one line goes down, another line can take on the workload. Looking back at my real-world story of what can go wrong with a traditional MPLS connection, if SD-WAN was being used, the branch office that went down when the cleaning person bumped the connection cable between offices could have automatically failed over to another line such as a much less expensive DLS line, keeping the network up until operations returned to work at a more reasonable hour. SD-WAN removes the need for an on-call technician to arrive onsite late at night! Figure 11-4 represents the transformative change SD-WAN makes in how the branch, datacenter, and roaming user access cloud resources. This approach removes the dreadful bottleneck by not forcing traffic through a centralized gateway.

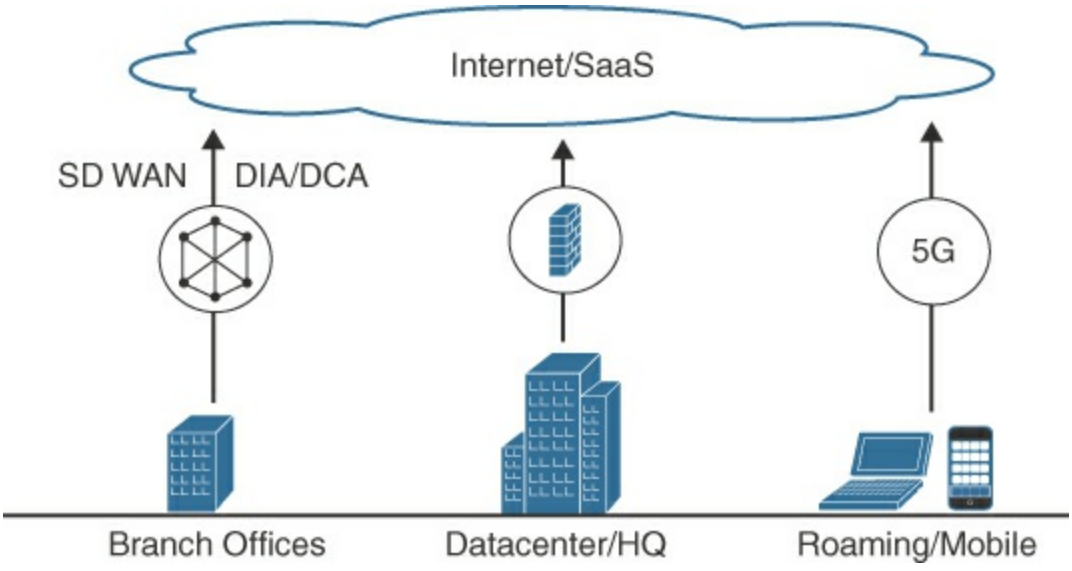

FIGURE 11-4 SD-WAN Approach to Connectivity

SD-WAN Benefits

SD WAN offers much better management technology than MPLS, including options to measure performance and improve response time when there are potential issues with connectivity between sites. Having more visibility means IT has a much better understanding of the problem. Looking back at my example, if a connection issue occurs, IT can identify whether the issue is a problem at the application layer, a problem on the branch side of the connection, a problem with the traffic traveling over a WAN connection, or a user problem, whereas in the real-world MPLS scenario I described I had to be physically onsite to figure out the problem. As stated earlier, SD-WAN reduces the likelihood that an on-call engineer will be engaged because the technology is able to better understand the problem and offer automated corrections when a problem occurs. Figure 11-5 shows an example of the Cisco vManage SD-WAN management dashboard summarizing various WAN performance statistics. SD-WAN dashboards such as this provide a ton more information than traditional WAN management platforms.

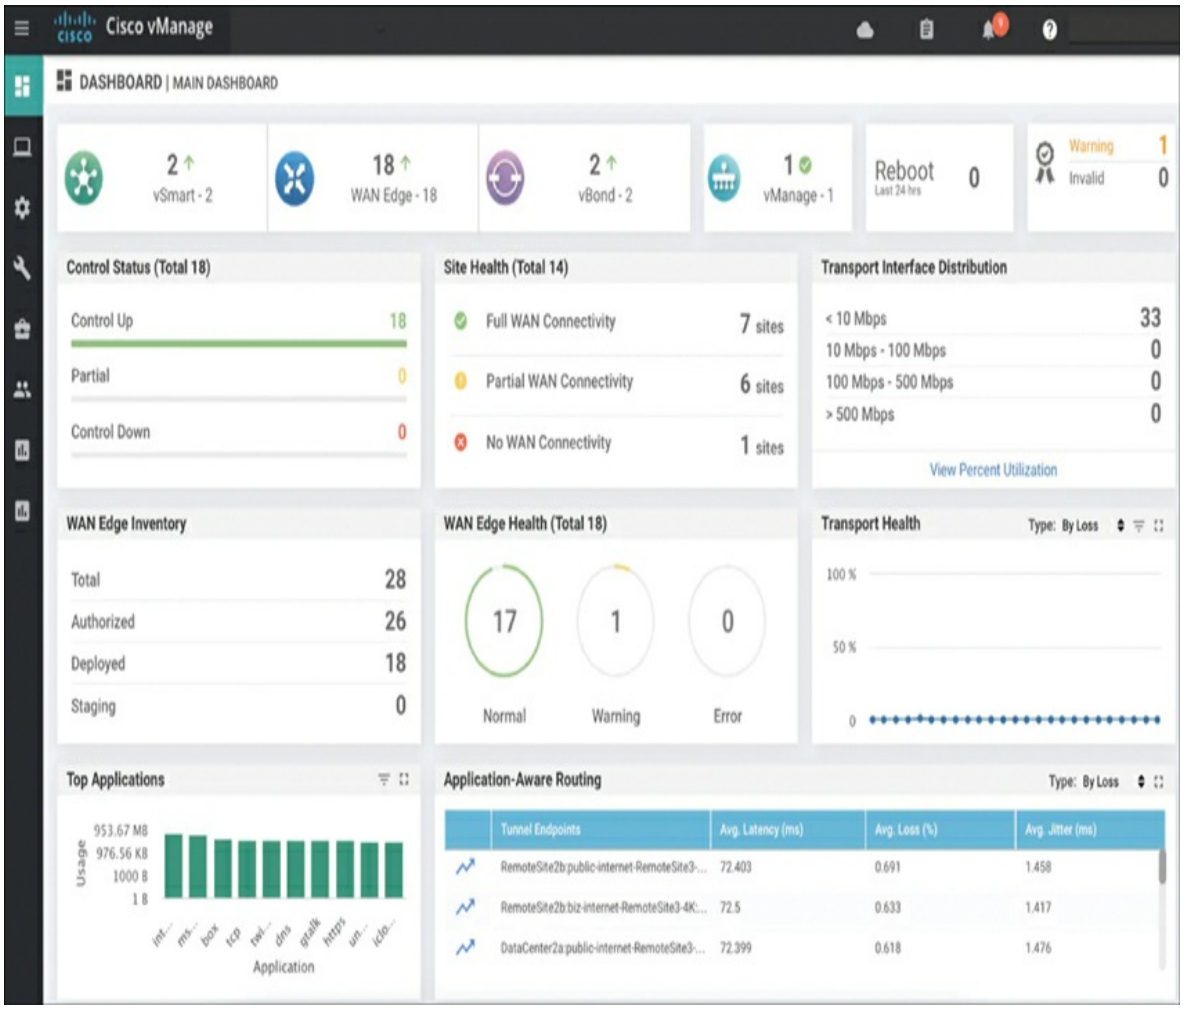

FIGURE 11-5 Cisco SD-WAN Dashboard Example

To summarize the value of SD-WAN, the following are five core values gained by investing in SD-WAN:

- Software intelligence exists within a network virtual environment providing agility and flexibility using a web portal configuration, versus managing routers and traffic monitoring taps.
- SD-WAN leverages an Internet connection, which is much cheaper than an MPLS service line; however, MPLS can be used when dedicated lines are needed.
- SD-WAN is an agnostic technology, allowing you to select the right connection for your business, whether it is using the Internet, MPLS, or other service.

- SD-WAN is able to terminate any connection from any source and apply security, load balancing, and local traffic shaping depending on network conditions.
- SD-WAN builds on the traditional WAN router and is revolutionizing the telecoms market.

If you compare why organizations have invested or will invest in VoIP technology and SD-WAN technology, you will find that both technologies offer similar value points and both are considered major disruptions to how they have impacted IT services. Both technologies provide huge cost savings by reducing what organizations must pay to an external service provider and both offer modern technology as part of the migration plan, which includes better management, troubleshooting, and an overall better experience for end users and IT managers. Unlike VoIP, however, a major catch to operationalizing SD-WAN is causing many organizations to hold back from making a full SD-WAN investment. That major roadblock to the success of SD-WAN is security!

SASE Solves SD-WAN Problems

Figure 11-4 showed a general concept of the SD-WAN architecture, which essentially involves connecting branch networkers over the Internet without using a service provider's network. Any security professional will point out that a major concern with this approach to connecting your offices is that SD-WAN exposes to the public Internet what in the past was somewhat hidden through a private service provider network, putting SD-WAN traffic at risk of attack from anywhere in the world. Imagine an organization that traditionally sent all traffic from multiple locations around the world through a service provider network, yet only had to focus on traffic hitting the Internet at a few gateway points now having every branch office represent a network gateway that must be secured. When I meet with organizations excited about SD-WAN and I mention the concept of securing every network point as part of the deployment, common responses are "Wait, now we need a firewall and IPS at every branch office?"; "Wait, now every branch office can be attacked by malicious parties somewhere on the Internet?"; and "Wait, now we're exposed to exploitation in a way we can't monitor and are not ready to

protect against?" These concerns are real, which is why every SD-WAN conversation must also be a security conversation. Security must be included in order for SD-WAN to be operationally feasible, which is why Gartner coined the term as *Secure* Access Service Edge.

SASE combines network security features with WAN capabilities to support the dynamic secure access needs of organizations exposing every branch and its users to the risks associated with the Internet. Security for organizations also needs to extend beyond the branch, and so does Gartner's vision of SASE, meaning all aspects of cloud security must be included. SASE includes connectivity between clouds, the use of Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and even the remote worker accessing the Internet from a mobile device over a cellular network. For WAN capabilities, Gartner's SASE model leverages the value provided by SD-WAN, including performance, troubleshooting, and managing networks, combined with security best practices recommended in popular industry guidelines such as those published by ISO and NIST. SASE considers how the future employee will access applications from their office or from the cloud, meaning traffic might never cross an employer-owned connection yet still needs to be secure. A true end-to-end SASE offering as defined by Gartner represents all of this as a unified service.

Note

I have been asked if secure SD-WAN and SASE are the same thing. My answer is no. Securing SD-WAN is only part of the SASE story. Securing SaaS usage (such as Dropbox or Box) is an example of a cloud problem that has nothing to do with SD-WAN, yet is part of the Gartner SASE vision. Another example is that Gartner's SASE model includes data loss prevention concepts. I consider securing SD-WAN to be a subset of SASE, but more is required than just securing SD-WAN to qualify as a SASE offering.

Gartner Shoots the SASE Flare Gun!

Gartner's SASE predictions represent a warning to the IT industry of a major

disruption that will occur. Cloud adoption is going to happen at a pace that hasn't previously happened. SD-WAN will become as common as VoIP is today in most organizations. Solving the security challenges that SASE solves will be a normal expectation from IT providers. Those that don't see the warning flare from Gartner will miss a huge market shift and likely not survive. Gartner is throwing huge numbers behind its SASE prediction with the following statements:

- By 2023, 20% of enterprises will have adopted secure web gateway (SWG), cloud access security broker (CASB), zero trust network access (ZTNA), and branch Firewall as a Service (FWaaS) capabilities from the same vendor, up from less than 5% in 2019.
- By 2024, at least 40% of enterprises will have explicit strategies to adopt SASE, up from less than 1% at year-end 2018.
- By 2025, at least one of the leading IaaS providers will offer a competitive suite of SASE capabilities.

To be clear on the numbers, 40% of enterprises represents billions of dollars of potential business for IT vendors. On top of that, these predictions show customers investing more with a single vendor, meaning the vendor with the best SASE offering has the potential to claim many areas of security and network business from firewall and branch service providers. This is why many experts are warning that if an IT company does not adapt to the SASE model, it will become obsolete and locked out of future business. As a result, multiple SASE-focused acquisitions occurred in 2020, including the following:

- Palo Alto Networks acquired CloudGenix
- Fortinet acquired OPAQ
- VMware acquired Nyansa
- McAfee acquired Light Point Security
- Zscaler acquired Edgewise Networks
- Cloudflare acquired S2 Systems

What are all of these organizations trying to achieve through their acquisitions? Read on to find out.

SASE Defined

Figure 11-6 represents a summary of different capabilities that are converging into the SASE concept from both the “connect it” side, meaning SD-WAN, and the “secure it” side via security capabilities. The challenge for this to occur, as pointed out by Gartner, is that many security solution providers do not offer SD-WAN or have experience with providing networking services, while many SD-WAN service providers have not built security into their offering as described by the Gartner SASE model and instead must partner with a third-party bolt-on security approach. This is also why Gartner stated in 2019 that no vendor was currently offering SASE as a single service. Some organizations were closer than others in 2020, and all security vendors who plan to stay in business while offering the capabilities listed in Figure 11-6 continue to adjust their strategies to accommodate the expected SASE disruption to the IT market. Gartner sees all of the components in Figure 11-6 as being a single service for a single cost, which Gartner believes is a true SASE service.

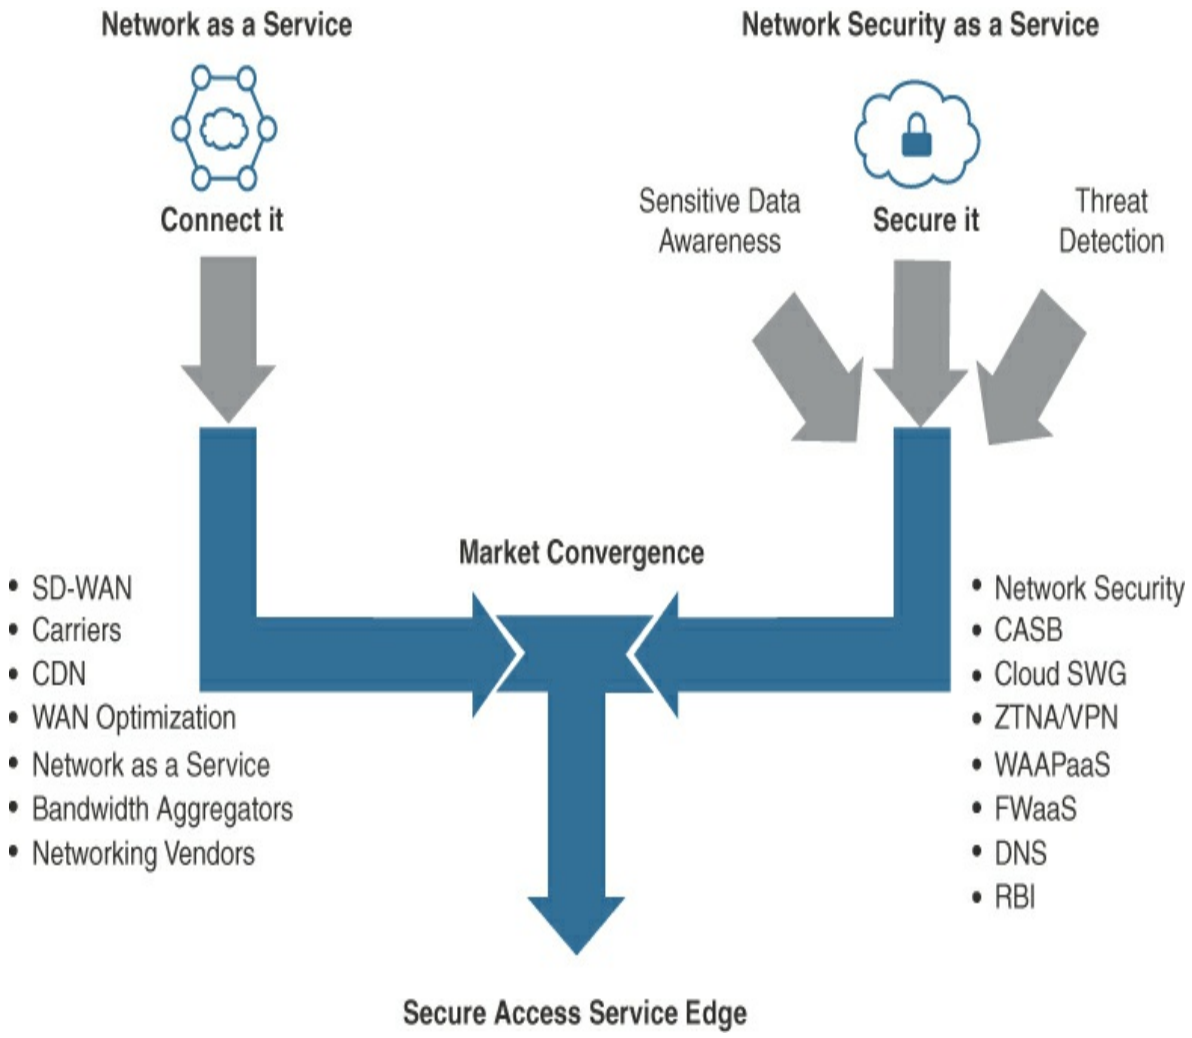

FIGURE 11-6 Services Combined to Form SASE

Future of SASE

Based on conversations I've had with executives of enterprises around the world, I've identified core fundamental SASE capabilities and methodologies that will exist in the SOCs of the future, some of which may already be present in your SOC. The data points in the sections that follow are lessons learned from the disruption of SASE and my predictions for how SASE will impact the future SOC. It is important to be forward thinking in your SOC planning by considering how cloud topics such as SASE will shape technology and how organizations will leverage IT services in the future.

Prediction: Integrated Network and Security

Organizations will look to acquire SD-WAN with security included versus keeping security as a dedicated expense and technology to manage. Organizations are already recognizing SD-WAN as a way to reduce the cost of services, but they will not want to spend more to secure it than they can save by switching to a SD-WAN architecture. Organizations are more likely

to invest in SD-WAN if it is secure (as described as SASE) and still provides cost savings after everything is acquired and deployed. Referring again to the example of VoIP adoption, migration to VoIP accelerated rapidly after the industry realized switching to VoIP had an extremely short return on investment, meaning the entire system would pay for itself within three to five years. If a SASE offering can demonstrate a similar quick return on investment with security included, organizations will rapidly jump on the offering. Organizations also will not want to spend more to secure an SD-WAN offering than they can save by switching to an SD-WAN architecture.

Note

In my role as a trusted advisor for organizations ranging from fortune 500 to federal government, I meet with customers interested in buying SD-WAN. I find much lower success rates when SD-WAN is covered by a trusted advisor without security being addressed since it leaves a huge gap in what needs to be addressed for the deployment to be successful. I tell my sales teams the likelihood of a customer adapting SD-WAN when security is addressed upfront increases by as much as 80%, according to previous sales history that is tracked based on win rate percentages. This is based on my personal experience with various types of organizations.

Prediction: SaaS Is the Future of Security

Software as a Service (SaaS) will continue to grow and it needs to be secured with technologies like cloud access security broker (CASB) and access control. Although SaaS means you are outsourcing most of the responsibilities to a service provider, you are still responsible for your data. I point out the concept of securing data in a SaaS environment because you must secure your data with the limited technologies that are available, while other aspects of security are covered by the SaaS provider. The future of SaaS will include security such as CASB capabilities as part of the offering, representing a flavor of SASE by combining a service with security for a single cost or as an option in a SaaS service-level agreement. This makes

sense as CASB is a security tool that acts as an intermediary between the users and cloud service provider, allowing the users' organization to enforce desired security policies and monitor how SaaS is used. I predict CASB offerings will become part of other services and will no longer be a dedicated solution, as it just makes sense to combine this with many SaaS services.

Prediction: Dynamic User/Device Fingerprint

Identifying a user and/or device is and will continue to be based on multiple factors, which includes where the user/device is located, what type of system the user is on, and what applications the user/device is permitted to access. Results of how a user/device is fingerprinted can dynamically determine what access rights they will be provisioned, regardless of where they are located, meaning the place of work can be anywhere in the world at any time. This type of technology exists today; however, it will continue to get better. Areas of improvement include the accuracy and details associated with a user/device, the potential risk, relating the user/device to historical data, comparing the user/device against behavior and attributes found with similar users/devices around the world, and continuously adapting to changes. I have seen movies depicting these types of futuristic features, such as a device scanning a user's eye from across the room to provision various services at an office or an automated voice greeting by name a customer walking into a retail store and then suggesting what the customer should purchase based on previous buying habits. I have seen aspects of these features already used in the real world, with collaboration technologies leveraging big data trends blended with scrapping social media accounts and facial recognition technology.

Prediction: OPEX Replaces CAPEX

Organizations will increase investment in subscription-based technology with the goal of converting CAPEX (capital expenses, meaning to buy something) to OPEX (operational expenses, meaning the cost to do something), reducing the need to acquire and manage technology. Many IT hardware providers have already made this shift, and the future of successful technology providers is dependent on reoccurring revenue business models. For example,

most vulnerability scanning providers in the past sold their scanners only as software you downloaded, installed, and managed (CAPEX). Today, it is more common for vulnerability scanning providers to offer the same technology from the cloud and as a subscription service (OPEX). From a customer viewpoint, the customer doesn't have to deal with setting up or updating the product, and management of their scanning is accessible from anywhere that has Internet access. From a vulnerability service provider viewpoint, the vendor can quickly add new features, onboard customers, and troubleshoot tickets, enabling them to offer a better experience to customers. SaaS also curtails pirating of software and allows for more control over how customers can test out their solution using a cloud-based proof of concept approach.

Prediction: The Office Is Everywhere!

User performance needs and requirements for flexibility will continue to increase, forcing organizations to move security to the application and host rather than requiring traffic to tunnel through a security stack. In older network designs, organizations would build a defense-in-depth design at the network gateway and put a VPN concentrator at the front end. Any VPN user requiring access to internal resources would have to send traffic through a VPN back to their organization's VPN concentrator, forcing traffic through the security stack within the gateway. Throughput requirements have already forced many organizations to leverage split tunneling—sending only those requests that need to go to the organization through the VPN while separating and sending all other traffic over the Internet. I previously explained how battling the bottleneck problem, meaning having too much traffic forced through a stack of security solutions, had led to a major driver for SD-WAN.

User traffic going through the gateway security stack is becoming more difficult to inspect due to SSL/TLS encryption, because data privacy rights require certain traffic to maintain its encryption even if decryption options are available. Secure Internet gateway (SIG) is a SASE concept that addresses this design problem by forcing user traffic through a cloud-based security stack, relieving organizations of the responsibility to force traffic through their corporate-owned security stack. SIG is growing in popularity and is a concept that will become the standard way to secure traffic from employees

who are not physically at a corporate office. Essentially, security will follow the users, devices, and data rather than everybody having to accommodate restrictions based on how the network is designed. The need for encryption will not go away, but VPN as a service will be the more common method to encrypt than traditional VPN hosted from a VPN concentrator owned by an organization. The COVID-19 pandemic has prompted many organizations to adopt a work-from-anywhere mentality to accommodate their workforce having to work from home. Many organizations I have spoken with since the pandemic began are developing a work-from-anywhere architecture not only to improve work efficiency but to prepare for any future event that would force their employees to work from home. Some organizations have even announced permanently using a work-from-home model as it has shown to work during the pandemic.

Prediction: Increase Data Loss Prevention Needs

Data loss and data at rest will continue to increase in demand as both technology options become easier to deploy and manage. Frameworks such as Zero Trust are extremely popular in concept but difficult to operationalize. In theory, the concept of encrypting important data is easy to state as a requirement; however, the biggest hurdle is determining what data is considered important and requires protection.

Predefined data such as credit card numbers and Social Security numbers are easy to set up for monitoring; however, automatically labeling an organization's data as important enough to be encrypted is much harder based on all the dynamics of what can be considered important and how fast data changes. Advances in technology based on machine learning and artificial intelligence are allowing for quicker and better decisions to be made regarding the importance of data, which can translate to a more accurate and automated data loss and data at rest technology. The future of the modern SOC will include automated data loss and data at rest technology as a backbone of its security practice similar to how antivirus is used today.

Prediction: Goodbye Tier One Support

Tier one support in most organizations represents the group responsible for

basic customer issues. They are seen as the first responder when an employee has technical needs. I predict tier one support for services will be fully automated in the future SOC. As SD-WAN adaption increases, organizations will have access to more details on how their services are functioning. These details will open the door to automated response when issues are found in a service and automated response when a potential security event is occurring.

By automating responses to these situations, tier one support will not be needed because the network will know when an issue or attack is occurring and either resolve it on its own or send event data directly to higher-tier support when its automated response fails to remediate the situation. Because SASE is a cloud-based service, support for tier one services will move to the service provider, saving the SOC the need to offer any low-level support services. Automated tier one support is in common use by businesses via automated assistance that attempts to solve your problem rather than sending you to a human.

Prediction: Automated Upgrades

With SASE, software upgrades will just happen rather than be a responsibility of the SOC. SASE means a service is provided by a cloud service provider, translating to the service provider always having access to the management tools. If new features are needed, the entire configuration and upgrade process can occur without the end user being notified because backup systems can be used while primary systems are being upgraded, leading to zero downtime during the upgrade process. This can and will all occur in the cloud without customers' knowledge, removing technology maintenance requirements from the SOC's list of responsibilities. Tesla provides a great example of this concept. Tesla owners agree to receiving updates, and their parked vehicles will inform drivers that an update is available and will be installed when they are not using the car. The Tesla owner doesn't have to do anything but wait for updates to be pushed from the cloud to their vehicle. SASE allows for the same update strategy to occur. I point out this example because technology continues to be part of all industries, leading to SASE impacting more than just our computers.

Another software development concept that will impact automated upgrades is agile software development methodologies. Agile software development

means smaller chunks of a program are developed and released more often. Using this approach also allows for pushing upgrades and patches along with software releases, providing a much smoother upgrade experience than waiting and combining multiple changes into one larger release. Cloud complements agile software development because the cloud service provider can continuously push smaller updates with little risk of impacting its customers.

SASE Predictions Summarized

The potential value of SASE to the SOC includes shifting the responsibilities of certain services to the cloud, which frees up time for SOC analysts to focus on different tasks. SASE technology also solves some challenges that SOCs deal with, reducing or even outright eliminating some risk that today's SOCs must account for.

Pay attention to the SASE trend as it impacts all of IT. There are many resources you can follow as SASE matures. I recommend using a blend of reading publications from industry firms such as Gartner, monitoring IT vendor acquisition trends, and speaking with other organizations regarding their cloud adaption roadmap. If you do not already have a SASE strategy today, now is a good time to start looking at how you plan to adopt cloud technology into your SOC's practice.

The next topic regarding the future of the SOC is how the services that it provides today will change in the future. Any forward-thinking SOC must consider not only what services the organization needs from the SOC today, but also how those and other services will change over time.

IT Services Provided by the SOC

Another concept that continues to change is how services are provisioned to customers. The term "customers" can mean the people buying something from an organization, the employees needing services from their organization to do their jobs, or the entire organization needing security services from the SOC. Regarding servicing employees, an organization is expected to deliver

a range of general services that enable employees to be productive. For many organizations, common services they provide to their employees include the following:

- Safe Internet connection
- Desktop and phone
- Collaboration tools for teams
- Access to corporate data
- Access to personal data
- Computer power to host applications
- Secure way to use third-party services
- Workflow and sales tracking technology

> **Note**
>
> This list is a general look at IT services employees expect from their employers. Many of these services are not related to what the SOC is responsible for, which is the point of my predictions to come!

IT Operations Defined

Organizations of the past designated all of the preceding services as the responsibility of "IT operations," a group within the organization operating out of its physical headquarters. As an organization of the past would grow, it would need more physical office space, it would need to provision more devices, and consequently it would require more staff responsible for supporting IT operations. If remote workers required external access to the corporate network to do their jobs, IT operations would configure tunneling technologies such as VPN to enable them to gain secure access to the services they needed. End-user equipment would be configured by IT operations at a corporate office and shipped to employees or made available in a corporate office to pick up. In the past, the SOC was not part of the team performing IT

operations' responsibilities, but sometimes the SOC would advise on security policies for assets that IT operations managed. Those policies included host security standards, requirements for VPN, password policies, and web policies.

If the SOC doesn't handle IP operations, how does the concept of IT services relate to a SOC? With the growing trend toward SASE, the focus is a convergence of "connect it" and "secure it" (as depicted earlier in Figure 11-6), meaning blending together IT operations and security. As this occurs, the SOC will need to be involved with IT operations because network connectivity and other IT services will have security as part of the package, meaning it will be one single service. If the SOC isn't involved with a SASE-based IT service, blind spots will emerge leading to dysfunctional SOC services.

The following are some examples of aligning IT operations to potential future security capabilities that the SOC would need to be responsible for:

- **Safe Internet connection:** Security visibility and controls will be included with many network services based on the SASE unified service concept. The SOC will need access to all security stats and will manage security policies such as what can be accessed and what to block as potential threats. Event data will be exported to the SOC's log management tools. All of these controls lead to providing end-users a more secure Internet connection.
- **Desktop and phone:** Collaboration technology continues to shift from hardware to cloud-provided services. A Chromebook is an example of hardware that is dependent on the cloud for many of its services. Most Chromebook capabilities are not available without Internet access. Another example is Cisco Meraki, which uses the cloud to manage the physical hardware. If the hardware is disconnected from the Internet, it can't be configured or updated, and sometimes might not even be functional. If desktops and phones become service-based technologies, the SOC will have to be involved to ensure those services are secured. The only way the SOC can accomplish applying security is to have involvement in the control of policies over such devices.
- **Collaboration tools for teams:** As a work-from-anywhere approach grows in popularity (or from necessity), collaboration tools will increase

in importance. This approach can also lead to serious data loss concerns based on the sensitivity of data that can be shared. Due to data loss risks, the SOC of the future will need to police the content being shared. Trends show collaboration tools being consumed in a SaaS format; therefore, the SOC of the future might own the responsibility of securing the data within collaboration tools while an external service provider handles all other management tasks.

- **Access to corporate and personal data:** Data privacy has become and will continue to be a huge concern for employees and customers. The challenge is that there is a constant battle for enforcing security (allowing security tools to evaluate data) versus maintaining privacy (not giving access to data), because maintaining privacy can lead to not having traffic evaluated by security, resulting in security gaps. This challenge justifies the continued involvement of the SOC with data loss and privacy policies. I predict data loss prevention will continue to grow in complexity and importance, eventually leading to a dedicated service offered by the SOC of the future.
- **Computer power to host applications and a secure way to use third-party services:** I pointed out in Chapter 10 that DevOps is a rapidly growing field of expertise based on the need to configure technology to work with other technology. Part of this demand is the need to host technology that can help accomplish an organization's mission. As organizations rely more on applications being stood up by employees, the risk of cybercriminals targeting those applications will increase. The SOC needs to include this risk as part of its risk management service. The same concept applies to leveraging third-party applications and services, which is when an organization outsources required applications.
- **Workflow and sales tracking technology:** The crown jewels of most organizations is their data. For organizations involved with sales and developing technology, tracking is critical to understanding how the business functions; therefore, the tracking system must be protected from cyberattack. The SOC needs to responsibly protect such tracking systems.

As you can see based on technology trends, the SOC of the future will have more IT operations responsibilities than the average SOC has today. Many organizations have kept IT operations and SOC teams siloed, but technology

trends such as cloud and SASE adoption are forcing a merger between compute, network, and security. I predict every organization will be forced to readjust how they view IT operations and the responsibility of the SOC as these technology trends continue. To be a forward-thinking SOC, you need to consider how to incorporate IT services, as that will eventually be part of the SOC responsibilities, if it isn't already assigned to the SOC.

The merging of IT operations and security will lead to new attack vectors. A few years ago, my good buddy Aamir Lakhani and I performed a penetration test targeting IT services. IT services in our targeted organization were not part of the SOC and were extremely vulnerable to social engineering. This use case demonstrates why it is critical to secure IT operations. The section that follows provides an overview of that research, which demonstrates the danger of keeping IT operations separate from security.

Hacking IT Services

Back in 2013, Aamir Lakhani and I spoke at the RSA Europe Conference about our research that showed flaws in the traditional model of provisioning IT services to employees. We created a fictitious person named Emily Williams using social media resources. She is not real; however, we used photos of a real person who allowed us to use the photos to create a completely new identity within social media circles. We created a Facebook account and a LinkedIn account for "Emily Williams" showing she was a new hire for an organization whose leadership authorized us to perform a penetration test of the organization's security capabilities and IT operations without the knowledge of the IT staff.

Hello World, Meet Emily Williams

Playing the role of Emily, we became connected via social media to key members of the organization we were penetration testing. Within a few days, we had over 170 employees of that organization associated as digital friends to our fake account. On LinkedIn, Emily Williams was receiving endorsements for her fake certifications. On Facebook, people were congratulating her on her new job and offering help as she started her fake

role in the organization. In a short amount of time, Emily Williams was welcomed in social media circles as a new hire looking to connect with fellow employees. Figure 11-7 shows LinkedIn endorsements for Emily's certifications. Some of these endorsements came within two hours after we created her account.

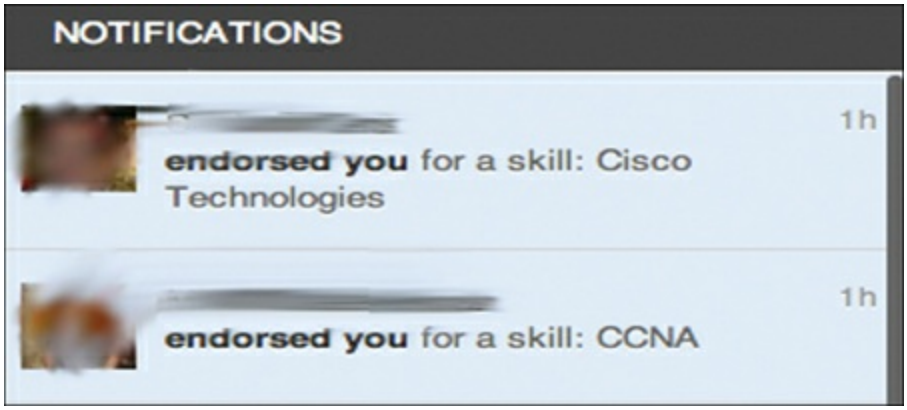

FIGURE 11-7 LinkedIn Endorsements for Emily Williams

During my initial outreach to targets using Facebook, one IT administrator at the organization being penetration tested asked "Emily" if he knew her after receiving her friend request that I sent, as shown in Figure 11-8. I quickly looked at his Facebook profile and saw that he had worked at Hungry Howie's Pizza 10 years prior to the conversation, so I immediately jumped on that topic and responded that I ("Emily") knew him from back then. My strategy was based on the assumption that it would be hard for this person to remember the girlfriends of people he knew 10 years ago. That strategy required finding another person he knew 10 years ago to be Emily's former boyfriend. I clicked the target's Facebook connections and saw he had a friend named Derrick that worked at Hungry Howie's during the same time period. Derrick's Facebook page showed that he currently was living in New York, allowing me to quickly make up a story about being Derrick's girlfriend 10 years ago and recently running into Derrick while in New York. I explained that Derrick told me the IT employee I was speaking with worked at the organization Emily was recently hired at; hence, the reason for the Facebook friend request. The person not only accepted my friend request but also provided me valuable information about the interworking of the organization's IT operations.

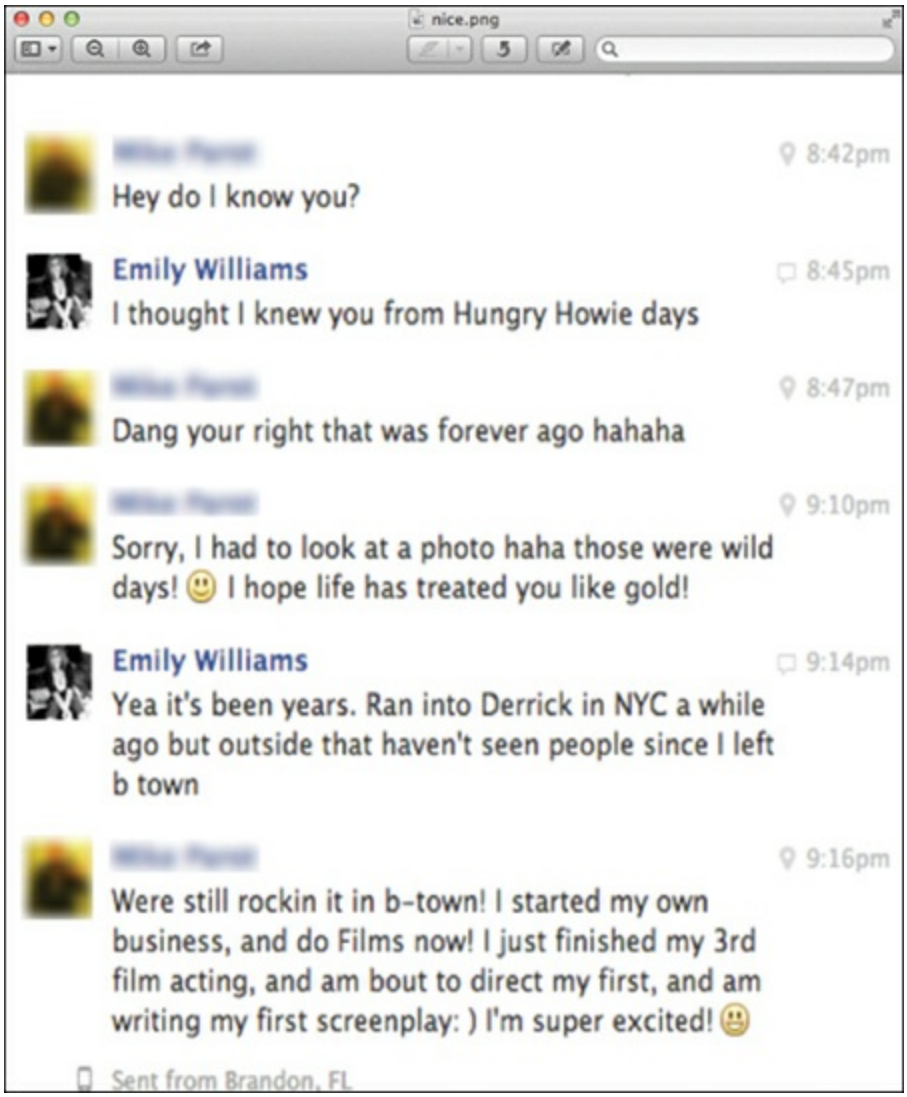

FIGURE 11-8 Facebook Social Engineering Attack Example

Impact of Emily Williams

The results of the penetration test enabled my team to conduct further social engineering that resulted in receiving a phone, a computer, and access to the target's network without stepping foot into the organization (all equipment was shipped to us and preconfigured). With leadership's explicit authorization, we were able to abuse vulnerabilities in how services were distributed, which led to a complete compromise of the organization.

Compromise not only occurred through receiving equipment from the organization, but also occurred through social engineering attacks launched

from Facebook that allowed my team to collect login credentials to remote worker VPN accounts using a social engineer penetration tool known as the Browser Exploitation Framework (BeEF). The goal of our presentation at RSA Europe 2013 was to show the vulnerabilities of social media as well as weakness in human trust, including those people responsible for delivering services to employees: IT operations. Figure 11-9 is an example of one of the many articles covering our research results.

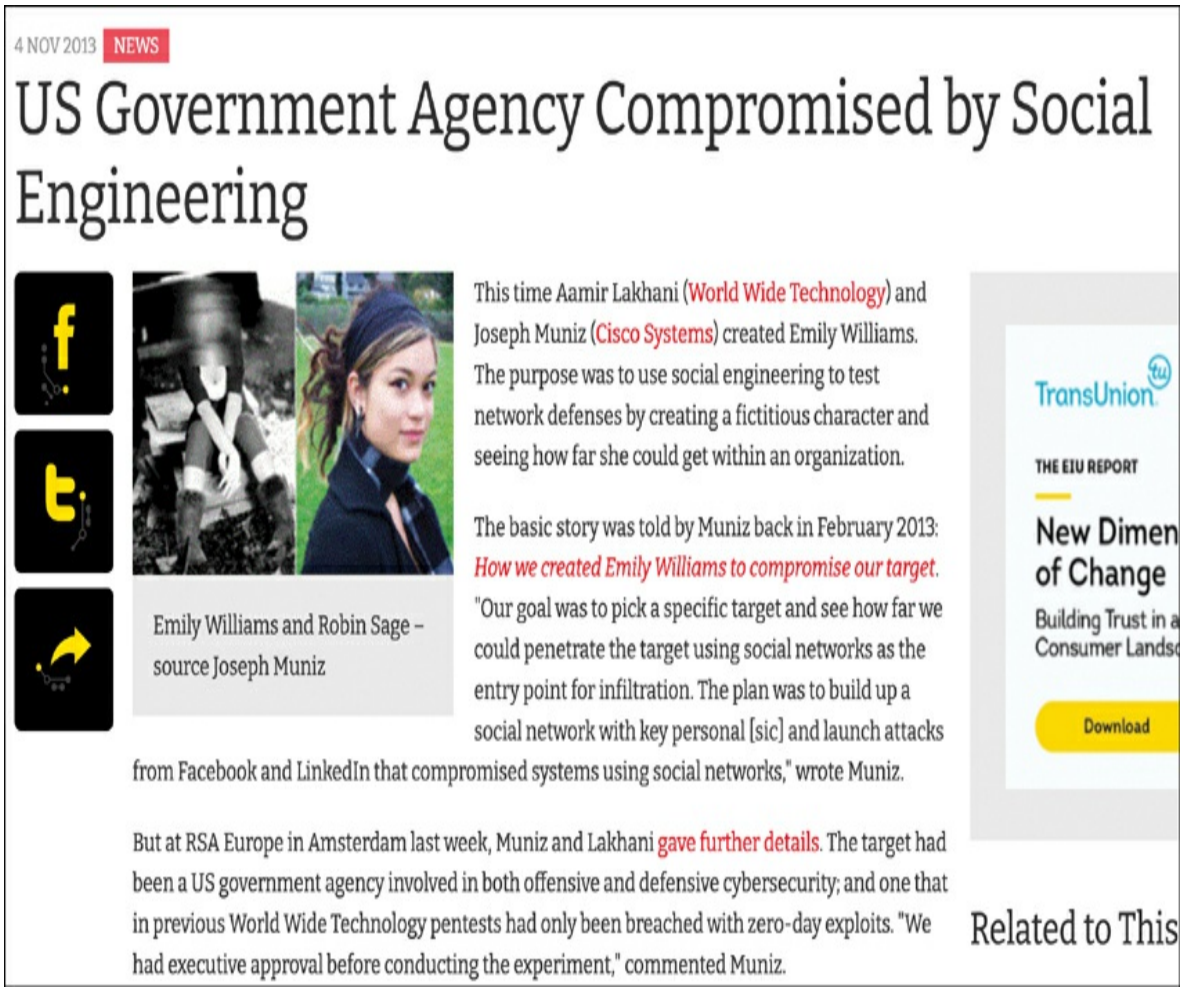

FIGURE 11-9 Article About the "Emily Williams" Penetration Test

The key takeaway from this research is the potential danger of not securing IT operations. I predict IT operations and security will merge and be part of the SOC due to the types of risk exposed in our research project. Every forward-thinking SOC must consider securing IT services, which implies the

SOC being involved with IT service security.

IT operations has changed since we conducted our Emily Williams research, and events such as COVID-19 have escalated certain trends such as the work-from-anywhere concept. Next, let's look at recent IT services trends.

IT Services Evolving

Over recent years, the traditional model for providing IT services to employees has been falling apart. The most drastic change is how organizations and SOCs are removing physical boundary requirements. In 2020 I spoke with dozens of CEOs during the COVID-19 pandemic and heard the same phrase over and over again: "We thought if everybody had to work from home, our business would shut down. We were surprised to be wrong, and even more surprised to find that many parts of our business were actually more productive!" Some organizations such as Google announced short sprints of required remote work programs for their employees, expiring when COVID-19 concerns were no longer relevant. Other organizations such as Twitter and Facebook announced a commitment that their employees could work remotely for as long as they choose. The removal of physical boundaries makes the old way to deliver services operationally challenging and dramatically increases the risk of exploitation, such as what my team showed back in 2013 with our Emily Williams penetration test.

SASE and IT Services

As I stated earlier in this chapter, the future of SOC services will include SASE components. A major driver for this shift is the need to adapt to a work-from-anywhere architecture. Using cloud services means that technology doesn't have to be shipped to an employee and that the risk of compromise can be shared or altogether handed off from the SOC to a service provider—other than securing the data being used. Data security will continue to be a fundamental responsibility of the SOC even after SASE becomes the norm.

Cloud technology allows for more flexibility, enabling employees to obtain the services they require rather than receiving tools they might or might not

use. I personally have worked remotely for the past 15 years and rarely use an employer-provided VoIP phone. In order to use the physical VoIP phone, I have to plug it into my home network and use it as a desk phone; however, I don't require that type of service for the work I do. Instead, I prefer to use a softphone or my mobile phone, making the company's investment in my physical VoIP phone a waste of IT resources. In a large company, if you multiply the number of remote employees that don't use their company-provisioned VoIP phone by the expense of provisioning and managing the unused VoIP phones, the sum will be a large amount of wasted money. One huge benefit of cloud computing is the ability to convert CAPEX to OPEX, which would be ideal versus always automatically provisioning a physical phone when a softphone is all that is needed.

Future of IT Services

Once again, I have some predictions that impact the future SOC. This time my predictions are for how IT services will be offered in the future. Many of these predictions are related to services that are already being offered, but for different industries and for specific use cases rather than for the average SOC. Because the SOC will have IT operations responsibilities in the future, you should be aware of these trends as they become part of normal IT operations within the SOC. Every forward-thinking SOC needs to have a plan for incorporating responsibilities for securing IT services.

Prediction: 3D Printing Resources

The IT operations function within the SOC of the future will offer a catalog of services that employees can request as needed rather than being provisioned as a set of services. As an example, 3D printing technology introduces the interesting concept of enabling employees to print something that was created by an outside party. I predict that in the future, if an employee needs a VoIP phone, they will be emailed a 3D image and be authorized to print their phone, removing the need to ship equipment. The services for the phone will automatically be set up, creating an extremely efficient process to provision such services.

When the asset is no longer needed, the service to the device will be automatically terminated and the user will be instructed to destroy the asset or mail it to a third party responsible for asset destruction. In the example of a VoIP phone, the future 3D printer could also offer the option to recycle the material, meaning melt the phone back into material that can be used to print future assets. Some manufacturers are already offering 3D sketches of replacement parts for their products. Rather than having to purchase a replacement part, a customer can simply download the 3D diagram and print what is needed.

Prediction: Virtualized Computers

Another service model that is shifting to the cloud is how compute power is provided to employees. Organizations today provide employees with computers for daily work. If higher compute power is needed, organizations either upgrade an employee’s computer or provision access to high-end

servers in a corporate-owned datacenter. This approach to provision resources requires interaction with multiple departments, including the datacenter administrators to create the account and service, network administrator to connect the system to the network, and security administrator to open required services in existing security tools. Due to the multiple parties being involved, adjustment to services is tedious to enforce, causing security exceptions to exist beyond the life of the service and wasted services based on requests for more than what is needed to avoid having to go through the request process in the future.

The future of provisioning compute is to offer it from the cloud based on business need. When a user needs compute resources of a certain type or size, that user will be able to access a system that can quickly make the desired resources available. When the resource is no longer needed or is not actively being used for a certain time period, it will dynamically be reduced or be deleted. I predict even personal computers will function in this manner. Everybody will have template hardware similar to the Chromebook model; however, the organization will determine what resources that hardware will have access to. The future might even allow for template hardware to be vendor agnostic, meaning an employee could be using Windows while employed with a company, but later have the software switch to Linux on their personally owned hardware template based on a new company sending it different compute services. The employee would simply provide their MAC address and their employer would be able to send everything from the operating system to what is installed.

Another huge value of using cloud computing is the flexibility to adjust services when demand increases. Microsoft defines the ability to dynamically deal with peaks in IT demand as *cloud bursting*. According to Microsoft, Azure always has additional resources available to the customer via cloud bursting, and the customer is charged only for additional services used. This approach to provisioning services allows organizations to offer an almost limitless amount of services (setting aside concerns about associated costs to the organization). The concept of cloud bursting is the future of provisioning services, eliminating the traditional approaches of buying more services than needed just to accommodate future requirements or limited high-usage scheduling. I can imagine a future SOC analyst needing to run a task requiring heavy compute power and his/her laptop being able to pull down

the needed resources dynamically. An example could be needing to decrypt packets using a brute-force technique, which would be much quicker if a supercomputer could be accessed for this purpose for a short period of time.

Prediction: Cloud Management Platforms

Management platforms for IT operations and security products are shifting to the cloud. This makes sense because having a hardware platform for management of technology affords no benefit, outside of supporting network environments that do not permit Internet access. The benefits of moving management platforms to the cloud include the following:

- Easy access from anywhere
- No hardware maintenance, power cost, or labor for updates
- Easy vendor access for support and troubleshooting
- Simplified integration with other tools because management is already in the cloud
- CAPEX converted to OPEX
- Much more robust high availability

Many vendors are already moving to this model. Certain customers, such as those that do not permit Internet access to their management tools, will need hardware options; however, I predict the majority of all IT and security tools used by the future SOC will be managed from the cloud. This includes firewalls, IPS, routers, switches, honeypots, and sandboxes.

IT Services Predictions Summarized

Of all my predictions, the biggest prediction I'm making regarding IT operations is that it will become part of security and therefore be a responsibility of future SOCs. Other predictions for IT operations fall in line with the predictions I made regarding the impact of SASE to the SOC. Tier one support will go away. Updates will be automatic. The office will be anywhere the employee can work. Cloud services will be the standard for delivering IT operations. Many of these predictions are starting to happen

today and others might have already happened within your organization. Make sure your SOC has a plan for rolling in IT services, as all signs point to this becoming a normal practice within the SOC.

The third topic to dive into regarding future SOCs is how employees will be trained. Remember that security involves people, process, and technology. I predict there will be major changes in how people are trained as well as what training will be needed within the future SOC. Every organization leader I meet with has training as part of their agenda for maturing their SOC.

Future of Training

People are an organization's greatest asset. The most skilled and best performers need to be provided a competitive work environment, including the opportunity for career advancement or they will go somewhere else that offers their desired future role. Unskilled and newer employees need to be trained to support the skilled and best performers so that they eventually can take over the tasks of other skilled employees. Doing so opens up the possibility for veteran employees to move to a higher rank within the organization with new and more challenging responsibilities. This entire lifecycle of employee development requires offering quality training to employees so that they can learn new tasks and eventually increase their responsibilities. This section focuses on trends in training today as well as what the future of training might look like for your SOC. Since training is such a critical element of improving the SOC, any forward-thinking SOC will be looking at training trends to take advantage of the latest training offerings.

Training Challenges

A common goal of leadership teams across organizations is to improve how training is delivered in their organization. I continually hear from leaders I meet with from various organizations that a top concern for their organization is finding and maintaining the right people.

According to Karen Greenbaum, president and CEO of the Association of

Executive Search and Leadership Consultants (AESC), in her article "The Five Top Talent Challenges of Today's C-Level Executives" (*Forbes*, 2017), a survey of senior executives concluded the top talent challenges are as follows:

- Lack of diversity
- Lack of key leadership successors
- Competition for talent
- Mismatch of current talent and future strategies
- Need for digital experience

Training is a critical component to address all five of these challenges. Lack of diversity means organizations are not supporting certain groups of people to grow their careers into leadership roles. There are many reasons why this can occur in an organization. One cause is not providing the right training for people at lower levels to be able to perform at a higher level of leadership. The right training not only includes how to be a leader but also provides the opportunity to interact with leadership, as business connections are part of what leads to people being promoted to executive roles.

The same concept regarding lack of training applies to the challenge of not having leadership successors. If you don't train people to be able to perform at a leadership level, then you won't have successors at hand when a leader takes on another role or leaves the organization. Top-performing employees want to improve their abilities and advance to higher ranking roles in the organization, which requires a history of strong performance along with training. Not providing training will allow another organization to poach your employees by offering higher ranking roles with better pay. People leave an organization for a reason, and it is not always purely for an increase in pay. People want career growth opportunities, and if they don't find them in your organization, they will find them somewhere else.

Technology continues to change, and training is required to keep up with it. Without training, your SOC won't have the people to run the latest technologies, leading to a SOC with limited capabilities. An example of a skill needed for running security orchestration, automation, and response is DevOps, and I find some SOCs lack this skill within their staff. Knowing DevOps is just one of the skills that represents the digital experience needed

to manage current SOC technologies.

Note

Some organizations are making huge strides to address the challenge of a lack of diversity in leadership roles. While working at Cisco, I was identified as being a top performer and having Latino heritage. I was invited to join a group within Cisco that met with the goal of gaining exposure to senior leadership and being provided special coaching/leadership training to help top employees with a Latino heritage obtain leadership roles. Other similar groups within Cisco have been formed, such as women in IT. I believe these types of investments not only show that an organization's leadership is invested in addressing diversity problems, but also build employee loyalty to the organization.

Training Today

There are various ways organizations offer training to their employees. One option is to contract with an external training service to run a class, which can be at a discount per employee if enough employees sign up. Another option is to have the organization cover the expense of external training for employees who request it. This often entails attending certification bootcamp training, vendor product training, or specific skill training such as a programming language. Another training option is online on-demand courses, which tend to be less expensive than live classes because the class size is not limited.

Note

Some organizations offer to pay for employees' external training with the precondition that the employees sign a contract specifying that if they do not continue to work for the organization for a specific period of time, they must reimburse the organization for the training. Organizations do this to retain talent and avoid the situation of

investing in an expensive training for an employee that will leave before the organization benefits from providing the training.

SOCs can also develop their own “over-the-shoulder” training options, which allows employees to share skills. The SOC’s situational and security awareness responsibilities can include running security training as well as offering options for non-SOC employees to shadow different SOC team members. This enables the SOC to not only identify employees within the organization who are capable of moving into SOC roles but also identify SOC members who show leadership skills. Sometimes organizations ask employees who have taken external training to build a course based on what they learned and train fellow employees. Many of these internal training options are advantageous because they do not have a cost outside of employee time to allow shadowing or develop training material.

Training Today Summarized

To summarize training today, I refer to the article “Top 10 Types of Employee Training Methods” by Corey Bleich (CEO of EdgePoint Learning, an employee training service provider), which lists the following top 10 ways organizations provide training:

1. Instructor-led training
2. Learning
3. Simulation employee training
4. Hands-on training
5. Coaching or mentoring
6. Lectures
7. Group discussion and activities
8. Role-playing
9. Management-specific activities
10. Case studies or other required reading

When I ask an organization’s leadership about how they deliver training to

their employees, the typical answer is that they use a blend of these options. My experience is that the most common options used by organizations are instructor-led training, eLearning, and hands-on training, all three of which tend to be outsourced to either the vendor of a certification program or a third party that is aligned with a certification program.

Next, I'll briefly describe which of the preceding list of methods I use when I require training.

Case Study: Training I Use Today

In my personal experience, I tend to use a blend of training resources. My training includes eLearning from free resources, lectures from industry conferences, case study research, conversations with experts on a specific topic, and lots of hands-on training using labs provided by my employer as well as labs I build on my own. When I need to aggressively shorten the time to learn required material for an industry certification exam, I take an instructor-led, bootcamp-style course after reading a book on the topic; I also run through tons of practice exam questions to ensure I mastered all required concepts. If I need to figure out how to make something work, I always start by searching the Internet to see if there is a short video or article that gives me enough details to accomplish my goal. If I need to learn about industry concepts or recent threats, I start with specific blogs and bookmarked resources rather than searching the Internet, because I trust the opinions of the experts I follow over anonymous Internet blog posts. I also have friends whose technical knowledge I trust and I sometimes call them to ask questions about concepts I need to better understand. In short, I leverage a lot of different resources for my training.

The following is a summary of the resources I personally use for training:

- **Instructor-led training:** Bootcamps I take when preparing for a specific certification
- **eLearning:** I use Google and YouTube to research concepts
- **Simulation employee training:** When my employer asks me to do mandated training such as ethics training

- **Hands-on training:** Labs I built, or employer provides
- **Lectures:** YouTube or bootcamps
- **Group discussion and activities:** Industry friends
- **Case studies or other required reading:** Internet research

Everybody has their own way to learn, and my personal case study on resources I use shows a blend of different training options. I can claim seven out of ten from Corey Bleich's top ten training methods. Some of these resources cost as much as a few thousand dollars, while others are free online options (or freely provided advice by people I trust to contact with questions, which I reciprocate). Sometimes I lean on a free resource, while other times investing in a professional resource makes sense regarding time saved and quality of content. You should expect a similar blend of training needs from all of your SOC employees.

Free Training

One life hack that can help save your organization's training budget is to ask product vendors for free training. Vendors want customers to try out their products with the intent of eventually purchasing from them. This helps overcome customer concerns that a large learning curve is required to use a vendor's product. I find many vendors offer really good free training to help encourage a sale, on par with third-party training that has a cost. Consider including training with a proof of concept before the sale or bundle training with a purchase to avoid having to invest in future training.

Online free resources can be very solid as well. Many universities are posting free versions of an entire education program, such as a master's degree in computer science or cryptology, for anybody to access and learn. Resources such as YouTube continue to grow their libraries of on-demand content, allowing anyone who is interested in learning a topic to instantly pull up material. I can't count the number of times I needed to figure something out and immediately searched YouTube and followed an expert's instructions. Consider all of this as you validate what your personal case study for training requirements or SOC staff requirements will include.

Note

I have an important tip regarding sharing your knowledge about your job role. Early in my career, the CEO of my employer gave me valuable advice that I live by: share your knowledge, and become so good at your job that your employer could replace you because you are no longer needed. Doing so allows you to be promoted as well as name your successors during the interview for a better role within the organization. People who hoard knowledge and are required for a position will be stuck in that position until they can be replaced by technology or new processes and will never be promoted because their employer needs them to continue in their current position. Try this tactic the next time you ask for a promotion. As you request new responsibilities, not only explain that the people you have trained can do your job, but specifically name a few employees who are ready for your role. Doing so will improve your chances of the promotion since your manager doesn't have to worry about backfilling your role. You also demonstrate leadership by bringing a solution to backfilling your current role.

Gamifying Learning

One interesting learning style that is growing in popularity is gamifying content. Learning is not always fun and can be overwhelming when combined with daily work and personal responsibilities. There is a growing field in learning management applications, known as learning management systems (LMSs), that aims to drive interest in learning through interaction by turning learning into a game. A paper by Juan Burguillo titled "Using Game Theory and Competition-based Learning to Stimulate Student Motivation and Performance" described a study that compared using a gamified approach to teaching topics versus using traditional methods. Results showed students using the gamified approach were much more motivated, were satisfied with the experience, and had a better understanding of the topic.

Successful video games tap into satisfying basic human needs by providing

small accomplishments as the player works through the game. The concept of "leveling up" is addictive and encourages the player to proceed through the game even if leveling up requires very tedious and repetitive tasks. By tapping into the desire to level up, training becomes a challenge that doesn't seem hard to the student if they can see success as a result of their time spent.

Learning Management Systems

Learning management systems apply gaming concepts to learning complex content with the goal of motivating students to complete tasks. An example of a popular LMS is Moodle (https://moodle.org/). I run a lab at Cisco that has a lab guide containing more than 300 pages of steps to perform all tasks. The older versions of my lab required a student to download the lab guide and perform all of the tasks as listed in a huge PDF file. I had no way to know how far students got in the lab, where they encountered problems (other than when they reached out to me), what areas they found the most and least interesting (unless I asked using a survey), and where they had to spend the most time. Therefore, I had trouble determining how and where to improve the lab to help encourage students to want to complete the lab. I also had no way to motivate students to overcome challenging tasks. Essentially, I was blind to how my lab was being used.

LMS technology (Moodle in my case) changed the way my team delivered the lab by allowing my team to associate points with completing steps. I could see how much time was spent on each step, what topics were considered interesting, and how far the average person would get in the lab. I was able to adjust the lab instructions based on real feedback as displayed in Moodle as it tracked people's progress. I found that when students ran into problems, they were more likely to push through when they were close to scoring enough points for a digital badge. Student feedback from the new approach to delivering the lab was very positive. Figure 11-10 is a screenshot from my lab guide converted into the Moodle platform. LMS didn't replace anything that had to do with the actual lab configuration; LMS replaced the lab guide.

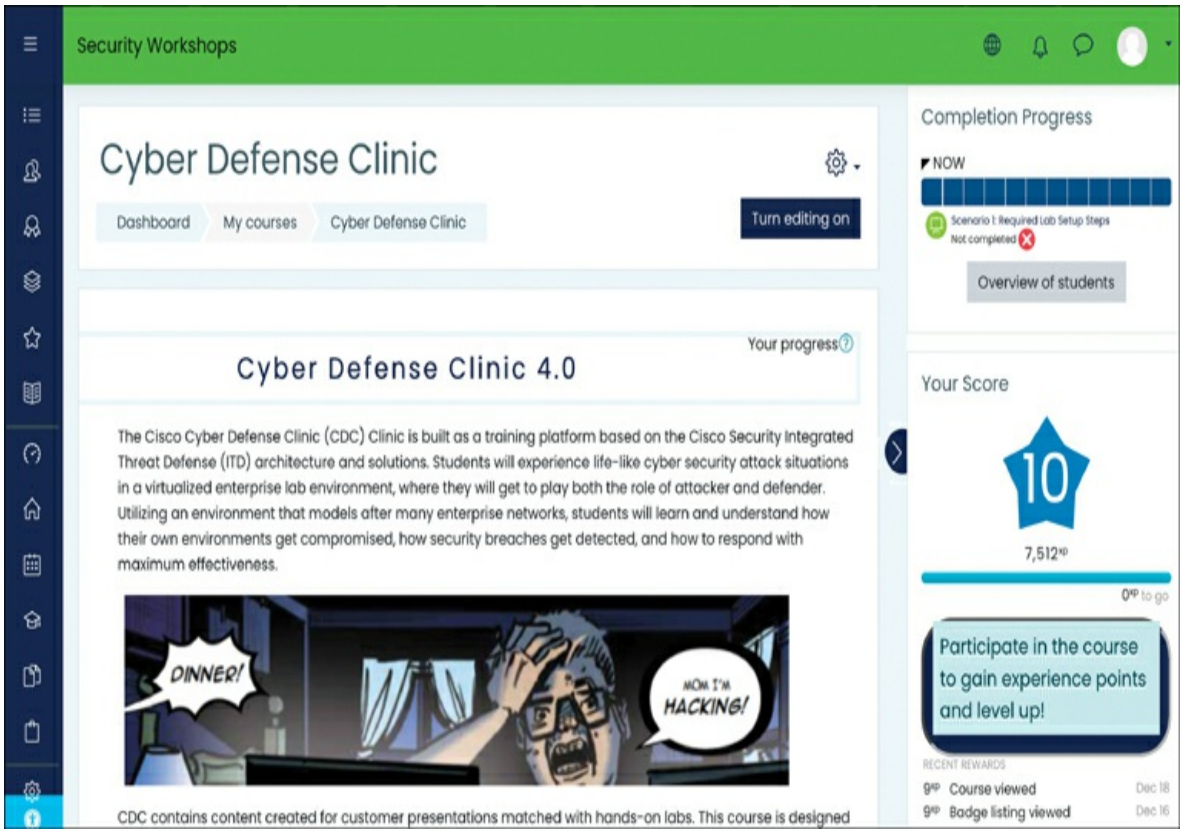

FIGURE 11-10 Lab Guide Converted to Moodle

The first key point to note in Figure 11-10 is the Completion Progress bar at the top right. LMSs include this type of progress indicator to show the student how far they are into the learning objectives as well as offer a click and access approach to picking up where a student left off when they need to come back to a task at a later time by clicking the progress bar step. Notice on the right in Figure 11-10 the Your Score area showing a ranking system, leader board, and other gaming tactics included to encourage a student to continue learning. The lab guides are posted on an Internet-based platform that offers language translation plugins to open up content to more people without the steps needed to translate the lab guide.

I predict LMSs will be the future for learning technical content. The LMS approach connects the developers with the students in a way that allows for more interactive collaboration between what is delivered and how it is perceived by the end user, leading to better results. The LMS approach also drives students to complete content and encourages better performance through various reward options and ranking systems. Don’t be surprised if

you start to find more learning platforms moving away from static PDFs and becoming more interactive.

On-Demand and Personalized Learning

Another trend I predict will become the industry standard is on-demand learning. Just like the “instant expert” concept using YouTube that I previously mentioned, on-demand learning is a way to get quick access to education. YouTube is great for small knowledge chunks; however, there are much more-defined learning curriculums that allow the student to proceed at his or her own pace. I used one version of this learning style to obtain my master’s degree in cybersecurity and information assurance while being employed and having a family. The on-demand approach contradicts the traditional education model, which is based on a calendar of specific times for classes and testing. The challenges to the traditional approach include having the time to be available when teaching occurs, the costs associated with providing live teaching, and limited capabilities to adjust to an individual student’s personal learning style. The traditional approach to learning is based on an adapt or die delivery.

According to https://news.fit.edu/archive/the-7-benefits-of-delivering-on-demand-training/, the following are key benefits from switching to an on-demand teaching approach:

- Convenience leads to compliance
- Consistency of quality
- Easy updating
- Lower costs
- Fosters independence
- Allows partnerships with experts
- Improved data collection

Looking closer at this list, the concept of convenience is critical for many people who have other obligations. Also, the flexibility enables individual students to take class at the time of day that is their peak learning point

mentally. Consistency of quality is another interesting point. The quality of a teacher will make or break a class. My 11-year-old daughter recently told me her favorite class is chemistry. When I questioned her further about why, her answer was that she enjoys the teacher, which has led to her performing very well in that class. Easy updating is yet another huge benefit from on-demand learning. Looking back at the last section's example of the lab I deliver using Moodle, I found that every time I released new content, I would have several bugs in the instructions that students would point out. On-demand systems allow the teacher to quickly make adjustments to the content stored in the cloud so that students not only benefit from the changes but also see their feedback quickly converted into actions. I also pointed out in the Moodle example the value of collecting statistics regarding how students use the learning platform.

The final point from the news.fit.edu article I want to focus on is cost savings from on-demand learning. On demand means top talented teachers can record their content and share it with hundreds of thousands of students. The teacher can make much better profits by expanding to a much larger audience, and students can learn from the highest rated teachers rather than being limited to whoever is available to teach locally or during a specific time period. Questions and labs can also be injected at any point, allowing for a change of teaching style to keep content interesting. Education that would cost thousands of dollars can now be obtained from anywhere at a far lower cost using this on-demand approach.

On-Demand Learning Example: Khan Academy

One fantastic example of on-demand learning is Khan Academy, which is a free version of Kindergarten-College education. One specific feature of on-demand learning that Khan Academy uses is personalized learning. Personalized learning means the content adjusts to the student; as the student masters topics, the content will adjust what is covered next based on what concepts the student needs to learn. Also, if the evaluation of a subject shows that more improvement is needed, the content will adjust and keep focusing on a topic until the student shows they have a firm understanding of the content. Personalized learning mixed with 24/7 access means students can learn when they want to learn and maximize their time, making this approach

extremely effective. Figure 11-11 shows an example of a Khan Academy dashboard offering computer programming, early math, and grammar classes.

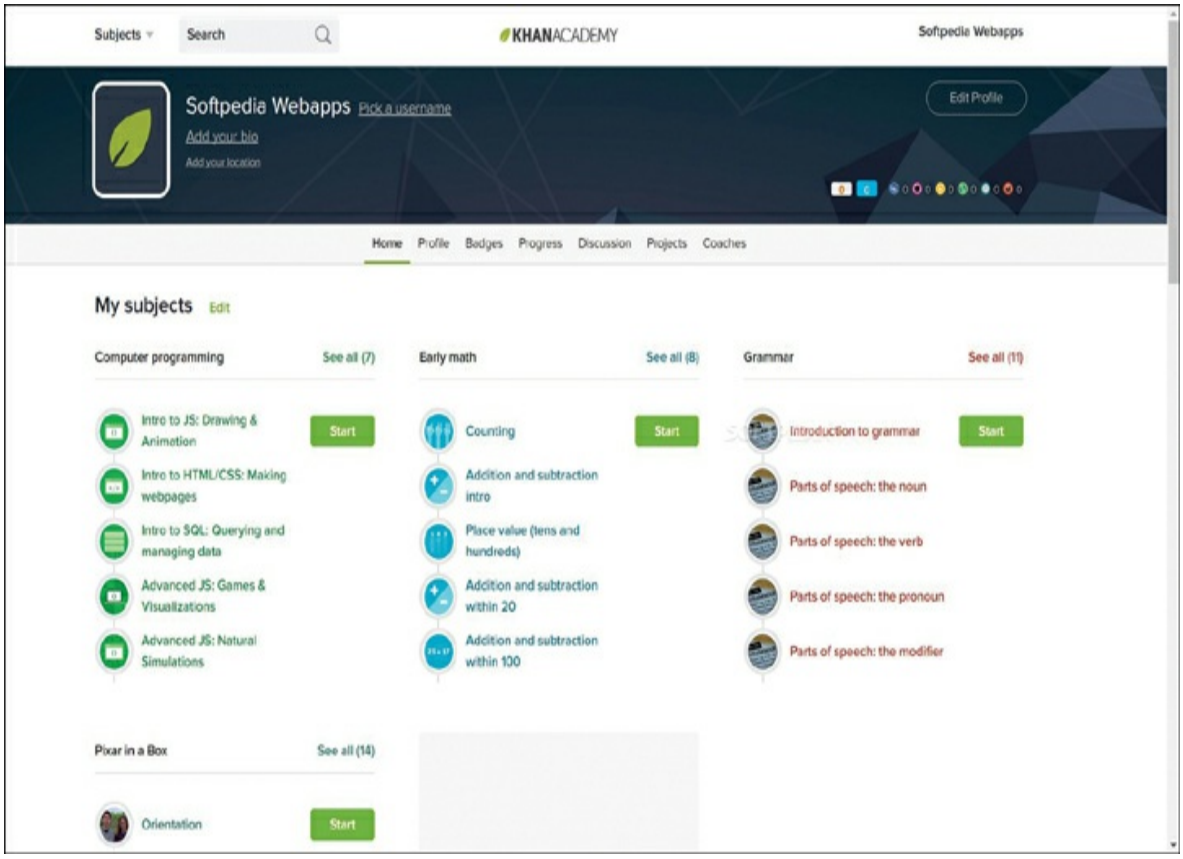

FIGURE 11-11 Khan Academy Dashboard

Future of Training

I predict the future of training will move to an on-demand platform. When a SOC analyst needs to master a skill such as programming, they will simply register for a course and complete it when time permits. Regarding other predictions for the future of training, this section represents a few forward-looking views of training. These predictions are based on technology trends and expected future training requirements. I encourage any forward-thinking SOC to change its training strategy from classic in-person training to on-demand options that allow for a "learn from anywhere" delivery, incorporation of external industry experts from anywhere in the world, and more interaction with LMS-type platforms. The sections that follow offer some predictions that further drive my case for this change.

Future of Training: On-Demand Experts

I am a landlord, but I have never gone through formal training to learn how to repair things in my rental units. Instead, YouTube provides me access to very experienced advisors. If I need to install a new sink, I can buy the materials as listed in a YouTube video, follow the video's steps to install it, and

complete the project without any prior experience. I predict on-demand training will continue to grow and become the standard for learning how to complete technical projects just like it has helped me become my own handyman. This trend will lead to engineers not needing training to accomplish certain tasks, such as setting something up or performing a specific task, since they can learn in real time.

The SOC of the future will rely on a blend of wizards and on-demand learning to accomplish tasks, with expectations that the tools will teach the people rather than the people being expected to know how to use the tools. For example, a new analyst will be able to verbally ask a tool how to do something and the tool's wizard will explain and do the work along with the analyst. This will also lead to a different skillset requirement such as DevOps and programming, meaning understanding how to create or build things rather than needing skillsets for managing existing tools, because those tasks can be done by following a wizard.

I experienced a version of a configuration wizard that understands verbal commands and explains complex tasks when setting up an HP Pavilion Windows 10 laptop. Upon startup, a robotic voice asked me questions with a detailed explanation of what was needed to get my laptop ready for use. I did not need a desktop support expert to help me with the setup because the on-demand virtual wizard clearly explained what I needed. I also saw a link to open a chat with a live on-demand expert if the wizard did not clearly explain what I needed to accomplish. The wizard not only explained features such as Windows multifactor authentication, but also explained why I would want to use them and how to set them up. I predict on-demand support will eventually only offer artificial support, similar to wikis, that learns and adapts based on its usage.

Future of Training: Universal Language

I predict that future technology will remove the need for people to speak the same language to communicate, meaning the future SOC will allow all employees to speak their native language and language translation will just occur in real time. In January of 2020, Apple added language translation to its software, allowing for quick translation between languages. Movies depicting the future show much more advanced versions of this technology where a

simple earpiece can understand and translate any language on demand, which isn't too far from becoming a reality. The future SOC applications will offer displays to support any language and voice will be translated in real time as conversations occur through Voice over IP technology. The breakdown of language barriers will allow the SOC to hire from all parts of the world and focus on what somebody knows rather than if they can communicate effectively with other team members. I predict that language training will be replaced with culture training, helping people understand other cultures, because that is a much harder concept to address with technology.

Future of Training: DevOps

I introduced the concept of DevOps in Chapter 10, which I pointed out is the future of the IT professional in the SOC. I explained how today organizations are throwing tons of money at technology that can produce outcomes, including SOAR and XDR. Outcome-oriented programming such as automation and orchestration continues to grow in popularity and will become a required skillset of the future SOC analyst. The future Internet will offer hundreds of services that can be leveraged using DevOps concepts replacing the need to start any project from scratch. Rather than buying a tool or building something new, the future SOC will simply obtain a generic template and choose from available configurations what they need and adjust the configuration as they see fit. If services or data are needed, such as threat data, that data will be pulled from other systems using APIs. I predict the future SIEM, SOAR, and XDR will simply be a platform that pulls in features and data from various sources, representing a true vendor-agnostic central point for all SOC data.

Cue Robot by Wonder Workshop is an example of a toy that incorporates the concept of using prebuilt programmable templates. My daughter has a Wonder Workshop robot for which she can download prebuilt templates that are customizable to make the robot do different things. My daughter's role in the robot's operation is to choose from the different available applications and make modifications until she is happy with how the robot responds. She doesn't have to build the robot or configure the software for the robot. Instead, she just chooses existing software templates such as "move forward" or "turn the headlamp red" and modifies how that software works, saving

time and the need to re-create what others in the Wonder Workshop community have posted online for her to leverage. For example, she can download a program named “Joey’s cool dance” and the robot will start moving around and making sounds based on how Joey configured and saved the program.

Many video games such as Minecraft are designed with the focus of modifying the code the game functions on. Players can copy the code of an environment they like and modify it so others can enjoy a new world built by a fellow player. Object-oriented programming (OOP) employs smaller programs, or “objects,” that contain code or data that can be called upon, and OOP is popular for teaching children programming. Both Minecraft and the Wonder Workshop community use the “move” command to move something, which represents a set of commands hidden within the move object. I predict the typical SOC analyst of the future will rely heavily on OOP rather than having to understand complex code. DevOps uses lots of OOP by calling APIs to enable prebuilt tasks to function. DevOps engineers rarely build things from scratch; rather, they learn how to create code that can make things work together or update existing code.

IT Training Predictions Summarized

People are a SOC’s most critical asset, and employee retention will continue to be challenging as SOCs around the world fight for the best people. Good training programs can not only help retain your best SOC employees but also help other employees become your best people. Leadership within organizations realize this and will continue to invest heavily in training options that are easy to obtain and deliver as well as have the best impact. This will fuel many of the predictions I have made about IT training, including training becoming a service that anybody can access from any part of the world. I highly recommend ensuring that your SOC keeps up with the latest training trends and considers migrating from classic in-person training to on-demand options that include the latest training technology.

My next focus for future SOC technology is automation. Chapter 10 hit this topic pretty hard; however, I believe most SOCs have barely scratched the surface regarding the potential of automation and orchestration concepts, making this topic a futuristic concept. Any future-looking SOC will be

adopting automation that includes machine learning concepts.

Full Automation with Machine Learning

One vision that has been popular over recent years is automating mundane and repetitive tasks with technology like SOAR. There are technology limitations that don't allow for full automation; however, technology is only improving in the area of understanding human behavior. Recent challenges in automation are related to very complex tasks that only humans can perform. For example, identifying and preventing phishing attacks is extremely complex and something all security vendors are warning is not a valid use case for automation today. I believe the future holds a fix for this challenge and I predict full automation of even the most complex tasks in the future SOC. Who would have thought cars would be able to drive themselves? Elon Musk is one person that believed this would happen, and today, I can use my smartphone to summon my Tesla car to drive for me and park itself. Musk predicts his cars will become a driverless taxi in the near future. If cars can drive themselves, why not the SOC?

Before I dive deeper into what the future holds for automation, I must stress that automation will never fully replace the need for people in the SOC. Any industry security professional and certification program will hammer home that the fundamental elements for security are people, process, and technology. Automation has and will continue to improve how these three elements function. People will handle more complex tasks while repeatable processes will be automatically executed. Automation also means more efficient use of technology such as automating responses and sharing data. Chapter 10 covered how automation works today.

Where does the future of automation take the SOC? Let's first look at one foundational concept that is changing how automation works, which is machine learning.

Machine Learning

Machine learning is based on computers learning patterns from a large amount of data. Consuming and analyzing a large amount of data requires constantly tracking and correlating millions of external and internal data points across a number of endpoints, which is not feasible to do manually and therefore requires automation.

Note

When a vendor claims its research team reviews thousands or millions of events, it is referring to the use of machine learning and automation dealing with big data. Analysts are seeing only the output of consolidated results and predictions based on that data.

Machine learning works in the context of security by recognizing patterns and predicting threat behavior in a massive data set at machine speed. Machine learning is responsible for viewing the raw data analysis of extremely large datasets and simplifying that data into a smaller, digestible dataset summarizing what is important for a human analyst to view. This enables the analyst to make a better decision based on what was deemed important by the machine learning algorithm.

According to TK Keanini, a distinguished engineer at Cisco, machine learning can use the following learning techniques to provide value to the SOC:

- **Detecting surreptitious attackers on networks:** Machine learning can detect behavioral anomalies to find attackers on the inside or logged in with stolen credentials.
- **Predicting "bad neighborhoods" online:** By learning from Internet activity patterns, machine learning can automatically identify attacker infrastructure being staged to launch the next threat.
- **Detecting attacks through novelty and outliers:** Machine learning finds attack patterns humans cannot readily detect, like a new peer relationship on the network with hosts communicating that can't or shouldn't be.
- **Finding suspicious cloud user behavior:** Analytical techniques uncover

suspicious user behavior indicative of cloud account compromise to extract data or perform malicious operations.

- **Leveraging modern malware detection:** Machine learning is valuable in detecting polymorphic malware, breaking down threat attributes to better stop new and reengineered polymorphic threats.

Machine Learning Hurdles

When machine learning works well, it provides a lot of value. There are, however, a few major hurdles that must be addressed to obtain value from machine learning. First, it must provide quality results, or you will be responding to bad recommendations. This can be a huge problem if automation is tied to the results from poor feedback provided by machine learning, leading to a complete breakdown of the SOC's processes. From a high-level viewpoint, bad recommendations mean creating false positives or false negatives causing the SOC to either miss critical events or add work for the analyst that doesn't lead to any real value.

Another major hurdle for machine learning is accounting for change within technology and the SOC. I've pointed out repeatedly in this book that the threat landscape is a continuously changing battleground. Part of that change is attackers adjusting their tactics to circumvent whatever has been preventing them from accomplishing their goals. If a machine learning tactic is found to be extremely successful, malicious parties will attempt to reverse engineer how the successful defense works and adjust their strategy until they go undetected. The best way to overcome this hurdle while using machine learning is to continuously retrain what was learned during the machine learning training cycle to ensure new data elements are accounted for.

One final hurdle of machine learning is explaining how an outcome was found to the average person. Answering HOW is critical to support any recommendations made by the machine learning algorithm, and it is critical to know how an outcome was found when automation is applied to the response. Simply saying "true" or "pattern matched" will not provide enough evidence for why a conclusion was made. Overcoming this hurdle is much harder than it might seem because the reason a decision was made by a

machine learning algorithm can be based on millions of smaller events. This could mean that a proper explanation of how the algorithm came up with its decision requires pages of data based on events that are constantly changing. By the time an analyst reviews why a decision was made, the associated factors could have changed!

Machine Learning Applied

When should a SOC use machine learning today? If you find that a domain is static, well understood, and has limited variability, machine learning would not be a good fit; in fact, it would be overkill. This same concept applies to tactical threat intelligence feeds that provide domains that should be blocked. You don't need complicated decision making when you know that something is bad and should be blocked. Simply use pattern matching techniques and block when a match occurs.

Machine learning is more appropriate for domains that are evolving with large variability and are not well understood. The goal of using machine learning is to better understand the domain of interest so a decision can be made about its risk just like we can do with the static domains, but machine learning fills in the missing data needed to make the best decision. This makes machine learning a complementary technology rather than a specific tool used for security. In short, machine learning is ideal for the following conditions:

- Advanced threat is evolving rather than static
- Associated data sets are very large, including the 1% that tends to be overlooked
- Many opinions need to be consolidated into a more understood view
- A need to support other security tools so they can make better decisions

The best way to develop a plan to apply machine learning is to work through the following three steps:

Step 1. Determine what outcomes you wish to see. The outcomes should be adjusted for the intended party; for example, is it an incident responder, a threat hunter, a forensic specialist, or a general network

engineer that will be using the results?

Step 2. Determine what telemetry is used to obtain the desired outcomes. Machine learning should be just one of the multiple analytical approaches that is used, since machine learning is a complementary technology. Think of machine learning as *N*+1, where *N* represents all the possible ways to analyze data and 1 represents machine learning. The following are potential datasets that could be used for common security-related outcomes:

- Entity modeling
- Statistical analysis
- Predictive analysis
- Machine learning
- Unsupervised learning
- Supervised learning
- Reinforced learning
- Artificial intelligence

Step 3. Determine what data will be analyzed, also known as *telemetry*. Chapter 5, "Centralizing Data," covered the various types of possible datasets you can leverage to generate a desired outcome.

Figure 11-12 shows a simple diagram explaining the three key concepts to leveraging machine learning. The best way to address these three key points is to work left to right.

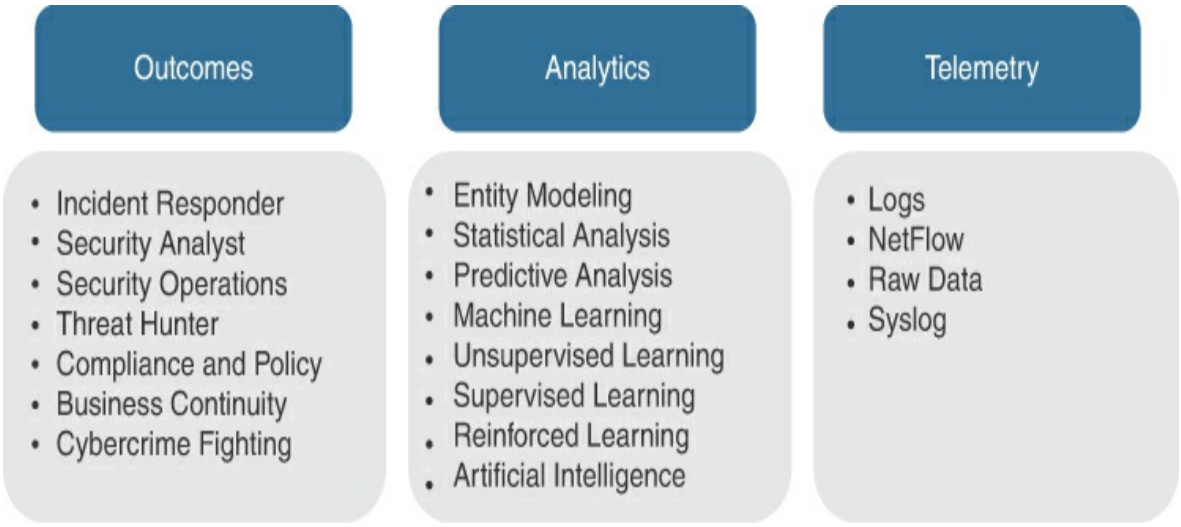

FIGURE 11-12 Key Elements to Using Machine Learning

Training Machine Learning

I've spoken about teaching and retraining machine learning; however, what does that really mean and how is it applied? There are three general approaches to training a machine in regard to security:

- **Supervised learning:** This approach is used to teach a machine the question you want answered when you know the question. To use supervised learning, you must be able to say "We know this concept is right but want machine learning to answer more questions about what we know is right."
- **Unsupervised learning:** This approach uses less structure based on not knowing how the structure should look to find what you are looking for. Simply put, you don't know the question to ask but know you have a problem to solve. Unsupervised learning is ideal when you don't know the answer or fully understand the questions you should be asking. Essentially, unsupervised learning is about understanding how elements of a state relate to each other rather than answering a specific question.
- **Reinforced learning:** The third approach is using a trial and error tactic to develop desired outcomes. Reinforced learning allows for the machine to help formalize the problem being addressed so conclusions can be made regarding how to solve what you don't know needs to be solved. In

short, this approach addresses when you know what the question is and attempt to answer it using different approaches, hence trial and error.

Training a machine can be broken down into a few steps. First, you have training data, which you apply against the machine learning algorithm. As the machine learns how to respond, you move from the learning phase to the live phase where real data is applied. The machine learning engine will classify things of interest and develop a prediction. Training data is continuously applied, allowing for adjustments to change. Figure 11-13 shows a simple diagram representing this concept.

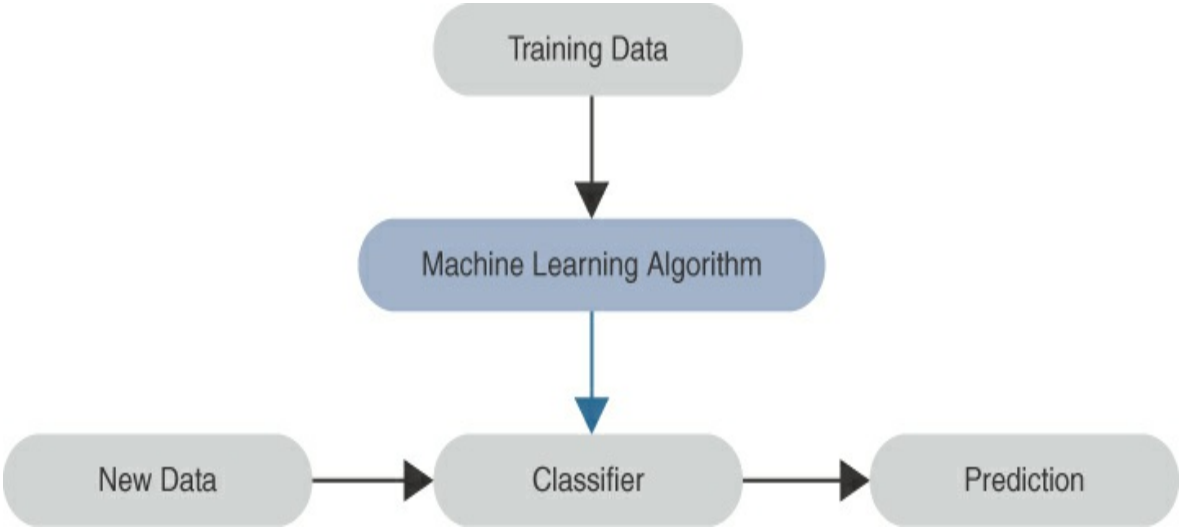

FIGURE 11-13 How to Train Machine Learning Classification Tools

Future of Machine Learning

Machine learning and automation offer a ton of potential benefits as these areas of technology improve. There are a number of trends and indicators explaining how the future of machine learning will impact the SOC. The sections that follow outline some of my predictions based on such trends. Every security technology vendor I speak with points out these concepts as the future of IT security.

Future of Machine Learning: Auto Adapting to Threats

Human behavior is complex, but with enough data, it can be predictable. Daniel Negreanu is a famous world champion poker player who has been ascribed the ability to predict what cards other players are holding. He has been filmed yelling out “You have king queen not suited!” and everybody is shocked to discover he is correct (sometimes). Mr. Negreanu has explained that what he does isn’t magic but based on various datapoints leading to a conclusion. He also explains that the slightest human gestures can determine how somebody will respond to his moves, allowing him to control the other

players based on what he says and how he bets.

Strategies from poker champions such as Daniel Negreanu are being collected and automated with poker software. Such software uses a combination of behavior learning and math to determine what cards should or should not be played as well as the possibility of a certain outcome when cards are played. I believe the same concepts apply to security, which is why security concepts such as artificial intelligence are becoming more relevant in how a tool predicts attack behavior. I predict the security tools of the future will rely heavily on various big data learning mechanisms that take into consideration not only how threats function but also general human decision patterns, enabling the tools to both predict an attack and identify the attacker and their intent. Today, many security tools can identify an attack, but few can determine who is perpetrating the attack or why. By knowing those additional details, more holistic responses can be performed, including preventing future attacks by addressing the threat directly.

One amusing example of this concept is a web application firewall (WAF) I saw at a security conference a few years back. One feature being demonstrated was that when the WAF identified a certain type of web application attack based on matching attack behavior, the WAF would display to the attacker the classic Clippy paperclip stating “It looks like you are trying to deliver a certain type of attack. You can learn more about what you are doing here.” I predict the future of machine learning will include this level of response, such as “Hi Joey Muniz located at (*home address*), it looks like you are attempting a cross-site scripting attack against my client. This is illegal and I have recorded what you are doing. This is your last warning before I send the authorities to your location.” In order for this type of response to occur, there would need to be a lot of evidence with clear accuracy; however, I believe we are not too far from the technology needed for this level of detail to be collected.

Future of Machine Learning: Chatbots Take Over

A chatbot is an automated response to a conversation with a human. You likely have seen chatbots on websites asking if you have questions. When you ask a question, the chatbot will attempt to control the conversation by limiting possible responses and asking you to work within its response

catalog. If the chatbot can't answer your question, you are offered alterative options including contacting a real support representative.

I predict technology of the future will allow for hyper-personalized conversations between humans and machines. I will be able to instruct a chatbot such as Amazon Alexa "I see these details; explain to me what type of attack is occurring" and obtain a very specific answer identifying what the attack is and how it is impacting my organization. I use the example of Amazon Alexa because I also forecast that the type of security data used exclusively in SOCs today will be readily available for anybody in the future.

The best example I've seen of this type of conversation is in the *Iron Man* movies, in which Iron Man (Robert Downey, Jr.) speaks with Just A Rather Very Intelligent System, shortened to J.A.R.V.I.S, regarding real-time events. Iron Man asks J.A.R.V.I.S to predict what is the best action to take based on hundreds of factors. The idea and technology needed to allow J.A.R.V.I.S to understand what a threat is and cause a change in J.A.R.V.I.S's outcome doesn't exist today, but it is coming, as computer AI is designed to mimic human behavior. I predict the future SOC will have a J.A.R.V.I.S-like system that everybody will speak with as if it is part of the SOC's staff. This AI to employee collaboration is also very common in movies featuring a spaceship crew interacting with the spaceship computer system.

Future of Machine Learning: Self-Configuration

Automation in security means machines automatically making changes a human would make. Many organizations are extremely cautious about deploying automated responses that have the potential to lead to drastic negative results. When the technology proves there is a very low risk of negative impact, I predict automatic security will become common in the SOC, meaning automated response will become part of all SOC procedures. Automated response will include adjusting the configuration of tools, logging all actions taken, and giving a post-incident explanation of what occurred both to the SOC and to whoever was impacted by the incident. I mentioned earlier in this chapter that I predict management tasks will be automated. This will occur due to improvements with machine learning, leading to a trust in the automated response from artificial intelligence.

An interesting depiction of futuristic trust in machines appears in the movie *I, Robot,* in the scene in which Spooner (Will Smith) manually drives the car. In the future world of *I, Robot,* machines drive all the cars. Manually driving is seen as a huge risk and a violation of insurance policies since the world trusts machines more than humans with driving. I predict this mindset will eventually apply to decisions made by AI tools in the SOC. Imagine a scenario in which an analyst makes a prediction but is dismissed by the rest of the SOC because they believe a different prediction that came from an AI tool. In my opinion, this is the future of AI in the SOC!

Self-configuration and response will lead to fewer job positions related to performing those types of tasks. I stated it before and will make the prediction again that jobs involving asset management tasks will be completely automated in the future SOC. The good news is that there will be an increase in demand for engineers that specialize in machine learning. Tasks for these future engineers include developing training data, adjusting classifiers in machine learning algorithms, and tuning how self-configuration tasks are handled.

Future of Machine Learning: Battle of the Bots!

As indicated in my previous predictions, I foresee automated responses becoming a normal activity in a SOC. I predict attackers will also leverage the same technology, leading to an AI versus AI cyberthreat landscape. Some future networks will become so hostile that manual interaction will not be able to keep up with response, leading to only being able to use AI to have tables stake performance, meaning only AI can keep up with AI-based attack. Future environments that become too hostile to monitor will become the top concerns for organizations based on the lack of controls and unknowns that are associated within the automated AI chaos. I can’t predict how big of a problem AI versus AI will become, but my gut tells me it will become a leading problem that won’t be solved with more technology. More technology will just fuel the problem.

Machine Learning Predictions Summarized

There continues to be a lot of hype in the security industry about machine

learning, and for good reason. As the number of events a SOC needs to deal with increases, automation will be needed to keep up with the workload. Machine learning will also become a requirement to match the future speed of change in the threat landscape, which is why leading security vendors are heavily investing in machine learning capabilities. This is why many of my machine learning predictions are focused on how security tools will incorporate machine learning to improve capabilities and response time. Every forward-looking SOC will include automation and orchestration as part of their maturity roadmap. The future of automation and orchestration is dependent on machine learning.

The final topic to address is how to predict how your SOC will function in the future. I will attempt to apply many of the concepts from this book with the hope that the concepts positively influence the outlook for your future SOC.

Future of *Your* SOC: Bringing It All Together

The final future topic to review in this chapter is where you should expect your SOC to be one, three, five or even ten years from now. For new to existing SOCs, it all starts with having an accurate and meaningful mission statement. Chapter 2, "Developing a Security Operations Center," covered what needs to be done to convert a security team into a formal SOC. Every SOC needs a mission statement, which has to be aligned with the organization's goals. If the organization adjusts its focus, the SOC needs to also validate that the current needs match the SOC mission statement and, if not, adjust accordingly.

Every SOC needs executive sponsorship, which can and will change over time but always must be identified to ensure the proper support is provided to the SOC's mission. The SOC's scope will also change as more services are introduced or adjustments are required to meet the organization's needs. My recommendation is to periodically evaluate your SOC's mission statement and scope statement with your executive sponsor. You can also draft new statements and have them reviewed by your organization's leadership as you plan for new services, to ensure you continue to have the proper support and aligned vision with leadership. Anytime there is a change in leadership or

SOC sponsor, you must validate your SOC's mission statement to ensure alignment with the executive sponsor. Do not assume your executive sponsor has bought into the SOC. This can lead to losing critical support, resulting in lack of power and funding for your SOC.

Future SOC Key Point

Review and update your mission statement and scope statement with your SOC's executive sponsor. If that sponsor ever changes, you must perform a review with the new sponsor.

Your Future Facilities and Capabilities

Because your facilities will change, the way you deliver security capabilities will change as well. Chapter 1, "Introducing Security Operations and the SOC," covered different approaches to assessing your security capabilities. For technical capabilities, you learned how to perform a SOC capabilities assessment, which documents the capabilities you have in a vendor-agnostic approach. I also suggested that you leverage industry guidelines to verify what are considered best practices and include those capabilities even if they are missing. I recommended creating a heat map that identifies what you have and what is missing regarding industry best practices. Figure 11-14 is an example of performing this type of assessment against endpoint and network security capabilities.

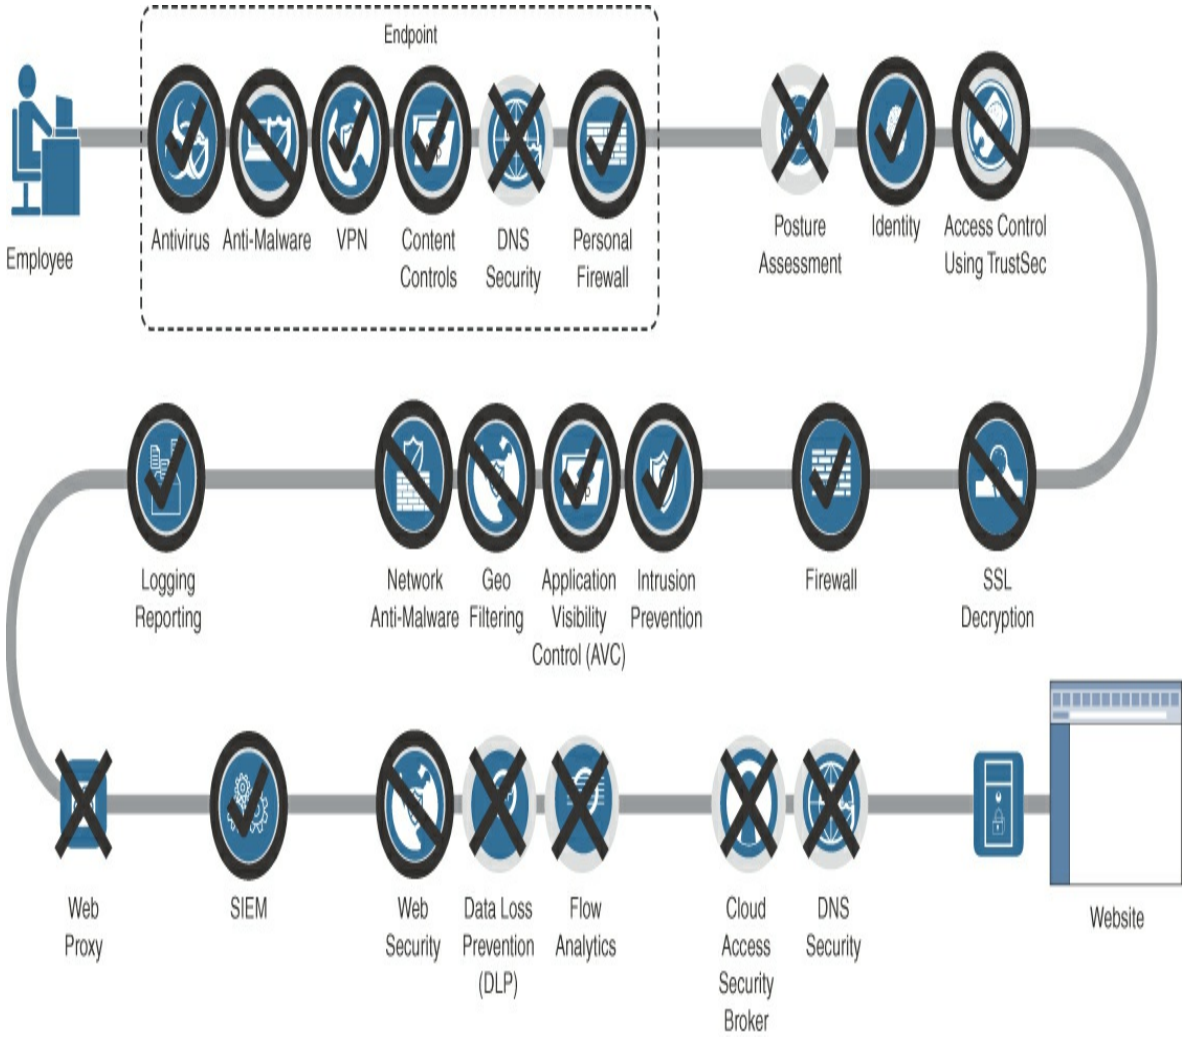

FIGURE 11-14 Capabilities Evaluation Example

Assessing Accessing the Datacenter

As I've covered in this chapter, many facility services will shift to the cloud in the future. You will need to adjust your SOC capabilities assessment to meet those changes. Figure 11-15 is an example of a SOC capabilities assessment that starts with the endpoint and works through security capabilities up to a server in the datacenter. Rather than including the endpoint, this diagram focuses on the security capabilities in the network and the servers in the datacenter.

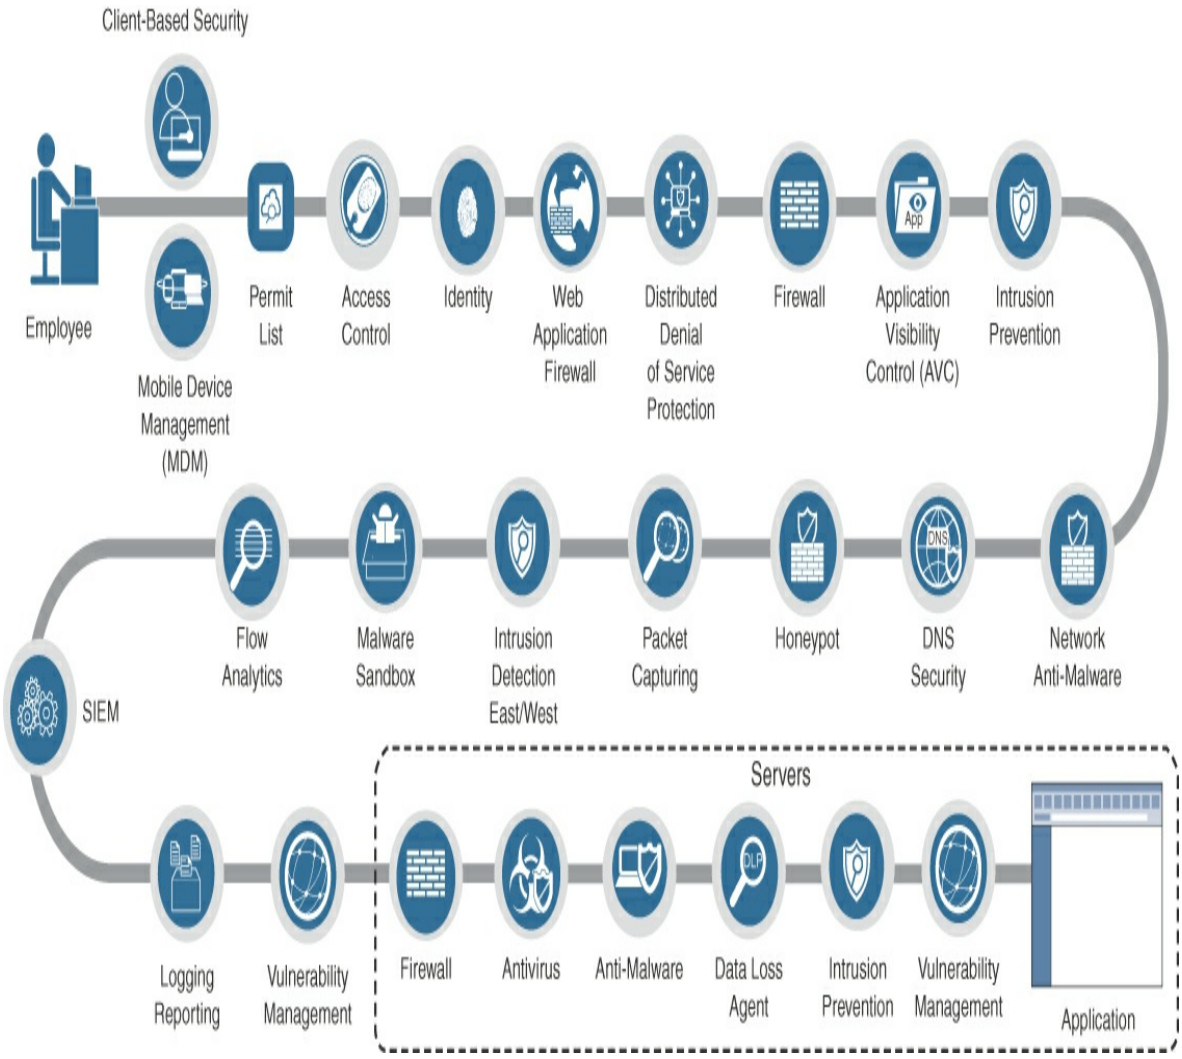

FIGURE 11-15 Server-Focused Capabilities Evaluation Example

Assessing Remote Users

Another scenario that exists today and will exist in the future is protecting remote users. VPN technology is not going away anytime soon since data privacy is a top concern and encryption is the industry standard for protecting data in motion. In the future, I predict remote-access VPN as a Service (VPNaaS) will be more common than traditional VPNs deployed through hardware known as VPN concentrators; however, the concepts are the same regarding security capabilities. Figure 11-16 is an example of creating a SOC capabilities assessment focused on VPN. This example uses the traditional VPN concentrator to provide VPN services, but how the VPN service is

provisioned is not a security capability, which is why the VPN service provider is shown as an unlabeled icon because it doesn't matter.

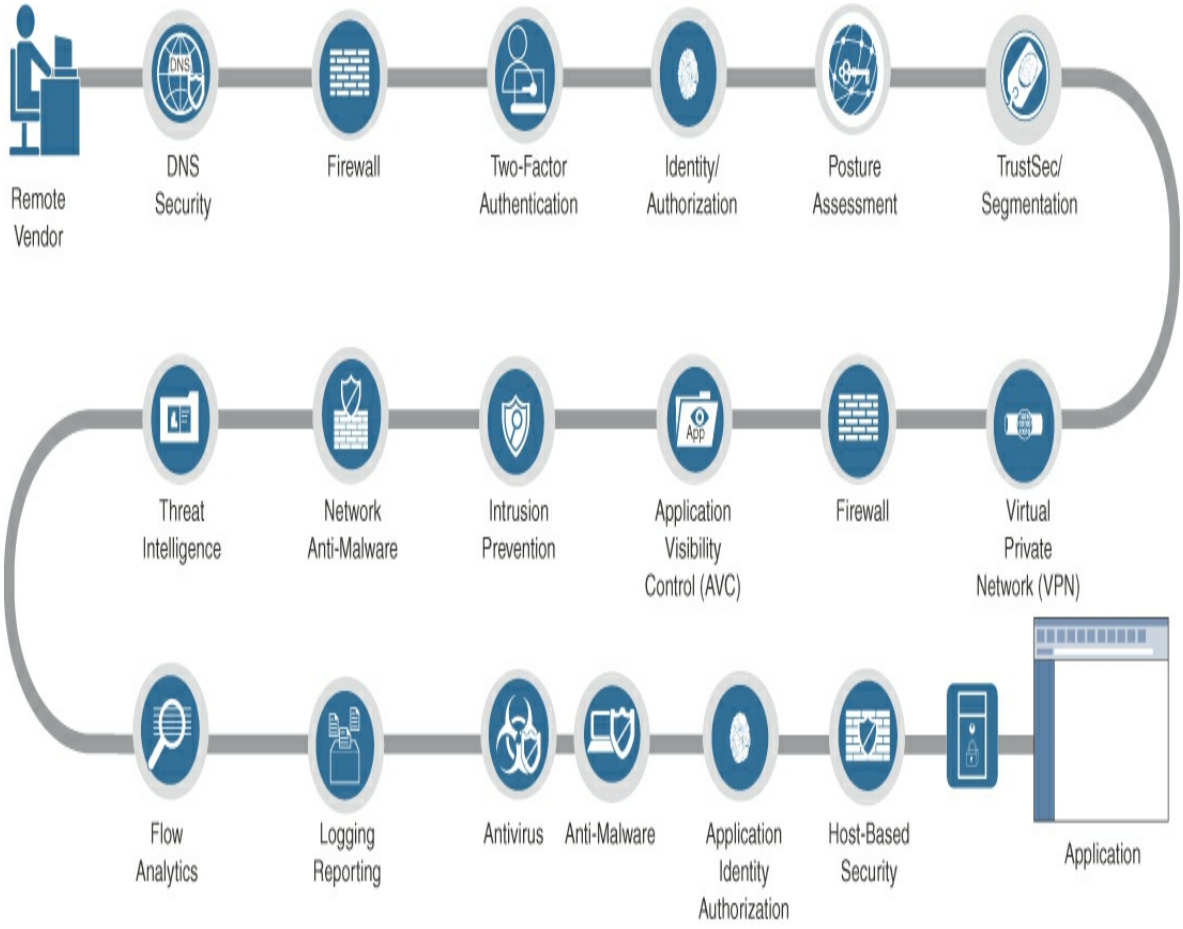

FIGURE 11-16 VPN-Focused Capabilities Evaluation Example

Future SOC Key Point

Review and update your security capabilities for all parts of the organization your SOC is responsible to protect. This will include physical, virtual, and cloud technologies and be made up of multiple SOC security capability maps.

Based on changes in technology trends, your future SOC's facility needs will be different than they are today. Earlier in this chapter, I spoke about how the

concept of the office is shifting away from a physical location to one that allows for SOC members to work from anywhere. Many of the facility planning concepts covered in Chapter 2 will change to services and technology that are delivered to wherever the SOC team member is located. Collaboration technology will be critical to keep a physically separated SOC in synch. The savings from reduced facility expenses can be used to fuel a more advanced collaboration technology as your SOC makes the transition away from a dedicated facility.

Figure 11-16 has one specific change regarding how network access control is performed compared to the server-focused capabilities assessment shown in Figure 11-15. I replaced NAC with TrustSec/segmentation for provisioning access to the network. TrustSec classifies traffic based on contextual identity versus using an IP address. TrustSec represents the concept of group tags, which leads us to the next topic regarding the future of your SOC. That topic is the focus of using segmentation concepts that will replace traditional VLANs and ACLs.

Group Tags

The way your SOC applies segmentation today will change as trends show a shift from a network-focused segmentation strategy to an application-focused segmentation strategy. Figure 11-17 shows a general look at the different types of segmentation that can be used (as covered in Chapter 2). Application-focused segmentation functions at the packet level, meaning packets are either allowed or denied per application rather than controlling where they can go on the network, which is shown at the bottom of Figure 11-17.

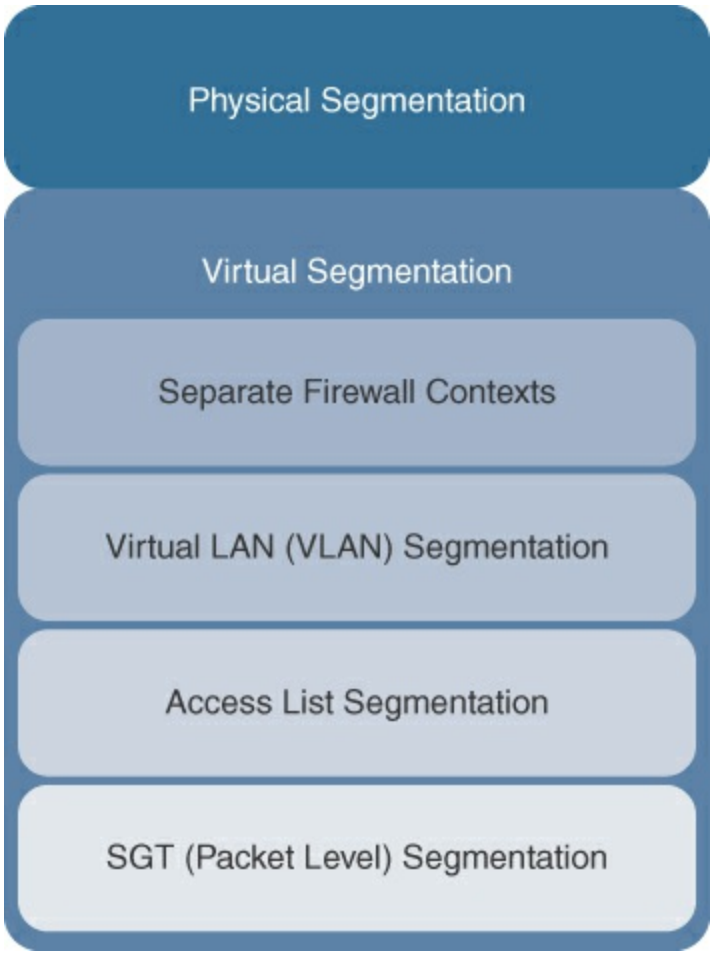

FIGURE 11-17 Different Segmentation Options

I did not cover the concept of group tagging earlier in this book, as it is not as commonly deployed as VLANs and ACLs for segmentation; however, I predict concepts such as group tagging will become the industry standard in the near future. The main reason for this prediction is that packet-level segmentation allows the removal of different responsibilities when provisioning access to a resource. For example, if a new server needs to be stood up in an average datacenter today, a datacenter specialist provisions the server, a network specialist assigns the network services, such as which IP address, VLAN, or ACL should be used, and a security specialist opens only the ports in the firewall that the server requires to function properly. Group tags are understood by network and security tools so all devices know when a person has rights to access a server; hence, no changes to the network are needed. Only a group tag is required to inform the entire network of who or what can access what. The future SOC can place a server on the network and

assign it an access policy, which the rest of the network and security tools would understand.

Group Tags Explained

A group tag works by being assigned when a user authenticates to the network. That user is assigned a group such as the PCI group, which in this example is configured to allow access to the PCI resources. No access is granted to any packets if they do not have the PCI group tag assigned. This Layer 2 tag is held as the user accesses resources on the network, which means the entire network must support group tagging for this model to work.

Today, most modern network and security tools support group tagging, which has been one of the major holdbacks for organizations supporting older equipment and unable to invest in a network-wide upgrade. As organizations update their old network equipment, they will eventually be able to support group tag type technologies, making the case for switching from VLANs and ACLs to tagging a much easier conversation.

Figure 11-18 shows an example of a user connecting to the network and Cisco's NAC solution, called Identity Services Engine (ISE), providing the secure group tag (SGT) classification of 5. When the user attempts to access the HR servers, she is denied based on not having the right access (SGT level 10). The firewall and switch both deny this access. This diagram also shows a financial server with an SGT of 4 attempting to access the HR servers and, once again, being denied by the switch due to not having the proper authorization. The key takeaway is that the entire network becomes the enforcement point of SGTs, allowing a much simpler security policy to be enforced network wide. I predict the current model of having thousands of ACLs and VLANs will be replaced by a centralized group tag policy matrix that is understood network wide.

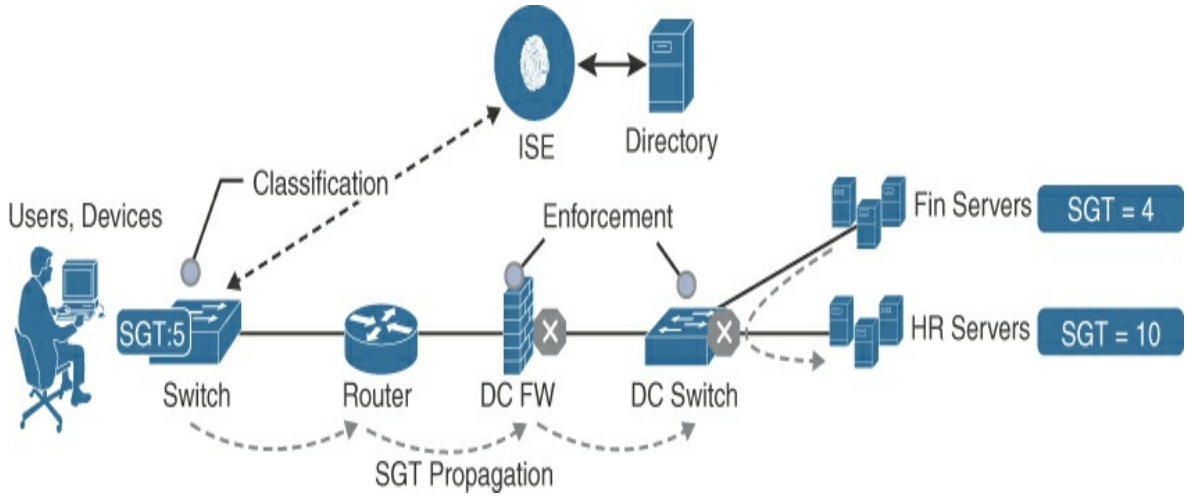

FIGURE 11-18 Different Segmentation Options

Future SOC Key Point

Review how your organization is implementing segmentation today and consider moving to an application-focused Layer 2 segmentation style if (or when) your technology supports it.

Your Future SOC Staff

Earlier in this chapter, I offered my predictions for how technology will change the job roles in your future SOC. I surmised product configuration and technology management roles will be automated and replaced by DevOps and data modeling roles. Tactical tasks will be handled using wizards or outright automated. Regardless of what changes come about, the end results of your SOC services must continue to provide value and improve overtime.

Chapter 4, "People and Process," covered different job roles found in SOCs today as well as how to recruit and maintain talent. As the job roles and processes change in your SOC, you will need to continuously evaluate how effective the people and processes are at accomplishing the SOC goals. The best approach to this is evaluating your SOC people and process using a

tabletop exercise followed by a penetration test. NIST SP 800-84 is one of the many industry guidelines that covers steps to deliver a tabletop exercise, which was also a topic covered in Chapter 6.

NIST Tabletop Exercises

NIST SP 800-84, *Guide to Test, Training, and Exercise (TT&E) Programs for IT Plans and Capabilities*, points out that before performing a tabletop exercise, your organization needs to consider your organization's overall objectives. The guideline also assumes your organization has a TT&E program. Those objectives can be evaluated using the following questions:

- Have the personnel who would participate in the tabletop exercise been trained on their roles and responsibilities within the plan? If the personnel have not yet been trained, the TT&E (test, training, and exercise) program coordinator should consider conducting a training event before the tabletop exercise so that the personnel can participate more effectively in the tabletop exercise, increasing its benefits.
- When was the last time the organization conducted a tabletop exercise for the plan?
- Have recent organizational changes been made that could impact the content of the plan?
- Has new TT&E guidance been issued that could impact the content of the plan?

Your SOC should periodically perform tabletop exercises to accommodate change and to evaluate your SOC's current maturity level. Chapter 1 covered different approaches to evaluating each SOC service's maturity as well as the overall SOC's maturity. Figure 11-19 is an example of ISACA's approach to evaluating the maturity level of the people, process, and technology in a SOC. It is critical to establish an understanding of the maturity of your SOC and associated services in order to develop a roadmap for what is needed to improve your capabilities.

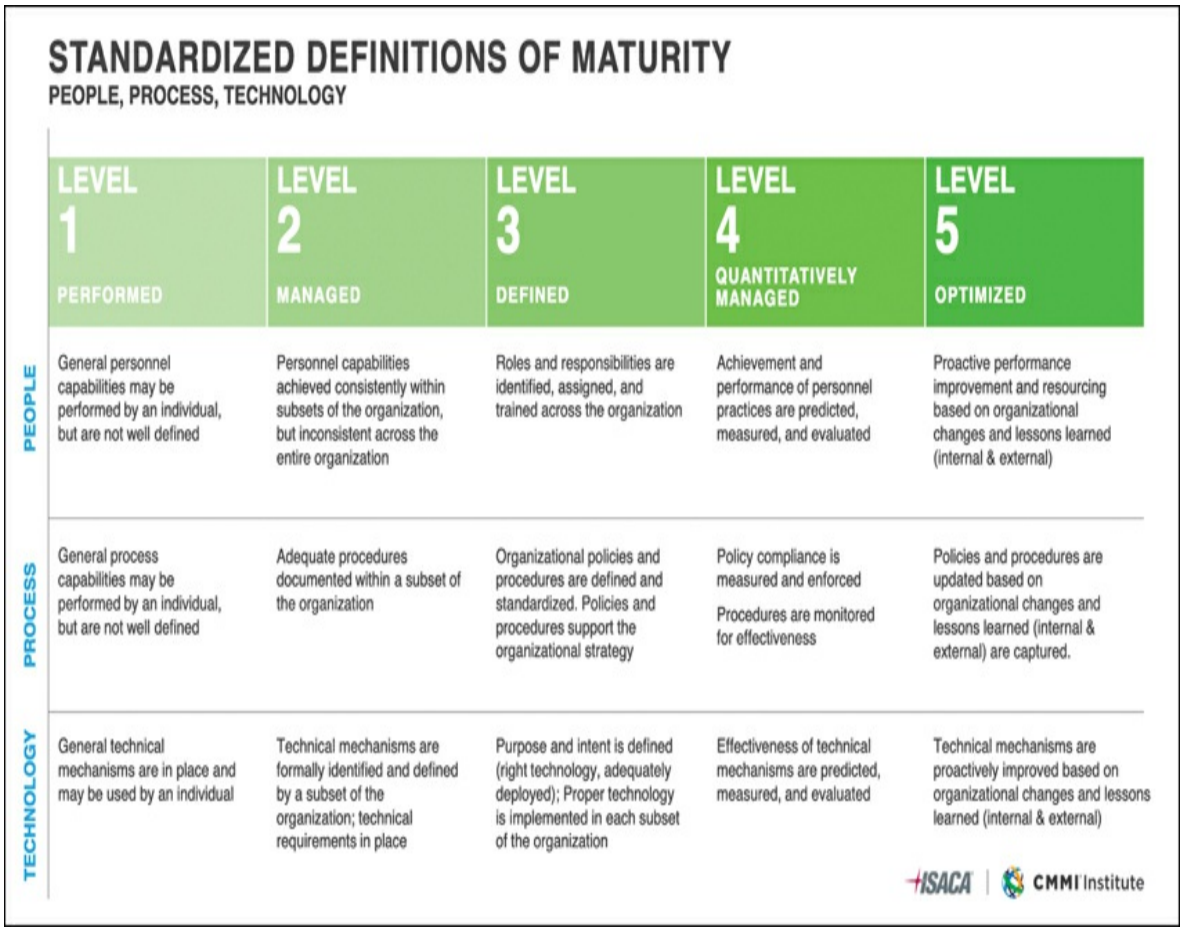

FIGURE 11-19 Standardized Definitions of Maturity from ISACA

Future SOC Key Point

Periodically perform tabletop exercises to evaluate your existing SOC services and maturity level.

Audits, Assessments, and Penetration Testing

Common services that are used to evaluate the risk associated with an organization's network and endpoints are audits, assessments, and penetration testing. Chapter 3, "SOC Services," explained the fundamental differences between each of these services and how they relate to what

services a mature SOC provides. A key point from Chapter 3 is that all of these services are a subset of risk management, which takes into consideration all risk associated with the organization. Audits, assessments, and penetration testing are not SOC services but processes that are part of core SOC services.

Chapter 6, "Reducing Risk and Exceeding Compliance," went into more details regarding where audits, assessments, and penetration relate to SOC services. It is important to remember that *audits do not make your organization more secure*. Audits are built around ensuring compliance is met for some specific regulation such as PCI DSS or HIPAA but do not consider risk outside of what is being evaluated. Audits can align with industry guidelines such as those covered in Chapter 6, but they are limited in scope and more focused on compliance than on security.

Which Service to Use?

Assessments are a general evaluation of potential risk. Chapter 9, "Vulnerability Management," focused on identifying and responding to IT vulnerabilities. I covered various approaches to performing assessments of IT systems using network- and client-based techniques that result in potential vulnerabilities. I also pointed out the risk of false positives being found, since assessments do not attempt to exploit the target to validate the risk is real. Chapter 6 explained how penetration testing is a more thorough but also higher risk and heavier service engagement for evaluating people, process, and technology for vulnerabilities. Penetration testing attempts to exploit the target to provide a more realistic assessment of a potential risk, leading to more accurate results. Due to the increase in risk and expected services, I recommended properly planning how you will execute a penetration test. As a refresher from Chapter 6, Figure 11-20 is a high-level diagram of the NIST SP 800-15 four-stage penetration testing methodology.

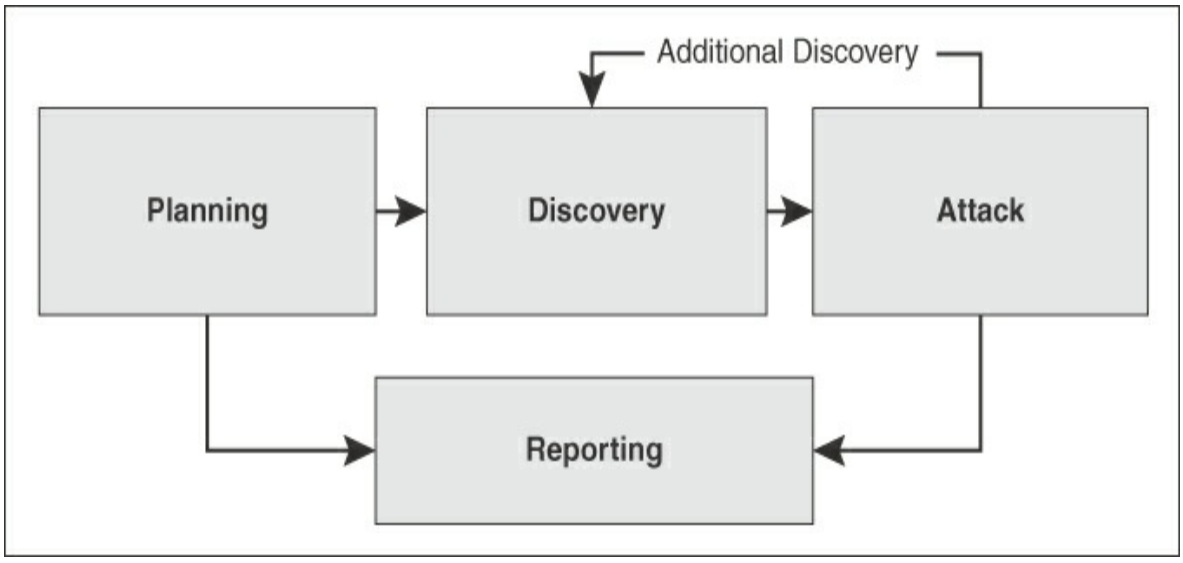

FIGURE 11-20 NIST SP 800-115 Four-Stage Penetration Testing Methodology

Your SOC services today and in the future must adopt audit, assessment, and penetration testing services based on your business needs. All three services are fundamental to maintaining a mature risk management service, and all three are used for different use cases. I pointed out earlier in this book that you do not want to perform a penetration test when you know you are vulnerable, as that is a waste of resources to inform you what you already know. You want to first perform an assessment and apply the best controls to reduce risk of exploitation before engaging in a penetration test, meaning the penetration test will evaluate how well you responded to what was found during the assessment. Audits should be assigned to required compliance and separate from how you perform assessment and penetration testing since audits are not designed to make your organization more secure.

Future SOC Key Point

Periodically perform audits, assessments, and penetration testing as part of your risk management program. Know when to use an audit, assessment, and penetration test, as each has different values and associated cost.

Future Impact to Your Services

Some concepts in this book will not exist within your SOC today, but they will be part of your future SOC. One common missing concept in many SOCs is the use of the various types of threat intelligence. Chapter 7, "Threat Intelligence," broke down threat intelligence into four categories:

- **Strategic intelligence:** Nontechnical intelligence that is heavily risk-based and used by high-level decision makers
- **Tactical intelligence:** Provides details of threat actor tactics, techniques, and procedures (TTPs)
- **Operational intelligence:** Reveals actionable information about specific incoming attacks
- **Technical intelligence:** Technical details about threat indicators such as malicious IP addresses and hashes of malicious artifacts

It is common for SOCs to leverage technical threat intelligence today, but leveraging tactical and operational intelligence will become just as common in the future SOC as big data–based technology such as machine learning and artificial intelligence increase in popularity. In Chapter 7, I covered how to utilize different forms of threat intelligence, including scrapping social media websites and converting security news into actionable intelligence. I highly recommend that your SOC start using publicly available resources to your advantage now so that later it will be a smoother transition to leveraging big data tools.

Chapter 10 explained the differences between SIEM, SOAR, and XDR. Many SOCs around the world are using these technologies but few are leveraging their full capabilities for automation and orchestration. Your future SOC will become more automated, requiring SOC analysts to have DevOps skills and rely on big data technology to make decisions and respond with automated blocking or configuration changes. I highly recommend your SOC start adopting the concepts from Chapter 10, as they will become standard practice.

The first step to adopting concepts from Chapter 10 is to develop playbooks. If your SOC has not formally converted processes into playbooks, now is a great time to invest in that effort. Playbooks lead to more mature SOC

services, as they provide repeatable steps that result in an expected outcome. As your SOC formalizes and leverages playbooks, it can automate steps in the playbooks to reduce the number of mundane and repetitive tasks your SOC analysts need to perform. I highly recommend using the industry-standard workflow symbols shown in Figure 11-21 as you convert all of your SOC processes into playbooks.

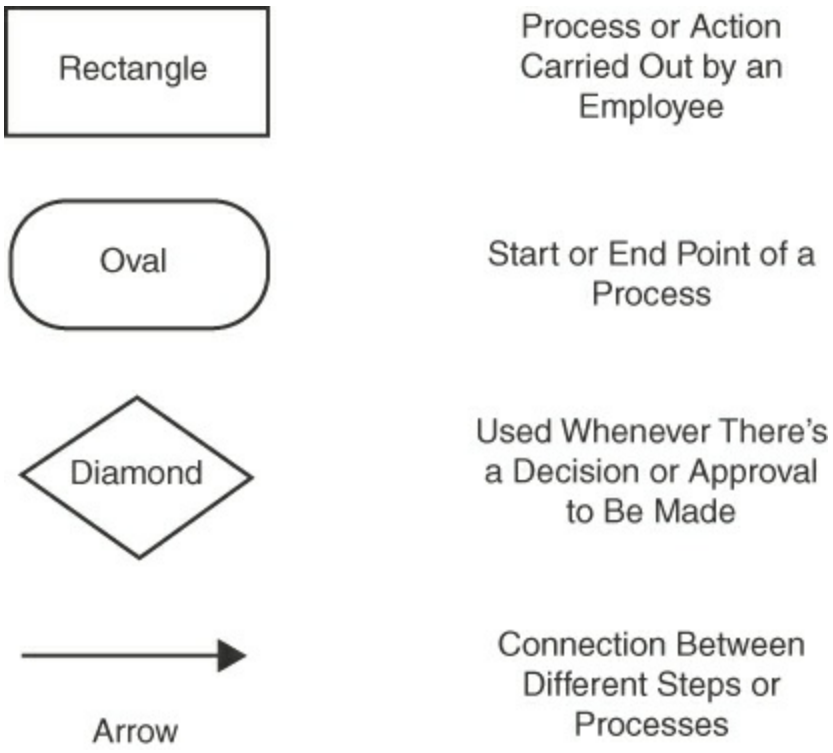

FIGURE 11-21 Industry Symbols for Playbooks

Learning DevOps

The best way to make the jump from manual processes to automation and orchestration in your playbooks is to gain experience with DevOps. I covered many core DevOps concepts in Chapter 10, but mastering DevOps will require much more reading and hands-on experience. I recommended Cisco's free labs found at (https://developer.cisco.com/site/sandbox) as a great starting point for SOC analysts to get hands-on DevOps training now at no cost. There are also industry certification programs that SOC analysts can work toward, which will help them gain the skills and confidence to convert

steps in your SOC's playbooks into automated and higher-functioning processes, leading to a much better SOC service. Figure 11-22 shows a screenshot of some of the many free labs offered in the Cisco DevNet Sandbox labs catalog that focus on Security.

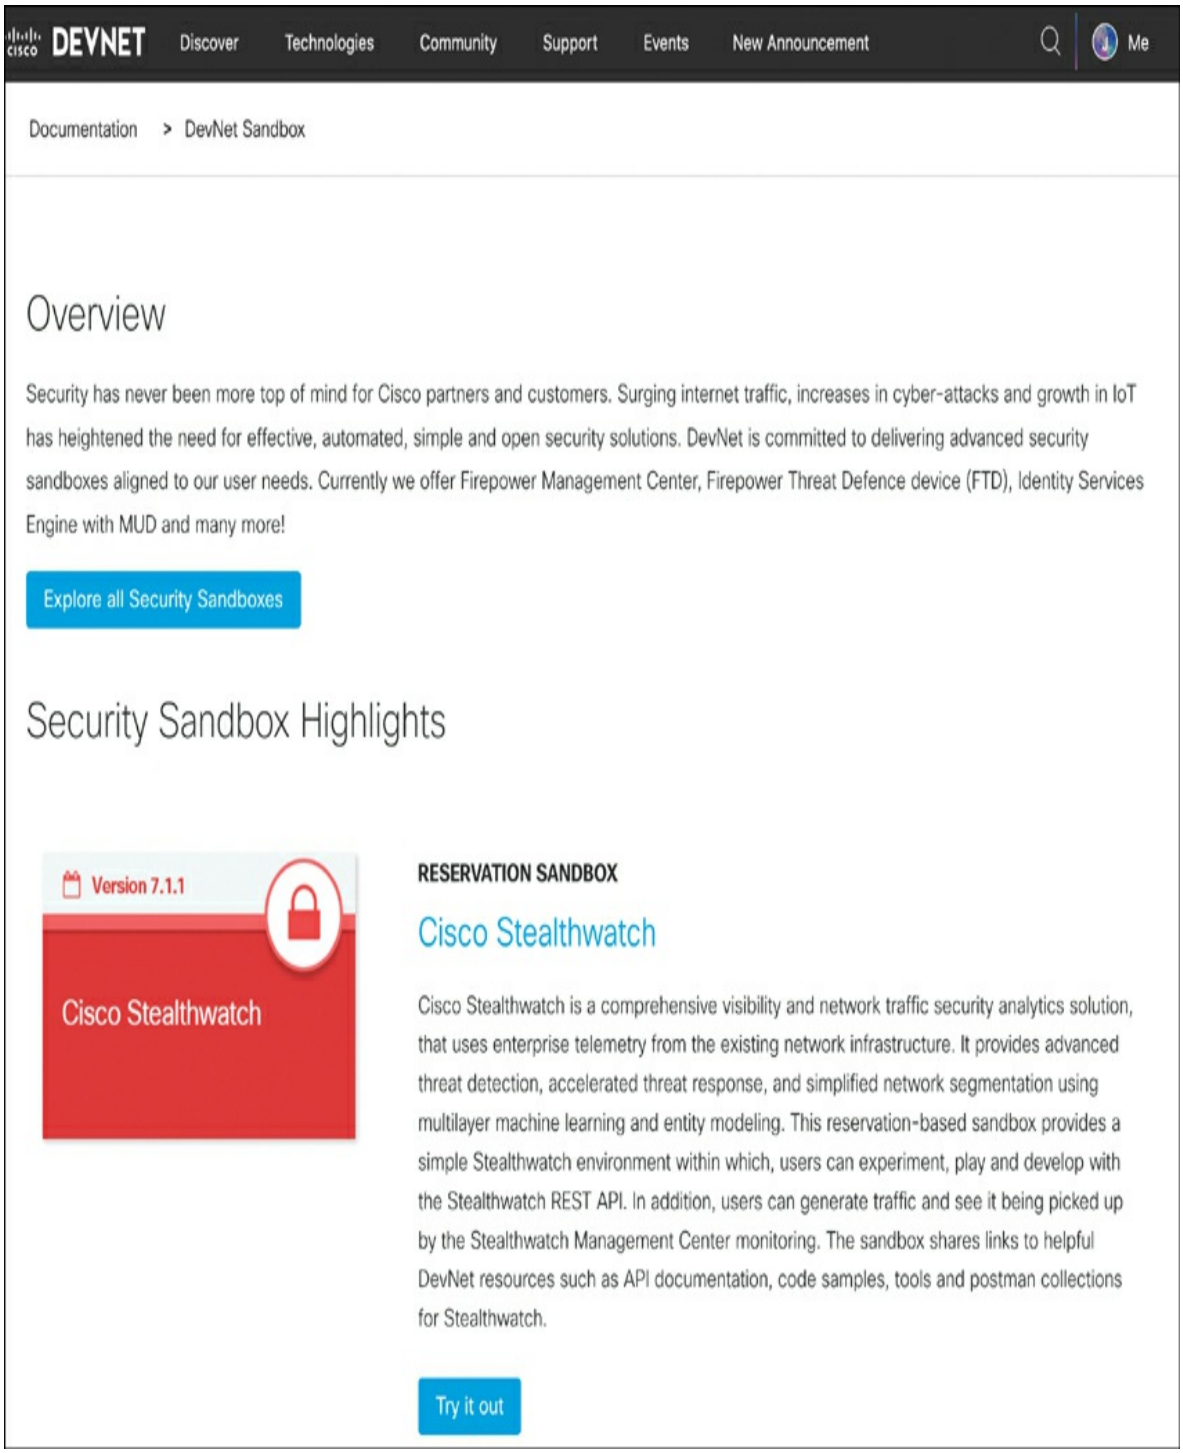

FIGURE 11-22 Cisco DevNet Sandbox Catalog of Free DevOps Labs

Future SOC Key Point

Convert all of your SOC processes into playbooks. Test your playbooks with tabletop exercises. Invest in DevOps to learn how to apply automation and orchestration to your playbooks. This is the future of all SOC services.

Hunting for Tomorrow's Threats

The final future SOC concept to consider is responding to the threats of the future. Chapter 8, "Threat Hunting and Incident Response," provided a framework for identifying and responding to security incidents as well as how to develop a hypothesis and hunt for it. Some of the concepts covered will change as SASE and other IT trends become part of the future SOC's services. The fundamental concepts, however, remain the same. Your SOC must adapt how it detects threats using the three core security capabilities in a defense-in-depth approach to ensure that any threat, regardless of tactics being used, is detected within a reasonable timeframe to avoid long-term exposure to malicious parties. Figure 11-23 is a diagram representing the three core security detection capabilities used by all security tools.

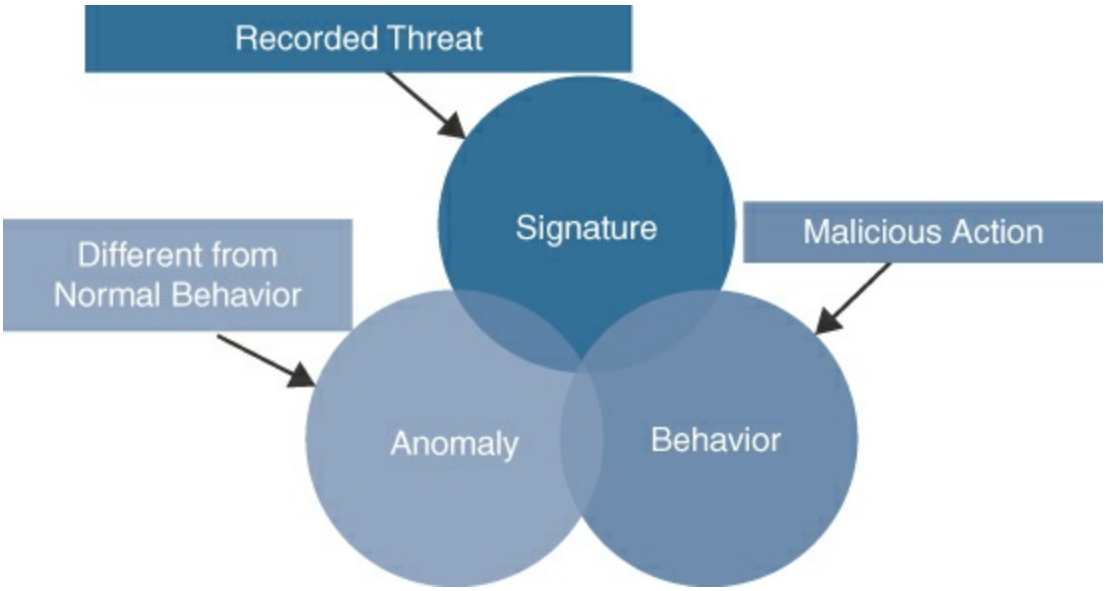

FIGURE 11-23 Security Detection Capabilities

These three core detection capabilities must be part of your SOC's defense-in-depth approach to identify threats. Signature detection represents pattern matching. I have seen statements refer to signature-based security as a fading technology, but I completely disagree with that assessment. Signatures are a quick and simple way to identify the majority of known threats, dramatically reducing the volume of risk that would need to be evaluated by the other detection capabilities. I would agree that an organization can get by without signature-based detection, but that approach will put a lot of stress on other capabilities. Also, many threat intelligence resources that complement signature tools would not be useful. I predict your future SOC will use signature detection and it will be fundamental to big data–based security that offers predictions of potential threats. Those predictions will be converted into signatures for quick evaluation.

The second detection capability is behavior-based detection. I believe of the three detection capabilities, the future SOC will see the most changes and usage of this form of detection. As technology improves and artificial intelligence becomes part of all security tools, technology will be able to determine if something is a risk and automate a response, leading to my previous prediction of future AI versus AI battles.

The final detection capability, anomaly detection, baselines normal behavior and determines if there is an outlier. I predict the future SOC will heavily

leverage big data as part of its anomaly detection capabilities. As an example of this capability, consider McAfee, which has a huge market share of consumer products, including TVs and endpoints, using McAfee antivirus solutions. This huge network of antivirus agents opens up the potential to evaluate a file against how the same file format is seen on millions of other endpoints around the world to identify when a file functions in an unusual manner based on a baseline of how this file typically behaves. I believe the future of security will include the capability of comparing an artifact against millions of others, evaluating its behavior along with comparing it against a known deny list making these three capabilities the formula to the industry standard for future SOCs detecting threats.

Responding to Threats

Once a threat is detected, incident response begins. Chapter 8 covered best practices for performing incident response. How your SOC responds to incidents today will be automated or outsourced to cloud services in the future. Guidelines such as NIST SP 800-61 provide recommendations for how to handle an incident along with who should be involved. I also pointed out in Chapter 8 that organizations such as FIRST.org have well-documented guidelines for PSIRT and CSIRT, all of which can help you develop mature incident response playbooks.

The technology may change but the fundamental approach to handing threats will remain the same. One key point from Chapter 8 that complements this thought is that you must plan for what you know and why you don’t know. You should prepare for attack techniques you have never seen and understand potential risk, legal concerns, and costs as you develop and launch your response. In certain cases, you will need to leverage external resources such as legal authorities or forensic specialists. Other times, you will depend on your SOC tools and analysts to return the impacted systems to normal operation. Once again, I highly recommend performing the different assessment services, ranging from tabletop exercises to penetration testing, to ensure your SOC analysts and the playbooks you have developed are capable of handling the threats encountered today and ones that will attempt to attack your organization tomorrow.

Future SOC Key Point

Ensure your SOC leverages the three core security detection capabilities across all parts of your business. Formulate incident response playbooks based on common attack categories and continuously test the effectiveness of your incident response program.

Summary

This final chapter provided predictions for how the future SOC will operate. I first covered secure access service edge (SASE) and its impact on technology. SASE continues to be a transformational technology that will change how your future SOC will deliver its services. Next, I made the bold prediction that IT operations will become part of the SOC's responsibility, based on trends from the impact of SASE and cloud services. The third was the future of training, which will change how the SOC will hire and maintain talent. I followed that topic with machine learning and automation concepts, which will heavily influence the future of every SOC. Finally, I closed this chapter by summarizing the concepts from this book designed to help you improve your future SOC. Every forward-looking SOC will have one or more of these concepts built into its maturity roadmap. I believe it is important for every SOC to be aware of these trends.

I stated in Chapter 1 that security is a journey, not a destination. That also applies to your SOC services and how your SOC delivers those services. Do not fear change. Embrace it, as there are a lot of exciting things coming to your future SOC. Also be aware that those same technology advances will be used by your adversaries. You must continue to adapt or you will fall behind the capabilities of the threat actors, leading to a cyberbreach.

The reason I wrote this book is to share the experience and knowledge I have gathered while working with various organizations around the world to improve their security. The SOC is a huge topic, and concepts applicable to the SOC, such as penetration testing and compliance, have their own dedicated books and certifications. My hope is that this book provides a

primer for all of the concepts I covered and ties together many technologies and concepts to the SOC. I don't expect this book will completely change how you run your SOC today, but I hope your SOC will apply some of the concepts to improve your SOC and its services. Even if those improvements are minor, over time, a collection of smaller improvements will lead to huge positive impact. Tony D'Amato (Al Pacino) said it best in the movie *Any Given Sunday*, "You find out life's this game of inches, so is football. Because in either game—life or football—the margin for error is so small. I mean, one half a step too late or too early and you don't quite make it." Security is also a game of inches. Every small improvement matters.

I hope you enjoyed this book. It was a journey to develop it, and the only reason I did so was to help you, the SOC professional!

References

Bleich, C. (n.d.). Top 10 Types of Employee Training Methods. EdgePoint Learning. https://www.edgepointlearning.com/blog/top-10-types-of-employee-training/

Burguillo, J. (2010, September). Using Game Theory and Competition-based Learning to Stimulate Student Motivation and Performance. Computers & Education, 55(2), 566–575. https://doi.org/10.1016/j.compedu.2010.02.018

Grance, T., et al. (2006, September). NIST Special Publication 800-84: Guide to Test, Training, and Exercise Programs for IT Plans and Capabilities. NIST. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-84.pdf

Greenbaum, K. (2017, July 3). The Five Top Talent Challenges of Today's C-Level Executives. Forbes. https://www.forbes.com/sites/forbeshumanresourcescouncil/2017/07/03/the-five-top-talent-challenges-of-todays-c-level-executives/?sh=7611aa25f185

Keanini, TK. (n.d.). Machine Learning and Security: Hope or Hype? Marsh & McLennan. https://www.mmc.com/insights/publications/2018/sep/machine-learning-and-security-hope-or-hype.html

Lerner, A. (December 30, 2019). Say Hello to SASE (Secure Access Service

Edge). Gartner. https://blogs.gartner.com/andrew-lerner/2019/12/23/say-hello-sase-secure-access-service-edge/

Marsala, F. (2019, September 13). The Future of Network Security Is in the Cloud. Gartner. https://www.gartner.com/en/documents/3957375/invest-implications-the-future-of-network-security-is-in

Microsoft Azure. (n.d.). What Is Cloud Bursting? Microsoft. https://azure.microsoft.com/en-us/overview/what-is-cloud-bursting/

SD-WAN-Experts. (n.d.). Welcome to the Ultimate SD-WAN Guide. SD-WAN-Experts. https://www.sd-wan-experts.com/the-ultimate-sd-wan-guide/

Florida Tech. The 7 Benefits of Delivering On-Demand Training. https://news.fit.edu/archive/the-7-benefits-of-delivering-on-demand-training/

Index

Numbers

3D printing, 638

A

Abuse.ch Feodo Tracker, 412

access

- ACL, segmentation, 117
- computer rooms, access control, 113
- datacenters, 661–662
- NAC
  - automated NAC, 501
  - manual NAC, 501
  - profiling, 128
  - SOC development, 92, 128–130
  - values, 129–130
- privileges, 265
- RBAC, 140

accreditation policies, 331–332

ACL, segmentation, 117

acoustics, facility design, 104

actionable intelligence, 378, 392

flowcharts, 414
processing data, 414
active vulnerability scanning, 86–87, 515–516
activity-attack graphs, 34–35
activity threads, 33
actors, threat, 5
cyberterrorists, 7
hacktivists, 5–6
insider threats, 7
motivations of, 7
state-sponsored actors, 6–7
AD, segmentation, 119–120
addressing risk, 172–173
business contingency planning, 173
risk heat mapping, 173–174
advanced static analysis, 448–451
adware, 456
aesthetics, SOC interior design, 105
AI (Artificial Intelligence), 315
airflow, computer rooms, 108–109
aisles, hot/cold design, 108–109
alerting levels in Cisco products, 142–143
AlienVault OTX (Open Threat Exchange), 412–413
AM (Account Managers), 214
Amazon DevOps, 612–613
analysis services, 45, 151
dynamic analysis, 200
hidden extensions diagrams, 197

job roles, 240
static analysis, 197–200
TrIDNET, 197
analytic pivoting, 30–31
anomaly detection, 15–16
Ansible
automated DevOps, 596
DevOps labs, 596–598
hosts files, 597–598
installing, 597
playbooks, 598–600
setting up, 597–598
antivirus data assessment example, 267–269
API (Application Programming Interfaces), 303–304
architectures, 304–305
examples of, 305–307
event-driven/streams, 305
IBM QRadar dashboard, 303–306
leveraging, 303–304
network programmability, NetDevOps, 605
Rapid7 Nexpose, 303–307
REST, 304
RPC, 305
applications
event logs, 273
firewalls, 534
NBAR and SOC development, 93
architectures, API, 304–305

artifacts, incident response
analyzing, 442–443
file identification tools, 445
identifying artifact types, 443–444
file identification tools, 445
magic numbers, 443–444
magic numbers, 443–444
ASHRAE, temperature/humidity in computer rooms, 108–109
assessment officers, 220–221
assessments, 355
capabilities assessments, 60–65
data, 267
antivirus data assessment example, 267–269
SOC services, 270–272
executive summaries, 357–360
FedRAMP security assessment reports, 356
future of, 667–668
goal assessments, 53
defining goals, 54–55
ranking goals, 56–58
summary of, 60
impact assessments, 356
results of, 357
risk assessment phase, vulnerability management, 504
risk assessments, 356
templates, 357–360
threat assessments, 355
types of, 355–356

vulnerabilities
assessments, 355–356, 505
scanning, 360–361
weaknesses of, 361
asset inventory phase, vulnerability management, 500–502
assets, 265
vulnerability evaluation
asset collection, 529–532
prioritizing assets, 536
vulnerability management, 522, 527
assigning tasks to incident response playbooks, 427–430
assurance of information, 9
Atomic Red Team, penetration testing, 182–185
ATT&CK Model, 35
chaining together attack behaviors, 36–38
PRE-ATT&CK research, 36–37
using, 38
attack graphs
activity-attack graphs, 34–35
Diamond Model attack graphs, 34–35
attack vectors, tactical threat intelligence, 394–395
audits, 351
compliance services, 188–189
example of, 351–352
external auditors, 353–354
firewall audit example, 351–352
future of, 667–668
internal audits, 352–353

PDCA cycle, 188–189
tools, 354–355
authenticated scanning, 86
automation
DevOps, 595–596
ML, 651
NAC, 501
playbooks, 575–578
upgrades, SASE, 630
avoiding risk, 542

B

backdoors, 456
baseline security, establishing, 11, 94, 133–135
behavior detection, 15
best-of-breed capabilities, 17
big data, centralized data management, 307–308
Hadoop, 308
challenges, 309–311
securing, 311–312
threat feeds, 312
black-box testing, 181
block pages, reputation security, 89–90
Blocklist.de, 412
blueprinting, 600–601
blue teaming. See threat hunting

boolean data type, 265
botnets, 457
branch networks, capability maps, 64–65
breaches
defense tools, 439–440
impact of, 9–10
Verizon 2020 Data Breach Investigations Report, 189–190
business challenges, SOC, 40–41
business contingency planning, 173
bytes, 264

C

capability assessments, 60
capability maps, 61, 68–69
branch networks, 64–65
endpoint security, 61–63
gap analysis, 66–68
network security, 63–64
gap analysis, 66–68
NIST CSF, 344–345
capacity planning, SOC development, 95–96, 99
careers vs. jobs, 210–211
case management, Phantom, 562–563
CEF format, logs, 278
centralized data management, 144–146, 260–261, 263
API, 303–307

architectures, 304–305
leveraging, 303–304
big data, 307–308
Hadoop, 308–312
threat feeds, 312
data assessments, 267
antivirus data assessment example, 267–269
SOC services, 270–272
data context, 265–267
access privileges, 265
asset information, 265
identity context, 265
network maps and geolocation, 266
nontechnical feeds, 266
process and operational context, 266
social and online context, 266
vulnerability context, 266
data types
booleans, 265
bytes, 264
chars, 265
doubles, 264
floats, 264
int, 264
longs, 264
primitive data types, 263–265
shorts, 264
logs, 272, 279

application event logs, 273
CEF format, 278
common log format, 278
directory service logs, 273
DNS server logs, 273
ELF, 278
endpoint logs, 272
formats of, 274–279
IoT logs, 273
JSON, 276
network device logs, 273
replication logs, 273
security tool logs, 273
syslog, 275
types of, 272–274
Windows event logs, 277
ML, 313
AI, 315
cross-validation models, 316–317
cybersecurity, 314
hold-out models, 316
models of, 315–317
semi-structured data, 263
SIEM, 279
dat digest flows, 283
data correlation, 281–282
data enrichment, 283
data processing, 280–281

IBM QRadar dashboard, 299–302
solution planning, 284–285
Splunk dashboard, 291–300, 311–312
troubleshooting, 287–301
tuning, 285–287
strategic data, 262
structured data, 263
tactical data, 262
threat mapping, 270
unstructured data, 263
Certero dashboard, vulnerability management, 522
certifications, 255–256, 331–332
chain of custody, digital forensics, 470–474
chaining together attack behaviors, ATT&CK Model, 36–37
challenges for services, 152
lack of experience, 154
limited tools, 153
low maturity, 153
people, 152
change
as cyberthreat, 8
impact of, 11–13
management, SOC development, 135–136
char data type, 265
chatbots, 657
ChatOps tools, 594–595
checklists
content quality, 390–391

threat intelligence, 389–390
Chef, automated DevOps, 596
choosing
segmentation, 117–118
threat models, 38–39
CINS Score, 412
CIS Controls, 347–349
Cisco products, alerting levels, 142–143
Cisco Webex Teams, ChatOps, 595
CISO (Chief Information Security Officers), 231–233
clean rooms, facility design, 106
client/server segmentation, 118–119
cloning Gmail, 203–204
cloud programmability
DevOps, 609–612
IT services, 639
orchestration in, 611–612
cloud/database engineers, 215
COBIT (Control Objects for Information and Related Technology)
capability scoring, 49–51
ISACA COBIT 5 Process Assessment Model, 49–51
ISACA COBIT 2019, 349
severity model, impact of incidents, 195
collaboration tools, SOC development, 138–140
collecting/processing threat intelligence, 399–400
actionable intelligence, 414
operational threat intelligence data, 402
Google Alerts, 402–403

scrapers, 403–404
social media, 404–407
strategic threat intelligence data, 400–402
technical threat intelligence data, 407
Abuse.ch Feodo Tracker, 412
AlienVault OTX, 412–413
Blocklist.de, 412
CINS Score, 412
CSV, 411
Cyber Threat System from FortiGuard Labs, 413
Dan.me.uk, 412
Emerging Threats Rule Server, 412
FBI InfraGard, 412
IBM X-Force Exchange, 413
JSON, 407–408
OpenIOC, 408
Regex, 411
SSH Bruteforce logs, 412–413
STIX, 408–409
TAXII, 409–411
XML, 407
common log format, 278
company cultures, 257
competitive workplaces, 252
compliance, 316–317
assessments, 355
executive summaries, 357–360
FedRAMP security assessment reports, 356

impact assessments, 356
results of, 357
risk assessments, 356
templates, 357–360
threat assessments, 355
types of, 355–356
vulnerability assessments, 355–356
vulnerability scanning, 360–361
weaknesses of, 361
audits, 351
example of, 351–352
external auditors, 353–354
firewall audit example, 351–352
internal audits, 352–353
tools, 354–355
CIS Controls, 347–349
exceeding compliance, 321, 350–351
FIRST CSIRT services framework, 350
frameworks, 340–350
guidelines, 340–350
industry compliance, 371–375
ISACA COBIT 2019, 349
ISO/IEC 27005, 345–347
NIST CSF, 342
capability assessments, 344–345
mapping Cisco security products to CSF, 354
tiers, 343–344
officers, 214

penetration testing, 361–362
known environments, 367
NIST Special Publication 800-115, 362–367
partially known environments, 367
planning, 368–371
scope statements, 369–371
types of, 367
unknown environments, 367
policies, 322, 327
accreditation, 331–332
certifications, 331–332
definitions and terms, 327
enforcing, 330–331
history of, 328
launching, 328–329
overview, 322–324
procedures, 332–333
purpose of, 324
scope of, 325
statements, 325–327
tabletop exercises, 334–340
services, 45, 151, 187–188
audits, 188–189
job roles, 240
SOC design considerations, 127–128
standards, 340–350
tools, vulnerability management, 522
Compromise (IOC), Indicators of, 382

computer rooms, 107
access control, 113
airflow, 108–109
equipment racks, 109
fire safety, 112
flood protection, 112
grounding, 111
hot/cold aisle design, 108–109
humidity/temperature, 108–109
lighting, 110
locks, 113
monitoring, 112
power requirements, 107–108
power-dense equipment, 109
raised floors, 111
redundancy planning, 110–111
temperature/humidity, 108–109
video surveillance, 113
connectivity (inline), network considerations, 123
containment
eradication and recovery phase, 455–483
incident response, threat hunting, 455–456
example of, 460–462
grouping, 455–456
maturity models, 460–462
performing, 459–460
stack counting, 459
techniques, 458–459

content quality, threat intelligence, 390
checklists, 390–391
key factors, 390
context
data, 265, 266–267
access privileges, 265
asset information, 265
identity context, 265
network maps and geolocation, 266
nontechnical feeds, 266
process and operational context, 266
social and online context, 266
vulnerability context, 266
threat intelligence, 379, 385–388
contingency planning, business, 173
contracted job roles, services, 165
corrective actions, vulnerability management, 539
correlating data, SIEM, 281–282
cross-validation models, ML, 316–317
CrowdStrike Falcon dashboard, EDR, 566–569
cryptographers/cryptologists, 229–230
CSF (NIST Cybersecurity Framework), 20–21, 342
capability assessments, 344–345
Framework Core, 21–22
mapping Cisco security products to CSF, 354
tiers, 343–344
CSIRT (Computer Security Incident Response Teams), 23, 350, 493–494
CSV, processing technical threat intelligence data, 411

Cuckoo sandboxes, dynamic analysis, 454
cultures of companies, 257
custody (digital forensics), chain of, 470–474
CVSS (Common Vulnerabilities Scoring System), 86, 507–508
CVSS v2, 508–512
CVSS v3, 512–514
cyber insurance, 544–547
Cyber Kill Chains, 25–29, 132
Cyber Threat System from FortiGuard Labs, 413
cybercriminals, 5
cybersecurity, ML, 314
cyberterrorists, 7
cyberthreats, 4–8
change as cyberthreat, 8
hacktivists, 5–6
insider threats, 7
motivations of, 7

D

Dan.me.uk, 412
dashboards
Certero dashboard, vulnerability management, 522
CrowdStrike Falcon dashboard, EDR, 566–569
IBM QRadar dashboard
API, 303–306
SIEM troubleshooting, 299–302

Khan Academy, 647–648
QRadar dashboard, centralized data management, 144–145
SD-WAN, 622–623
SOC development, 140–141
Splunk dashboard
centralized data management, 144–145
Hadoop, 311–312
SIEM troubleshooting, 291–300
data
assessments, 267
antivirus data assessment example, 267–269
SOC services, 270–272
at rest/in motion, SOC development, 92–93
breaches
impact of, 9–10
Verizon 2020 Data Breach Investigations Report, 189–190
context of, 265, 266–267
access privileges, 265
asset information, 265
identity context, 265
network maps and geolocation, 266
nontechnical feeds, 266
process and operational context, 266
social and online context, 266
vulnerability context, 266
correlating, SIEM, 281–282
digest flows, SIEM, 283
logs, 272, 279

application event logs, 273
CEF format, 278
common log format, 278
directory service logs, 273
DNS server logs, 273
ELF, 278
endpoint logs, 272
formats of, 274–279
IoT logs, 273
JSON, 276
network device logs, 273
replication logs, 273
security tool logs, 273
syslog, 275
types of, 272–274
Windows event logs, 277
modeling, DevOps, 589–590
processing, SIEM, 280–281
SIEM, 279
data digest flows, 283
data correlation, 281–282
data enrichment, 283
data processing, 280–281
IBM QRadar dashboard, 299–302
solution planning, 284–285
Splunk dashboard, 291–300, 311–312
troubleshooting, 287
tuning, 285–287

structures of
semi-structured data, 263
structured data, 263
unstructured data, 263
threat mapping, 270
types of
booleans, 265
bytes, 264
chars, 265
doubles, 264
floats, 264
int, 264
longs, 264
primitive data types, 263–265
shorts, 264
data management (centralized), 144–146, 260–261, 263
API, 303–307
architectures, 304–305
leveraging, 303–304
big data, 307–308
Hadoop, 308–312
threat feeds, 312
data assessments, 267
antivirus data assessment example, 267–269
SOC services, 270–272
data context, 265–267
access privileges, 265
asset information, 265

identity context, 265
network maps and geolocation, 266
nontechnical feeds, 266
process and operational context, 266
social and online context, 266
vulnerability context, 266
data types
booleans, 265
bytes, 264
chars, 265
doubles, 264
floats, 264
int, 264
longs, 264
primitive data types, 263–265
shorts, 264
logs, 272, 279
application event logs, 273
CEF format, 278
common log format, 278
directory service logs, 273
DNS server logs, 273
ELF, 278
endpoint logs, 272
formats of, 274–279
IoT logs, 273
JSON, 276
network device logs, 273

replication logs, 273
security tool logs, 273
syslog, 275
types of, 272–274
Windows event logs, 277
ML, 313
AI, 315
cross-validation models, 316–317
cybersecurity, 314
hold-out models, 316
models of, 315–317
recovery, digital forensics, 479–480
semi-structured data, 263
SIEM, 279
dat digest flows, 283
data correlation, 281–282
data enrichment, 283
data processing, 280–281
IBM QRadar dashboard, 299–302
solution planning, 284–285
Splunk dashboard, 291–300, 311–312
troubleshooting, 287–301
tuning, 285–287
sovereignty laws, 374
stealing software/keyloggers, 457
strategic data, 262
structured data, 263
tactical data, 262

threat mapping, 270
unstructured data, 263
data orchestration
blueprinting, 600–601
DevOps, 582
Amazon DevOps, 612–613
Ansible and DevOps labs, 596–598
automated DevOps, 595–596
cloud programmability, 609–612
common data formats, 585–589
data management, 583–584
data modeling, 589–590
IaaS DevOps, 610
JSON, 586
manual DevOps, 592–595
NETCONF, 590–591
NetDevOps, 604–609
PaaS DevOps, 610
RESTCONF, 591
SaaS DevOps, 610, 613–614
targets, 592
text-file formats, 584–585
tools, 591
XML, 585–586
YAML, 586–589
YANG serializers, 589–590
EDR
CrowdStrike Falcon dashboard, 566–569

NISTIR 8011 Attack Methodologies, 566
network programmability, NetDevOps, 604–605
API, 605
examples of, 606–609
OODA loop diagrams, 557–558
playbooks, 569
automation, 575–578
components of, 569–570
IRC, 571–572
malware outbreak playbooks, 196, 572–575
workflows, 570–571
workflows, examples, 579–582
SIEM, SOAR comparisons, 558
SOAR, 556–558, 560–561
Phantom, case management, 562–563
Phantom, DevOps usage example, 564–566
Phantom, example of, 561–562
Phantom, playbooks, 563–564
SIEM comparisons, 558
XDR, 559–560
database/cloud engineers, 215
datacenters, accessing, 661–662
defense-in-depth strategies, 9, 17, 136–137
defining goals, SOC goal assessments, 54–55
designing
interior design of SOC, 103–105
procedures, 83–84
SOC facilities

computer rooms, 107–113
in-house services vs. outsourcing, 102–103
interior design, 103–105
layouts, 113–114
locating, 103
physical vs. virtual SOC, 102–103
rooms, 106–113
WBDG, 101–102
desktop support, IT job roles, 215
detecting/preventing
detection and analysis phase, incident response lifecycle, 438–454
detection, 13–14
anomaly detection, 15–16
baselines, 94
behavior detection, 15
best-of-breed capabilities, 17
defense-in-depth strategies, 17
evaluating security technologies, 17–18
honeypots, 94
intrusions, 133
NBAR, 93
NetFlow, 93, 133–134
researching security technologies, 18–19
signature detection, 14
SOC development, 93–94
developing SOC
baseline tools, 133–135
centralized data management, 144–146

change management, 135–136
compliance, 127–128
dashboards, 140–141
data retention and, 143–144
detection technologies, 93
    baselines, 94
    honeypots, 94
    NBAR, 93
    NetFlow, 93, 133–134
encryption, 130–131
evaluating vulnerabilities
    active vulnerability scanning, 86–87
    CVSS, 86
    passive vulnerability scanning, 87–88
facility design
    computer rooms, 107–113
    in-house services vs. outsourcing, 102–103
    interior design, 103–105
    layouts, 113–114
    locating, 103
    physical vs. virtual SOC, 102–103
    rooms, 106–113
    WBDG, 101–102
host systems, 136–137
internal security tools, 132
intrusion detection/prevention, 133
mobile device security concerns, 94–95
NAC, 128–130

NetFlow, 133–134
network considerations, 114–115
    disaster recovery, 125–126
    inline connectivity, 123
    redundancy, risks reduction, 124–125
    segmentation, 115–120
    throughput, 120–121
network security guidelines, 137–138
packet capturing, 133–134
phases of development, 80–82
planning, 95
    capacity planning, 95–96, 99
    goal alignment, 96
    growth planning, 96–97
    redundancy planning, 98
    resource planning, 98
    technology planning, 97–98
preventive technologies, 88–89
    data at rest/in motion, 92–93
    firewalls, 89
    NAC, 92, 128–130
    reputation security, 89–91
    VPN, 91–92
procedures, 83–85
reporting, 140–141
security
    considerations, 126–127
tools, 85

storage
data retention and, 143–144
throughput and, 141–144
throughput, 141–144
tool collaboration, 138–140
development milestones, SOC, 69–70
device fingerprints, SASE, 628
DevOps, 582
Amazon DevOps, 612–613
Ansible and DevOps labs, 596–598
automated DevOps, 595–596
cloud programmability, 609–612
common data formats, 585–589
data management, 583–584
data modeling, 589–590
IaaS DevOps, 610
JSON, 586
learning, 670–671
manual DevOps, 592–593
ChatOps tools, 594–595
wikis, 593–594
NETCONF, 590–591
NetDevOps, 604–609
PaaS DevOps, 610
Phantom usage example, 564–566
RESTCONF, 591
SaaS DevOps, 610, 613–614
targets, 592

text-file formats, 584–585
tools, 591
training, future of, 650
XML, 585–586
YAML, 586–589
YANG serializers, 589–590
Diamond Model, 30–31
attack graphs, 34–35
Diamond Model for Incident Management, 32–33
Extended Diamond Model, 31
digital forensics
incident response, 467–468, 482–483
chain of custody, 470–474
data recovery, 479–480
dynamic analysis, 480–482
evidence, 474–476
first responders, 470
hashing, 476–478
process of, 468–469
static analysis, 478–479
volatile data, 480–482
labs, facility design, 106
services, 46, 151, 200–202, 240–241
directory service logs, 273
disaster recovery, network considerations, 125–126
disposal (secure), facility design, 104
disassemblers, static analysis, 199–200
distance, networks, 534–535

DLP, SASE, 629
DMZ, IDS/IPS, 534–535
DNS server logs, 273
documentation, risk documentation, 171–172
double data type, 264
downloaders, 456
DRP (Disaster Recovery Planning), 125–126
duplicating evidence, digital forensics, 474–476
dynamic analysis
analysis services, 200, 452
isolated systems, 453
sandboxes, 453–454
forensic dynamic analysis, 480–482
dynamic users/device fingerprints, SASE, 628
dysfunctional SOC, factors of, 3–4

E

EDR (Endpoint Detection and Response)
CrowdStrike Falcon dashboard, 566–569
NISTIR 8011 Attack Methodologies, 566
ELF (Extended Log Format), 278
email
ESA, 420–421
threat intelligence security, 420
deploying email security, 421
ESA, 420–421

Emerging Threats Rule Server, 412
Emily Williams hacking example, IT services, 633–636
employees
certifications, 255–256
company cultures, 247
job roles, 165
managing, 250–252
onboarding, 249–250
training, 253–255
EMV (Expected Monetary Value), 170–171
encoding files, malware, 14
encryption
LAN, 131
SOC development, 130–131
endpoint logs, 272
endpoint security
capability maps, 61–63
defense in depth strategy, 136–137
enforcing policies, 330–331
enriching data, SIEM, 283
EPS (Events Per Second)
digesting by a monitoring system, 141–142
reducing, 142–143
equipment racks, computer rooms, 109
eradication phase, incident response, 462
eradication playbooks, 464–465
system order, 463
ESA (Email Security Appliance), 420–421

evaluating
security technologies, 17–18
soft skills, 242–243
threat intelligence, 388–389
Three Pillars of Foundational SOC Support Services, The, 159
vulnerabilities, SOC development
active vulnerability scanning, 86–87
CVSS, 86
passive vulnerability scanning, 87–88
evaluation procedures, vulnerability management, 528–539
asset collection, 529–532
choosing corrective actions, 539
launch scanning, 537–539
event-driven/streams, API, 305
evidence, digital forensics, 474–476
exceeding compliance, 321, 350–351
exceptions, vulnerability management, 552–553
executive summaries, assessment template, 357–360
experience (lack of), challenges for services, 154
exploitation tools, vulnerability management, 520–521
Extended Diamond Model, 31
extensions diagrams, hidden, 197
external auditors, 353–354
external SOC services, 164
external threat intelligence, 385–386

F

Facebook, Emily Williams social engineering attack example, 634–635
facility design
computer rooms, 107–113
future of, 659–661
in-house services vs. outsourcing, 102–103
interior design, 103–105
layouts, 113–114
locating, 103
physical vs. virtual SOC, 102–103
rooms, 106–113
WBDG, 101–102
Falcon dashboard (CrowdStrike), EDR, 566–569
false positives, anomaly detection, 16
FBI InfraGard, 412
FedRAMP (Federal Risk and Authorization Management Program)
industry compliance, 374
security assessment reports, 356
feedback, threat intelligence, 421–422
file identification tools, artifact identification, 445
finding people for services, 152, 157
fingerprints
device fingerprints, SASE, 628
Nmap, 503
fire safety, computer rooms, 112
Firepower passive vulnerability scanning, 87–88, 306–307
firewalls
application-layer firewalls, 534
audit example, 351–352

SOC development, 89
first-generation SOC, 51
first responders, digital forensics, 470
FIRST service frameworks, 493
CSIRT, 23, 160–161, 350, 493–494
PSIRT, 23–24, 493
FISMA (Federal Information Security Modernization Act), 373–374
float data type, 264
flood protection, computer rooms, 112
floor layouts, facility design, 113–114
Foremost data recovery, 479–480
forensics (digital)
incident response, 467–468, 482–483
chain of custody, 470–474
data recovery, 479–480
dynamic analysis, 480–482
evidence, 474–476
first responders, 470
hashing, 476–478
process of, 468–469
static analysis, 478–479
volatile data, 480–482
labs, facility design, 106
services, 46, 151, 200–202, 240–241
forensic dynamic analysis, 480–482
forensic engineers, 230–231
forensic static analysis, 478–479
formalizing pay scales, 212–213

Foundational SOC Support Services, 154–155
evaluating, 159
people, 156–157
technology, 158–159
fourth-generation SOC, 52
Framework Core, CSF, 21–22
frameworks
compliance/risk reduction, 340–350
NIST CSF, 342–344
security, 19–20
applying, 24–25
CSF, 11, 20–22
FIRST service frameworks, 23–24, 350
free training, 644
fundamental security capabilities, 13
anomaly detection, 15–16
behavior detection, 15
best-of-breed capabilities, 17
defense-in-depth strategies, 17
evaluating security technologies, 17–18
researching security technologies, 18–19
signature detection, 14
fundamental SOC services, 150–152

G

gamifying learning, 644–645

gaps in SOC capabilities, analyzing, 66–68
geolocation and network maps, 266
Gmail, cloning, 203–204
goals
alignment, SOC development, 96
assessments, SOC, 53
defining goals, 54–55
ranking goals, 56–58
ranking threats, 58–59
summary of, 60
service job roles, 165–166
Google
Google Alerts, operational threat intelligence data, 402–403
reputation warning banners, 90–91
governance references, SOC scope statements, 80
gray-box testing, 181
grounding, computer rooms, 111
group tags, 664–665
grouping, threat hunting, 459
growth planning, SOC development, 96–97
GS pay scales, 211–213
guidelines
compliance/risk reduction, 340–350
security, 19–20
ISO 3100:2018, 22–23
NIST, 22
SOC network security, 137–138

H

hacktivists, 5–6
Hadoop, 308
challenges, 309–311
securing, 311–312
hash matches, 458
hashing, digital forensics, 476–478
heat mapping, risk, 173–174
helpdesks, IT job roles, 215
hidden extensions diagrams, 197
HIPAA (Health Insurance Portability and Accountability Act), 373
HipChat, ChatOps, 595
hold-out models, ML, 316
honeypots, 29, 94
host scanning, 516, 534
host systems, SOC development, 136–137
hot/cold aisle design, computer rooms, 108–109
humidity/temperature, computer rooms, 108–109
hunting threats, incident response, 424, 455–456
consortium playbooks, 196
example of, 460–462
grouping, 455–456
incidents, defining, 425
lifecycle of, 425–426
containment, eradication and recovery phase, 426–438
detection and analysis phase, 438–454
post-incident activity phase, 484–492

preparation phase, 426–438
maturity models, 460–462
performing, 459–460
planning, 194
SOC job roles, 221–222
stack counting, 459
techniques, 458–459
hybrid services, 44

I

IaaS, DevOps, 610
IBM QRadar dashboard
API, 303–306
SIEM troubleshooting, 299–302
IBM X-Force Exchange, 413
identity context, 265
IDS/IPS (Intrusion Detection/Prevention Systems), 534
impact assessments, 356
impact of incidents, incident management services, 194–195
incident management
Diamond Model for Incident Management, 32–33
services, 45, 151
COBIT severity model, 195
impact of incidents, 194–195
incident response planning, 194
job roles, 239–240

NIST Special Publication 800–61 Revision 2, 190–193
playbooks, 195
Verizon 2020 Data Breach Investigations Report, 189–190
incident response, 424
artifacts
analyzing, 442–443
identifying artifact types, 443–445
breach defense tools, 439–440
communication, 430–431
containment phase, threat hunting
example of, 460–462
grouping, 455–456
maturity models, 460–462
performing, 459–460
stack counting, 459
techniques, 458–459
core security capabilities, 439–440
detecting malware behavior, 441
digital forensics, 467–468, 482–483
chain of custody, 470–474
data recovery, 479–480
dynamic analysis, 480–482
evidence, 474–476
first responders, 470
hashing, 476–478
process of, 468–469
static analysis, 478–479
volatile data, 480–482

dynamic analysis, 452
isolated systems, 453
sandboxes, 453–454
eradication phase, 462
eradication playbooks, 464–465
system order, 463
FIRST service frameworks, 493
CSIRT, 493–494
PSIRT, 493
guidelines, 492–494
incident detection, 438–439
incidents, defining, 425
infected systems, 441–442
law enforcement, 432–435
Lessons Learned reports, 489–492
lifecycle of, 425–426
containment, eradication and recovery phase, 426–438
detection and analysis phase, 438–454
post-incident activity phase, 484–492
preparation phase, 426–438
malware
categories of, 456–457
threat hunting, 455–456, 458–462
packing files, 445–447
planning, 194
planning templates, 437
playbooks
consortium playbooks, 196

eradication playbooks, 464–465
recovery playbooks, 466
task assignments, 427–430
recovery phase, 466
SOC job roles, 221–222
static analysis, 446–447
advanced static analysis, 448–451
Pframe, 448
WannaCry kill switch malware analysis, 451–452
third-party interactions, 431–432
threat analysis, 440
threat hunting, 455–456
example of, 460–462
grouping, 455–456
maturity models, 460–462
performing, 459–460
stack counting, 459
techniques, 458–459
ticketing systems, 435–436
industry compliance, 371–372
data sovereignty laws, 374
FedRAMP, 374
FISMA, 373–374
HIPAA, 373
SOX, 373
industry threat models, 25
ATT&CK Model, 35–38
chaining together attack behaviors, 38

PRE-ATT&CK research, 36–37
using, 38
choosing, 38–39
Cyber Kill Chain model, 25–29
Diamond Model, 30–31
attack graphs, 34–35
Diamond Model for Incident Management, 32–33
Extended Diamond Model, 31
social-political meta-features, 31
technology meta-features, 31
infected systems, incident response, 441–442
information assurance, 9
information management phase, vulnerability management, 502–503
ingesting log data from security devices, service areas, 162–163
in-house SOC services, 42, 102–103, 164
advantages of, 42–43
disadvantages of, 43–44
inline connectivity, network considerations, 123
insider threats, 7
installation/post-sales engineers, 214
int data type, 264
interior design of SOC, 103–105
internal audits, 352–353
internal security tools
Cyber Kill Chains, 132
SOC development, 132
internal threat intelligence, 385–386
interviewing, job roles, 247

interview prompters, 247–248
post interview process, 249
intrusion detection/prevention, SOC development, 133
investing in security
defense-in-depth strategies, 9
information assurance, 9
NSA Information Assurance and Defense-in-Depth Strategy, 8–9
Investment (ROI), Return on, 421–422
IOC (Indicators of Compromise), 382, 408
IoT logs, 273
IRC playbooks, 571–572
ISACA COBIT 5 Process Assessment Model, 49–51
ISACA COBIT 2019, 349
ISO (International Organization for Standardization)
ISO 3100:2018, 22–23
ISO/IEC 27005, 345–347
isolated systems, dynamic analysis, 453
IT job roles, 213–214, 216
AM, 214
compliance officers, 214
database/cloud engineers, 215
desktop support, 215
helpdesks, 215
installation/post-sales engineers, 214
managers, 215
marketing engineers, 214
network engineers, 215
SE, 214

software engineers, 215
IT services, 631, 639–640
3D printing, 638
cloud programmability, 639
hacking, Emily Williams example, 633–636
IT operations, defined, 631–633
IT services, IT operations defined, 631–633
SASE, 637
training, 640–651
virtualized computers, 638–639
IT teams, vulnerability management, 527

J

Jenkins, automated DevOps, 596
job retention, 252–253
job roles, 206, 210–211
analysis services, 240
careers vs. jobs, 210–211
certifications, 255–256
company cultures, 247
competitive workplaces, 252
compliance services, 240
developing, 211–213
digital forensics services, 240–241
incident management services, 239–240
interviewing, 247

interview prompters, 247–248
post interview process, 249
IT job roles, 213–214, 216
AM, 214
compliance officers, 214
database/cloud engineers, 215
desktop support, 215
helpdesks, 215
installation/post-sales engineers, 214
managers, 215
marketing engineers, 214
network engineers, 215
SE, 214
software engineers, 215
managing employees, 250–252
NICE Framework, 233–237
onboarding employees, 249–250
pay scales
formalizing, 212–213
GS pay scales, 211–213
pre-interviewing, 246–247
research and development services, 241
retaining jobs, 252–253
risk management services, 239
security clearances, 244–245
services
contracted vs. employee job roles, 165
goals, 165–166

resource planning, 166–167
situational and security awareness services, 241
SOC job roles, 216–217, 231–233
assessment officers, 220–221
cryptographers/cryptologists, 229–230
forensic engineers, 230–231
incident responders, 221–222
penetration testers, 218–219
security administrators, 224–225
security analysts, 217–218
security architects, 227–229
security engineers, 225–226
security trainers, 227
systems analysts, 222–224
SOC services and associated job roles, 238–241
soft skills, 241–242
evaluating, 242–243
SOC soft skills, 243–244
tiers, 237–238
training employees, 253–255
vulnerability management services, 239
Joe sandbox, dynamic analysis, 453–454
JSON (JavaScript Object Notation), 276
DevOps, 586
processing technical threat intelligence data, 407–408

K

Kali Linux, penetration testing, 186
keyloggers/data stealing software, 457
Khan Academy, on-demand/personalized learning, 647–648
known environment penetration testing, 367

L

lack of experience, challenges for services, 154
LAN, encryption, 131
launchers, 456
launching policies, 328–329
law enforcement, incident response, 432–435
layouts, facility design, 113–114
learning
DevOps, 670–671
gamifying, 644–645
LMS, 645
on-demand learning, 646–648
personalized learning, 646–648
Lessons Learned reports, 489–492
lighting
computer rooms, 110
facility design, 104
limited tools, challenges for services, 153
LinkedIn, Emily Williams hacking example, 634
Linux (Kali), penetration testing, 186
LMS (Learning Management Systems), 645

locating SOC facilities, 103
lockers, facility design, 105
locks, computer rooms, 113
logical segmentation, 116–118
logs, 272, 279
application event logs, 273
CEF format, 278
common log format, 278
data (security devices), ingesting for service areas, 162–163
directory service logs, 273
DNS server logs, 273
ELF, 278
endpoint logs, 272
formats of, 274–279
IoT logs, 273
JSON, 276
network device logs, 273
replication logs, 273
security tool logs, 273
SSH Bruteforce logs, 412–413
syslog, 275
types of, 272–274
Windows event logs, 277
long data type, 264
low maturity, services, 153

M

magic numbers, 443–444
malware
adware, 456
backdoors, 456
botnets, 457
categories of, 456–457
detecting behavior, 441
downloaders, 456
encoding files, 14
keyloggers/data stealing software, 457
launchers, 456
matching hashes, 458
outbreak playbooks, 196, 572–575
packing files, analysis services, 445–447
phoning home, 457
port scanning, 457–458
ransomware, 457
rootkits, 456
scareware, 457
signature detection, 14
spam, 457
threat hunting, 455–456
example of, 460–462
grouping, 455–456
maturity models, 460–462
performing, 459–460
stack counting, 459
techniques, 458–459

viruses, 457
WannaCry kill switch malware analysis, 451–452
worms, 457
managers, IT job roles, 215
manager's office, facility design, 106
managing
analysis services, job roles, 240
asset management, vulnerabilities, 522
change, SOC development, 135–136
compliance services, job roles, 240
data management (centralized), 144–146, 260–261
API, 303–307
big data, 307–313
data assessments, 267–272
data context, 265–267
data structures, 263
data types, 263–265
Hadoop, 308–312
logs, 272–279
ML, 314–317
semi-structured data, 263
SIEM, 279–302
strategic data, 262
structured data, 263
tactical data, 262
threat mapping data, 270
unstructured data, 263
digital forensics services, job roles, 240–241

incident management services, 45, 151
COBIT severity model, 195
impact of incidents, 194–195
incident response planning, 194
job roles, 239–240
NIST Special Publication 800–61 Revision 2, 190–193
playbooks, 195
Verizon 2020 Data Breach Investigations Report, 189–190
information management phase, vulnerability management, 502–503
MDM, 94–95
Nmap scanning, 501–502
people, 250–252
power
power-dense equipment, computer rooms, 109
UPS, computer rooms, 110–111
research and development services, job roles, 241
risk management services, 45, 150, 169
addressing risk, 172–174
four responses to risk, 169–170
job roles, 239
reducing risk, 169–172
situational and security awareness services, job roles, 241
vulnerability management, 498–499, 501
accuracy, 540–541
asset access, 535
asset inventory phase, 500–502
asset management, 522
best practices, 499–500

Certero dashboard, 522
CVSS, 507–514
cyber insurance, 544–547
deployment example, 535
evaluation procedures, 528–539
exceptions, 552–553
exploitation tools, 520–521
host scanning, 516
information management phase, 502–503
measuring vulnerabilities, 506
NAC, 501, 522–524
network scanners, 501–502, 515
patching systems, 547–549
process summary, 554–555
program diagrams, 527–528
remediation approval, 550–551
report and remediate phase, 505
reporting, 552
respond and repeat phase, 506
responses, 540, 542–544
risk assessment phase, 504
shorthand, 511–512
Struts vulnerability example, 507, 512–514
temporal/environmental metrics, 511
threat detection tools, 524–525
vulnerability assessments, 505
vulnerability scanning, 515–520
vulnerability management services, 45, 150, 175, 525

best practices, 175–176
job roles, 239
OpenVAS, 178
penetration testing, 179–187
roles, 527–528
scannng services, 525–527
Tenable.sc vulnerability tracking, 177
vulnerability tracking, 179
manual DevOps, 592–593
ChatOps tools, 594–595
wikis, 593–594
manual NAC (Network Access Control), 501
maps
capability maps, 61, 68–69
branch networks, 64–65
endpoint security, 61–63
gap analysis, 66–68
network security, 63–64
data, threats, 270
risk heat maps, 173–174
marketing engineers, 214
matching hashes, 458
maturity (low), services, 153
maturity models, 47
assessments, 47–48
ISACA COBIT 5 Process Assessment Model, 49–51
program maturity, 51–53
services, 167–168

SOC-CMM Model, 49
threat hunting, incident response, 460–462
MDM (Mobile Device Management), 94–95
measuring vulnerabilities, 506
Metasploit, penetration testing, 14, 186–187
Microsoft Teams, ChatOps, 595
mission statements, 74–75
developing, 75–76
sample statements, 76–77
MITRE ATT&CK Model, 35–38
chaining together attack behaviors, 36–37
penetration testing, 182
PRE-ATT&CK research, 36–37
using, 38
ML (Machine Learning), 313, 651–652
AI, 315
applied, 653–654
automation, 651
chatbots, 657
cross-validation models, 316–317
cybersecurity, 314
future of, 656–659
hold-out models, 316
hurdles of, 652–653
models of, 315–317
training, 655
mobile devices
MDM, 94–95

security concerns, SOC development, 94–95
modified waterfall model, processing threat intelligence, 400–402
monitoring, computer rooms, 112
monitoring systems, EPS, digesting, 141–142
Moodle, LMS, 645
motivations of threat actors, 7

N

NAC (Network Access Control), 12
automated NAC, 501
profiling, 128
SOC development, 92, 128–130
values, 129–130
vulnerability management, 522–524
name servers, rogue, 282
NAT (Network Address Translation), 534
NBAR (Network-Based Application Recognition), 93
NERC CIP (North American Electric Reliability Corporation, Critical Infrastructure Protection), 375
NETCONF, 590–591
NetDevOps, 604–605
API, 605
examples of, 606–609
NetFlow, 93, 133–134
network scanners, 501–502
application-layer firewalls, 534
distance, networks, 534–535

IDS/IPS, 534
NAC, 522–524
network scanners, 515
perimeter networks (DMZ), 535
segmentation, 534–535
templates, 534, 536
VPN, 534
networks
branch networks, capability maps, 64–65
connectivity, inline connectivity, 123
device logs, 273
disaster recovery, 125–126
distance, 534–535
engineers, 215
LAN, encryption, 131
maps and geolocation, 266
perimeter networks (DMZ), 535
programmability, NetDevOps, 601–604
API, 605
examples of, 606–609
redundancy, risks reduction, 124–125
SD-WAN, 618–622
benefits of, 622–623
dashboard example, 622–623
DLP, 629
tier one support, 629–630
security, capability maps, 63–64
segmentation, 115–116

ACL, 117
AD segmentation, 119–120
choosing, 117–118
client/server segmentation, 118–119
logical segmentation, 116–118
server segmentation, 118–119
SOC design considerations, 114–115
network security guidelines, 137–138
segmentation, 115–120
throughput, 120–121
throughput, 120–121
requirements, 121–123
VPN, 534
SASE, 628–629
SOC development, 91–92
WAN, 618–620
Nexpose vulnerability scanner, 86–87
NICE Framework, 233–237
NIST (National Institute of Standards and Technology)
CSF, 11, 20–21, 342
capability assessments, 344–345
Framework Core, 21–22
mapping Cisco security products to CSF, 354
tiers, 343–344
guidelines, 22
SP 800–61 Rev. 2 Incident Response Lifecycle, 425–426
containment, eradication and recovery phase, 426–438
incident management, 190–193

preparation phase, 426–454, 484–492
SP 800-84, future of SOC staff, 666–667
SP 800-86, digital forensics services, 201–202
SP 800-115, penetration testing, 180–182, 362–367
NISTIR 8011 Attack Methodologies, 566
Nmap
fingerprinting, 503
scanning, 501–502
nontechnical feeds, 266
nontechnical intelligence. See strategic threat intelligence
NSA Information Assurance and Defense-in-Depth Strategy, 8–9

O

onboarding employees, 249–250
on-demand experts, future of training, 649
on-demand learning, 646–648
online and social data context, 266
OODA loop diagrams, 557–558
OpenIOC, processing technical threat intelligence data, 408
OpenVAS, vulnerability scanning, 178
operational threat intelligence, 205, 382, 384–385
data expectations, 396–397
processing data, 402
Google Alerts, 403–404
scrapers, 403–404
social media, 404–407

operations rooms, facility design, 106
OPEX (Operating Expenses), 628
orchestrating data
blueprinting, 600–601
DevOps, 582
Amazon DevOps, 612–613
Ansible and DevOps labs, 596–598
automated DevOps, 595–596
cloud programmability, 609–612
common data formats, 585–589
data management, 583–584
data modeling, 589–590
IaaS DevOps, 610
JSON, 586
manual DevOps, 592–595
NETCONF, 590–591
NetDevOps, 604–609
PaaS DevOps, 610
RESTCONF, 591
SaaS DevOps, 610, 613–614
targets, 592
text-file formats, 584–585
tools, 591
XML, 585–586
YAML, 586–589
YANG serializers, 589–590
EDR
CrowdStrike Falcon dashboard, 566–569

NISTIR 8011 Attack Methodologies, 566
network programmability, NetDevOps, 604–605
API, 605
examples of, 606–609
OODA loop diagrams, 557–558
playbooks, 569
automation, 575–578
components of, 569–570
IRC, 571–572
malware outbreak playbooks, 196, 572–575
workflows, 570–571, 579–582
SIEM, SOAR comparisons, 558
SOAR, 556–558, 560–561
Phantom, case management, 562–563
Phantom, DevOps usage example, 564–566
Phantom, example of, 561–562
Phantom, playbooks, 563–564
SIEM comparisons, 558
XDR, 559–560
Osquery
blueprinting, 600–601
running, 601–604
outsourcing services, 42, 102–103

P

PaaS, DevOps, 610

packed files
Peframe, 198–199
static analysis, 197–199
packet capturing, SOC development, 135
packing files, analysis services, 445–447
partially known environment penetration testing, 367
passive vulnerability scanning, 87–88, 516–517
patching systems, vulnerability management, 547–549
pay scales
formalizing, 212–213
GS pay scales, 211–213
PDCA cycle, audits, 188–189
Peframe packed file analysis, 198–199
penetration testing, 179, 361–362
Atomic Red Team, 182–185
black-box testing, 181
Emily Williams example, hacking IT services, 635–636
future of, 667–668
gray-box testing, 181
Kali Linux, 186
known environments, 367
Metasploit, 14, 186–187
MITRE ATT&CK Model, 182
NIST SP 800–115, 180–182, 362–367
partially known environments, 367
planning, 368–371
scope statements, 369–371
SOC job roles, 218–219

Surveyor, 185
types of, 367
unknown environments, 367
people
finding for services, 152, 157
managing, 250–252
Three Pillars of Foundational SOC Support Services, The, 156–157
perimeter networks (DMZ), 535
personalized learning, 646–648
Pframe, static analysis, 448
Phantom
case management, 562–563
DevOps usage example, 564–566
playbooks, 563–564
SOAR example, 561–562
phases of SOC development, 80–82
phoning home, malware, 457
physical SOC, facility design, 102–103
pivoting, analytic, 30–31
planning
business contingency planning, 173
DRP, 125–126
incident response planning, 194
incident response planning templates, 437
penetration testing, 368–371
redundancy planning, computer rooms, 110–111
resource planning, service job roles, 166–167
SOC, 95

capacity planning, 95–96, 99
goal alignment, 96
growth planning, 96–97
redundancy planning, 98
resource planning, 98
technology planning, 97–98
solution planning, SIEM, 284–285
threat intelligence, 393–398
vulnerability evaluation procedures, planning, 532–537
work environments, 155–156
playbooks, 569
Ansible, 598–600
automation, 575–578
components of, 569–570
eradication playbooks, 464–465
incident management services, 195
incident response
consortium playbooks, 196
eradication playbooks, 464–465
IRC, 571–572
malware outbreak playbooks, 196, 572–575
Phantom usage example, 563–564
recovery playbooks, 466
workflows
examples, 579–582
sample workflow, 570–571
symbols, 570
policies, 322

accreditation, 331–332
certifications, 331–332
compliance, 327
definitions and terms, 327
enforcing, 330–331
history of, 328
launching, 328–329
overview, 322–324
procedures, 332–333
purpose of, 324
scope of, 325
statements, 325–327
tabletop exercises, 334–335
example of, 337–340
executing, 336–337
format of, 337–338
options, 334–335
port scanning, 457–458
post interview process, 249
post-incident activity phase, incident response lifecycle, 484–492
post-sales/installation engineers, 214
power management
power-dense equipment, computer rooms, 109
UPS, computer rooms, 110–111
power requirements, computer rooms, 107–108
power-dense equipment, computer rooms, 109
PRE-ATT&CK research, 36–37
pre-interviewing, job roles, 246–247

preparation phase, incident response lifecycle, 426–438
prevalence, threat intelligence, 387
preventing intrusions, SOC development, 133
preventive technologies
data at rest/in motion, SOC development, 2–93
firewalls, SOC development, 89
NAC
profiling, 128
SOC development, 92, 128–130
values, 129–130
reputation security, SOC development, 89–91
SOC development, 88–93
VPN, SOC development, 91–92
primitive data types, 263–265
prioritizing assets, vulnerability evaluation, 536
procedures, 82
designing, 83–84
examples of, 84–85
policies, 332–333
process and operational context, 266
processing data, SIEM, 280–281
processing threat intelligence, 399–400
actionable intelligence, 414
operational threat intelligence data, 402
Google Alerts, 402–403
scrapers, 403–404
social media, 404–407
strategic threat intelligence data, 400–402

technical threat intelligence data, 407
Abuse.ch Feodo Tracker, 412
AlienVault OTX, 412–413
Blocklist.de, 412
CINS Score, 412
CSV, 411
Cyber Threat System from FortiGuard Labs, 413
Dan.me.uk, 412
Emerging Threats Rule Server, 412
FBI InfraGard, 412
IBM X-Force Exchange, 413
JSON, 407–408
OpenIOC, 408
Regex, 411
SSH Bruteforce logs, 412–413
STIX, 408–409
TAXII, 409–411
XML, 407
profiling NAC, 128
proxy servers, rogue, 282
PSIRT (Product Incident Response Teams), 23–24, 493
Puppet, automated DevOps, 596

Q

QRadar dashboard, centralized data management, 144–145
quality of content, threat intelligence, 390

checklists, 390–391
key factors, 390

R

raised floors, computer rooms, 111
ranking
SOC goals, 56–58
threats, 58–59
ransomware, 457
Rapid7 Nexpose
API, 303–304, 305–307
Struts vulnerability example, 514
RBAC (Role-Based Access Control), 140
recovering data, digital forensics, 479–480
recovery phase, incident response, 466
reducing EPS, 142–143
reducing risk, 169, 316–317
assessments, 355
executive summaries, 357–360
FedRAMP security assessment reports, 356
impact assessments, 356
results of, 357
risk assessments, 356
templates, 357–360
threat assessments, 355
types of, 355–356

vulnerability assessments, 355–356
vulnerability scanning, 360–361
weaknesses of, 361
audits, 351
example of, 351–352
external auditors, 353–354
firewall audit example, 351–352
internal audits, 352–353
tools, 354–355
CIS Controls, 347–349
EMV approach, 170–171
FIRST CSIRT services framework, 350
frameworks, 340–350
guidelines, 340–350
industry compliance, 371–375
ISACA COBIT 2019, 349
ISO/IEC 27005, 345–347
NIST CSF, 342
capability assessments, 344–345
mapping Cisco security products to CSF, 354
tiers, 343–344
penetration testing, 361–362
known environments, 367
NIST Special Publication 800–115, 362–367
partially known environments, 367
planning, 368–371
scope statements, 369–371
types of, 367

unknown environments, 367
policies, 322
accreditation, 331–332
certifications, 331–332
compliance, 327
definitions and terms, 327
enforcing, 330–331
history of, 328
launching, 328–329
overview, 322–324
procedures, 332–333
purpose of, 324
scope of, 325
statements, 325–327
tabletop exercises, 334–340
redundancy, 124–125
risk documentation, 171–172
risk register systems, 172
standards, 340–350
redundancy
planning
computer rooms, 110–111
SOC development, 98
reducing risk, 124–125
Regex (Regular Expressions), 411
remediation approval, vulnerability management, 550–551
remote users, 661
replication logs, 273

report and remediate phase, vulnerability management, 505
reporting
SOC development, 140–141
vulnerability management, 552
reputation security
block pages, 89–90
Google reputation warning banners, 90–91
SOC development, 89–91
reputation warning banners, Google, 90–91
research and development services, 46, 151, 205–206, 241
researching security technologies, 18–19
residual risk, 550
resource planning
service job roles, 166–167
SOC development, 98
respond and repeat phase, vulnerability management, 506
REST (Representational State Transfer), 304
RESTCONF, 591
retaining jobs, 252–253
reverse engineering files, static analysis, 199–200
risk, 39–40
assessment phase, vulnerability management,
assessments, 356, 504
avoidance, 542
contingency, 171
flowcharts, 542–543
heat mapping, 173–174
modifying, 542

reducing, redundancy, 124–125
register systems, 172
retention, 542
scope statements, managing risk, 80
transfer/sharing, 542
risk management services, 45, 150, 169
addressing risk, 172–173
business contingency planning, 173
risk heat mapping, 173–174
four responses to risk, 169–170
job roles, 239
reducing risk, 169
EMV approach, 170–171
risk documentation, 171–172
risk register systems, 172
risk reduction, 316–317
assessments, 355
executive summaries, 357–360
FedRAMP security assessment reports, 356
impact assessments, 356
results of, 357
risk assessments, 356
templates, 357–360
threat assessments, 355
types of, 355–356
vulnerability assessments, 355–356
vulnerability scanning, 360–361
weaknesses of, 361

audits, 351
example of, 351–352
external auditors, 353–354
firewall audit example, 351–352
internal audits, 352–353
tools, 354–355
CIS Controls, 347–349
FIRST CSIRT services framework, 350
frameworks, 340–350
guidelines, 340–350
industry compliance, 371–375
ISACA COBIT 2019, 349
ISO/IEC 27005, 345–347
NIST CSF, 342
capability assessments, 344–345
mapping Cisco security products to CSF, 354
tiers, 343–344
penetration testing, 361–362
known environments, 367
NIST Special Publication 800–115, 362–367
partially known environments, 367
planning, 368–371
scope statements, 369–371
types of, 367
unknown environments, 367
policies, 322
accreditation, 331–332
certifications, 331–332

compliance, 327
definitions and terms, 327
enforcing, 330–331
history of, 328
launching, 328–329
overview, 322–324
procedures, 332–333
purpose of, 324
scope of, 325
statements, 325–327
tabletop exercises, 334–340
standards, 340–350
rogue name servers, 282
rogue proxy servers, 282
ROI, threat intelligence feedback, 421–422
rootkits, 456
RPC (Remote Procedure Calls), 305

S

SaaS (Software as a Service)
DevOps, 610, 613–614
future of, 627
SaltStack, automated DevOps, 596
sandboxes, dynamic analysis, 453–454
SANS, vulnerability management best practices, 12
SASE (Secure Access Service Edge), 616–617, 623–625

automated upgrades, 630
defined, 625–626
dynamic users/device fingerprints, 628
future of, 627–631
IT services, 637
OPEX, 628
SaaS, 627
VPN, 628–629
scanning for vulnerabilities, 12, 176–177
active vulnerability scanning, 86–87
assessments, 360–361
authenticated scanning, 86
Firepower, 87–88, 306–307
Nexpose vulnerability scanner, 86–87
passive vulnerability scanning, 87–88
unauthenticated scanning, 86
scanning services, vulnerability management, 525–527
scareware, 457
SCIF (Sensitive Compartmented Information Facilities), 106
scope of policies, 325
scope statements, 74–75
challenges of, 79–80
developing, 77–78
governance references, 80
penetration testing, 369–371
risk management references, 80
sample statements, 78–79
scrapers, operational threat intelligence data, 403–404

SD-WAN (Software-Defined Wide-Area Networks), 618–622
benefits of, 622–623
dashboard example, 622–623
DLP, 629
tier one support, 629–630
SE (Sales Engineers), 214
second-generation SOC, 51
secure disposal, facility design, 104
security
administrators, 224–225
analysts, 217–218
architects, 227–229
baselines, establishing, 11, 94
breaches, impact of, 9–10
change, impact of, 11–13
clearances, job roles, 244–245
detection capabilities, 13–14
anomaly detection, 15–16
behavior detection, 15
best-of-breed capabilities, 17
defense-in-depth strategies, 17
evaluating security technologies, 17–18
researching security technologies, 18–19
signature detection, 14
email, threat intelligence security, 420
deploying email security, 421
ESA, 420–421
endpoint security, defense in depth strategy, 136–137

engineers, 225–226, 527
evaluating security technologies, 17–18
facility design, 104
frameworks, 19–20
  applying, 24–25
  CSF, 11, 20–22
  CSIRT, 23
  FIRST service frameworks, 23–24
  PSIRT, 23–24
fundamental security capabilities, 13
  anomaly detection, 15–16
  behavior detection, 15
  best-of-breed capabilities, 17
  defense-in-depth strategies, 17
  evaluating security technologies, 17–18
  researching security technologies, 18–19
  signature detection, 14
guidelines, 19–20
  ISO 3100:2018, 22–23
  NIST, 22
incident response, 424
  artifact analysis, 442–443
  breach defense tools, 439–440
  communication, 430–431
  consortium playbooks, 196
  core security capabilities, 439–440
  detecting malware behavior, 441
  identifying artifact types, 443–445

incidents, defining, 425
incidents, detecting, 438–439
infected systems, 441–442
law enforcement, 432–435
lifecycle of, 425–426
lifecycle of, containment, eradication and recovery phase, 426–438
lifecycle of, detection and analysis phase, 438–454
lifecycle of, post-incident activity phase, 484–492
lifecycle of, preparation phase, 426–438
packing files, 445–447
planning, 194
planning templates, 437
playbooks, 196, 427–430
SOC job roles, 221–222
static analysis, 446–448
third-party interactions, 431–432
threat analysis, 440
ticketing systems, 435–436
internal security tools
Cyber Kill Chains, 132
SOC development, 132
investing in
defense-in-depth strategies, 9
information assurance, 9
NSA Information Assurance and Defense-in-Depth Strategy, 8–9
log data from security devices, ingesting for service areas, 162–163
mobile devices, SOC development, 94–95
officers, vulnerability management, 527

reputation security, 89–91
block pages, 89–90
Google reputation warning banners, 90–91
researching security technologies, 18–19
SOC design considerations, 126–127
SOC technology, 158–159
standards, 19–20
threat intelligence security tools, 414–416
email security, 420–421
SIEM, 416–419
tools
logs, 273
SOC development, 85
trainers, 227
segmentation, 115–116, 534
ACL, 117
AD segmentation, 119–120
choosing, 117–118
client/server segmentation, 118–119
group tags, 664–665
logical segmentation, 116–118
server segmentation, 118–119
semi-structured data, 263
servers
compromise, 282
rogue name servers, 282
rogue proxy servers, 282
segmentation, 118–119

service areas, 160
developing, 161–163
FIRST CSIRT services/service areas, 160–161
log data from security devices, ingesting, 162–163
services, 46, 150
analysis services, 45, 151
dynamic analysis, 200
hidden extensions diagrams, 197
job roles, 240
static analysis, 197–200
TrIDNET, 197
challenges, 152
lack of experience, 154
limited tools, 153
low maturity, 153
people, 152
compliance services, 45, 151, 187–188
audits, 188–189
job roles, 240
SOC design considerations, 127–128
data assessments, 270–272
digital forensics services, 46, 151, 200–202, 240–241
external SOC services, 164
FIRST CSIRT services/service areas, 160–161
fundamental services, 150–152
future impact of, 669–671
in-house services, 42, 102–103, 164
advantages of, 42–43

disadvantages of, 43–44
incident management services, 45, 151
COBIT severity model, 195
impact of incidents, 194–195
incident response planning, 194
job roles, 239–240
NIST Special Publication 800–61 Revision 2, 190–193
playbooks, 195
Verizon 2020 Data Breach Investigations Report, 189–190
IT services, 631, 639–640
3D printing, 638
cloud programmability, 639
hacking, Emily Williams example, 633–636
SASE, 637
training, 640–651
virtualized computers, 638–639
job roles
contracted vs. employee job roles, 165
goals, 165–166
resource planning, 166–167
SOC services and associated job roles, 238–241
tiers, 237–238
maturity models, 167–168
outsourcing services, 42, 102–103
research and development services, 46, 151, 205–206, 241
risk management services, 45, 150, 169
addressing risk, 172–174
four responses to risk, 169–170

job roles, 239
reducing risk, 169, 170–172
scanning services, vulnerability management, 525–527
situational and security awareness services, 46, 151, 202–203
cloning Gmail, SET, 203–205
job roles, 241
user training, 203–205
Three Pillars of Foundational SOC Support Services, The, 154–155
evaluating, 159
people, 156–157
technology, 158–159
vulnerability management services, 45, 150, 175, 525
best practices, 175–176
job roles, 239
OpenVAS, 178
penetration testing, 179–187
roles, 527–528
Tenable.sc vulnerability tracking, 177
vulnerability tracking, 179
SET, cloning Gmail, 203–204
short data type, 264
SIEM (Security Information and Event Management), 279
dat digest flows, 283
data correlation, 281–282
data enrichment, 283
data processing, 280–281
IBM QRadar dashboard, 299–306
SOAR comparisons, 558

solution planning, 284–285
Splunk dashboard, 291–300, 311–312
threat intelligence security, 416–419
troubleshooting, 287, 291
actionable intelligence, 300–301
data input, 288, 293–299
data processing, 289–291
data storage, 291–293
IBM QRadar dashboard, 299–302
Splunk dashboard, 291–300, 311–312
validating results, 299–300
tuning, 285–287
signature detection, 14
situation rooms, facility design, 106
situational and security awareness services, 46, 151, 202–203
cloning Gmail, SET, 203–205
job roles, 241
user training, 203–205
Slack, ChatOps, 595
SOAR (Security Orchestration, Automation and Response), 557–558, 560–561
Phantom
case management, 562–563
DevOps usage example, 564–566
example of, 561–562
playbooks, 563–564
SIEM comparisons, 558
SOC (Security Operations Center), 2–3

business challenges, 40–41
capabilities assessments, 60
    capability maps, 61–65
    gap analysis, 68–69
developing
    baseline tools, 133–135
    centralized data management, 144–146
    change management, 135–136
    compliance, 127–128
    dashboards, 140–141
    data retention and, 143–144
    detection technologies, 93–94
    encryption, 130–131
    evaluating vulnerabilities, 86–88
    facility design, 101–114
    host systems, 136–137
    internal security tools, 132
    intrusion detection/prevention, 133
    mobile device security concerns, 94–95
    NAC, 128–130
    NetFlow, 133–134
    network considerations, 114–125
    network security guidelines, 137–138
    packet capturing, 133–134
    phases of development, 80–82
    planning SOC, 95–99
    preventive technologies, 88–93
    procedures, 83–85

reporting, 140–141
security tools, 85
throughput, 141–144
tool collaboration, 138–140
development milestones, 69–70
dysfunctional SOC, factors of, 3–4
facility design
computer rooms, 107–113
future of, 659–661
in-house services vs. outsourcing, 102–103
interior design, 103–105
layouts, 113–114
locating, 103
physical vs. virtual SOC, 102–103
rooms, 106–113
WBDG, 101–102
first-generation SOC, 51
fourth-generation SOC, 52
future of, 659
goal assessments, 53
defining goals, 54–55
ranking goals, 56–58
ranking threats, 58–59
summary of, 60
job roles, 216–217, 231–233
analysis services, 240
assessment officers, 220–221
certifications, 255–256

company cultures, 247
competitive workplaces, 252
compliance services, 240
cryptographers/cryptologists, 229–230
digital forensics services, 240–241
forensic engineers, 230–231
incident management services, 239–240
incident responders, 221–222
interviewing, 247–249
managing employees, 250–252
onboarding employees, 249–250
penetration testers, 218–219
pre-interviewing, 246–247
research and development services, 241
retaining jobs, 252–253
risk management services, 239
security administrators, 224–225
security analysts, 217–218
security architects, 227–229
security clearances, 244–245
security engineers, 225–226
security trainers, 227
situational and security awareness services, 241
SOC services and associated job roles, 238–241
soft skills, 241–244
systems analysts, 222–224
tiers, 237–238
training employees, 253–255

vulnerability management services, 239
maturity models, 47
assessments, 47–48
ISACA COBIT 5 Process Assessment Model, 49–51
program maturity, 51–53
SOC-CMM Model, 49
mission statements, 74–75
developing, 75–76
sample statements, 76–77
network considerations, 114–115
disaster recovery, 125–126
inline connectivity, 123
redundancy, risks reduction, 124–125
segmentation, 115–120
throughput, 120–121
phases of development, 80–82
physical vs. virtual SOC, 102–103
planning, 95
capacity planning, 95–96, 99
goal alignment, 96
growth planning, 96–97
redundancy planning, 98
resource planning, 98
technology planning, 97–98
procedures, 82
designing, 83–84
examples of, 84–85
risk, 39–40

scope statements, 74–75
challenges of, 79–80
developing, 77–78
governance references, 80
risk management references, 80
sample statements, 78–79
second-generation SOC, 51
security considerations, 126–127
service areas, 160
developing, 161–163
FIRST CSIRT services/service areas, 160–161
ingesting log data from security devices, 162–163
services, 46
analysis services, 45, 151, 197–200, 240
associated job roles, 238–241
challenges, 152–154
compliance services, 45, 151, 187–189, 240
data assessments, 270–272
digital forensics services, 46, 151, 200–202, 240–241
external SOC services, 164
FIRST CSIRT services/service areas, 160–161
fundamental services, 150–152
in-house services, 42–44, 102–103
in-house SOC services, 164
hybrid services, 44
incident management services, 45, 151, 189–195, 239–240
job roles, tiers, 237–238
maturity models, 167–168

outsourcing services, 42, 102–103
research and development services, 46, 151, 205–206, 241
risk management services, 45, 150, 169–174, 239
situational and security awareness services, 46, 151, 202–205, 241
Three Pillars of Foundational SOC Support Services, The, 154–159
vulnerability management services, 45, 150, 175–187, 239
staff, future of, 666–667
third-generation SOC, 52
vulnerabilities, 39–40
SOC-CMM maturity model, 49
social and online data context, 266
social engineering
attack example, hacking IT services, 634–635
SET, cloning Gmail, 203–204
social media, operational threat intelligence data, 404–407
social-political meta-features, 31
soft skills, job roles, 241–242
evaluating, 242–243
SOC soft skills, 243–244
software
engineers, 215
SaaS, 610, 613–614, 627
solution planning, SIEM, 284–285
sovereignty of data, 374
SOX (Sarbanes-Oxley Act), 373
spam bots, 282
spam malware, 457
Splunk

dashboard
centralized data management, 144–145
Hadoop, 311–312
SIEM troubleshooting, 291–300
Phantom
case management, 562–563
DevOps usage example, 564–566
playbooks, 563–564
SOAR example, 561–562
SSH Bruteforce logs, 412–413
stack counting, threat hunting, 459
standards
compliance/ risk reduction, 340–350
security, 19–20
state-sponsored actors, 6–7
static analysis
analysis services, 446–447
advanced static analysis, 448–451
disassemblers, 199–200
packed files, 197–199
Pframe, 448
reverse engineering files, 199–200
WannaCry kill switch malware analysis, 451–452
forensic dynamic analysis, 480–482
forensic static analysis, 478–479
stealth strategies, tactical threat intelligence, 395
STIX, processing technical threat intelligence data, 408–409
storage

data retention and, 143–144
facility design, 104
SOC development
data retention and, 143–144
throughput and, 141–144
throughput and, 141–144
strategic data, 262
strategic threat intelligence, 205, 382, 383
data expectations, 393
processing data, 400–402
strike packs, 18–19
structures of data, 263
semi-structured data, 263
structured data, 263
unstructured data, 263
surveillance (video), computer rooms, 113
Surveyor, penetration testing, 185
syslog, 275
system order, eradication phase (incident response), 463
systems analysts, 222–224

T

tabletop exercises, policies, 334–335
example of, 337–340
executing, 336–337
format of, 337–338

options, 335
tactical data, 262
tactical threat intelligence, 205, 382–384
attack vectors, 394–395
data expectations, 394–396
infrastructures, 395
stealth strategies, 395
tools, 395
task assignments to incident response playbooks, 427–430
TAXII, processing technical threat intelligence data, 409–411
technical threat intelligence, 206, 382, 385
Abuse.ch Feodo Tracker, 412
AlienVault OTX, 412–413
Blocklist.de, 412
CINS Score, 412
Cyber Threat System from FortiGuard Labs, 413
Dan.me.uk, 412
data expectations, 397–398
Emerging Threats Rule Server, 412
FBI InfraGard, 412
IBM X-Force Exchange, 413
processing data, 407
CSV, 411
JSON, 407–408
OpenIOC, 408
Regex, 411
STIX, 408–409
TAXII, 409–411

XML, 407
SSH Bruteforce logs, 412–413
technology
domains, 35
meta-features, 31
planning, SOC development, 97–98
securing SOC technology, 158–159
Three Pillars of Foundational SOC Support Services, The, 158–159
temperature/humidity, computer rooms, 108–109
Tenable.sc vulnerability tracking, 177
testing, threat intelligence, 392
text-file formats, DevOps, 584–585
third-generation SOC, 52
threat actors, 4–5, 6–7
cyberterrorists, 7
hacktivists, 5–6
insider threats, 7
motivations of, 7
threat hunting, incident response, 424, 455–456
consortium playbooks, 196
example of, 460–462
grouping, 455–456
incidents, defining, 425
lifecycle of, 425–426
containment, eradication and recovery phase, 426–438
detection and analysis phase, 438–454
post-incident activity phase, 484–492
preparation phase, 426–438

maturity models, 460–462
performing, 459–460
planning, 194
SOC job roles, 221–222
stack counting, 459
techniques, 458–459
threat intelligence, 205, 262, 378–379
actionable intelligence, 378, 392
flowcharts, 414
processing data, 414
categories of, 382–385
checklists, 389–390
collecting/processing, 399–400
operational threat intelligence data, 402–407
strategic threat intelligence data, 400–402
content quality, 390
checklists, 390–391
key factors, 390
context, 379, 385–388
evaluating, 388–389
external threat intelligence, 385–386
feedback, 421–422
internal threat intelligence, 385–386
IOC, 382
nontechnical intelligence. *See* strategic threat intelligence
operational threat intelligence, 205, 382, 384–385
data expectations, 396–397
processing data, 402–407

overview, 379
planning, 393–398
prevalence, 387
ROI, 421–422
security tools, 414–416
email security, 420–421
SIEM, 416–419
strategic threat intelligence, 205, 382–383
data expectations, 393
processing data, 400–402
tactical threat intelligence, 205, 382–384
attack vectors, 394–395
data expectations, 394–396
infrastructures, 395
stealth strategies, 395
tools, 395
technical threat intelligence, 206, 382, 385
data expectations, 397–398
processing data, 407–413
testing, 392
threat data, 380
example of, 380
limitations, 381–382
value of, 380–381
threat models, 25
ATT&CK Model, 35–38
chaining together attack behaviors, 38
PRE-ATT&CK research, 36–37

using, 38
choosing, 38–39
Cyber Kill Chain model, 25–29
Diamond Model, 30–31
attack graphs, 34–35
Diamond Model for Incident Management, 32–33
Extended Diamond Model, 31
social-political meta-features, 31
technology meta-features, 31
threats
assessments, 355
data, 380
example of, 380
limitations, 381–382
value of, 380–381
detection tools, vulnerability, 524–525
feeds, big data, 312
future of, 671–673
mapping data, 270
ranking, 58–59
response to future threats, 673
zero-day threats, 7
Three Pillars of Foundational SOC Support Services, The, 154–155
evaluating, 159
people, 156–157
technology, 158–159
throughput, 120–121
requirements, 121–123

SOC development, 141–144
storage and, 141–144
ticketing systems, incident response, 435–436
tools
collaboration, SOC development, 138–140
limited tools, challenges for services, 153
tracking vulnerabilities, 179
training, 640
case study, 643–644
challenges of, 640–641
DevOps, 650
employees, 253–255
free training, 644
future of
on-demand experts, 649
universal language/language translation, 649
learning
on-demand learning, 646–648
gamifying, 644–645
LMS, 645
personalized learning, 646–648
ML, 655
today's training, 641–643
TrIDNET analysis service, 197
troubleshooting SIEM, 287, 291
actionable intelligence, 300–301
data input, 288, 293–299
data processing, 289–291

data storage, 291–293
validating results, 299–300
tuning SIEM, 285–287
types of data
booleans, 265
bytes, 264
chars, 265
doubles, 264
floats, 264
int, 264
longs, 264
primitive data types, 263–265
shorts, 264

U

unauthenticated scanning, 86
unknown environment penetration testing, 367
unstructured data, 263
upgrades (automated), SASE, 630
UPS, computer rooms, 110–111

V

Verizon 2020 Data Breach Investigations Report, 189–190
video surveillance, computer rooms, 113
video walls, facility design, 104–105

virtualized computers, 638–639
viruses, 457
VirusTotal, 14
VoIP (Voice over IP), 617–618
volatile data, digital forensics, 480–482
VPN (Virtual Private Networks), 534
SASE, 628–629
SOC development, 91–92
vulnerabilities, 39–40
active vulnerability scanning, 86–87
assessments, 355–356, 505
authenticated scanning, 86
context, 266
CVSS, 86
evaluating, SOC development
active vulnerability scanning, 86–87
CVSS, 86
passive vulnerability scanning, 87–88
Nexpose vulnerability scanner, 86–87
passive vulnerability scanning, 87–88
SANS vulnerability management, best practices, 12
scanning, 12, 176–177
tracking, 179
unauthenticated scanning, 86
vulnerability management, 498–499
accuracy, 540–541
assessments, 505
assets

access, 535
inventory phase, 500–502
management, 522
best practices, 499–500
Certero dashboard, 522
compliance tools, 522
CVSS, 507–508
CVSS v2, 508–512
CVSS v3, 508–512
cyber insurance, 544–547
deployment example, 535
evaluation procedures, 528–529
asset collection, 529–532
choosing corrective actions, 539
launch scanning, 537–539
planning, 532–537
prioritizing assets, 536
exceptions, 552–553
exploitation tools, 520–521
host scanning, 516
information management phase, 502–503
management services, 45, 150, 175
best practices, 175–176
job roles, 239
OpenVAS, 178
penetration testing, 179–187
roles, 527–528
Tenable.sc vulnerability tracking, 177

vulnerability tracking, 179, 525
measuring vulnerabilities, 506
NAC, 522–524
automated NAC, 501
manual NAC, 501
network scanners, 501–502, 515
Nmap
fingerprinting, 503
scanning, 501–502
patching systems, 547–549
planning, 532–537
process summary, 554–555
program diagrams, 527–528
remediation approval, 550–551
report and remediate phase, 505
reporting, 552
respond and repeat phase, 506
responses, 540, 542–544
risk assessment phase, 504
scanning, 515–520
active scanning, 515–516
assessments, 360–361
Firepower, 87–88, 306–307
passive scanning, 516–517
services, 525–527
shorthand, 511–512
Struts vulnerability example, 507
CVSS v2, 512

CVSS v3, 513–514
temporal/environmental metrics, 511
threat detection tools, 524–525
tracking, 179, 525

W

WAN (Wide-Area Networks), 618–620. *See* also SD-WAN
WannaCry kill switch malware analysis, 451–452
war rooms, facility design, 106
waterfall model (modified), processing threat intelligence, 400–402
WBDG (Whole Building Design Guide), SOC facility design, 101–102
Webex Teams, ChatOps, 595
wikis, manual DevOps, 593–594
Windows event logs, 277
work environments
planning, 155–156
Three Pillars of Foundational SOC Support Services, The, 155–156
workflows, playbooks
examples, 579–582
sample workflow, 570–571
symbols, 570
workplaces, competitive, 252
workstations, facility design, 105
worms, 457

X

XDR (Cross-layered Detection and Response), 559–560
XML (Extensible Markup Language)
DevOps, 585–586
processing technical threat intelligence data, 407

Y

YAML, DevOps, 586–589
YANG serializers, DevOps, 589–590

Z

Zenmap, NAC, 523–524
zero-day threats, 7

VIDEO TRAINING FOR THE **IT PROFESSIONAL**

LEARN QUICKLY
Learn a new technology in just hours. Video training can teach more in less time, and material is generally easier to absorb and remember.

WATCH AND LEARN
Instructors demonstrate concepts so you see technology in action.

TEST YOURSELF
Our Complete Video Courses offer self-assessment quizzes throughout.

CONVENIENT
Most videos are streaming with an option to download lessons for offline viewing.

Learn more, browse our store, and watch free, sample lessons at
informit.com/video

Save 50%* off the list price of video courses with discount code **VIDBOB**

Pearson

informIT®
the trusted technology learning source

*Discount code VIDBOB confers a 50% discount off the list price of eligible titles purchased on informit.com. Eligible titles include most full-course video titles. Book + eBook bundles, book/eBook + video bundles, individual video lessons, Rough Cuts, Safari Books Online, non-discountable titles, titles on promotion with our retail partners, and any title featured as eBook Deal of the Day or Video Deal of the Week is not eligible for discount. Discount may not be combined with any other offer and is not redeemable for cash. Offer subject to change.

Code Snippets

Many titles include programming code or configuration examples. To optimize the presentation of these elements, view the eBook in single-column, landscape mode and adjust the font size to the smallest setting. In addition to presenting code and configurations in the reflowable text format, we have included images of the code that mimic the presentation found in the print book; therefore, where the reflowable format may compromise the presentation of the code listing, you will see a “Click here to view code image” link. Click the link to view the print-fidelity code image. To return to the previous page viewed, click the Back button on your device or app.

```
# Options to syslogd
# -m 0 disables 'MARK' messages.
# -r enables logging from remote machines
# -x disables DNS lookups on messages received with -r
# See syslogd(8) for more details
SYSLOGD_OPTIONS="-m 0"
# Options to klogd
# -2 prints all kernel oops messages twice; once for klogd to decode, and
#    once for processing with 'ksymoops'
# -x disables all klogd processing of oops messages entirely
# See klogd(8) for more details
KLOGD_OPTIONS="-x"
#
SYSLOG_UMASK=077
# set this to a umask value to use for all log files as in umask(1).
# By default, all permissions are removed for "group" and "other".
```

```
{ "timestamp": "2020-05-24 22:19:41", "id": 0, "class": "audit", "event":
"startup", ...
{ "timestamp": "2020-05-24 22:19:45", "id": 0, "class": "connection", "event":
"connect", ...
{ "timestamp": "2020-05-24 22:19:45", "id": 1, "class": "connection", "event":
"disconnect", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 0, "class": "connection", "event":
"connect", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 1, "class": "general", "event":
"status", ...
{ "timestamp": "2020-05-24 22:19:47", "id": 2, "class": "connection", "event":
"disconnect", ...
{ "timestamp": "2020-05-24 22:19:53", "id": 0, "class": "audit", "event":
"shutdown", ...
```

```
#Version: 1.0
#Date: 20-Jan-2020 00:00:00
#Fields: time cs-method cs-uri
00:34:23 GET /foo/bar.html
11:21:16 GET /foo/bar.html
11:45:52 GET /foo/bar.html
11:58:34 GET /foo/bar.htm
```

```
192.168.2.1 user-identifier joey [15/Oct/2019:21:54:12 -0700] "GET /apache_pb.gif
HTTP/1.0" 200 4782
```

```
<?xml version="1.0" standalone="yes"?>
<conversation>
<greeting>Hello, world!</greeting>
<response>Stop the planet, I want to get off!</response>
</conversation>
```

{Date:03/01/2020, DayofTheWeek:Tuesday, Event:Phishing}

```
{
   "type": "campaign",
   "id": "campaign--8e2e2d2b-17d4-4cbf-938f-98ee46b3cd3f",
   "spec_version": "2.1",
   "created": "2016-04-06T20:03:00.000Z",
   "modified": "2016-04-06T20:03:23.000Z",
   "name": "Green Group Attacks Against Finance",
   "description": "Campaign by Green Group against targets in the
      financial services sector."
}
```

2020-01-03,03.06.41,Ad-Aware,trojan.genericKD.3164632
2020-01-03,02:23:26,AegisLab,Uds.Dangerousobject

```
<? xml version="1.0" encoding="UTF-8" ?>
   <books>
      <book bookID="1">
      <lead>Joey Muniz</lead>
      <title>The Modern Security Operations Center</title>
      <genre>Technology</genre>
      <desc>A technical journey to understanding a SOC</desc>
   </book>
</books>
```

```
{
   "books":{
      "book":[
      "author":"Joey Muniz",
      "title":"The Modern Security Operations Center",
      "genre":"Technology",
      "desc":"A technical journey to understanding a SOC",
      "_bookID":"1"
      ]
   }
}
```

```
---
- books:
      book:
- author:Joey Muniz
- title:The Modern Security Operations Center
- genre:Technology
- desc:A technical journey to understanding a SOC
- bookID:1
```

```
#This program will divide one number by another
- divider:
    do:
      divide:
        - dividend: ${input1}
        - divisor: ${input2}
    publish:
      - answer: ${quotient} #Answer represents the result of the division
    navigate: [{ILLEGAL: FAILURE}, {SUCCESS: printer}]
```

```
---
- name: Update web servers
  hosts: webservers
  remote_user: root

  tasks:
  - name: Ensure apache is at the latest version
    ansible.builtin.yum:
      name: httpd
      state: latest
  - name: Write the apache config file
    ansible.builtin.template:
      src: /srv/httpd.j2
      dest: /etc/httpd.conf

- name: Update db servers
  hosts: databases
  remote_user: root

  tasks:
  - name: Ensure postgresql is at the latest version
    ansible.builtin.yum:
      name: postgresql
      state: latest

- name: Ensure that postgresql is started
  ansible.builtin.service:
    name: postgresql
    state: started
```

```
sudo app-get update -y && sudo apt-get upgrade -y
ansible -version
sudo apt-get install software-properties-common
sudo apt-add-repository ppa:ansible/ansible
sudo apt-get install ansible
ansible -version
```

```
---
- hosts: Router1
  gather_facts: true
  connection: local
  tasks:
   - name: show run
   ios_command:
   commands:
     - show run
     host: "{{ ansible_host }}"
     username: admin
     password: cisco
     register: config

  - name: save output to /etc/ansible/router_backups/
  copy:
  content: "{{ config.stdout[0] }}"
  dest: " /etc/ansible/router_backups/ show_run_{{ inventory_hostname }}.txt"
```

Contents

Cover Page
About This eBook
Title Page
Copyright Page
Dedication
Table of Contents
Preface

1. Vision
2. Who Should Read This Book?
3. How This Book Is Organized
4. Book Structure

We Want to Hear from You!
Reader Services
Acknowledgments
About the Author
Figure Credits
Chapter 1. Introducing Security Operations and the SOC

1. Introducing the SOC
2. Factors Leading to a Dysfunctional SOC
3. Cyberthreats
4. Investing in Security
5. The Impact of a Breach
6. Establishing a Baseline
7. Fundamental Security Capabilities
8. Standards, Guidelines, and Frameworks
9. Industry Threat Models
10. Vulnerabilities and Risk
11. Business Challenges
12. In-House vs. Outsourcing
13. SOC Services

14. SOC Maturity Models
15. SOC Goals Assessment
16. SOC Capabilities Assessment
17. SOC Development Milestones
18. Summary
19. References

Chapter 2. Developing a Security Operations Center

1. Mission Statement and Scope Statement
2. Developing a SOC
3. SOC Procedures
4. Security Tools
5. Planning a SOC
6. Designing a SOC Facility
7. Network Considerations
8. Disaster Recovery
9. Security Considerations
10. Internal Security Tools
11. Guidelines and Recommendations for Securing Your SOC Network
12. SOC Tools
13. Summary
14. References

Chapter 3. SOC Services

1. Fundamental SOC Services
2. The Three Pillars of Foundational SOC Support Services
3. SOC Service Areas
4. SOC Service Job Goals
5. Service Maturity: If You Build It, They Will Come
6. SOC Service 1: Risk Management
7. SOC Service 2: Vulnerability Management
8. SOC Service 3: Compliance
9. SOC Service 4: Incident Management
10. SOC Service 5: Analysis
11. SOC Service 6: Digital Forensics

12. SOC Service 7: Situational and Security Awareness
13. SOC Service 8: Research and Development
14. Summary
15. References

Chapter 4. People and Process

1. Career vs. Job
2. Developing Job Roles
3. SOC Job Roles
4. NICE Cybersecurity Workforce Framework
5. Role Tiers
6. SOC Services and Associated Job Roles
7. Soft Skills
8. Security Clearance Requirements
9. Pre-Interviewing
10. Interviewing
11. Onboarding Employees
12. Managing People
13. Job Retention
14. Training
15. Certifications
16. Company Culture
17. Summary
18. References

Chapter 5. Centralizing Data

1. Data in the SOC
2. Data-Focused Assessment
3. Logs
4. Security Information and Event Management
5. Troubleshooting SIEM Logging
6. APIs
7. Big Data
8. Machine Learning
9. Summary
10. References

Chapter 6. Reducing Risk and Exceeding Compliance

1. Why Exceeding Compliance
2. Policies
3. Launching a New Policy
4. Policy Enforcement
5. Procedures
6. Tabletop Exercise
7. Standards, Guidelines, and Frameworks
8. Audits
9. Assessments
10. Penetration Test
11. Industry Compliance
12. Summary
13. References

Chapter 7. Threat Intelligence

1. Threat Intelligence Overview
2. Threat Intelligence Categories
3. Threat Intelligence Context
4. Evaluating Threat Intelligence
5. Planning a Threat Intelligence Project
6. Collecting and Processing Intelligence
7. Actionable Intelligence
8. Feedback
9. Summary
10. References

Chapter 8. Threat Hunting and Incident Response

1. Security Incidents
2. Incident Response Lifecycle
3. Phase 1: Preparation
4. Phase 2: Detection and Analysis
5. Phase 3: Containment, Eradication, and Recovery
6. Digital Forensics
7. Phase 4: Post-Incident Activity

8. Incident Response Guidelines
9. Summary
10. References

Chapter 9. Vulnerability Management

1. Vulnerability Management
2. Measuring Vulnerabilities
3. Vulnerability Technology
4. Vulnerability Management Service
5. Vulnerability Response
6. Vulnerability Management Process Summarized
7. Summary
8. References

Chapter 10. Data Orchestration

1. Introduction to Data Orchestration
2. Security Orchestration, Automation, and Response
3. Endpoint Detection and Response
4. Playbooks
5. Automation
6. DevOps Programming
7. DevOps Tools
8. Blueprinting with Osquery
9. Network Programmability
10. Cloud Programmability
11. Summary
12. References

Chapter 11. Future of the SOC

1. All Eyes on SD-WAN and SASE
2. IT Services Provided by the SOC
3. Future of Training
4. Full Automation with Machine Learning
5. Future of Your SOC: Bringing It All Together
6. Summary

7. References

Index

i
ii
iii
iv
v
vi
vii
viii
ix
x
xi
xii
xiii
xiv
xv
xvi
xvii
xviii
xix
xx
xxi
xxii
xxiii
xxiv
xxv
xxvi
xxvii
xxviii
xxix
xxx
xxxi
xxxii
xxxiii
xxxiv

xxxv
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37

38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74

75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111

112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148

149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185

186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222

223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259

260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296

297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333

334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370

371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407

408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444

445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481

482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518

519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555

556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592

593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629

630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666

667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703

704
705
706
707
708
709
710
711
712
713
714
715
716
717
718