

About ISACA®

With more than 100,000 constituents in 180 countries, ISACA (www.isaca.org) is a leading global provider of knowledge, certifications, community, advocacy and education on information systems (IS) assurance and security, enterprise governance and management of IT, and IT-related risk and compliance. Founded in 1969, the nonprofit, independent ISACA hosts international conferences, publishes the *ISACA® Journal*, and develops international IS auditing and control standards, which help its constituents ensure trust in, and value from, information systems. It also advances and attests IT skills and knowledge through the globally respected Certified Information Systems Auditor® (CISA®), Certified Information Security Manager® (CISM®), Certified in the Governance of Enterprise IT® (CGEIT®) and Certified in Risk and Information Systems Control™ (CRISC™) designations.

ISACA continually updates and expands the practical guidance and product family based on the COBIT® framework. COBIT helps IT professionals and enterprise leaders fulfill their IT governance and management responsibilities, particularly in the areas of assurance, security, risk and control, and deliver value to the business.

Disclaimer

ISACA has designed and created *Responding to Targeted Cyberattacks* (the “Work”) primarily as an educational resource for security, governance and assurance professionals. ISACA makes no claim that use of any of the Work will assure a successful outcome. The Work should not be considered inclusive of all proper information, procedures and tests or exclusive of other information, procedures and tests that are reasonably directed to obtaining the same results. In determining the propriety of any specific information, procedure or test, security governance and assurance professionals should apply their own professional judgment to the specific circumstances presented by the particular systems or information technology environment.

Reservation of Rights

© 2013 ISACA. All rights reserved. No part of this publication may be used, copied, reproduced, modified, distributed, displayed, stored in a retrieval system or transmitted in any form by any means (electronic, mechanical, photocopying, recording or otherwise) without the prior written authorization of ISACA. Reproduction and use of all or portions of this publication are permitted solely for academic, internal and noncommercial use and for consulting/advisory engagements, and must include full attribution of the material’s source. No other right or permission is granted with respect to this work.

ISACA

3701 Algonquin Road, Suite 1010
Rolling Meadows, IL 60008 USA
Phone: +1.847.253.1545
Fax: +1.847.253.1443
Email: info@isaca.org
Web site: www.isaca.org

Provide Feedback: www.isaca.org/Cyberattacks

Participate in the ISACA Knowledge Center: www.isaca.org/knowledge-center

Follow ISACA on Twitter: <https://twitter.com/ISACANews>

Join ISACA on LinkedIn: ISACA (Official), <http://linkd.in/ISACAOOfficial>

Like ISACA on Facebook: www.facebook.com/ISACAHQ

ISBN: 978-1-60420-337-0

Responding to Targeted Cyberattacks

Acknowledgments

Development Team:

Chris Crevits, CISSP, CCNA, Ernst & Young LLP, USA
 Jose Granado, CISSP, America's Practice Leader, Information Security Services,
 Ernst & Young LLP, USA
 James O. Holley, CCE, NSA ISSO, NSA ISSP, Ernst & Young LLP, USA
 Patrick J. Hynes, CISA, CISM, CCNA, CISSP, NSAIAM, Ernst & Young LLP, USA
 David C. Kovar, CCE, CISSP, EnCE, GCFA, GCIH, GREM, Ernst & Young LLP, USA
 Mark G. Manglicmot, CEH, CISSP, GCIH, Security+, Ernst & Young LLP, USA
 Austin J. Murphy, GCFA, Security+, Ernst & Young LLP, USA
 Daniel J. Quealy, CAMS, CIWM, GSEC, HTCIA (past president), NSAIAM, Ernst & Young LLP, USA
 Joshua Theimer, CISA, CISSP, Security+, Ernst & Young LLP, USA
 Bryan C. York, CEH, CISSP, Linux+, Security+, Ernst & Young LLP, USA

Subject Matter Expert Reviewers:

Vilius Benetis, Ph.D., CISA, CRISC, BAIP, Lithuania
 Jeimy Cano, Ph.D., CFE, CMAS, Ecopetrol, Colombia
 Christos K. Dimitriadis, Ph.D., CISA, CISM, CRISC, INTRALOT S.A., Greece
 Jo Stewart-Rattray, CISA, CISM, CGEIT, CRISC, CSEPS, BRM Holdich, Australia, Director
 Rolf M. von Roessing, CISA, CISM, CGEIT, CISSP, FBCI, FORFA AG, Switzerland

ISACA Board of Directors

Gregory T. Grocholski, CISA, The Dow Chemical Co., USA, International President
 Allan Boardman, CISA, CISM, CGEIT, CRISC, ACA, CA (SA), CISSP, Morgan Stanley, UK,
 Vice President
 Juan Luis Carselle, CISA, CGEIT, CRISC, Wal-Mart, Mexico, Vice President
 Christos K. Dimitriadis, Ph.D., CISA, CISM, CRISC, INTRALOT S.A., Greece, Vice President
 Ramses Gallego, CISM, CGEIT, CCSK, CISSP, SCPM, Six Sigma Black Belt, Dell, Spain,
 Vice President
 Tony Hayes, CGEIT, AFCHSE, CHE, FACS, FCPA, FIIA, Queensland Government, Australia,
 Vice President
 Jeff Spivey, CRISC, CPP, PSP, Security Risk Management Inc., USA, Vice President
 Marc Vael, Ph.D., CISA, CISM, CGEIT, CRISC, CISSP, Valuendo, Belgium, Vice President
 Kenneth L. Vander Wal, CISA, CPA, Ernst & Young LLP (retired), USA,
 Past International President
 Emil D'Angelo, CISA, CISM, Bank of Tokyo-Mitsubishi UFJ Ltd., (retired), USA,
 Past International President
 John Ho Chi, CISA, CISM, CRISC, CBCP, CFE, Ernst & Young LLP, Singapore, Director
 Krysten McCabe, CISA, The Home Depot, USA, Director
 Jo Stewart-Rattray, CISA, CISM, CGEIT, CRISC, CSEPS, BRM Holdich, Australia, Director

Knowledge Board

Marc Vael, Ph.D., CISA, CISM, CGEIT, CRISC, CISSP, Valuendo, Belgium, Chairman
 Rosemary M. Amato, CISA, CMA, CPA, Deloitte Touche Tohmatsu Ltd., The Netherlands
 Steven A. Babb, CGEIT, CRISC, Betfair, UK
 Thomas E. Borton, CISA, CISM, CRISC, CISSP, Cost Plus, USA
 Phil J. Lageschulte, CGEIT, CPA, KPMG LLP, USA
 Jamie Pasfield, CGEIT, ITIL V3, MSP, PRINCE2, Pfizer, UK
 Salomon Rico, CISA, CISM, CGEIT, Deloitte, Mexico

Acknowledgments *(cont.)*

Guidance and Practices Committee

Phil J. Lageschulte, CGEIT, CPA, KPMG LLP, USA, Chairman
Dan Haley, CISA, CGEIT, CRISC, MCP, Johnson & Johnson, USA
Yves Marcel Le Roux, CISM, CISSP, CA Technologies, France
Aureo Monteiro Tavares Da Silva, CISM, CGEIT, Brazil
Jotham Nyamari, CISA, Deloitte, USA
Connie Lynn Spinelli, CISA, CRISC, CFE, CGMA, CIA, CISSP, CMA, CPA, BKD LLP, USA
Siang Jun Julia Yeo, CISA, CPA (Australia), MasterCard Asia/Pacific Pte. Ltd., Singapore
Nikolaos Zacharopoulos, CISA, CISSP, DeutschePost–DHL, Germany

ISACA and IT Governance Institute® (ITGI®) Affiliates and Sponsors

Information Security Forum
Institute of Management Accountants Inc.
ISACA chapters
ITGI France
ITGI Japan
Norwich University
Socitum Performance Management Group
Solvay Brussels School of Economics and Management
Strategic Technology Management Institute (STMI) of the National University of Singapore
University of Antwerp Management School

ASIS International
Hewlett-Packard
IBM
Symantec Corp.

ISACA thanks Ernst & Young for its generous donation of services to develop this publication.



Table of Contents

Chapter 1. Introduction	9
1.1 A Case Study of the Need for Change	9
1.2 Evolution of the Threat Landscape	10
1.3 Adaptive Attack Vectors	13
1.4 A Watershed Event	15
1.5 The APT Life Cycle	16
1.6 What Are Others Doing?	20
1.7 Summary	21
Chapter 2. Preparation	23
2.1 Build a Team, Make a Plan	23
2.2 Establish Key Relationships	23
2.2.1 External Relationships	23
2.2.2 Internal Relationships	24
2.3 Determine Authorities	24
2.4 Inventory Existing Technologies	25
2.5 Standardize the Investigation Process	26
2.6 Training and Governance	28
2.6.1 Exercises	28
2.6.2 Security Program and Response Plan Reviews	29
2.7 Establish Critical Capabilities	29
2.7.1 Host-level Activity Awareness	31
2.7.2 Network-level Activity Awareness	33
2.7.3 Search	35
2.7.4 Computer Forensic Analysis	36
2.7.5 Malware Analysis	37
2.7.6 Threat Intelligence	37
2.7.7 Vulnerability Identification	41
Chapter 3. Investigation	43
3.1 Conducting a Security Breach Investigation	43
3.2 Who Attacked Us?	47
3.3 What Was Targeted?	47
3.4 When Did Various Events Occur?	48
3.5 From Where Did the Attacks Come?	48
3.6 Why Did They Attack?	49
3.7 How Did They Get In, Stay In and Get the Data Out?	50
3.8 Other Important Areas to Consider	50
3.9 On the Quality of Intelligence	51

3.10	Evidence Handling	51
3.10.1	Preservation and Collection Memorandum	52
3.10.2	Chain of Custody	52
3.10.3	MD5 Hashing	53
3.10.4	Write Blockers	53
3.10.5	Reconcile Record Counts	53
3.10.6	Attorney-client Privilege or Attorney Work Product Privilege	54
3.10.7	Insurance Claims	54
3.11	Investigating Anonymously	54
3.12	Safeguarding the Investigative Actions	55
3.12.1	Data	55
3.12.2	Data in Motion	55
3.13	Protecting the Investigation	55
3.13.1	Credential Protection	56
Chapter 4.	Eradication	57
4.1	Plan for Eradication	57
4.1.1	Create the Eradication Event Team	57
4.1.2	Develop the Eradication Event Plan	58
4.1.3	Determine the Eradication Event Date	59
4.1.4	Know the Attacker’s Techniques, Tactics and Procedures	59
4.1.5	Establish Communication Protocols	60
4.1.6	Establish a “War Room”	62
4.1.7	Establish Secure Communications and Information Sharing Mechanism(s)	62
4.2	Execute the Plan	63
4.2.1	Execute a Password Change	64
4.2.2	Block Attacker Command and Control	66
4.2.3	Rebuild Compromised Systems	67
4.2.4	Submit Malware to Antivirus Vendors	68
4.3	Monitor for Attempted Reentry	68
Chapter 5.	Post-eradication	71
5.1	Validate Eradication Activities	71
5.1.1	Maintain a Heightened State of Alert	71
5.1.2	Validate Controls	73
5.2	Brief Stakeholders	74
5.3	Lessons Learned	75
5.4	Strategic Change—Cybersecurity Transformation	77

Chapter 6. Conclusion	81
Appendix A. Other Questions the Investigation Team Will Address	83
Appendix B. Investigative Tools	87
List of Figures	
Figure 01—Ten Assessment Scenarios	10
Figure 02—Evolution of the Threat Landscape	11
Figure 03—Adaptive Attack Vectors	14
Figure 04—The APT Life Cycle.....	16
Figure 05—The APT Life Cycle: Another View	16
Figure 06—APT <i>Modus Operandi</i>	19
Figure 07—What Are Others Doing?	20
Figure 08—Basic Incident Triage Process.....	27
Figure 09—Alert Decision Process	28
Figure 10—CSIRT Capability Requirements	31
Figure 11—Potential Impact of Threat Intelligence on the Attack Life Cycle	38
Figure 12—The Incident Response Process	44
Figure 13—The Incident Response Life Cycle From NIST SP 800-61	45
Figure 14— <i>xkcd Comic</i> : “Password Strength”	66
Figure 15—War Room After-action Report Template	76

Page intentionally left blank

Chapter 1. Introduction

Never tell people how to do things. Tell them what to do, and they will surprise you with their ingenuity.

War As I Knew It (1947) by General George S. Patton

1.1 A Case Study of the Need for Change

Imagine that you are the chief information officer (CIO) of a Fortune 100 company and you are preparing to give the following briefing to both the chief executive officer (CEO) and the board of directors:


It took the attackers only six minutes to circumvent the perimeter defenses. From there, they achieved domain administrator privileges in less than 12 hours. In less than a week they fully compromised all 30 of our global domains. They harvested more than 200,000 credentials, giving them the ability to log in to the network masquerading as any of us—they could even change our investment elections in our 401(k)s or transfer money out. There was no place on our global network they could not go and only a handful of computers they did not have easy access to—only 10 percent of our manufacturing facilities are behind firewalls, segregating them from our network. The attackers were in a position to electronically transfer millions of dollars out of our bank accounts through our accounts payable system. Their tools did not set off any alarms—our antivirus software did not trigger any alerts. They had direct access to our manufacturing environment and could affect both the quality of our production processes and safety on our shop floors. They had access to our most sensitive intellectual property, including our past, current and future plans for major acquisitions and divestitures as well as the results of the billions of dollars we have invested in a decade of research and development. And, in the end, they were able to steal all the data. We were not able to stop them, or even see them in our network.

That is a chilling story! But that is the story the CIO at one of the largest clients of Ernst and Young (EY) relayed to the CEO and the board of directors after the EY attack and penetration team finished a complex cybersecurity posture assessment.

The engagement encompassed the design of a broad series of assessments—not intended to measure compliance with policy or conformance to leading practice/technology controls, but rather to focus on assessing the company’s cybersecurity posture in the face of attacks other Fortune 100 companies had recently suffered. Essentially, the questions were, “What if someone does this to us? How would we fare? How are we positioned to make an attacker’s tasks difficult, to detect that an attack scenario is underway, and to respond to attacks we detect?”

Ten assessment scenarios¹ were developed and executed, using a wide range of attackers with differing motives and differing accesses to their client’s technology infrastructure (**figure 1**).

FIGURE 01 Ten Assessment Scenarios



The results, as described by the CIO to the CEO and the board, reflected a current and urgent need to fundamentally alter the way the enterprise approaches cybersecurity. This need is driven primarily by one key fact: The threats that enterprises face by being connected to the Internet are evolving at a much faster pace than the information security architectures, technologies and processes they have deployed.

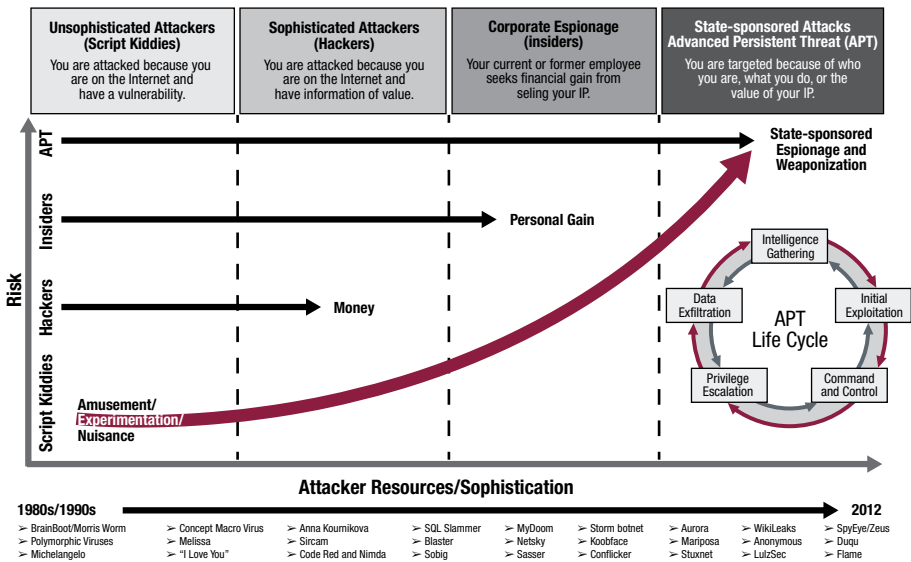
The threats that enterprises face by being connected to the Internet are evolving at a much faster pace than the information security architectures, technologies and processes they have deployed.

1.2 Evolution of the Threat Landscape

The progression from experimentation, harassment, monetization, industrial espionage and weaponization threats is illustrated in **figure 2**.

¹ The scenarios were: 1) a malicious insider gains access to trade secrets, 2) a distributed denial-of-service (DDoS) attack is made against a key revenue-generating web application, 3) a primary web site is defaced by hacktivists, 4) an organized crime targets treasury workstations for illicit wire transfer, 5) a disgruntled employee modifies an automated equipment control system at a remote site, 6) a shop floor Programmable Logic Controller is modified or shut down, resulting in a production stoppage or quality control issues, 7) an insider accesses pre-release financial data, 8) an Internet-facing web application is attacked, allowing access to customer personally identifiable information (PII) in a back-end database, 9) a focused cyberespionage attack is made by an advanced persistent threat (APT) to steal critical intellectual property (IP), and 10) an attack is made against the corporate intranet and facilitated by a joint venture connection.

FIGURE 02 Evolution of the Threat Landscape



In November 1983, when Fred Cohen, then a graduate student at the University of Southern California, USA, conducted experiments to demonstrate the security risk of self-replicating computer code, no one had yet heard of the term “computer virus.” Len Adelman (Cohen’s academic advisor and the “A” in RSA) had only recently coined the term.² Thirty years later, hundreds of millions of people know what computer viruses are and hundreds of millions of computers have been infected with them. While Cohen conducted his virus research on a mainframe virtual address eXtension (VAX) machine running a version of UNIX®, it was not long before the concepts of self-replicating code that he demonstrated were also implemented in the PC world.

The Brain—a boot sector virus designed to infect floppy disks—was the first IBM® PC-compatible virus.³ It appeared in January 1986 and was followed by the Morris worm in November 1988.⁴ By 1990, virus developers integrated polymorphism⁵ techniques into their code as a means to evade nascent antivirus technologies.

² Cohen, Fred, “Experiments With Computer Viruses,” 1984, <http://all.net/books/virus/part5.html>

³ Leyden, John, “PC Virus Celebrates 20th Birthday,” *The Register*, UK, 19 January 2006,

http://www.theregister.co.uk/2006/01/19/pc_virus_at_20/

⁴ The Morris worm, or Internet worm of 1988, was one of the first computer worms distributed via the Internet. It is considered the first worm and was certainly the first to gain significant mainstream media attention. It also resulted in the first conviction in the United States under the 1986 Computer Fraud and Abuse Act. It was written by a student at Cornell University, Robert Tappan Morris, and launched on 2 November 1988 from MIT. Kehoe, Brendan P.; Zen and the Art of the Internet: A Beginner’s Guide to the Internet, Prentice Hall, USA, 1992, <http://groups.csail.mit.edu/mac/classes/6.805/articles/morris-worm.html>

⁵ Polymorphic code is code that uses a polymorphic engine to mutate while keeping the original algorithm intact. The code changes itself each time it runs, but the function of the code (its semantics) will not change. Virus writers adapted polymorphic code methods in an effort to evade signature-based antivirus detection.

In the late 1980s and early 1990s, with relatively simple functional applications running on PCs that were often turned off when not in use, the real power of the emerging Internet was centered on servers. Major exploits of the day typically targeted operating system (OS) vulnerabilities, and hackers' targets were generally educational institutions, government and military systems, and corporate data centers. Hackers required deep technical skills, knowledge of multiple programming languages, knowledge of server OS internal technical functionality and networking protocol stacks. Hackers generally worked to discover vulnerabilities on their own and developed their own exploits. There was great disdain for "script kiddies" who were perceived as not having great technical skills and used vulnerabilities, exploits and scripts developed by others to accomplish their goals. Hackers also had few resources to learn new techniques from other hackers.

Each attack typically compromised only a few systems because so few systems were connected to the Internet. But even at this very early stage of the battle in cyberspace, governments targeted other governments using technical espionage. Clifford Stoll's *The Cuckoo's Egg*, published in 1989, profiles a West German hacker working for the Soviet KGB who, in 1986, broke into the Advanced Research Projects Agency Network (ARPANET) and Military Network (MILNET) computers to steal secrets about the original "Star Wars" missile defense system.

From those initial experimental stages in the 1980s, computer systems, applications and OSs have dramatically evolved. Much has changed; in fact, everything has changed. Today's PCs and servers are often Internet-connected 24 hours a day, seven days a week for extended periods. Technical exploits target not only vulnerabilities in OSs, but also significant vulnerabilities in extremely complex, highly interactive, web-aware user applications (e.g., Adobe® Reader®, Microsoft® Office® applications, web browsers, JAVA™). Attackers also seek to take advantage of computer users (e.g., spear phishing or other social engineering) by deploying commonly available web-enabled, user-friendly hacking tool kits (Zeus is currently the most popular). Today, just about anyone who can use a web browser and a mouse can deploy and control a botnet.

Access points into an enterprise also have dramatically expanded. The perimeter once existed at the true boundary of the enterprise, where Internet-facing servers sat in the demilitarized zone (DMZ) behind a firewall. Today, cloud services, social media, mobile devices and bring your own device (BYOD) policies have moved the real perimeter through the firewall, beyond the laptop/desktop/server, right down to the most sensitive data that must be protected. As a result of the redefined perimeter, enterprises could find that their most critical data have been copied to personally owned tablets, uploaded to a file-hosting service or emailed to personal email accounts. Unfortunately, protection capabilities have not kept pace with changes to technology.

While cybersecurity professionals find themselves struggling to keep up with technical innovation, they must also deal with the fact that learning resources for would-be hackers have increased and are often freely available online. The Metasploit® framework has revolutionized vulnerability testing, making powerful vulnerability scanners freely available to anyone who calls himself/herself a penetration tester. While many trusted penetration testers have made use of Metasploit's open source and commercial testing tools, they have also been leveraged by attackers to achieve less-than-honorable objectives.

Today, each attack can potentially compromise hundreds of thousands of computers around the globe using automated tools. For example, the US Federal Bureau of Investigation (FBI) reported that the Coreflood botnet, created by a group of Russian hackers in 2010 and designed to steal web banking credentials, contained as many as 2.3 million computers worldwide, infecting systems of “approximately 17 state or local government agencies, including one police department, three airports, two defense contractors, five banks or financial institutions, approximately 30 colleges or universities, approximately 20 hospital or health care companies, and hundreds of businesses.”⁶

Attackers have used exploitable vulnerabilities in server, desktop/workstation, and laptop class systems and application software to harass and embarrass their targets; have monetized the capabilities of vast botnets to harvest banking credentials; and have used sophisticated malware to steal intellectual property from major enterprises. Government-sponsored attack groups have stolen hundreds of millions of US dollars of intellectual property to enable state-owned enterprises to leapfrog their competition. And today, malicious software with deceptive and destructive capabilities has been weaponized: Governments now use cyberweapons such as Stuxnet, the first cyber smart bomb, as an extension of politics to create effects in the physical world.⁷

1.3 Adaptive Attack Vectors

The threat landscape will continue to evolve as attackers adapt new and innovative attack methods to existing or adaptive attack vectors while defenders deploy new defense strategies. The concept of adaptive attack vectors (i.e., attacking company A as a means to enable an attack on company B, the real target) is illustrated in **figure 3** with three recent examples that have been in the news.

⁶ Russo, Tracy, “Coordinated Law Enforcement Action Leads to Massive Reduction in Size of International Botnet,” The United States Department of Justice, 27 April 2011, <http://blogs.justice.gov/main/archives/1320>

⁷ Stuxnet was allegedly designed by the US and Israeli governments to slow production of uranium at Iran's main nuclear enrichment facility in Natanz, Iran. See Sanger, David, “Obama Order Sped Up Wave of Cyberattacks Against Iran,” *The New York Times*, 1 June 2012, www.nytimes.com/2012/06/01/world/middleeast/obama-ordered-wave-of-cyberattacks-against-iran.html?pagewanted=all.

FIGURE 03 Adaptive Attack Vectors

Security Issue	Security Solution	Adaptive Attack Vector
Single-factor authentication (e.g., something you know, such as a user ID and a password) is too weak; passwords are easily compromised or guessed.	Multifactor authentication—something you know (user ID/ password) plus something you have (password retrievable from a token that changes at regular intervals based on strong encryption algorithms)	Break into the token vendor (RSA, March 2011) and steal the encryption keys that are used by the real target (Lockheed Martin, May 2011).
There are thousands of malware writers, some of whom masquerade their code as being from a trusted developer.	Digital certificates used to “sign” code from a vendor so that the code can be trusted	Break into a credible vendor whose software is run on almost every computer (Adobe) and use its code-signing infrastructure to sign the malicious code (September 2012).
The antivirus approach (defining what “bad” software is and blacklisting or quarantining it) is not able to keep pace with malware writers. There are more than 200,000 new blacklist signatures each day.	Application whitelisting (defining what is “good” and assuming everything else is “bad”)	Break into the application whitelisting vendor (Bit9) and have its code-signing infrastructure sign the malicious code so that it is effectively on the whitelist (February 2013).

Creative, talented and aggressive attackers continue to drive the threat world into new areas. They adapt and change; enterprises must adapt and change as well. Enterprises must take “prevent” out of the cybersecurity dictionary because there are no universal solutions to prevent a sophisticated attacker from infiltrating any environment (even the Natanz uranium enrichment plant, which was “air-gapped” from other networks, was penetrated). If a well-funded and sophisticated attacker targets a specific environment, he/she will get in. The key outlook for information security professionals must be: The network is compromised, or soon will be. How do we protect the most important data in a compromised environment? How do we make it difficult for attackers to be successful? How do we detect that an attack is underway? How do we respond to today’s sophisticated attacks?

The network is compromised, or soon will be. How do we protect the most important data in a compromised environment? How do we make it difficult for attackers to be successful? How do we detect that an attack is underway? How do we respond to today’s sophisticated attacks?

1.4 A Watershed Event

In January 2010, Google disclosed that it had been the target of a sophisticated computer network attack emanating from China and apparently seeking to access the Gmail™ accounts of Chinese human rights activists.⁸ Google also indicated that as many as 20 other large companies may have been similarly attacked. For many in the corporate security community, this was the beginning of the advanced persistent threat (APT) era.

However, US federal law enforcement, intelligence and counterintelligence communities were already very familiar with “the APT” by the time Google made its disclosure. As Richard Bejtlich pointed out in July 2010⁹:

The United States Air Force (USAF) coined the phrase advanced persistent threat in 2006 because teams working within the service needed a way to communicate with counterparts in the unclassified public world. Department of Defense and intelligence community members typically assign classified names to specific threat actors, and use the term intrusion set to describe activities by those threat actors. If the USAF wanted to talk about a certain intrusion set with uncleared personnel, they could not use the classified threat actor name. Therefore, the USAF developed the term APT as an unclassified moniker.

The term APT was not originally intended to be the generic term that marketers have transformed it into. It was developed to refer to specific, known state-sponsored groups in the Asia-Pacific region that conducted attacks against specific targets at the direction of their government. The public announcement by Google of the APT attack against its environment brought into the public light for the first time a group of attackers who had actually been at work for many years prior, but whose existence had not been acknowledged outside of tightly controlled circles.

In this third year of the APT era, many other enterprises have disclosed what they termed “sophisticated” attacks. And while the term APT no longer means exclusively what it once meant, it does represent a new breed of attacker: one who specifically targets a person or enterprise for attack to achieve a specific purpose. Today and tomorrow, other attackers will adopt the APT methodology because it works. In the future, it will be harder to attribute an attack to a specific APT group. But the fact remains that sophisticated attackers with incredible resources will be successful in attacking enterprises to steal information of value.

⁸ Eunjung Cha, Ariana; Ellen Nakashima; “Google China Cyberattack Part of Vast Espionage Campaign, Experts Say,” *The Washington Post*, 14 January 2010, www.washingtonpost.com/wp-dyn/content/article/2010/01/13/AR2010011300359.html

⁹ Bejtlich, Richard, “What APT Is (and What It Isn’t),” *Information Security*, July/August 2010

1.5 The APT Life Cycle

History shows that most sophisticated attackers, regardless of their motives, funding or control, tend to operate in a certain cycle (figures 4 and 5) and are extremely effective at attacking their targets.

FIGURE 04 The APT Life Cycle





FIGURE 05 The APT Life Cycle: Another View



The APT life cycle includes:

- **Conduct background research.** The APT conduct detailed research on their targets to identify very specific avenues of attack. Imagine that you are a senior executive for a Fortune 100 company just returning from an industry conference. In your inbox is an email message apparently from a conference speaker thanking you for attending his talk. The email contains an updated copy of the speaker’s presentation deck. Would you open it? What if you were the general counsel of a large enterprise and you received an email message apparently from your outside counsel informing you of new litigation filed in district court and containing a copy of the complaint. Would you open it? These scenarios represent the kind of detailed research the APT do to maximize the chances that their target will act as desired.

- **Execute initial attack.** Generally, the initial attack targets one or more specific individuals through some form of social engineering: embedding a link to malicious content into an email message, an instant message, a social media posting or another attack vector, and enticing the target to act by opening an attachment or clicking on a link to infect one or more devices with malicious software.
- **Establish foothold.** Once a user acts, the APT establish an initial foothold into the target environment using some version of customized malicious software. In almost every case, that custom software does not trigger any antivirus alert, but it does beacon to the APT to inform them of the successful attack. The initial infection tool, sometimes called first-stage malware, may have very little malicious functionality, but it generally is able to beacon home and download additional functionality, sometimes called second-stage malware.
- **Enable persistence.** One of the primary objectives of the APT is to establish persistent command and control over compromised computers in the target environment—meaning control and access that will survive a reboot of the targeted device and provide the APT with regular connectivity to the target environment. In most cases, this persistence is established simply by installing new services (including the attacker’s command-and-control software) on the target computer that automatically start when the computer boots. Users generally do not interact with services and often do not know they are running. In other cases, the attackers establish persistence by modifying the system registry to automatically run their malicious applications each time the computer starts. This is more risky for the attacker because a user might notice a running application and might terminate it. And, in some cases, the attackers establish persistence by replacing legitimate OS or application software components with compromised components that include additional functionality for their command-and-control requirements.
- **Conduct enterprise reconnaissance.** After establishing persistent access to the target environment, the APT generally conduct enterprise reconnaissance in an effort to find the computers, servers or storage areas holding the information they have been instructed to steal. In most cases, they conduct the reconnaissance using the tools available to them on the compromised computers. In some cases, they upload scanning tools to search for specific types of systems (e.g., identity and access management, authentication, virtual private network (VPN), database or email servers).
- **Move laterally to new systems.** Part of enterprise reconnaissance necessarily includes moving laterally to new systems to explore their contents and understand to what new parts of the enterprise they might gain access from the new systems. They also directly install their command-and-control software on new systems to expand their persistent access to the environment.
- **Escalate privileges.** As the attackers conduct reconnaissance and move around the network using the compromised credentials of their first few targets, they inevitably seek to escalate from local user to local administrator to higher levels of privilege in the environment so that they are not constrained to any specific part of

the environment. In enterprises where access to information is tightly controlled, compromising all the credentials in the environment (typically, but not always, an Active Directory® domain) allows the attackers to masquerade as anyone in the environment and access any resource they desire.

- **Gather and encrypt data of interest.** Having found the data of interest to them, the APT generally gather the data into an archive and then compress and encrypt the archive. This enables them to hide the contents of the archive from technologies that include deep packet inspection capabilities and from data loss prevention (DLP) at the enterprise boundary.
- **Exfiltrate data from victim systems.** The ultimate objective of an APT attack is to steal something of value. The APT use traditional file transfer protocol (FTP) applications if the enterprise allows unhindered FTP outbound to anywhere. But they also use custom data transfer technologies operating on standard and nonstandard ports if FTP is not allowed.
- **Maintain persistent presence.** Finally, the APT seek to attain what they have been tasked by their controllers to do: maintain access to the target environment. It is not uncommon for the APT to sit undetected in an enterprise network for lengthy periods of time, waiting to complete their directive. The APT do not attack “targets of opportunity.” If they are in the network, it is because they were told to be there—and they will seek to stay!

This APT life cycle can also be used to describe attacks by other sophisticated attackers, but the essence of an APT attack is its targeted nature. Take, for instance, the Coreflood botnet, a very sophisticated malware application. For more than 10 years the botnet controllers evaded antivirus software by updating the malware on a regular basis. The botnet controllers generally delivered the malware through social engineering attacks. The software self-replicated to move laterally through an infected network seeking specific computers with access to accounts payable systems. Although the software did not include functionality to escalate privileges to domain administrator level, it operated with the privileges of its targeted users. It captured keystrokes from accounts payable systems and relayed those keystrokes to the attackers, including user IDs and passwords necessary to access online banking systems and initiate electronic funds transfers. Then the attackers, masquerading as enterprise employees, used the stolen credentials to move hundreds of millions of US dollars to overseas accounts.

But the Coreflood botnet was not “the APT” and should not be described as “an APT.” The targets of the Coreflood controllers were selected only because they represented the potential for financial gain. They were targets of opportunity attacked for no reason other than they have computers connected to the Internet and might use web-based banking systems. That certainly does not lessen the impact that the Coreflood botnet controllers had on the enterprises they attacked, but defending against attacks like the Coreflood botnet is quite different from attempting to defend against an APT.

The APT have adapted their tactics, techniques and procedures to the typical information security architecture they find deployed (**figure 6**).

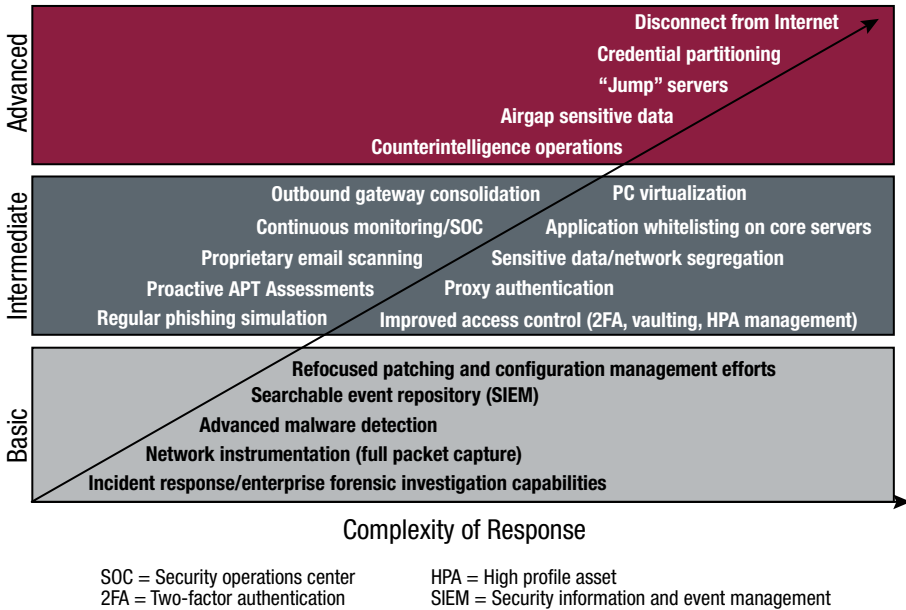
FIGURE 06 APT *Modus Operandi*

Traditional Security Practice	APT <i>Modus Operandi</i>
Network boundary/perimeter devices inspecting traffic content	Secure Sockets Layer (SSL), custom encryption, and password-protected/encrypted container files make packet content inspection difficult or impossible.
Network firewalls monitoring and assessing traffic metadata	Communication is initiated from within the network using standard ports and protocols (Hypertext Transfer Protocol [HTTP], Domain Name System [DNS], SSL, Simple Mail Transport Protocol [SMTP], etc.).
Host firewalls monitoring and assessing local traffic metadata	Initial infection tool adds malware to the host firewall whitelist.
Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) with real-time assessment and alerting running on servers and workstations	Communication uses common ports and protocols; the APT hide in plain sight within obvious/allowed traffic.
Antivirus software on every server and workstation, with multiple daily automatic signature updates	Malicious code is compiled immediately before use and tested on the latest antivirus definitions. Custom code is created per target. Protection is gained with multiple anti-reverse engineering techniques (i.e., kernel drivers, packers, other obfuscation).
Monthly or quarterly vulnerability assessments	Attacks are not based on server OS vulnerabilities, but on host/application vulnerabilities, zero-day attacks and user action.
Two-factor authentication	Custom malware is installed with user privileges. The hacker authenticates to the malware, circumventing two-factor requirements.
Hypertext Markup Language (HTML)-formatted email disallowed and real-time email content filtering	Links to malware (instead of malware itself) are embedded in very well-crafted, specifically targeted phishing email.
Real-time malicious web site and URL filtering	Third-party sites are compromised to host malware: a different site for each victim and a different site for each wave of attacks on a victim.

1.6 What Are Others Doing?

Because the APT and other advanced attackers have adapted to the environments they seek to attack, targeted enterprises can respond in a variety of ways—some of which are very advanced (figure 7).

FIGURE 07 What Are Others Doing?



Not every enterprise will consider the advanced mitigations. For instance, disconnecting from the Internet might break critical business processes that rely on connectivity with partners, suppliers or customers. Many of the other exercises are effective, yet might be determined to be overkill by an enterprise because, in almost any global enterprise, the level of effort required is extreme.

Many enterprises implement some of the intermediate-level concepts. Because the APT and other advanced, sophisticated attackers have such a high success rate, it is recommended that every enterprise implement all of the basic concepts.

1.7 Summary

The threat environment has radically changed over the last 10 years. Most enterprises have not kept pace. The remainder of this publication covers a few of the basic concepts depicted in **figure 7**, which will help answer the key questions posed by a new outlook: that a breach WILL eventually occur.

Enterprises need information on the fundamentals of conducting investigations of advanced, targeted cyberattacks, some of which are sponsored by government entities and conducted using methods designed to maintain a persistent foothold in the enterprise's environment. The attackers might be "the APT" or they might be some other entity that has adopted the APT approach and methodology. In any event, the enterprise must prepare for an attack executed by an advanced, sophisticated, organized, well-funded and persistent adversary. The enterprise must be able to conduct an investigation of a breach to feed threat intelligence into a detailed remediation/eradication plan, and it must be able to execute the remediation/eradication plan to successfully expel the attackers from the environment.

Page intentionally left blank

Chapter 2. Preparation

2.1 Build a Team, Make a Plan

Enterprises that invest time and resources into preparing for a breach will fare much better in their response and eradication efforts.

Objective: *Establish the people, plan and capabilities to enable a thorough and efficient breach investigation and eradication event.*

It can be difficult to convince enterprise leadership of the need to have a well-thought-out and practiced cybersecurity breach investigation capability. It is not uncommon for management to balk and question, “Why should I expend resources on response preparation when we can just work to keep attackers out altogether?”

The truth is that no matter the strength of enterprise defenses, sophisticated attackers with advanced capabilities have the means and determination to adapt and defeat the most complex prevention and detection measures the enterprise might deploy, regardless of size. In short, targeted, determined attackers are highly likely to achieve some degree of success. In addition, targeted attackers often establish strong footholds within their victim’s network, making it very difficult to eradicate them.

Prior preparation for a cyberattack, with the perspective that an attack will be successful, allows an enterprise to be both efficient and comprehensive in its investigation and eradication activities. So what can the enterprise do to prepare? The following section identifies six activities that enterprises can undertake to position themselves to manage a cybersecurity breach efficiently and effectively.

2.2 Establish Key Relationships

2.2.1 External Relationships

To operate effectively and efficiently during a breach investigation, enterprises should establish relationships with important third parties in advance of a breach. These third parties may include business relationships, joint ventures, anyone with a link into the enterprise’s network, on- and offsite contractors, and anyone else who would be impacted in the event the enterprise must operate in a degraded capacity. Once these parties are identified, their contact information should be retained and kept easily accessible by the appropriate individuals, including a management representative and a network security team leader.

Enterprises that do not establish these relationships prior to attacks end up scrambling to get the correct security companies and professionals to help them respond quickly once a breach has occurred. Conversely, enterprises that establish these relationships ahead of time and keep third parties on a retainer for times

of need can move much more smoothly to full-fledged breach investigation and eradication. The size of an enterprise and the maturity of its security operations play significant roles in determining the level of third-party involvement during a breach investigation and eradication event. Enterprises with mature security programs may conduct most of their operations in-house, while those with less mature security programs may depend entirely on third parties. Regardless of where the enterprise falls on this spectrum, it must consider many possible third-party relationships prior to an attack including:

- Breach investigation teams
- Managed services for security monitoring
- Denial-of-service (DoS) response services
- Malware analysis support
- Forensics support
- Eradication event teams

2.2.2 Internal Relationships

Network administrators in medium-sized or large enterprises commonly do not know about all of the departments within their enterprise. A comprehensive contact list of all internal relationships will significantly speed investigation, containment and eradication efforts. Knowing who has the authority to take systems and users offline or otherwise degrade operations in each department will prove invaluable. There is nothing more confusing or frustrating to the people within an enterprise that is under attack than not knowing who is responsible for what and who has the authority to make the decisions required during the investigation and eradication.

Once these individuals have been identified, they should be notified that they may be contacted by security team leaders to make a quick containment decision. Creating a responsibility assignment matrix (e.g., responsible, accountable, consulted, informed [RACI] matrix) helps to properly outline roles, responsibilities and authorities. Establishing documented processes and procedures and subsequently conducting tabletop exercises will open lines of communication among the parties involved in investigations and eradication events. Having all of these internal contacts identified and prepared to react to an advanced attacker allows the enterprise to fight through the attack with business continuity, rather than simple survival, in mind.

2.3 Determine Authorities

After the team is established, it must be empowered to act. Some of its actions may prove to be unpopular with others in the enterprise, but senior management must commit to supporting the security team in those cases. As long as business operations are not impacted beyond acceptable levels, the security team must be allowed to act in the best long-term interest of the enterprise. This empowerment is generally best accomplished with a very clear team charter signed by the CIO and outlining

the team's mission, goals, objectives, authorities and responsibilities. A copy of this charter should be given to every direct report of the CIO as well as to key business unit leaders. All executives must understand the response team's mission, goals, objectives, authorities and responsibilities.

2.4 Inventory Existing Technologies

A well-maintained, actionable inventory of existing technologies and assets aids in conducting an efficient investigation and eradication event. Such an inventory is a baseline step of preparation. This extends beyond the obvious IDSs and IPSs to all devices that touch the network because any device could be a connection point for lateral network movement in an attack. Moreover, the devices that serve business operation functions could be compromised, and the investigation team needs to have a comprehensive inventory to understand each device's purpose, criticality and point of contact. Examples of these types of devices include supervisory control and data acquisition (SCADA)-type industrial control systems (ICSs), payment card industry (PCI) equipment, or health care technology connected to a network. If these systems are compromised, it can be catastrophic to an enterprise; therefore, a complete inventory is imperative.

In an advanced attack, having an inventory of all devices (with their major attributes) allows the investigation team to follow each of the possible leads (rabbit holes) and examine each indicator of compromise (IOC) thoroughly. Following the leads means taking whatever initial IOCs have been discovered and following each attribute until an investigative end is reached. Each of the following attributes is significant and should be completely explored:

- Date/time
- Internet protocol (IP) address (internal or external)
- Port (source or destination)
- Domain
- File (e.g., .exe, .dll)
- System (hardware vendor, OS, applications, purpose, location)

The date and time give the security team a point on which to focus its log and SIEM analysis. The analysis should start with a small window prior to and following the specific time provided (hopefully, it is at least at the hour level) and look for other activity matching the initial indicator. The time of the event detected could be the initial attack, but it is more likely a mistake by the attacker during a later phase of its efforts. Once the initial time window has been examined, the parameters should be broadened systematically to look for other events matching the initial and subsequently discovered indicators. Even if it is possible to discern easily what occurred from these data, it is wise to hold off on cleaning and restoration actions until the eradication event.

If an external IP address or port is provided, all communication to that IP address or through that port should be examined, starting with the traffic from the initial time window. Finding all of the internal hosts that were talking to that IP address and then discovering everything that was occurring on those systems in a window around that event can be critical. Next steps should include seeing who else these systems were communicating with around the event (internally or externally) and looking into those hosts as well.

In cases where an internal IP address or port is provided, one can follow the previous steps, but in reverse order (inward-to-outward instead of outward-to-inward). If a repeated pattern of traffic is discovered (e.g., one packet being sent out every hour on the hour), this is likely a sign of a command-and-control channel between the two hosts. This traffic is called “beaconing” and lets the attacker know that its access to the system still exists and is ready to be exploited.

If the attacker is utilizing a specific domain, it is again critical to find all traffic to the domain from the network and investigate those systems. Domains can switch IP addresses at the will of the owner, so one should be wary of relying on old data connecting these two attributes.

When the initial IOC is tied back to a specific system, such as a server type or domain controller, all similar systems of that type should be examined. Because they likely have the same configurations, patches and vulnerabilities, there is a high likelihood that these also could be compromised. Advanced attackers are very careful and thorough, so they have redundancy built into their campaign-style attacks.

Without a thorough inventory of all network assets, the investigation team will be operating at a disadvantage. The up-front investment in establishing enterprise visibility will pay large dividends in maintaining security compliance and, more important, responding quickly in the event of a breach.

2.5 Standardize the Investigation Process

While every network situation is unique, there are commonalities that allow for a standardized plan that an enterprise can proactively implement and adapt as needed. The plan must be sufficiently comprehensive and agile to cover, and adapt to, many different scenarios. This means that the plan needs to be written at a higher level. An enterprise should begin by building a comprehensive investigation and eradication plan in preparation for a security breach. This plan can be adapted to the situation in light of new information and discoveries.

A solid plan outlines standard processes of the incident identification phase. It breaks down important sections and outlines overarching actions to take during each step of the process. This process typically includes the following major categories:

- **Determine that an incident has occurred.** Determine whether a reported event constitutes a security incident, a nonmalicious policy violation or another nonmalicious event.
- **Analyze and categorize the incident.** Assign a category and potential subcategory to an incident to reference the appropriate preestablished response actions.
- **Rate and prioritize the incident.** Analyze the actual or anticipated impact of the incident to determine priority of response actions and assign a severity level.
- **Track the incident.** Assign the incident a unique identifier in the tracking system that will follow the incident through resolution, communications, recovery and lessons learned.

The plan may include a basic triage process such as that in **figures 8 and 9**. If the team can determine, based on malware and threat research, that the event does not pose a great threat, the team can follow standard remediation procedures. Otherwise, the team will elevate the incident to management and execute a more in-depth investigation and cleanup.

FIGURE 08 Basic Incident Triage Process

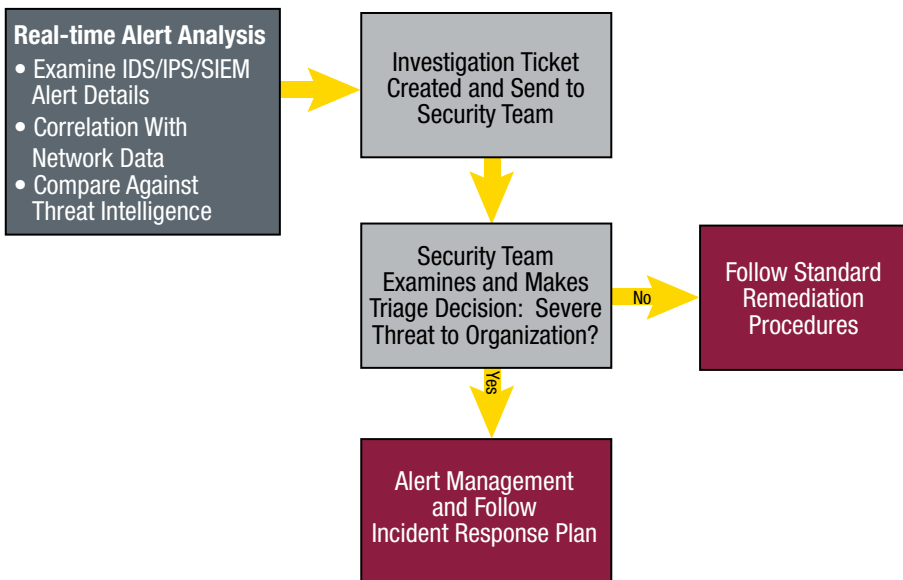
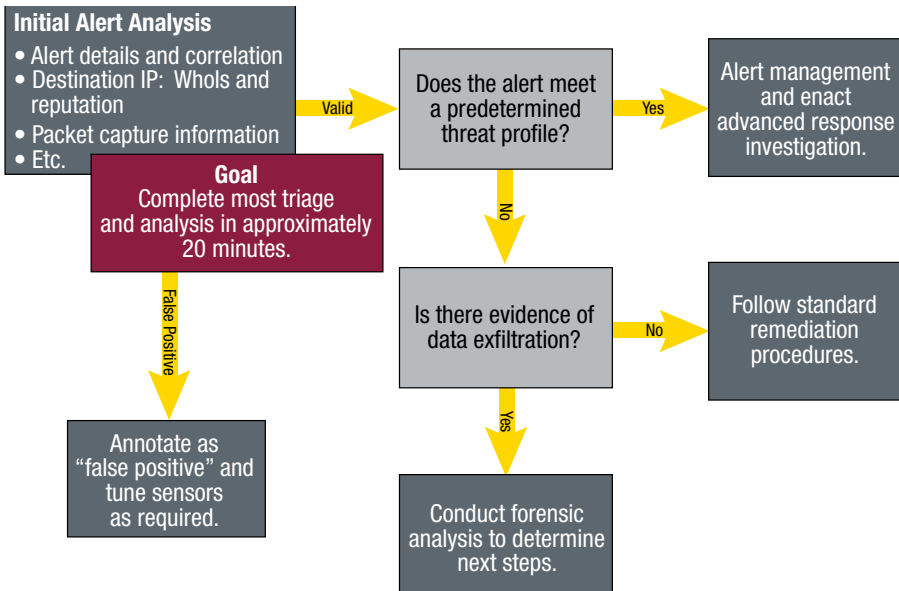


FIGURE 09 Alert Decision Process



The investigation, eradication and post-eradication sections in this publication highlight more granular details regarding investigation execution.

2.6 Training and Governance

2.6.1 Exercises

The proficiency of the investigation and eradication teams should be tested at least annually. A comprehensive exercise should involve all of the key factors: communications, coordination, resource availability and response. Management must be involved in all phases, from planning and execution through lessons learned and policy and security posture improvement.

Exercises may follow a variety of formats:

- Full-scale
- Tabletop simulation
- Seminar

A full-scale exercise takes a specific scenario and plays out all or most of the actions necessary to deal with the situation. The security team may not have prior knowledge of the scenario or its scope. It may not even be told that the situation is an exercise. This “no-notice” exercise is the most robust and advanced form. Most often, there is a large exercise planning group to coordinate most of the actions necessary in

advance. Evaluators may be appointed or brought in to help identify lessons learned and areas of strength. The smaller the planning team, the more the element of surprise is maintained for the players, and the more accurate the depiction will be to an actual response.

Records should be maintained for all exercises. The records should include the date range of the exercise (if covering multiple days); the intended and actual scope of the exercise (these may be different); and a description of all activities, participants, observations and lessons learned, as well as noted improvement areas.

2.6.2 Security Program and Response Plan Reviews

It is a good idea to review and update the comprehensive security program and subsequent investigation and eradication response plan annually to fit changing organizational needs. This includes all of the program performance targets, which should be reviewed, evaluated for success and feasibility, and adjusted as necessary. Compliance with training requirements should be measured, and technology requirements for the upcoming year considered and planned. Advances in technology, evolving attacks and linked vulnerabilities, and trends in user behaviors should all be explored when reviewing the security program.

The best time to conduct these reviews is prior to annual budget planning and the general financial management process. An enterprise should consider the funding needs of the program for the following year and its program objectives, needs and planned investments (e.g., training activities, awareness activities, equipment and technology, personnel and operating expenses).

2.7 Establish Critical Capabilities

To conduct a thorough and efficient investigation and perform an effective eradication event, critical capabilities should be established.

Determining the scope of the breach is critical to any investigation. This requires a certain set of capabilities, especially with highly complex attacks. As part of breach preparation activities, it is important to perform a gap analysis of existing capabilities and develop a plan to close the gaps so the enterprise can better manage a security breach when it occurs.

A critical set of capabilities consists of appropriately skilled people guided by well-designed processes that enable the effective use of relevant technologies. The right set of capabilities is key to conducting a thorough investigation and successfully eradicating adversaries deeply embedded and persistent in the environment. The capabilities required to conduct an effective, efficient investigation and eradication are provided in **figure 10**. The capabilities are divided into minimum and preferred

levels. The minimum capabilities are those that can enable a team to investigate a breach with some reasonable level of success, depending on the complexity of the breach. The more complex the breach (e.g., scope of systems compromised, sophistication of the malware), the more difficult it will be to operate effectively without the preferred capabilities.

Technologies or tools tend to get most of the attention and often are the focus of incident response (IR) improvement projects. Technical capabilities can be acquired by procuring expensive enterprise appliances and software, or they can be addressed in a more cost-effective manner by first maximizing the use of existing tools and, where appropriate, utilizing open source solutions to fill the technical capability gaps. Before procuring a new solution, an enterprise should first make a concerted effort to understand the technical capabilities of high-performing IR functions and determine which are in place. Where gaps are found, enterprises should first look to open source possibilities before purchasing products. Taking a disciplined approach to managing costs related to tool acquisition will make more funds available to acquire highly skilled people and invest in training and education to improve the existing team's skills.

It is true that effective and efficient IR is highly dependent on having certain technological capabilities. However, more often than not, a sustained investment in people and process capabilities will yield a greater return on the enterprise's investment in technology.

Figure 10 summarizes the people, process and technology capabilities that a computer security incident response team (CSIRT) needs to effectively and efficiently manage a cybersecurity incident. Subsequent sections elaborate on each capability.

As outlined in **figure 10**, an investigation team needs to have visibility into both host-level and network-level activity. This could be described as understanding data in motion, data at rest and data in use on an enterprise network of hosts. There are basic requirements needed to conduct an analysis, and then there are more mature capabilities that will increase the efficacy of an incident-handling team. If an enterprise stops at the basics and fails to think about a more robust capability, the ability to find IOCs and apply them in a systematic manner to search for the full scope of compromised systems will be very limited, and the ability to proactively detect a compromise before significant damage occurs will be challenging and cost-prohibitive. This section outlines the capabilities that are basic, common and required, and the ones that go beyond those levels and should be considered only after all of the basic capabilities are met.

FIGURE 10 CSIRT Capability Requirements

Capability (People, Process, Technology)	Minimum	Preferred
Host-level activity awareness	<ul style="list-style-type: none"> • Logs from end-point software agents (e.g., antivirus) • Native OS logging (e.g., Microsoft Windows®) 	<ul style="list-style-type: none"> • Host-based intrusion detection • Remote enterprise forensic analysis • Agent-based, live memory analysis
Network-level activity awareness	<ul style="list-style-type: none"> • Network flow data (e.g., layer 3) • Proxy logs • Firewall logs 	<ul style="list-style-type: none"> • Network intrusion detection logs • Full packet capture at all egress points • SSL inspection
Search	Decentralized log searches on a per-system basis: <ul style="list-style-type: none"> • Local logging • Manual retrieval • Limited automation 	<ul style="list-style-type: none"> • Centralized aggregation of searchable log data • Event correlation (e.g., SIEM)
Digital forensics	<i>Ad hoc</i> , local	<ul style="list-style-type: none"> • Remote enterprise (acquisition) • Case management systems
Malware analysis	<ul style="list-style-type: none"> • Dynamic malware analysis • Basic static and automated analysis 	<ul style="list-style-type: none"> • In-depth static code analysis • Reverse engineering
Threat intelligence	<i>Ad hoc</i> , open source research	<ul style="list-style-type: none"> • Subscription-based • Business partner information sharing • Repeatable, automated integration
Vulnerability identification	Enterprise application inventory	Enterprise vulnerability identification

2.7.1 Host-level Activity Awareness

Each host on an enterprise network contributes to the attack surface of an organization. Each server, desktop workstation, laptop, printer and network appliance can be used by an attacker to gain a foothold or expand access into a network. To detect this activity, it is important to have a good understanding of the activities that are happening within each host. Depending on the size of the enterprise, it may not be possible for its security monitoring function to analyze all activities involving all hosts in real time. However, these data should be centrally collected so they can be parsed, searched and understood in the event an investigation is needed.

The most basic task of a digital forensics examiner is to re-create a timeline of host activity. This can help the team understand the capabilities of malware on a machine, or possibly help in the building of IOCs that can be searched for across the enterprise to find other compromised hosts. The team can also use this information to understand an attacker's intent. Is the enterprise dealing with hacktivists who are looking for information that might be used to defame or damage the enterprise if released publicly? Is it dealing with an APT that is interested in research and development servers that contain sensitive proprietary information? Or, are the attackers financially motivated and simply interested in collecting personal data, including credit card or Social Security numbers? An enterprise must ask itself why the attackers compromised its network. What were they after? Why was the enterprise a target?

Local logging of activity on a machine is vital to achieving a sense of host-based awareness. At a minimum, hosts should be configured so they log the following activities both locally and to a centralized repository:

- Local and network logon/logoff activities
- Antivirus alerts
- Process execution
- Account creation/deletion or group membership modifications
- Service creation/start/stop
- Web browser history
- For network security appliances, ALLOW and DENY records

It is important to confirm that the maximum log file is sufficient to capture a meaningful amount of local logging time. The default configurations of some systems or applications may be sufficient in certain cases, but not in others. The Windows Event Log file size maximum (20 MB before rolling over) may be sufficient to track several weeks of activity on a workstation, depending on its function. However, a log file of that same size may capture less than a day of activity on a heavily used server or domain controller. If an incident is not discovered until days or weeks after it happened, valuable information about an attacker's activity may already be overwritten.

Some attackers may attempt to clean the system of log files after they compromise a server, but a skilled digital forensics practitioner can look for a number of other artifacts to correlate the events on a host. Attempts to clean artifacts off of a system often generate more artifacts that look like a bright red flag of malicious activity to a skilled analyst.

Along with log files, other system artifacts that contain very meaningful and relevant time stamps are:

- File system metadata layer time stamps (MACB value)
- Linkfile generation and associated time stamps
- Registry modification time stamps
- Prefetch file creation and modification time stamps (point to process execution)
- Shellbags
- Recovery of \$MFT entries marked for deletion
- Recovery of registry entries marked for deletion
- Scheduled tasks, AT jobs or cron jobs

When analyzed in order, these time stamps can provide a detailed picture of attacker activity and generate a list of powerful IOCs that can be used to identify malicious activity elsewhere in an enterprise.

These artifacts can be analyzed on a host-by-host basis by a skilled analyst who knows how to parse all of these data structures with open source tools and look for correlated events, but this is not feasible on a large scale. Many host-based monitoring systems can perform this type of analysis at a higher, centralized level. The most detailed, meaningful IOCs typically are generated after conducting in-depth forensics on a given compromised host; a centralized solution then can be leveraged to search for some of these IOCs elsewhere in the enterprise to gain a more complete understanding of the scope of the intrusion and identify other servers that should receive attention during an investigation.

The recommended level of host-based awareness involves the implementation of a host-based intrusion detection system (HIDS), an enterprise forensic analysis capability and a live memory analysis capability. The HIDS would extend the capability of SIEM by allowing a distributed search for artifacts on hosts themselves, on top of central log management. An HIDS gives powerful distributed search functionality to an investigator by being able to search through more than just logged data. Among the interesting things an investigator could search for with an HIDS are:

- The presence of a malicious file on disk, by file path, name or hash
- Services
- Running processes
- Network connections or listening agents

2.7.2 Network-level Activity Awareness

It is equally important to gain visibility into data in motion in an enterprise via network-level awareness. Often, the initial awareness of an incident comes from a network-based indicator. Most network infrastructure devices have the ability to generate NetFlow data and forward that activity to a central server. One way to

achieve the most basic required level of network awareness is to generate NetFlow data that can be queried. If NetFlow data are not available, another option for basic-level capability is monitoring firewall and proxy logs. Since most enterprises log proxy data and at least certain rules on the firewall, these data can prove valuable to an investigation into attackers who attempted to communicate outside the network through the proxy or firewall. Other capabilities may also meet the minimum required needs if they provide the following network-level data:

- Time stamp
- Source IP address
- Destination IP address
- Port number
- Packet size

Similar information can be gathered from the firewall and proxy. While this level of information by itself may not be sufficient, it can lead to very powerful conclusions regarding an attack when aggregated in a way that can be queried. Capabilities such as these can help establish network baselines. Developing extreme familiarity with enterprise policies and standards and monitoring network-level behavior over a period of time can help identify deviations from these baselines that may warrant a closer look. These baselines are specific to each enterprise, and they require a detailed understanding of the roles of the IP space and hosts on the enterprise's network and the standard methods in which they interact with each other. Depending on the enterprise, some examples of network-level activity that may deviate from the baseline are:

- A new host in an operations department segment communicating high levels of network activity between a series of hosts in the research and development subnet
- Web servers initiating web traffic
- Specific insecure methods for file transfer not typically permitted by the enterprise (e.g., FTP, Trivial FTP [TFTP])
- Remote administration of devices from outside the network (e.g., Secure Shell [SSH], Remote Desktop Protocol [RDP], Telnet)
- Port 80 or 443 traffic where more data are sent than received
- If the enterprise does not do business with certain countries, overlaying geolocation data to NetFlow data and searching for traffic to/from other countries

Many network administration teams could brainstorm a list of baseline indicators that is pages long, and an analyst could generate searches and alerts to look for these indicators in the minimum network-level monitoring capability.

The preferred level in network analysis involves an investment in network awareness appliances that include protocol decoders. This category includes a combination of firewall logs, proxy logs, network intrusion detection systems (NIDS), NetFlow data

and full packet capture. These can be configured to look for anomalous traffic that does not conform to standards or deviates from established thresholds/baselines set by administrators. For deeper analysis, alerts from these appliances can be fed into a centralized repository for search and/or correlation with other data.

The most comprehensive level of network awareness is full packet capture. Enterprises that opt for full packet capture should give thought to how it is designed and implemented. Full packet capture solutions can collect an incredible amount of data quickly, which may lead to a lack of storage if the solution is not properly designed. Full packet capture allows teams to replay sessions and extract suspicious files that were transmitted on the network for analysis. It can help the enterprise determine what data may have been exfiltrated or help recover the source installation file of a malware infection, allowing malware analysts to develop signature-based and heuristic indicators of compromise.

2.7.3 Search

Critical log data must be readily available and searchable. The minimum capability required to conduct effective searching during investigation and eradication includes sending log files to a central syslog server or another server that stores these data. Basic search capabilities can be developed through custom scripting, using open source tools or purchasing technologies designed to enable log data searching. This provides the enterprise quick visibility into its network and preserves data for responding to incidents.

At a minimum, enterprises should focus first on developing processes and procedures to regularly monitor individual log sources. Good candidates for this type of monitoring include firewall, proxy and antivirus logs. For example, the enterprise should be monitoring and responding to antivirus logs that alert to malware found on a system that the antivirus program could not delete or quarantine.

A more mature enterprise tends to have a scalable architecture to manage and search logs that enables correlation and alerting. This can be accomplished by the many SIEM solutions in the market today. SIEM serves two detection functions: as a repository and correlation platform for alerts generated by other means, and as an alerting capability based on network traffic flows and past logs. While this solution is not initially required to provide enterprisewide coverage for every log source, it should be scalable to meet future demands. Several log management and search products on the market can meet this requirement. Enterprises that have a limited number of high-priority log sources available in a centralized environment can get quick answers to on-demand searches. This sets the foundation for more advanced correlation and behavior-based detection.

When implementing SIEM, the enterprise should consider incorporating logs in a phased approach. This phased approach should integrate logs to meet the demands of a particular threat actor's tactics, techniques and procedures; they should not be phased into SIEM simply because someone thinks it is a good idea or because they may be useful someday. Integrating unnecessary log sources may cause system performance hits and cost more money in licensing fees. An enterprise that has SIEM needs to remove unnecessary data from the system and seek to home in on the most relevant alerts. An optimized enterprise understands the most relevant alerts, has well-defined processes to respond to those alerts, and experiences very low incident detection and response times.

2.7.4 Computer Forensic Analysis

All user activity leaves an enormous number of footprints on a host. The challenge of analysts is to sift through data to determine which data are useful, which artifacts point toward user activity and which data are results of system activity. Some data are recorded as a result of system activity that is caused by user interaction and can be leveraged to learn more about user activity. This undertaking is part art and part science and is commonly referred to as digital forensics.

Forensics is traditionally thought of as the application of scientific methods and technologies to establish facts of interest in relation to criminal or civil law. Enterprises may conduct investigations to determine the scope of an infection without pursuing legal recourse, so the term "forensics" may not necessarily apply. However, the same processes are still followed. At a basic level, an enterprise can conduct in-depth forensic reviews of a host by manually visiting the host and running an application that will dump the contents of memory and then image the full disk. This information should be dumped to a removable device or a network location to prevent overwriting unallocated disk space on the host.

There is no "minimum vs. preferred" level of conducting forensic analysis; however, some investment can be made in the logistics behind evidence acquisition and case management. There are enterprise agents that can be deployed remotely to collect images of volatile memory and full disk, hash the evidence, and store it on a remote location. This can save valuable time because it can be time-consuming for an enterprise to deploy an individual to a physical location. Also, it may not be operationally possible to shut down a server and do physical drive collection due to uptime requirements. Additionally, some enterprises find it helpful to organize the efforts of their forensic analysts by investing in and using a case management system that tracks workflow and standardizes the documentation of findings for evidence that has been analyzed.

2.7.5 Malware Analysis

There are two basic types of malware analysis: dynamic and static. Both provide information that aids an analyst or a reverse engineer in understanding the purpose, methods, impacts and potential indicators of the malware's function and origin. However, dynamic and static analysis techniques leverage different means to reach those conclusions. While an automated form of dynamic analysis is the minimum requirement for effective investigations and eradications, the recommended approach is to leverage the benefits of both techniques to help discover as much as possible during the investigation.

Dynamic malware analysis functions by executing possible malware in a controlled and contained environment—typically achieved through virtualization—to observe its behaviors. This controlled environment contains sensors that monitor the behavior of the executable, such as a file written to disk, network traffic, modified registry keys and spawned processes. Dynamic analysis can be completed much faster than static analysis through automation and typically can be conducted by a less technical resource; but dynamic analysis is not sufficiently thorough because it does not reduce the malware to its smallest possible parts.

It is better to combine dynamic analysis with static malware analysis. Static analysis involves reverse engineering or taking apart the software's code to analyze its parts and functions. Disassemblers, debuggers and decompilers are common tools used to perform these functions. The end goal is to access the source code, or the most basic code, of the malware. As author Dennis Distler wrote, "While performing code analysis, antivirus software will run on the malware, string searches will be performed, and files such as shell scripts will be analyzed."¹⁰ In the event of a targeted and advanced attack, the intentions and identities of the attacker will become more apparent. Because most malware involves some sort of command-and-control function, it allows the attacker access to control the infected system or communicates back to the attacker, and attribution can be gained. Malware continues to evolve, however. Some malware has functions built in to detect tampering, in which case the malware will perform functions to look benign or hide its purpose.

2.7.6 Threat Intelligence


As part of its network security preparations, an enterprise needs to look at existing external threats. By profiling threats, the enterprise can evaluate those it believes have both a high interest in attacking them and the highest likelihood of succeeding. Looking inward and outward in a balanced approach is necessary to make intelligent and agile defensive decisions.¹¹

¹⁰ Distler, Dennis; *Malware Analysis: An Introduction*, SANS Institute, Information Security Reading Room, 2007, page 20, www.sans.org/reading_room/whitepapers/malicious/malware-analysis-introduction_210

¹¹ The Mitre Corporation. "Standardizing Cyber Threat Intelligence Information with the Structured Threat Information eXpression (STIX™)," 2012, <http://makingsecuritymeasurable.mitre.org/docs/STIX-Whitepaper.pdf>

Cyberthreat intelligence can be viewed as research into malicious actors to determine their capabilities, motives and likely actions. When a security team conducts and applies cyberthreat intelligence, the team will more clearly understand the tactics, techniques and procedures of the attackers and can defeat weaker actors by disrupting or degrading their efforts. Moreover, threat intelligence potentially allows a security team to push the “front lines” back from the point of typical notification (generally well after the real damage is done) to somewhere in the attacker’s preparation or initial attack stages, as depicted in **figure 11**.

FIGURE 11 Potential Impact of Threat Intelligence on the Attack Life Cycle



Threat intelligence can also be a key indicator that a compromise has already occurred. Threat intelligence may not arrive before an advanced attacker has established a foothold, but it can be used to search network SIEM and logs for evidence of a breach.

To determine which intelligence to focus on, the enterprise should answer a series: Who would want to attack us? Where are we vulnerable? What equipment do we have, and what weaknesses does it have? Answering questions like these can help the enterprise narrow the plethora of available intelligence to a reasonable amount to review on a regular basis.

Threat intelligence can be obtained in a number of ways. The most common way, but by no means the best way, is on an *ad hoc* basis. This means that a member of a security team may perhaps receive a random tip from an acquaintance, happen upon something useful while digging through technology web sites, or see a news story that may relate to a configuration or tool on the enterprise’s network.

The next step up from *ad hoc* is to make threat research a recurring task, a part-time duty or even a full-time research job. This collection method can have varying degrees of success depending on the skill of the employee, his/her knowledge of the threats and of where to look for data, and luck in research efforts.

For most enterprises, a better method is to subscribe to a regular feed of threat data. Threat intelligence subscriptions have many advantages that make them an attractive option. Specialty companies, antivirus vendors and other security providers offer many customizable services for enterprises. They take feeds for thousands of end points and sources in their globally-reaching web and analyze trends and anomalies for valuable data.

Once their data analysis is complete, they can offer their reports to subscribers in a variety of formats. Vulnerability reports address security holes discovered in specific software or hardware. Flash bulletins are meant to alert security teams to events occurring at that moment. Malicious code reports thoroughly break down the attributes, attack methods and impacts of malware, while threat reports similarly describe specific threat actor teams up to the APT level. Because these third-party vendors can reach into a broad spectrum of enterprises, they can establish more complete patterns than an enterprise acting alone could achieve.¹² These reports may even be tailored to specific industries to further aid security teams.

Many industries are sharing threat intelligence at levels from real-time analyst up to senior leaders. Chief information security officers (CISOs) meet periodically through conferences or summits to aggregate trend reporting, and analysts can contact peers when appropriate. There are a number of formal and informal sharing partnerships, such as the sector-centric information sharing and analysis centers (ISACs). A commercial example is the Red Sky Alliance,¹³ whose members can collaborate on conducting counter-APT research; exchange data; and subscribe to detailed threat indicator tables, which take disparate reports across multiple industries targeted by the APT and combine them into a single source for readers.

Industry partnerships are another potential source of threat intelligence. Enterprises that share attack indicators with potential competitors benefit by reinforcing a united front against common adversaries.¹⁴ Within the technology sector, a newly formed Cyber Security Research Alliance (CSRA)¹⁵ is uniting large technology companies to move beyond their already established sharing of data and reports to take on “grand

¹² Verizon, “2012 Data Breach Investigations Report,” 2012.

www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012-ebk_en_xg.pdf

¹³ Red Sky Alliance, www.redskyalliance.org

¹⁴ Lange, Captain Adam M., USAFR; Captain Mark G. Manglicmot, USAFR, “Merging Industry Best Practices, Intelligence Improvements, and Adversary Analysis to Enhance Cyber Defense,” *Air and Space Power Journal*, 2013

¹⁵ Shalal-Esa, Andrea, “Lockheed, Intel, Others Team up to Tackle Cyber Challenges,” Reuters, 24 October 2012, www.reuters.com/article/2012/10/25/net-us-cyber-companies-idUSBRE89O03120121025

challenges” by producing breakthrough technologies. This nonprofit alliance intends to work with its multiple affiliates to manufacture long-lasting, advanced network security solutions.¹⁶

For industries and enterprises addressing custom APT-targeted malware, it is valuable to share data in this way. Free antivirus service providers such as VirusTotal are a good start, but are insufficient for dealing with sophisticated actors.¹⁷ Even if one instance of a virus is defeated in this method, all of the other characteristics of an attack are still valid and available to an APT, such as the initial infection vector (i.e., phishing) or the command-and-control parameters. Furthermore, because all advanced actors test their malware against antivirus software before deploying it, they would be able to identify and recode the exposed version. After that, they could continue their campaign and the target would not be any more secure than before using the antivirus service provider.

The US Computer Emergency Readiness Team (US-CERT), a subcomponent of the Department of Homeland Security, has a history of sharing strategic and technical reports among individuals, businesses, government and ICS network professionals. The public and private organizational liaison of the US-CERT, The National Council of ISACs, specializes in actionable sector-centric threat intelligence.¹⁸ US-CERT encourages submissions through its web site of user reports of phishing, incidents and software vulnerabilities. In return, US-CERT releases daily threat alerts, vulnerabilities analyses and malware dossiers.

The military started sharing indicators at an analyst-to-analyst level a few years ago, and these personal relationships have evolved into an annual CERTCON (a grassroots conference for interservice computer network defense service providers). From this conference, the military has started formally sharing ideas, tools, processes and case studies. Other industries can model after this template, starting with trends and lessons learned, and then growing into more real-time indicator sharing.¹⁹

Understandably, sharing within an industry, with potential competitors or with any other security source, may be difficult to accept at first. However, if all parties participate equally, the benefits outweigh the costs. On the other hand, enterprises will want to balance self-reliance and overreliance on vendors. When a group is self-reliant, using only its own response measures without sharing indicators within a trusted network of business partners, it is essentially acting without any quality assurance. Enterprises that rely on outside vendor support may see previously private security events become widely known, even to APT actors.

¹⁶ *Ibid.*,

¹⁷ VirusTotal, www.virustotal.com/

¹⁸ Pelgrin, Will, “The Role of Information Sharing And Analysis Centers (ISACs) in Private/Public Sector Critical Infrastructure Protection,” January 2009, pp. 1-2, www.isaccouncil.org/index.php?option=com_docman&task=doc_download&gid=1&Itemid

¹⁹ *Op cit*, Lange and Manglicmot

Combining threat intelligence with reviewing recent logs allows for advanced traffic analysis. The analysis activities recommended by the US National Institute of Standards and Technology (NIST) in Special Publication 800-61 can be used to find not only trends in usage, but also advanced targeted actors.²⁰ A targeted attack may attempt to blend in with normal traffic to evade detection, and only advanced log analysis may be able to discover its presence. Targeted attackers can also attempt to keep their network footprints small enough to remain under any clipping levels on sensors. Advanced log analysis can seek out these “low and slow” actions to find outlier activities and correlate them to legitimate or illegitimate traffic.

Performing advanced traffic analysis can help to identify network trend deviations. When these occur, it is time for the analysts to take a deeper look. Some of the trends to consider are:

- Where are most security incidents occurring (business unit, location, equipment type)?
- What type of incident is occurring the most (specific type of malware, keylogging, incidents from phishing, social media-related attacks)?
- Are the incidents readily solvable (adequate logging, network visibility, acceptable use and investigation policies, user cooperation)?
- What are the latest IDS/IPS/SIEM alert creation trends? What are they in response to? What are the intelligence/compromise sources?

2.7.7 Vulnerability Identification

The security team should conduct a complete vulnerability scan of the enterprise regularly. The frequency depends on the size of the enterprise, and incremental scans of network segments may be necessary. A quarterly scan for small- or medium-sized enterprises may be adequate, while larger enterprises may decide on a monthly schedule.

No matter how the enterprise decides to schedule vulnerability scans, the general purpose is to discover what devices in the enterprise are open to well-known vulnerabilities. The key is to conduct the scan, analyze the results, and address the findings before an adversary can discover them and exploit the weaknesses.

There are both free and subscription scanners available for enterprise usage. While all scanners are designed for the same general purpose, they may be more tailored for specific data, OS or hardware types. Different scanners have unique attributes; when selecting the scanner type, the security team should weigh the impact the scanner will have on the devices evaluated, as well as the thoroughness of the assessment.

Once the scans have been completed, vulnerabilities should be systematically addressed. More often than not, vendors will have patches and updates available to resolve any issues presented. Business or political reasons may impede the

²⁰ NIST, “Computer Security Incident Handling Guide,” SP 800-61, 2012, <http://csrc.nist.gov/publications/drafts/800-61-rev2/draft-sp800-61rev2.pdf>

application of patches to certain systems. This is where regular updates to senior management are valuable and can be used to present the case for improved security.

If updates cannot be applied, it is critical to find a way to increase monitoring capabilities of those systems. The results from the scans should also lead to advanced traffic analysis on the systems with identified vulnerabilities—advanced attackers could have already breached these systems. While vulnerability scanning is primarily used as a prevention tool, it can also be a part of an enterprise’s greater detection capabilities. As a key information security motto states, “Prevention is ideal, but detection is a must.”

Chapter 3. Investigation

3.1 Conducting a Security Breach Investigation

The core component of any investigation is the collection and analysis of facts pertinent to the matter. Facts related to a cyberbreach can be gathered from multiple sources, including witness testimony, infrastructure logs, computer systems and third parties. In the case of a cyberattack by an APT, the primary purpose of the investigation is to provide intelligence to the eradication and remediation plan. The eradication plan is generally designed to remove the attackers from the environment, and this cannot be successfully executed without:

- Detailed information on the attacker's tactics, techniques, procedures and intent
- Knowledge of the systems the attackers have compromised, the credentials they have stolen, the data they are after, and the command-and-control systems they used or are using to maintain persistent access to the environment

Details of eradication planning are discussed in chapter 4.

The process of gathering and analyzing facts does not usually follow a perfectly linear path (e.g., interview first, then review and analyze electronic logs). Fact gathering is a dynamic, fluid, cyclical process that involves sound judgment and experience. As new information is discovered or certain initial assumptions prove right or wrong, the process may have to begin again. Accordingly, an investigation might start with an interview, move to log review and then go back to a follow-up interview based on new information.


The IR or investigative process can be broken down into five phases: identification, containment, eradication, recovery and follow-up, as shown in **figure 12**.

While the five phases are sequential, they are not necessarily discrete—they may overlap and recur as the investigation demands. But the fact remains: An incident that has not been identified cannot be contained, and an incident that has not been contained cannot be eradicated.

The major objectives of the IR process are to:

- **Identify** all potentially relevant data sources.
- **Preserve** those data with forensically sound methodologies.
- **Analyze** the data for facts related to the matter.
- **Report** findings.

FIGURE 12 The Incident Response Process

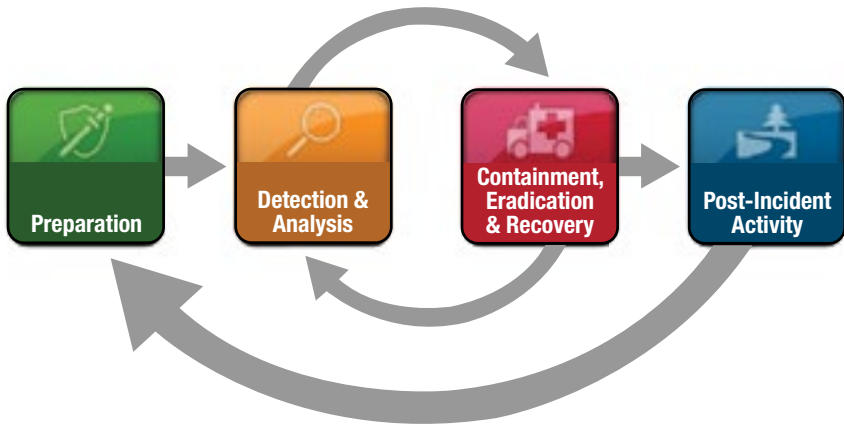


Other very good models exist, including models published by NIST and the SANS Institute. Some published resources and excellent guidance on conducting investigations of cyberattacks are available online, including:

- *Computer Security Incident Handling Guide* at <http://csrc.nist.gov/publications/drafts/800-61-rev2/draft-sp800-61rev2.pdf>
- *Guide to Industrial Control Systems (ICS) Security* at <http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf>
- *Guide to Malware Incident Prevention and Handling* at <http://csrc.nist.gov/publications/nistpubs/800-83/SP800-83.pdf>
- *Guide to Integrating Forensic Techniques into Incident Response* at <http://csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf>
- A wide range of resources at <http://computer-forensics.sans.org>, including *An Incident Handling Process for Small and Medium Businesses* also at http://www.sans.org/reading_room/whitepapers/incident/incident-handling-process-small-medium-businesses_1791

An acronym for the SANS model is PICERL: Preparation, Identification, Containment, Eradication, Recovery, Lessons learned. The NIST model combines multiple steps into four phases, including Preparation; Detection and Analysis; Containment, Eradication and Recovery; and Postincident Activity. The NIST model is depicted in **figure 13**.

FIGURE 13 The Incident Response Life Cycle From NIST SP 800-61



Source: Cichonski, Paul; Tom Millar; Tim Grance; Karen Scarfone; *Computer Security Incident Handling Guide*, National Institute of Standards and Technology (NIST) Special Publication (SP) 800-61, Revision 2, USA, August 2012.

As the NIST model depicts, an IR team might cycle between Detection and Analysis and Containment, Eradication and Recovery several times before the incident is resolved and Postincident Activity can begin. The core part of the investigation occurs in the Detection and Analysis phase, while eradication is primarily a function of the Containment, Eradication and Recovery phase.

Regardless of the model used as a framework for conducting investigations of cyberbreaches, the primary objectives remain the same: answering the basic “who, what, when, where, why and how” questions and, most important, feeding the answers as intelligence into the eradication plan.

The objectives of a cyberbreach investigation are to understand the:

- Attacker’s tactics, techniques and procedures, and, if possible, intentions and use that understanding to inform the eradication plan
- Scope, breadth and depth of the compromise and communicate the information to stakeholders (internal and external):
 - **Activity 1:** Collect a variety of electronic records relevant to the compromise.
 - **Activity 2:** Transform the collected data into information to help focus the investigation.
 - **Activity 3:** Analyze the information to determine the “who, what, when, where, why and how” details of the compromise.
 - **Activity 4:** Develop detailed reporting for stakeholders.

External security incidents or breaches facilitated with commodity malware are likely to be detected by existing enterprise security mechanisms. But it is often seen that enterprises that have been attacked by advanced, persistent and sophisticated attackers have generally been notified of the attack by a federal law enforcement or intelligence organization; enterprises have rarely independently detected the attacks.

Once an enterprise understands that it has been attacked, management will have many questions that the investigation will seek to answer. Chief among them will be: Who has attacked us? When did the attack occur? What did the attackers take from us? Why did they do it? Management will generally be less concerned with how or where the attack occurred, but will expect the investigation team to cover those questions as well to understand the scope, depth and breadth of the compromise.

The investigation team will seek to understand the attacker's tactics, techniques and procedures—or, more simply, how the attack happened—while the eradication team simultaneously plans the eradication event.

The investigation and eradication teams must be given proper guidance to understand their roles, responsibilities and expectations. The plan must include a risk analysis that empowers the teams to make certain decisions quickly.

It is a good idea to prepare detailed tracking processes before beginning the investigation. If the enterprise has case management or ticketing software that already performs the tracking function, this software can be used in the enterprise's plan. However, some investigations require more in-depth knowledge than these systems typically allow. While standard help desk ticket tracking systems are good for assigning various stages of a workflow, an investigation tracking process must allow for collaboration, storage of much technical detail and management of multiple tangential efforts at one time.

From the time alerts are received, no matter the method (IDS alert, user-reported, log analysis, system administrator-reported), they should be tracked in a uniform manner. This allows for long-term trend analysis of causes, attributes and response times. From these trends, a security team can advocate to management for the proper resources it needs to both secure its network and resolve incidents quickly.

There are many Internet resources from respected sources that describe in great detail how to conduct an investigation. For the most part, the approach is dependent on the technology tools available to the enterprise or that the enterprise can quickly acquire and deploy. For instance, an enterprise forensics capability will be needed; if it is not already in place, conducting an investigation will be much more difficult.

Regardless of how the inquiry is conducted, a focus on answering the following questions will lead the enterprise in the correct direction.

3.2 Who Attacked Us?

This is commonly referred to as attribution. The answer contributes to other facets of the investigation. Certain threat actors target specific industries or seek to acquire specific types of information. If Threat Actor X commonly targets financial data, then its presence in the environment guides the investigation toward financial systems. The associated questions that management and the board are likely to ask include:

- Are the attackers affiliated with a specific hacktivist or organized group?
- Are the attackers sponsored by a foreign government?
- Are the attackers sponsored by a competitor?
- Is there more than one group? If there is more than one group, are the groups collaborating?
- Are there multiple individuals collaborating?
- Are the attackers being monitored/watched by any US law enforcement or intelligence entity?
- Are the attackers using a third party (e.g., contractor, client, joint venture) as an attack vector?
- Do the attackers have insider assistance?
- Do they have physical access to our facilities or network? Do they have access in our overseas offices?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- Who notified us of the attack (law enforcement, other external third party or internal person/process)? When did they notify us? How did they notify us?
- Who knows details of the attack (employees, contractors, vendors, clients, law enforcement, others)? Who is responsible for the investigation and remediation activities (e.g., project management office [PMO])?
- Have we retained outside help? Who is aware of our investigation and remediation activities (employees, contractors, vendors, clients, others)?

3.3 What Was Targeted?

Initially, this is a list of systems that the attackers touched. At a minimum, these systems may need to be included in the remediation effort. Knowing which systems were touched may help answer other questions as well. The associated questions that management and the board are likely to ask include:

- What are the attackers after? What are their objectives?
- Have they already achieved their objectives? What data have they taken?
- Have our most sensitive data been compromised?
- What computers are infected with their malware?

- What computers have they accessed (but not compromised)? What accounts have the attackers compromised?
- What network devices (e.g., routers, firewalls, switches, printers) have they accessed or compromised?
- What networks/domains/IP addresses are they using for command and control?
- What third-party systems are they using to shield their command-and-control network?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- What is our liaison relationship with outside threat intelligence sources about the attackers?
- What has been released to the press about the attack?
- What is our plan for public statements if news of the attack leaks to the press or public?

3.4 When Did Various Events Occur?

One of the critical elements of an investigation is a time line, which helps put the event into perspective. An investigation often produces a time line for the entire event as well as granular time lines for each system that was touched. The associated questions that management and the board are likely to ask include:

- When was the first attack activity?
- How long have the attackers been here?
- When was the last attack activity?
- Is there any apparent pattern in the attacks (time of day/day of week/weekends/holidays)?
- Do the attacks coincide with any significant enterprise announcements?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- When will we conduct an initial detailed brief of the board of directors or other stakeholders?
- What is the schedule for follow-up briefings to the board/stakeholders?
- When will we execute our remediation plan to cut off the attackers' access?

3.5 From Where Did the Attacks Come?

The questions to determine attack origination include:

- Was the initial vector a legitimate host on the Internet that was compromised and turned into a "watering hole"?
- Was it a phishing email message supposedly from a legitimate organization with which the enterprise does business?
- Did it start with a potentially compromised host in some other enterprise's environment?

The associated questions that management and the board are likely to ask include:

- From where did the initial attack emanate?
- Where in our environment did the attackers focus the initial attack?
- Where in our environment do they have backdoor access?
- Where else in our environment have they explored (geographic locations, business units, offices, functions)?
- Where did they build repositories?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- Where are we most vulnerable to their attack?
- Do our vulnerabilities represent a heightened risk to our joint venture partners, clients, customers or other third-party business associates?
- Where are our most sensitive data? How well are they protected?
- When was the last update to the critical information inventory? This includes:
 - Executive communications and decision making
 - Human resources (HR) data/legal data
 - Patent/trademark/trade secret data
 - Merger/acquisition/divestiture plans and status
 - Research/engineering/manufacturing data
 - PCI-controlled data
 - Personally Identifiable Information (PII)
 - Other sensitive data

3.6 Why Did They Attack?

Questions to determine the reasons for the attack include:

- Did they seek intellectual property or financial records?
- Were they trying to influence enterprise decisions?
- Was it an effort to gather credentials to enable future attacks?
- Was it designed to facilitate a separate attack on someone else?

The associated questions that management and the board are likely to ask include:

- Do the data holdings of computers they have infected with malware suggest a motive?
- Do the data holdings of computers they have accessed suggest a motive?
- Do the accounts they have compromised suggest a motive?
- Does the timing of the attacks suggest a motive?
- Do the data they have taken suggest a motive?
- Does their network reconnaissance activity suggest a motive?
- Is there a pattern of lateral movement that suggests a motive?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- Why did we fail to detect the initial attack?
- Why did we fail to detect incremental attack events?

3.7 How Did They Get In, Stay In and Get the Data Out?

This is the most technically challenging part of the investigation and often requires experienced investigators with a wide range of skills: network analysis, forensic analysis, memory analysis and reverse engineering, to name a few. The associated questions that management and the board are likely to ask include:

- Did the attackers use email, cloud, social or mobile attack vectors?
- Did they target employee-owned devices or enterprise-owned/-managed assets?
- Did we publish data on our web site that facilitated the attack?
- How did they enable persistent access to our environment?
- How are they conducting reconnaissance?
- How did they escalate from local to domain administrative privileges?
- How are they moving laterally in our environment?
- How did they expand their foothold after the initial attack?
- How are they searching for and finding data of interest to them?
- How are they getting the data out of our environment?

During the investigation, it is important to keep management apprised of the answers to the following questions:

- How can we better defend ourselves from these and other attackers?
- What facts from this event can be used to enhance our employee awareness program?
- Do our employees or executives publish information on social media sites that might facilitate future attacks?

3.8 Other Important Areas to Consider

Besides the six basic questions on which the investigation should focus, there is a list of additional questions (appendix A) that should be answered. The answers to these questions should be provided to management. The investigation team may not be the correct people to answer these questions, but management and the board will most likely ask some or all of the questions eventually. Therefore, the team should find and be ready with the answers.

A good investigation and eradication plan outlines the people, processes and technologies involved in identifying initial indicators of compromise. While a variety of sources exist to identify security breaches, most breaches are initially identified by an external source and made known to the victim through some sort of notification process. Often, the notification can be vague, requiring further investigation to determine which systems have been compromised. For example, if the breach

notification comes from a federal agency, that agency may advise the victim only that there has been a large amount of data sent to a known malicious IP address(es) or domain name(s). Sometimes, the victim may not even know the destination countries, IP addresses or domains.

Other potential sources that may provide IOCs include:

- Outside sources (government, business partner, threat intelligence feeds)
- Security monitoring and response processes
- Automated network and end point analysis capabilities
- Vulnerability scanning follow-up analysis
- Penetration tests
- User-reported events
- Threat intelligence feeds

3.9 On the Quality of Intelligence

It is feasible that in the event of an attack, management may ask only some of the previously identified questions. However, it is also feasible that management may ask the investigation team all the questions in the previous sections. Being prepared to answer them can sometimes be a daunting task. The investigation team must be willing to hold blunt conversations with management. While an investigation of a cyberbreach should not be a “blame game,” management will naturally ask questions about who knew what, when they knew it, and what action they took or failed to take with that knowledge. These questions are inevitable.

The investigation team must be able to answer questions and provide intelligence to management. Leaders must make decisions, sometimes with imperfect information. They must be cognizant of the level of quality of the intelligence on which they are basing decisions.

Answers to the questions listed previously feed into the investigation and into many other security-related programs. Good documentation and coordination are required to capture all of the information, present it to groups inside and outside of the investigation, and, as fully as possible, understand all of the details.

That is a long list of things for the investigation team to address, but when facing today’s advanced, persistent attackers, any gap in the corporate armor can and will be exploited.

3.10 Evidence Handling

Proper handling of evidence is not only a good practice from an enterprise and management standpoint, it will also help to minimize the chances that evidence is excluded during a criminal or civil prosecution.

It is important to be able to demonstrate that each piece of source evidence collected, preserved or analyzed during the course of an investigation has not been altered. One technique of defense attorneys is to call into question whether the evidence presented is complete or has been altered, or if exculpatory evidence has been improperly and deliberately excluded.

There are several techniques and good practices that the investigator can use to minimize the risk of evidence being challenged. The following narrative provides a good summary, but, of course, does not supplant the instructions from the enterprise's legal counsel.

3.10.1 Preservation and Collection Memorandum

It is critical that contemporaneous preservation and collection documentation and memoranda be maintained from the start of the investigation. During the course of a complex investigation, dozens to hundreds of hard drives may be preserved, hundreds of log files may be copied, and thousands of system artifacts may be identified and extracted. Items that should be noted in preservation and collection documentation include:

- The host or person that was the source of each piece of evidence
- The date the evidence was preserved or collected
- The name of the investigator who collected the evidence
- The circumstances and methods of collection or preservation
- The tracking of collected evidence through a numbering or indexing scheme
- Any exceptions encountered during the collection or processing of evidence (e.g., sector read errors, gaps in log files)

3.10.2 Chain of Custody

It is important to secure evidence, particularly original source evidence, properly in a locked evidence room or in a locked drawer to which only the investigator has access. Procedures should be put in place to “check out” evidence each time someone other than the investigator needs it for analysis, or each time that custody of evidence needs to be transferred (e.g., from an enterprise investigator to law enforcement). Each of these custodial transfers should be tracked with the following information:

- Name of transferee
- Name of transferor
- Date and time of transfer
- Method of transfer (handoff vs. delivery service)
- Tracking number (if applicable)

3.10.3 MD5 Hashing

MD5 hashing of evidence is a generally acceptable method of authenticating that a forensic copy of data collected is a forensic match of the original data. This hash value can be generated at the physical device level (in the case of whole disk forensic imaging) or at the individual file level (e.g., copying individual log files). The hash value should be calculated for both the original evidence and the copy at the time of imaging or collection of the data, and any differences between the MD5 values of the original and copy should be addressed at the time of collection. Differences in MD5 hash values can occur if there are read or write errors during imaging, or if the source data change while collection is taking place (e.g., in the case of a “live” acquisition of a running system).

3.10.4 Write Blockers

Write blockers are generally utilized in traditional forensic preservation tasks to prevent the OS platform that the investigator is using from modifying the data on the original or evidence device. Not modifying the original device during imaging may be key if the evidence will later be presented during a court proceeding. Certain OS platforms, such as Microsoft Windows, create a recycle bin or change other metadata on a drive when it is connected to another system running Windows. A hardware write blocker prevents the investigator’s OS from making changes to the original evidence. Other OSs, such as the Helix™-modified version of Linux, do not automatically mount and write to a source/evidence drive. A properly configured Helix instance can be utilized as a software write blocker to prevent modifications to the original drive during acquisition.

In most IR investigations, the ability to collect data from live running systems, including memory dumps and disk images, is needed. In these situations, write blockers are not generally necessary, but it is critical to use a documented and tested methodology that does not change any data that might ultimately be relied on in court.

3.10.5 Reconcile Record Counts

Another method to test for completeness during collection is to reconcile record counts from the original source evidence and the preserved copy of evidence. This is the required method if a utility must be used to extract records such as log entries from a nontext format. Making note of the record count at the time of collection, and reconciling any inconsistencies during collection time (e.g., dips in record counts, unexplained gaps in dates and times of recorded events), can help the investigator become aware of and correct potential completeness problems at the time of collection. Tools such as Robocopy also output record counts, so the investigator can become aware whether the destination count of files copied is less than the source count, potentially due to permission or copy errors.

3.10.6 Attorney-client Privilege or Attorney Work Product Privilege

It should be established early on whether the investigation will be performed under attorney-client privilege with an enterprise's internal general counsel or external counsel from a contracted law firm, or will otherwise be protected as attorney work product. It is strongly recommended to seek legal advice on this matter early in the investigation.

3.10.7 Insurance Claims

Many insurance companies issue policies to enterprises to cover potential losses due to hacking or security breaches. The actual coverage amounts, deductibles and conditions of each policy vary, but, in addition to coverage for the monetary loss due to the breach, the policies often provide reimbursement to the enterprise for the costs incurred to engage external investigators and security consultants to investigate the breach and quantify the losses.

If the enterprise has insurance coverage, the investigator should keep the following points in mind:

- Each investigator should keep a detailed breakdown of hours incurred during the investigation, including the number of hours spent on each investigative task and a brief description of that task. Insurers often require this level of support to approve investigation reimbursement claims.
- It is necessary to quantify the losses incurred by the enterprise. For a treasury or credit card breach, the amount may be fairly straightforward and will include the fines imposed by credit card issuers. For a breach resulting in the theft of intellectual property, the loss amount may be more difficult to calculate because certain losses may not become apparent for months or years.

3.11 Investigating Anonymously

Sometimes the attacker's command-and-control servers host files or provide updates to their malware that expand their foothold or enhance their malware's capabilities to control the target's environment. The target enterprise might recover portions of the attacker's malware from the live memory of a running, compromised system, but might not be able to piece the file back together forensically. While it is tempting to pull a complete copy of the file from the command-and-control server or otherwise profile some of the command-and-control servers, it is important to remember to perform all network activity that directly touches an adversary command-and-control server (e.g., interfaces, scans, pings, curls, wget) from a machine that is not infected and is located outside and unrelated to the enterprise being attacked. If this is not handled carefully, it may alert the attackers that the enterprise is aware of their presence.

3.12 Safeguarding the Investigative Actions

There are defensive precautionary tactics that must be taken to prevent the investigation itself from eventually becoming part of the compromise.

3.12.1 Data

Investigations are often done on machines on which day-to-day work is performed. This can be an issue, depending on the level of compromise that has occurred. If the investigator is working off domain-joined machines and the threat actor has already extracted the password hashes of all of the employees in the environment, the box is at risk. Whenever possible, the attempt should be made to use software and hardware that function independently of the current environment.

3.12.2 Data in Motion

Most email in the current environment is reliant on domain-joined architecture for the authentication process. An attacker who compromises the domain has access to the entire enterprise's password hashes and, ultimately, everyone's email. As such, communication should take place outside of the realm of email. The team can leverage a content management system for the purposes of the investigation. If email is required, all communication should be encrypted, using at least AES-256 encryption.

A quick note about encryption: Not all password protection schemes are the same. For example, while the Microsoft Office® password protection options are seemingly secure, they have been susceptible to cracking attempts as well as to authentication bypass mechanisms, thus it is recommended that for the purpose of protecting sensitive investigative data, and use of PGP, TrueCrypt® or some other form of encryption should be considered.

3.13 Protecting the Investigation

After a cyberattack has occurred and it is clear that the network is compromised, it is understandable that there will be a strong sense of urgency to start the investigation. But an enterprisewide investigation is a formidable task and often requires input from a variety of teams across multiple geographic or business-unit lines. Communication, data collection and investigative tasks must be done quietly, without the adversary becoming aware of the activities or, even worse, undermining the entire investigation. If the adversary knows that the enterprise has been alerted to its presence before the enterprise is ready to eradicate it from the enterprise environment, the adversary may take action that is damaging to the enterprise. For example, the adversary may install additional command-and-control techniques that the enterprise has not yet detected so that when the enterprise blocks the command-and-control systems it has detected, the adversary can still maintain access. It is not uncommon for attackers install back-up command-and-control capabilities and to program those capabilities to “sleep” for a period of time to reduce the likelihood that the investigation team will find them. Attackers are smart and creative and should not be underestimated.

3.13.1 Credential Protection

The investigation team must log on to various machines in the environment at different points during the investigation. Whether collecting images, pulling log files or deploying agents, investigators must take care of their credentials.

It is important that any credentials used or created in the investigation environment differ from those that are normally used. If possible, reliance should not be placed on a pre-existing architecture unless there is 100 percent certainty that it is not compromised. Maintaining separation is a critical safety precaution.

Without in-depth monitoring and investigation controls, it is hard to determine the velocity with which the malicious actor is moving. Too often, in their haste to gather information or respond to an incident, enterprises neglect to protect domain account credentials. These accounts need to be powerful enough to move about the environment and, as such, are granted administrator access to a variety of hosts. A powerful account in the hands of an inexperienced or careless IR handler can result in the attacker gaining the IR handler's credentials. An attacker would need only an administrative account on the machine to access the hashes and, ultimately, the cached credentials of the investigator. In many enterprises, built-in administrator accounts may often share the same password. Simple issues like this can undermine the entire IR investigation.

Through all the investigative activity, investigators must keep management informed of the answers to the key questions. Also, more important, they must feed insight and intelligence to the eradication plan, which should be designed to remove the attackers from the environment.

Chapter 4. Eradication

It is important to understand that, while investigating a sophisticated cybersecurity breach may seem like a marathon, effective eradication must be a sprint. That is, efficiency in the eradication effort is critical to success. Many investigations take months to determine the complete scope of an attack, identify all the systems the attackers have accessed or compromised, list all the credentials they have stolen, find all the malicious software they have embedded in the enterprise, categorize all the sensitive data they have taken, and make preparations for an effective eradication event. But, once the investigation is finished, the eradication event must aim for completeness and speed.

Effective eradication plans must be executed with speed and precision because attackers often try to reestablish a base and then entrench themselves again into the network once they sense they have been discovered and eradication is underway. For example, during one eradication event, a threat actor realized that its access to a privileged account had been removed. The attacker then accessed a web shell it had installed on a demilitarized zone (DMZ) server nearly a year earlier that had gone undiscovered by the investigation team, and used the web shell to regain access to the environment. Attackers have also been known to simply register new IP addresses for their domain names once they suspect that eradication teams have blocked their IP addresses.

Too often, enterprises rush into eradication activities by blocking an infection vector or point of persistence as soon as it is discovered, before the full scope of the compromise is understood. This approach leaves the enterprise constantly chasing the adversary as the attacker's tactics change in answer to the enterprise's poorly planned response.

This can be both expensive and exhausting for the enterprise. Instead, eradication must be a coordinated effort to systematically remove the presence and persistent access of an advanced threat to the enterprise's environment. Investing the time and effort to carefully prepare the enterprise for an inevitable cybersecurity breach is a due diligence exercise that generates ample return on investment (ROI).

4.1 Plan for Eradication

A coordinated, well-planned eradication event must consider operational details (e.g., work schedule, timing).

4.1.1 Create the Eradication Event Team

Preparation for an eradication event starts with the enterprise identifying a team to conduct the various components of the effort. These preparation efforts should start

during the investigation period so that the eradication event's execution can begin soon after the investigation is completed. A dedicated team leader functions as the eradication plan coordinator and remains responsible for overseeing the execution of eradication activities. Team members should be appointed from among various groups within the enterprise. Each team member plays a critical role in one or more components of the eradication plan and may be responsible for certain aspects of the plan execution.

In the days leading up to eradication, resources need to be dedicated full-time in preparation for the event. On the day of eradication, members of the eradication team should focus their attention only on measuring the efficacy of the actions taken and react to any new activity that is generated by the eradication effort. The eradication team needs to plan its contingency scenarios and be prepared to adapt to the attacker's reactions to eradication. A comprehensive and methodical plan of engagement should be distributed to the team and key stakeholders at this time. It is a good idea to execute a tabletop exercise of the plan, including all of the steps from the beginning to the lessons learned session. This helps to develop answers to key questions and clarify ambiguities in the plan.

The eradication event often requires team members to adjust work schedules—providing 24/7 support—to achieve swift execution of the plan, beginning with cutting off an attacker's access. The approach an enterprise takes to achieve this depends heavily on the scope of the eradication event and the size and skill set of the team.

4.1.2 Develop the Eradication Event Plan

Consider the following when developing the eradication event plan:

- **The eradication plan must be comprehensive and synchronized with the investigation plan.** After the investigation and eradication teams have been assembled, each team needs to understand its roles and responsibilities along with the objectives and tasks that must be accomplished to achieve eradication. Leaders should develop a written investigation and eradication plan, obtain buy-in and socialize the plan with other key stakeholders. It is vital that the plan outline the relationship between the investigation team and the eradication team, as well as other business units. A good plan allows investigation and eradication teams to work in sync during each phase of the process, helping to achieve an effective, efficient and comprehensive investigation and eradication event.
- **Determine tasks that are necessary to complete before and after the eradication event.** A comprehensive plan includes a wide variety of items that need to be investigated; processes that need to be executed; technical fixes that need to be implemented; and functional capabilities that need to be optimized, developed or purchased. The various aspects of the plan must be prioritized based on the nature of the compromise. For example, one enterprise may determine that implementing application whitelisting on domain controllers is essential prior to eradication.

Another enterprise may instead focus on removing LAN manager password hashes from the environment prior to remediation because the attacker has been exploiting this encryption mechanism. Regardless of the situation, the enterprise must determine which items must be completed prior to eradication and which items can wait until after the eradication event. The enterprise should conduct a risk analysis based on the threat, the likelihood of exploitation and the enterprise's current security posture.

- **Extend eradication event activities beyond the initial event.** Planning should not end at eradication. Once the advanced attacker detects that eradication is underway, it will attempt to regain access. Vigilance should be maintained during the few weeks after eradication and elevated operations continued for at least a week after the eradication event. At that point, it is appropriate for the team to consider standing down all of the eradication teams.

4.1.3 Determine the Eradication Event Date

Selecting the appropriate date and time for an eradication event is important. It can be tempting to schedule the event around the eradication team's availability and/or network downtime. While these are important factors, it is also important to consider the attacker's behavior patterns. For example, analysis of the attacker's behavior may show that the attacker works according to a certain schedule, including taking weekends off and working only during standard business hours. Alternatively, attack patterns may demonstrate that the attacker is typically more active on weekends. In such a case, choosing Monday morning to begin the eradication event may give the eradication team the best possible chance of executing the full eradication plan before the attacker realizes that access has been lost. If eradication occurs when the attacker is known to be active, the attacker will likely actively fight to regain access.

Advanced, persistent attackers often fight to maintain their presence within an environment, whereas a stern warning of detection may thwart more novice attackers. When detected and blocked, advanced attackers may act quickly and aggressively to regain their access to a network that they have spent months or years working to penetrate. If an attacker has been predictable in access days/times, this can be used to the eradication team's advantage. Knowledge of the attacker should serve as important input when determining when to execute the eradication event.

4.1.4 Know the Attacker's Techniques, Tactics and Procedures

An enterprise must understand the attacker not only to improve defenses, but also to help enable efficient, comprehensive investigation and eradication. Targeted threat actors often create and maintain many points of presence and vectors of entry into an enterprise's network. Targeted threats have been known to compromise user and privileged accounts, infect hosts with backdoors, leverage web shells, and do much more to maintain persistence. These access vectors are cultivated for months or even

years. Once established, they may not be used until the attacker's primary access is cut off. For this reason, an investigation and eradication plan should be coordinated and holistic; if not, the risk that threat actors will use these alternate methods of access to reestablish their presence on the network is greatly increased.

4.1.5 Establish Communication Protocols

During the eradication event, it is important to establish clear lines of communication within the team to help promote timely attention to status updates and open issues. If the team is divided into around-the-clock shifts, the team and all relevant stakeholders should receive these schedules and contact information. Scheduled updates to leadership will help limit requests for *ad hoc* updates. Each team should keep a log to make the periodic updates easy to construct, keep team members on all shifts informed about previous actions, and simplify the process of creating reports on lessons learned and overall eradication.

Both top-down and bottom-up communication is essential to the success of the investigation and eradication efforts. The teams must keep key stakeholders and leadership informed of the plan far in advance so that authorization and approval are in place before the eradication event. Likewise, stakeholders and leadership must have insight into some details of the plans so that they can better understand the impact the necessary changes will have on the broader enterprise. Some of these changes may involve critical organizational processes or key technologies. Integrating all levels helps identify essential processes and technologies that may impact organizational objectives and operations throughout investigation and eradication.

Preparation to address any issues that may arise during investigation and eradication can be severely compromised if the investigation and eradication teams operate in isolation. Integrating the teams with the relevant IT department resources improves effectiveness and efficiency, especially when quick decisions and actions are critical to the investigation. The plan should develop processes on how the help desk, workstation, network and security teams will interact. It is worthwhile to define in advance the appropriate timing and procedures to turn over malware-infected systems to the workstation team for cleaning. The workstation should be kept informed during the investigation so that it can anticipate when its services are required.

Plans for large-scale investigation and eradication events should define how the security team integrates within IT on a daily basis. Having daily operational status meetings, during which the security team briefs IT teams and leaders, helps close out investigations quickly and return systems to production. Some questions the plans should answer include:

- Will the help desk get all IT-related calls, or will the security operations center (SOC) deal directly with users?

- When there is a disagreement between the security team and the network team regarding a course of action, who has the authority to decide the next step?

These are important questions that should be addressed in advance. Some enterprises choose to embed with the security team (on a part-time basis) members from the help desk, workstation team or network team to help with investigations. An *ad hoc* integration can bog down the time-critical security investigation and should be avoided when possible through clearly-defined integration roles and processes.

Part of planning for communication involves determining who needs to know about the investigation—a decision that cannot be taken lightly. At some point, not at the onset, it may become necessary to tell certain parts of the enterprise something about what is happening. A plan that keeps critical information limited to a select group is a leading practice. The plan should dictate who has the right to know and who does not.

External communication is a much more sensitive topic. If word of the attack leaks out, direct communication with media outlets may be unavoidable and should, therefore, be planned for in advance. While media publicity is typically a reactive, rather than proactive, step, a plan must be developed to consider how the enterprise will converse with the media should they become aware of the situation. Advance preparation of talking points for approved employees is a good first step.

In some cases, it may be beneficial or even essential to proactively reach out to key business constituents. Enterprises that are linked by open static connections should have an agreement in place that requires them to communicate openly and honestly about their network's health and well-being. This is not common, but it should be. When these linkages are established, the two networks essentially become one and are dependent on the other to stay clean. Additionally, if one enterprise's network becomes compromised in a way that may impact business or supply chain agreements, it may be prudent to alert business partners as early as possible. This can help in avoiding disruptions further down the line. Also, some business partners may have detection capabilities that can shed further light on the threat if it exists on their network as well. However, none of this intelligence sharing can occur without processes being put in place to communicate across enterprises.

Other third-party entities that may need to be contacted (if they do not reach out first) are the government and law enforcement. Every enterprise must be familiar with reporting requirements based on the type of compromise sustained.

4.1.6 Establish a “War Room”

It is a good idea to set up a “war room” during the advanced IR process. Because a custom response plan typically lasts for only a short time, the war room is a temporary measure and will be in use for perhaps no more than a week. It is the team’s primary meeting and collaboration space, where all relevant parties (incident responders, IT staff representatives, stakeholders and other leaders) assemble.

Using the eradication plan as the guide, the team leader serves as the conductor. He/she facilitates response discussion, coordination and selected actions; coordinates the timing of team members’ actions; and keeps everything moving in concert.

There may be a need for multiple war rooms, depending on the size of the team and the scope of the project. For example, a leadership room could be used to coordinate all of the actions, with smaller security monitoring, network and workstation teams implementing the plan in another room. All actions within each team should be logged to allow later shifts to see what happened and to feed into lessons learned after the project is complete.

If follow-up actions are required after the war room is disbanded, it can be reconvened, with the security team leader making that decision as necessary.

4.1.7 Establish Secure Communications and Information Sharing Mechanism(s)

Due to the composition of the eradication team and the nature of the eradication effort, secure communication plays a vital role. Team members should be briefed before, during and after the eradication efforts. Much communication in today’s enterprises flows through digital networks, which may be monitored by an attacker that has compromised the network. Therefore, it is important to establish and strictly adhere to operational security guidelines.

Operational security guidelines should dictate that information regarding the incident be held by only those who need to know. This may vary by enterprise, but often includes the eradication team, the investigation team, some aspects of IT operations and select management. Additionally, those who have the need to know must keep the information confidential, lest knowledge of the breach get into the wrong hands or inadvertently be released publicly. While legal requirements may stipulate that certain breaches must be disclosed, the eradication team should not make this determination without consulting senior management, legal counsel and any other pertinent players.

The teams must operate under the assumption that the attacker can monitor the network. Targeted attackers have been known to use this ability to search for signs that their presence has been discovered. If attackers sense they are being investigated, they

may change tactics or disconnect for a long period, knowing they can wait for weeks or months before reentering the enterprise through one of the many access vectors they previously put in place. To avoid discovery, the investigation and eradication project should be given a code word. Additionally, the eradication team must leverage strong encryption for in-band communication (e.g., email, file sharing). A separate set of encryption and passwords should be put in place for team communication, and team members should be frequently reminded of the sensitivity of the project and the care that should be taken with their communications. Whenever possible, the operations security guidelines should identify a preference for out-of-band communication (e.g., in-person meetings, phone calls, Short Message Service [SMS]).

To achieve swift, flawless execution of eradication, the team must keep key stakeholders and decision makers informed of the plan far in advance so that authorization and approval are in place before eradication. Management needs to be involved with the details of the plan so that it understands the impact of the changes that need to happen. Some of these changes may affect critical organizational processes or key technologies. Management's involvement helps identify essential processes and technologies that may impact organizational objectives and operations. While these conversations and decisions may be difficult, it is important that the plan be agreed on prior to eradication day.

Before an advanced investigation or eradication can occur, a well-thought-out communication plan is imperative. The team needs to decide when to communicate, with whom and how to do so securely. Notifications of advanced intrusions on the network are often carried out through a telephone call. The first step is to find a point of contact in the enterprise. The caller should not provide information directly, but should instead set up a face-to-face meeting. The data related to the intrusion itself are often fairly old; however, the time it takes to identify affected organizations and make initial contact is not necessarily short.

4.2 Execute the Plan

While executing the plan:

- **Keep team members focused.** During the eradication event, keep the dedicated team members isolated from the broader enterprise to minimize distractions. Some may want to join in and “help” the team, but because they have not been a part of the preparations, they could make a mistake that derails the operation. It is best to politely ask these employees to give the eradication team some privacy.
- **Appropriately label information in communications.** When concerns arise, they should be elevated quickly. It is imperative to label the information being passed along. Clear and appropriate information labels help decision makers make the right call. If someone's hunch is presented as a fact and later is found to be untrue, there could be negative consequences for the individual and the eradication effort.

- **Stay true to the eradication plan.** Team members may get anxious during eradication and want to do everything at once. This urge must be resisted. Team members must trust the plan. If the optimal time and energy have been invested in the preparations and investigation, eradication will be a concerted effort. Deviations from the plan are bound to happen, but every effort should be made to limit them. Any major deviation should be coordinated through the eradication team lead before being undertaken and then communicated across the team as it is being carried out.
- **Tight coordination matters.** Coordination is essential. For example, if an attacker compromises an Active Directory server and obtain domain administrator credentials for several accounts, changing those passwords must be coordinated to prevent attackers from simply using another account if only some of the compromised passwords are changed. If the attacker maintains access, the attacker may be able to capture the hash of the new passwords or possibly even create a different account that has the privileges needed.

If possible, the best approach is to isolate the environment during execution of eradication, but that is usually not feasible, especially if the attacker is well entrenched and has a large footprint in the enterprise's network. If the network can be disconnected while eradication occurs, the odds of a successful eradication increase.

4.2.1 Execute a Password Change

An enterprisewide password change is a standard component of any advanced breach response. Passwords are commonly stolen by attackers and used to gain access to assets while blending undetected into normal network traffic behaviors. Attackers also can use the passwords to regain access, often at the administrative level, when their other tools and malware are detected and removed. For these reasons, password management and comprehensive resets should be a major step in eradication of an advanced attacker's presence.

However, an enterprisewide password change might prove too challenging for business reasons. In that case, at a minimum, any account that has been identified as compromised or suspicious by the investigation team should be disabled or have the password changed simultaneously on eradication. Rules should be placed in the SIEM or other monitoring and search tools to identify failed logon attempts from these accounts. This rule may identify attempts to reenter the network after eradication and should be treated as high-priority alerts.

Prior to eradication day, it is important to establish and maintain a baseline level of security for the domain that has been compromised to prevent easy reentry.

Baseline-level activity should include the following activities:

- **Conduct a thorough account inventory.** All accounts that are used in the domain should be inventoried. Each account should be tied back to an individual who

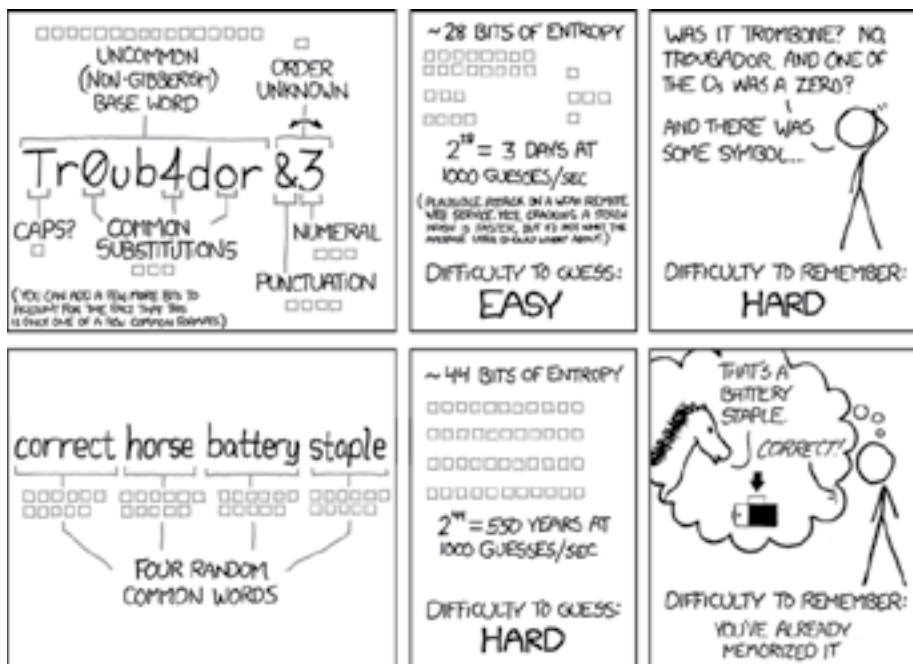
currently needs access. The inventory should also include checking for last changed password dates as well as the algorithm used to store the password.

At the conclusion of the account inventory, all of the accounts that need to be deleted or that require a password change should be incorporated into the eradication plan. This is one of the coordinated events that should happen simultaneously with other eradication plans. An unexpected password change event before the eradication could warn a threat actor that it has been detected and thus could complicate an ongoing investigation. Whenever possible, everything should be overhauled at once rather than addressing any part of the environment individually.

- **Review, rationalize and, where possible, reduce the number of domain administrator accounts.** Domain administrator accounts are the “keys to the castle” and should be guarded as such. The number of users with this privilege should be tightly restricted, and the activity of these users should be logged as thoroughly as possible.
- **Update password policy**—Password policy should be reviewed and updated, if needed. Windows passwords should be at least 15 characters to prevent LAN manager hashing from being used in the environment. There are a number of tools that can be used to evaluate the password hash database and identify passwords that are easily guessable or not long enough.
- **Enforce password strength**—After resetting the passwords using appropriate password strength (**figure 14**), a process should be implemented to maintain them properly. Passwords should be reset every quarter for standard users and more often for administrators. In addition, it is a good idea to establish account lockouts. For example, after users fail to enter their passwords a maximum of five times, their accounts should be locked until their passwords can be reset by the help desk. This will prevent an attacker from using a brute force attack to find the passwords on the network.

If room can be found within the budget, it is a leading practice to implement authentication that uses two or more factors. Two-factor authentication means having two of the following: something the user knows (e.g., password or personal identification number [PIN]), something the user has (e.g., some sort of token or smart card), or something the user is (e.g., fingerprint, retina scan). Multifactor authentication removes many of the weaknesses associated with passwords, but it can be expensive to implement and maintain.

FIGURE 14 *xkcd* Comic: “Password Strength”



THROUGH 20 YEARS OF EFFORT, WE'VE SUCCESSFULLY TRAINED EVERYONE TO USE PASSWORDS THAT ARE HARD FOR HUMANS TO REMEMBER, BUT EASY FOR COMPUTERS TO GUESS.

Source: *xkcd: A Webcomic of Romance, Sarcasm, Math, and Language*, <http://xkcd.com/936/>.
 Used with permission (<http://xkcd.com/license.html>).

4.2.2 Block Attacker Command and Control

During the investigation, the forensic keys section of the report should detail all of the communication paths that are being used by the threat actors. All IP addresses that receive or send command-and-control communications should be outlined, and all domain names and URLs used to resolve these IP addresses should be listed. Other information that is used to direct communication, such as Internet anonymizing proxy sources or URL obfuscating/shortening services, should be detailed. On eradication day, all of these URLs, networks and IP addresses must be blocked at egress points. The firewall should be used to block the IP addresses; however, if the attackers are utilizing name server resolution of their IP space, then “black holes” should be created for the URLs and domain names. They should be placed in the internal DNS resolver, pointing to one of the following:

- A loopback address or 0.0.0.0
- A local, internal web server configured to “answer” all image requests mapped to a single-pixel image file, and all text pages mapped to a warning message page
- A machine with a personal firewall

If the IP addresses are blocked but malware is still configured to talk to a fully qualified domain name (FQDN), then the attacker can change to a different IP address and update its name server to point to a different IP address, thus bypassing the block without having to update its malware.

The last two configurations—local, internal web server configured to answer, or machine with personal firewall—have the added advantage of generating log files for inspection as well as enabling the Snort® or other IDS system to continue to see traffic. To enforce successful DNS black holes, DNS traffic that is not being routed through the DNS resolver needs to be blocked. If DNS is not funneled through an enterprise’s owned and controlled name servers, then the malware can be programmed to use external, public DNS resolvers to beacon or exfiltrate data.

It is a good idea to block servers that provide services internally from communicating outside of the network, except to a specified whitelisted set of networks. For example, if the servers need to communicate with software vendors for patches or security definitions, the vendors should be placed on a whitelist and communication allowed.

4.2.3 Rebuild Compromised Systems

Finally, it is time to eradicate the threat. Threat intelligence can help identify the source of the breach, which can aid in the timing of the eradication. For example, if the threat is based in China, the attackers may not be very active during Chinese New Year celebrations, offering an opportune time to eradicate and implement preventive measures. If eradication occurs when the attacker is known to be active, the attacker will likely fight to regain access. There is no reason to give the attacker a fair fight, so it makes sense to act when the attacker’s guard is temporarily lowered.

If possible, all hosts that have been compromised should be taken off of the domain and recreated. Evidence showing that an advanced threat actor has compromised a box should be taken seriously. An advanced threat actor could have installed a rootkit that might not be detected by antivirus or system scanning. If possible, these systems should be reimaged to be sure that the attacker is not maintaining a foothold in the system. If resources allow, it is better to recreate the compromised hosts in advance so that the hosts can be simply “swapped out” on eradication day. This allows eradication to happen faster and in a more coordinated fashion.

Careful consideration should be taken in deciding whether to inform the user that his/her system has been compromised. While coordinating the logistics of cleaning the compromised systems, the eradication team should consider options to mitigate the risk that the user or others involved in the system cleaning process will interfere with the eradication effort.

4.2.4 Submit Malware to Antivirus Vendors

The investigation team should identify and archive all tools used by the threat actor as well as the malware that was deployed or stored on the network. The team should then analyze these malware artifacts, at least dynamically, to discern the attacker's intent. The eradication team should work with antivirus vendors and host-based intrusion detection vendors to build signatures for this code. On eradication day, definition files should be deployed throughout the enterprise to detect and quarantine any other existing copies/variants of this code.

4.3 Monitor for Attempted Reentry

Prepare to operate in a heightened level of monitoring. Once eradication day nears, it is time to take a step back and consider contingencies. It is entirely possible that the actions planned may not have the desired effect or will have an adverse second-order effect. The eradication team must think about the following:

- **What are the second-order effects of the eradication actions and the future prevention actions?** One way to address this is by making a conscious effort to consider these hidden consequences as planning nears an end. Another method is to bring on a new team member with a fresh perspective to evaluate the eradication plan.
- **What assumptions built into the plan could prove wrong? What would the consequences be?** A comprehensive overview can help clear up many of these assumptions. For the rest, it is best to have a contingency plan in place if needed to keep the eradication plan on track. The adversary may have new tactics in store. Threat intelligence may help to identify and anticipate such tactics.
- **Did the investigation team find enough IOCs to cover all of the attacker's activity? Was the coordinated effort done quickly enough to deny the attackers access before they could establish a second access method elsewhere in the enterprise?** The most challenging task of eradicating a threat is determining whether the effort was successful. It can be unsettling to realize that it may be impossible to say for certain that the network is secure.

Once the monitoring team is reasonably comfortable that all IOCs and artifacts have been eradicated, the SOC needs to brief management that eradication execution was completed successfully and note any exceptions and findings from the validation effort.

After an enterprise detects an advanced threat in its network, it should monitor and watch behavior to understand the attacker's tactics and identify mechanisms for persistence. The enterprise must make every attempt to discover all infection vectors, points of presence, persistence mechanisms, targeted vulnerabilities, and other methods that may allow the attacker back into the network. A single mistake

or missed item could allow the attacker back into the environment without the enterprise's knowledge. To help avoid such a catastrophic mistake, an enterprise must focus on:

- Developing quality investigation report documentation to feed eradication efforts
- Proper planning to achieve complete coordination and eradication
- Adherence to sound project management principles throughout the entire process

An effective eradication plan addresses all of the individual findings in the investigation report and has a clear start date with specific objective target dates.

The eradication team needs to be able to monitor all of the activities that are identified by the investigation report before execution of the eradication plan can begin. This monitoring should continue for several weeks after eradication day to watch for any recurring activity. Specific use cases should be built, tested and tuned to feed into an SIEM solution that alerts designated members of the eradication team so that the efficacy of the eradication can be logged, tracked and monitored.

Page intentionally left blank

Chapter 5. Post-eradication

The process of eradication can be an exhausting sprint. Often, the eradication event must take place after normal business hours or during a weekend maintenance window to minimize the impact on operations. Members of the SOC accomplish a challenging series of time-sensitive tasks cross-functionally with other IT functional areas. Immediately after the successful completion of the eradication checklist, the SOC must continue to operate at a heightened level of sensitivity because a renewed attack is most likely to occur in the weeks after the eradication event. An attacker will try to reestablish itself, and when the initial tactics are blocked, the attacker will rapidly change strategy. This will happen until the attacker is successful, or until the attacker decides that the effort to regain access exceeds the value of the information on the network or the value of the time spent building access. Depending on the severity of the initial infection or the value of an enterprise's secrets, this formula can be unsettling. SOC analysts need to validate the efficacy of the controls put in place during the eradication event, monitor for these events and react quickly.

5.1 Validate Eradication Activities

5.1.1 Maintain a Heightened State of Alert

If planned properly, eradication execution can be done very quickly, but monitoring needs to be in place and rolling forward for weeks or even months to maintain vigilance over the network against an advanced attacker. Network situational awareness is most important during this phase. NetFlow data should be generated and analyzed for changes in trends, new activity that is not aligned with the baseline and so on. All use cases for attacker tactics, techniques and procedures should be programmed into SIEM, and SIEM should be monitored actively for some time after eradication, specifically for alerts that match the attacker's threat profile.

As mentioned previously, advanced attackers will try to get back into the network through all of the methods at their disposal. They will also come back knowing that they are being hunted and their existing tactics, techniques and procedures have been discovered. This knowledge can lead an advanced attacker to become aggressive and unconcerned about being noticed. The attacker may have spent the previous months using very low-profile methods to evade detection, but now it knows that does not matter. In light of this, the attacker may try some desperate methods to get what it can before it is thrown out for good. Attackers may attempt to use a brute force attack to find passwords or scan network segments, tactics that generate many logged events and are easily detected as attacks by even default rule sets of SIEM solutions. If these techniques are detected, expeditious response is the most important thing.

Management must decide how long to leave active monitoring in place; this will likely depend on how aggressive the attacker is after eradication. Even after the active monitoring phase concludes, the rules that were put into SIEM should not be removed because advanced attackers are interested in the long-term strategic level of network access. They know that sometimes the most successful strategy, after they have been kicked off the network, is to disconnect for weeks or even months.

A major task for the SOC in a post-eradication environment is to rescan the environment for all of the IOCs identified in the investigation. The SOC team should be familiar with these IOC artifacts. The team should use its host-based awareness capabilities to check each of the infected machines to verify that they are free of malware on both disk and memory. It should use its network-based capabilities to identify any traffic that is going to or from any of the IP addresses used by any malicious actors in the incident. The team should also resolve again any FQDNs that were used from inside the network and make sure that they point to the designated black hole. Additionally, they should frequently resolve those FQDNs outside of the network to learn whether the IP addresses have been updated. Many attackers do not hard-code IP addresses into beaconing malware. Instead, they use URLs so that when an IP address is blocked, they can update the DNS entry and bypass the block. If the SOC periodically checks the FQDN resolution from outside the network and detects that these DNS records have been updated, it should add the new IP addresses to the threat intelligence database and block those as well.

At this point, the SOC is in monitoring mode, but now it is with special purpose. The SOC needs to remain in a heightened state of awareness for a time because as soon as the attackers realize that they have been kicked out, they are likely to try to regain access to the network until they are successful or exhausted. It is the SOC's job to exhaust them with successful monitoring and action. This can be a challenge because the SOC is in an interesting paradox with respect to intelligence:

- SOC analysts know that the attacker that has invested months or years into establishing, growing and maintaining access to the network that has just been disconnected, and they know how to detect its attacks because they have been watching and learning and have developed a series of IOCs.
- The attacker wants to regain access to the network, and it now knows that what it was doing before no longer works, so it must change its tactics, tools and procedures.

This puts an SOC in a difficult situation because it has to prepare for an imminent attack that likely will not match any of the alerts that the security monitoring systems have been set to detect. It is still important for the SOC to monitor its devices closely and prepare to act quickly for at least two weeks after eradication day.

The attacker will most likely notice first that its malware is no longer checking into the command-and-control server, and it may try to probe a bit to troubleshoot its access. It may then discover that its IP addresses are blocked at the firewall. It is important to study the firewall logs to understand the DENY events. These events will point to the location the attacker is trying to get back to, which could lead to the discovery of previously unidentified compromised systems with backdoors.

It is a common investment for advanced attackers to compromise systems and then treat them as “plan B,” to be used if they ever lose their primary activity. These “plan B” compromised systems can be extremely hard to find in an investigation. Because they are not being used by the attacker, network analysis and forensics do not point in that direction—but they should not be overlooked. For example, when compromised internal hosts stop checking in to the command-and-control server, the attacker may try to access a web shell that was placed on a web server months earlier. It will discover that it no longer has access to any part of the network because the firewall is tuned to block the attacker’s source IP address.

At that point, the SOC team is in a race with the attackers. If the team can understand this alert and determine that the web server is compromised before the attackers try to route their traffic through another location, the SOC wins. If the attackers route their traffic through another location or are in some way able to change their IP addresses to an unblocked addresses and they can get to their hidden backdoor, then the SOC may lose this battle. Depending on the architecture, the entire process may need to be restarted. The SOC analysts and the security team have a tremendous responsibility. Without a diligent, thoughtful, proactive approach to post-eradication monitoring, the attacker can quickly undo the large-scale and expensive eradication project.

5.1.2 Validate Controls

After the eradication and network hardening events have occurred, the enterprise should validate that the controls and blocks put into place enterprisewide will help alert the IR team to similar advanced attacks. An effective method to accomplish this is through an outside assessment, in which a team (with the permission of the network owner) attempts to hack into an enterprise’s network, exposing potentially unknown vulnerabilities. The enterprise should give the team a methodology that is similar to the recently repelled attacker’s as well as “flags” to capture on systems that the enterprise wants to evaluate explicitly. Even with these constraints, the team should be allowed the latitude to attempt other methods within the spirit of the assessment.

These tests on the newly implemented eradication controls are essential for a variety of reasons. They help the enterprise discover where its protection is weak. They show how well key assets and data are defended, and they also help present management with a realistic picture of the current state of the network. A side benefit is that the testing, depending on its thoroughness, may find evidence of a breach.

After conducting the initial assessment on the new controls, medium- and large-sized enterprises should have a controls audit or penetration test conducted on their network at least annually. This audit or test may be performed by a staff member (if he/she is capable) or, preferably, an outside entity. The time it takes to complete a penetration test can depend on the size of the enterprise, the goal of the test and the knowledge provided to the tester.

Possible goals of a security controls assessment in which the test team has been given no information about the network, except the enterprise's name, include:

- Determining whether testers can obtain access to a key system
- Determining whether testers can exfiltrate or steal a key piece of data
- Evaluating the enterprise's susceptibility to a social engineering attack
- Evaluating, in the words of Tom Kellerman, vice president of cybersecurity for Trend Micro, "... from the perspective of the database [or other system] that was [already] compromised to see where else the adversary could have moved laterally within the system and deposited back doors."²¹

Once the test report is available, the enterprise should make every effort to shore up the deficiencies. Because it is already proven that these weaknesses can be discovered and exploited, immediate action should be taken.

5.2 Brief Stakeholders

Once the team is reasonably certain that all IOCs and artifacts have been eradicated and the team has validated that enterprisewide eradication controls are in place, the SOC needs to brief management. The SOC should report that eradication execution was completed successfully and note any exceptions and findings from the validation effort.

Briefings to stakeholders about the results should be well planned and conducted soon after the event. An initial, high-level communication can be issued within one day of the event, followed by a deeper explanation of the activities that took place. Communication should be clear, concise and focused on problem resolution. It should clearly identify any gaps that remain and proposed efforts to mitigate them. It is also important that communications demonstrate lessons learned and identify any new processes created as a result of the effort that will make similar eradication events or activities more efficient and less taxing on the enterprise.

²¹ Homeland Security News Wire, "South Carolina Exploring Different Cybersecurity Plans," 13 November 2012, <http://www.homelandsecuritynewswire.com/dr20121113-south-carolina-exploring-different-cybersecurity-plans>

5.3 Lessons Learned

An essential part of post-eradication is conducting lessons-learned sessions. This should be viewed as an ongoing process through which team members can collaborate and learn from experiences. Establishing a formal lessons-learned program, with a clear leader who is responsible for conducting the meetings and following up on action items, will help the enterprise excel at learning from previous mistakes, incidents and experiences. Some enterprises choose to conduct lessons-learned sessions only during the post-eradication phase; others find it beneficial to conduct more regular sessions throughout the investigation and eradication as well.

To facilitate the lessons-learned program, all discussions and decisions conducted during the eradication event must be well documented. Documentation and collaboration enable lessons learned to be shared after the event has ended. It is important to gather postresponse feedback after situations that require quick decisions. Examples of points to document include:

- What went well?
- What could be improved? (Consider people, processes and technology related to security.)
- Could any unforeseen events have been prevented? If so, how?
- What are the immediate follow-up actions?

After eradication activities are complete, the participants should meet for a debriefing to share ideas on things that worked well and things that could be improved. The observations presented, along with lessons learned documented by the scribe during the war room activities, are captured in an after-action report. **Figure 15** illustrates an example after-action evaluation form.

While a developing a form does not solve the problems that a lessons-learned program aims to correct, it is a good step in establishing a predictable, repeatable lessons-learned process. All levels—from analysts to managers and direct reports—should be involved in this process, but the level of involvement will depend highly on the organizational culture and individual capabilities within the enterprise.

FIGURE 15 War Room After-action Report Template

Investigation time frame	
Investigation type	
Trigger	
War room commander	
Notification method	
Initial infection vector (What was targeted or compromised?)	
Actions During the War Room	
List actions here.	
Were notifications timely at all levels?	
Were notifications accurate?	
What actions that occurred were not related to the investigation at hand?	
Did team members use safe and secure methods so as not to inflict further damage?	
Were adequate facilities available? What was not readily available that will be needed for future war rooms? Did the delay in availability affect the timeliness of response actions?	
Were available checklists used?	
What checklists/processes need to be developed?	
What checklists/processes need to be improved? (Elaborate below.)	
Were there any special circumstances or serious unexpected problems?	
What other problems could have arisen? How would they have been handled?	
Were duties delegated to appropriate people? Were necessary adjustments made? How did they perform?	
Actions during the war room were complete when:	

FIGURE 15 War Room After-action Report Template (cont.)

Actions After the War Room	
List actions here.	
Lessons Learned	
Areas of Recognition	
List areas here.	
Areas for Improvement (Consider People, Processes, Technology)	
List areas here.	
How could the incident have been avoided? What technologies or processes could be added/alterd across the enterprise?	
Action Items	
List actions here.	

5.4 Strategic Change—Cybersecurity Transformation

One of the most important outcomes of the eradication event should be a road map describing how the enterprise will leverage lessons learned from the incident to become more resilient in the face of future attacks. The road map should include projects or initiatives, technical and nontechnical, that will improve the enterprise’s ability to reduce an attacker’s odds of success and more rapidly and effectively detect and respond to an attacker’s activities. Analysis of the security breach should consider whether technical capability gaps contributed to the attacker’s success or whether people or process gaps were the main culprit.

During the planning process, the team should have identified a series of “fix” actions to undertake prior to the eradication event. Other items should have been identified as longer-term initiatives that the enterprise may not be able to accomplish prior to eradication. The following list describes activities that some enterprises may determine to be longer-term strategic efforts. However, depending on the attacker’s tactics to gain access and maintain persistence, some of these activities may need to be accomplished prior to the eradication event. Possible long-term strategic efforts include:

- **Desktop environment**—Each host on a corporate network is a potential attack vector. Individual desktop hosts typically represent the largest attack surface due to the number of hosts, the challenge of monitoring/administering those hosts and the

technical education level of the users. If possible, a new security baseline should be established prior to eradication day. It may not be possible to prevent the inevitable compromise of desktop workstations given the unpredictable behavior of end users, but increasing the security baseline of the desktop environment can greatly mitigate the risk that a compromise poses to the network. Some leading practices are:

- Removing end users from the local administrators group
 - Confirming that application and system patches are being delivered promptly as required by policy
 - Maintaining up-to-date antivirus end point protection definitions
 - Providing administrators with nonadministrator accounts to use when escalated privileges are not required. Additionally, if deploying a more secure OS such as Windows 7[®] is not an option in the given time frame, it can be deployed to members of administrative groups.
- **Server environment**—Even if servers are not frequently the initial point of an attack on a network, they are typically the goal of any threat actor looking for proprietary, nonpublic information. Attackers target servers to increase their access in the environment by stealing more passwords, and they also target servers to harvest data to exfiltrate. The eradication team should assess the patch level of the server environment and generate a list of patches that need to be applied. Also, the team should conduct a service inventory to identify servers that are running unnecessary Disk And Execution MONitor tools (DAEMONS) and listeners. For example, if Common Internet File System (CIFS) file sharing is not required on a server, it should be disabled to decrease the attack surface of a host. All remaining required services should be brought to current patch levels.

Servers should not be accessed directly for administrative purposes, and enterprises should put network controls in place to limit this access. For administrative purposes, servers should be accessed via jump servers. A jump server is a hardened and monitored divide that spans two dissimilar security zones and provides a controlled means of access between them. User access is tightly controlled and monitored for these devices. This limits the available communication paths between compromised hosts and servers outside of specifically designated and controlled paths. Additionally, enterprises should put access controls in place to prevent servers from communicating with each other on unnecessary access channels. For example, if RDP is being used between a jump server and a domain controller, it should not be possible to move laterally from that domain controller and RDP into a different domain controller. If administrators need to move between servers, they should open a separate session from the jump server. This will control, log, monitor and limit administrative activity.

- **Other areas to address**—Besides hardening the desktop and server environment to make it more difficult for attackers to be successful, there are other areas of cybersecurity that must be addressed. Almost all of the following high-level topics

will necessitate multiple distinct projects, and many of those projects will have dependencies on other projects. This is why project management quickly becomes a core component of cybersecurity road maps. Some of the most topics with the most impact include:

- Establishing and running a PMO for cybersecurity
- Developing a security education and awareness program
- Conducting regular self-phishing exercises
- Building an enterprise electronic investigation capability
- Designing an enterprise password change management process
- Executing a rapid global password change
- Selecting and deploying a password vaulting solution
- Developing a vulnerability acceptance process
- Updating corporate policy on cybersecurity
- Developing an information classification program
- Disabling LAN manager in the enterprise
- Restricting anonymous connections in the environment
- Improving administrative and technical control over:
 - Domain administrator accounts
 - Local administrator accounts
 - Service and functional accounts
 - Accounts whose passwords are set to never expire
 - Dormant accounts
- Enhancing Active Directory security, including redesigning organizational unit structures
- Optimizing use of security tools including:
 - Encryption in the enterprise
 - Host-based/end point protection technologies
 - DLP and antivirus software
- Remediating default credentials
- Enhancing patch management
- Improving database security (Oracle, SQL, MySQL)
- Developing a high-value information protection program
- Deploying specialized technical information zones
- Enhancing technology controls in the enterprise
- Isolating shop floors from the intranet
- Enhancing egress filtering
- Reducing the number of points of presence
- Enhancing security of web mail
- Deploying multifactor authentication
- Implementing an SIEM system
- Selecting and deploying an application whitelisting solution

- Hardening the messaging infrastructure
- Scanning and securing web applications
- Deploying full packet capture and deep packet inspection at Internet points of presence
- Improving security in treasury and wire transfer functions
- Redesigning an information security function
- Enhancing security compliance
- Building a security governance capability
- Upgrading identity and access management capabilities
- Deploying information security governance, risk and compliance systems
- Establishing a vendor risk management program
- Deploying an insider threat program
- Building an SOC

Chapter 6. Conclusion

While following the phased approach to managing a cybersecurity incident will not prevent an event from occurring, it can be a key difference in the resilience of an enterprise. Well-prepared enterprises fare much better than those that are not prepared at all or are ill-prepared.

Following the phases of preparation, investigation and eradication, the post-eradication phase sets the stage for the future of the enterprise's security program. After the team has validated that eradication activities were successful, briefed stakeholders on eradication event results and conducted lessons-learned sessions, it can move on to more strategic changes that aim to transform the enterprise.

Transformation can help to elevate an enterprise so that it is better prepared to deal with any potential events. In order to assist security professionals with the challenges in transforming the enterprise, ISACA is creating additional guidance leveraging COBIT 5. The publication will be released in the third quarter of 2013 and will be available at www.isaca.org.

Page intentionally left blank

Appendix A. Other Questions the Investigation Team Will Address

Other questions on which the investigation team will focus beyond the “who, what, when, where, why and how” of the attack are:

- Are the enterprise’s basic hygiene efforts sufficient for today’s new threats?
- Does a reliable and continually updated computer asset inventory exist?
- Does a reliable and continually updated inventory of authorized and unauthorized software exist?
- Are accounts with elevated privileges handled differently from regular user accounts?
- Has the enterprise deployed a password vaulting system to control access to the most sensitive accounts?
- Have users been removed from the local administrators group on their computers?
- Does the enterprise have accounts whose passwords never expire or are changed infrequently?
- Does the enterprise have accounts with elevated privileges that work in multiple domains (i.e., same user ID and password)?
- Does the enterprise use nested groups that make account management difficult?
- Does the enterprise have good control of its domain administrator accounts, and are they prohibited from being members of any other groups?
- Are regular scans of the entire environment conducted for vulnerabilities, and are they addressed in a timely manner based on the risk they represent?
- Is full disk encryption capability deployed to protect against computer theft?
- Is the latest version of the enterprise antivirus solution deployed, and are all end points regularly updated with the latest signatures?
- Are regular antivirus scans of enterprise data repositories/file shares being conducted?
- Has the enterprise protected its public key infrastructure (PKI)?
- Are the most important keys and certificates separated from the network using an air gap?
- Do the technology controls appropriately harden the following technologies?
 - Devices, including firewalls, routers, printers and switches
 - Databases and applications
 - Laptops, workstations and servers
- Have all default accounts on all devices, databases and applications been changed?
- Have technology controls been reviewed and updated recently?
 - Does the enterprise have an automated way to continually measure compliance with its technology controls?

- Is a reliable patch management system in place to keep the following items continually updated/patched?
 - Microsoft software
 - Other software (e.g., Sun[®]/Oracle[®] Java, Adobe Acrobat/Flash, Apple QuickTime, Google Chrome, WinZip[™], other browsers)
- Has the enterprise secured access to its critical database servers and databases? Is that access logged, and are periodic appropriateness reviews conducted?
- Have all nonsupported OSs and applications (e.g., Windows NT[®], Windows 2000[®], or Microsoft Office 2000[®]) been removed from the environment?
- Does the enterprise disallow nonessential ports, protocols and services at the firewall and proxy servers?
- Are wireless access points tightly controlled, and are they configured in compliance with the enterprise's security policies?
- Are users required to authenticate to the proxy server to access the Internet?
 - Are uncategorized sites blocked at the proxy server?
 - Are dynamic DNS lookups disallowed?
- Is multifactor authentication required for all remote access, including web mail?
- Does the enterprise have a centralized logging and monitoring system that collects event logs from at least Active Directory, firewall, DNS, Dynamic Host Configuration Protocol (DHCP), VPN, proxy server, messaging system, antivirus software, IDS, IPS, DLP and web filter devices?
- Are application firewalls used to protect Internet-facing web applications?
- Is the code in Internet-facing web applications tested regularly for exploitable vulnerabilities?
- Have all application developers been sent to specialized training for secure coding?
- Have the accounts payable/treasury workstations been secured so that they are dedicated to this function and serve no other purpose?
- Are penetration tests using third-party experts conducted regularly?
- Has the enterprise refreshed the content in its security awareness program, and are all employees, contractors, onsite vendors and joint venture partners required to take the training annually?
- Are the enterprise's information security professionals sent to appropriate training events annually to keep their skills sharpened?
- Have more advanced steps been taken to protect the enterprise?
- Is the information security culture right for today's threat environment?
 - Is information considered an asset under active management?
 - Is trust specifically assigned based on need?
 - Does the enterprise generally view itself as a high-risk or high-value target?
 - Is there an understanding that prevention is impossible—that the network is already compromised or soon will be, and the enterprise must be able to protect its most sensitive information in a compromised environment?
 - Is there an understanding that people and data are the new perimeters?

- Are independent, standards-based reviews conducted?
- Is the enterprise externally aware of what is going on in the information security space?
- Does a risk-based, business-focused approach to information security exist?
- Has the enterprise replaced its policy exception process with a risk acceptance process?
- Does the enterprise actively manage vendors and third parties with access to its environment?
- Do the information security policies emanate from one central authority?
- Is information security considered a business enabler?
- Is the data center traffic fully managed?
- Are operations conducted to actively monitor for and alert on nonapproved or anomalous network traffic?
- Are the most sensitive environments segmented from the main intranet or disconnected from the intranet entirely?
 - Is access logged, and are periodic appropriateness reviews conducted?
- Does an SOC capability exist, either in-house or through a managed security service (MSS)?
- Have any malware detection capabilities been deployed that are not signature-based?
- Is all email and other network traffic reviewed for suspicious behaviors?
- Does the enterprise have a forensics investigation capability with a fully staffed and trained incident response team?
- Have jump servers been deployed to tightly control and monitor access to the member server environment?
- Are PC and server virtualization being used to reduce security exposure?
- Has application whitelisting been deployed to core servers and high-risk servers?
- Have credentials been partitioned to produce the following results:
 - Workstation administration accounts cannot be used on member servers or domain controllers.
 - Member server administration accounts cannot be used on workstations or domain controllers.
 - Domain administration accounts cannot be used on workstations or member servers.
 - Domain administration accounts can be used only on domain controllers and only from dedicated laptops and only after multifactor/smart card authentication.
 - Application accounts can be used only on the applications or systems for which they are specifically deployed.
- Does the enterprise have a robust full-packet capture capability on all of its egress points so that it can look backward in time and replay sessions that may contain actionable threat intelligence?
- Have self-phishing exercises been conducted as a way to measure employee awareness of this threat vector?
- Is a DLP solution being used on the network and on the end points to protect the most sensitive data?

- Have local firewalls been deployed to all end points?
 - Have IDS/IPS solutions been deployed to all end points?
- Has the enterprise limited the number of places its intranet touches the Internet (i.e., egress point consolidation)?
- Is there a current information security strategy driving the decision-making processes?
- Is the information security organization properly aligned, staffed and trained for today's threat environment?
- Has the enterprise fully developed social, mobile and cloud security programs?
- Has the enterprise developed and deployed an insider threat program?
- Do strong capabilities exist in the following areas:
 - Security governance team and organization
 - Governance, risk and compliance (GRC) program, including a security compliance program that regularly tests the enterprise
 - Identity and access management (IAM) program
 - Threat and vulnerability management (TVM) program
 - Information security risk management program embedded in the enterprise risk management (ERM) function
 - Vendor risk management program
- Has the enterprise refreshed all of its information security policies, standards, procedures, guidelines and recommendations?

Appendix B. Investigative Tools

There are some freely available virtual machine instances that come prebuilt with tools that are needed in an investigation. These virtual machines often contain default credentials that must be changed. The following is a list of some of the most common, open source, virtual machines and their associated passwords that would allow for administrative access to the virtual machine and possibly parts of the host, depending on the setup of the virtualized environment:

- *SANS SIFT Workstation: Investigative Forensic Toolkit*
 - Login: sansforensics
 - Password: forensics
 - <http://computer-forensics.sans.org/community/downloads#login>
- *REMnux: A Linux Distribution for Reverse-Engineering Malware*
 - Operate in REMnux as the user: remnux
 - Default password for this account: malware
 - <http://zeltser.com/remnux/remnux-malware-analysis-tips.html>
- *Backtrack: A Linux Security Distribution*
 - Default user name: root
 - Default password: toor
 - http://www.backtrack-linux.org/wiki/index.php/Basic_Usage#Changing_the_root_password

Page intentionally left blank