



REDEFINING HACKING

A COMPREHENSIVE GUIDE TO
RED TEAMING AND BUG BOUNTY HUNTING
IN AN AI-DRIVEN WORLD



OMAR SANTOS | SAVANNAH LAZZARA
WESLEY THURNER

FREE SAMPLE CHAPTER |



**REDEFINING HACKING:
A COMPREHENSIVE GUIDE TO
RED TEAMING AND BUG BOUNTY
HUNTING IN AN AI-DRIVEN WORLD**

This page intentionally left blank

REDEFINING HACKING: A COMPREHENSIVE GUIDE TO RED TEAMING AND BUG BOUNTY HUNTING IN AN AI-DRIVEN WORLD

Omar Santos, Savannah Lazzara, Wesley Thurner

◆◆ Addison-Wesley

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

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

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

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2025930851

Copyright © 2025 Pearson Education, Inc.

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

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

ISBN-13: 978-0-13-836361-1

ISBN-10: 0-13-836361-7

\$PrintCode

GM K12, Early Career and Professional Learning

Soo Kang

Director, ITP Product Management

Brett Bartow

Executive Editor

James Manly

Development Editor

Christopher Cleveland

Managing Editor

Sandra Schroeder

Senior Project Editor

Mandie Frank

Copy Editor

Chuck Hutchinson

Indexer

Timothy Wright

Proofreader

Donna E. Mulder

Technical Reviewers

Matthew Eidelberg

Petar Radanliev

Editorial Assistant

Cindy Teeters

Designer

Chuti Prasertsith

Compositor

codeMantra

I would like to dedicate this book to my lovely wife, Jeannette, and my two beautiful children, Hannah and Derek, who have inspired and supported me throughout the development of this book.

—Omar

I would like to dedicate this book to my devoted parents, Dannielle and Tom, as well as my amazing partner, Nickolas, who never stopped believing in me. I would also like to honor my loving dogs: Baxter, who brought joy to my life for 13 years and will forever be missed, and Caymus, who is the light of my life.

—Savannah

I would like to dedicate this book to my son, Westin, who put up with me throughout this entire endeavor and never stopped supporting me.

—Wes

This page intentionally left blank

Contents

1	The Evolution of Penetration Testing, Red Teaming, and Bug Bounties	1
	Into the Abyss	1
	The Uncharted Territory	2
	The AI Crash Course	2
	Into the Breach	3
	Culmination	5
	Next Steps	5
	Industry Overview	5
	What Are Penetration Testing, Red Teaming, and Bug Bounties?	6
	Why Are These Disciplines Important?	6
	How Have They Evolved Over Time?	6
	The Early Days and Evolution of Penetration Testing	7
	The Rise of Red Teaming	9
	The Emergence of Bug Bounty Programs	10
	The Convergence of Penetration Testing, Red Teaming, and Bug Bounties	11
	The Future of Penetration Testing, Red Teaming, and Bug Bounties	12
	The Increasing Use of AI and ML in Penetration Testing, Red Teaming, and Bug Bounties	13
	Test Your Skills	15
2	Introduction to Red Teaming	17
	What Is Red Teaming?	17
	The Importance of Red Teaming	18
	Identifying Critical Vulnerabilities	18
	Challenging Assumptions	18
	Emulating Industry-Specific Threats	19
	Red Team Frameworks and Methodologies	20

	MITRE ATT&CK Framework	20
	Unified Kill Chain	20
	TIBER-EU	21
	CBEST	21
	Red Team Engagements	22
	Full-Scope	22
	Objective-Based	23
	Adversary Emulation (Scenario-Based)	24
	Purple Team	24
	Tabletop	25
	Red Team Challenges	26
	Leveling Up Other Teams	26
	Lack of Findings	27
	References	28
	Test Your Skills	30
3	Red Team Infrastructure	33
	An Overview of Red Team Infrastructure	34
	Command and Control	34
	C2 Communication Channels	36
	Redirectors	37
	Command and Control Setup	37
	Test Your Skills	60
4	Modern Red Team Methodology and Tools	63
	Planning	64
	Rules of Engagement	64
	Kickoff Call	65
	Communication	65
	Reconnaissance	65
	Autonomous System Number (ASN)	65
	Certificate Transparency	66
	Proxying Traffic	66

Domain Name Service	71
Metadata	78
Username Collection	79
Software Discovery	81
Leaked Data	81
Initial Access	81
Password Spraying	82
Mobile Device Management (MDM)	87
Social Engineering and Physical Access	88
Payload Preparation	88
Web Application Vulnerabilities	95
Persistence	96
Lateral Movement	97
Obtaining Credentials	97
Gathering Domain Information	97
Kerberoasting	99
Active Directory Certificate Services	100
SCCM Abuse	104
Post-Exploitation	105
Reporting	105
Test Your Skills	107
5 Social Engineering and Physical Assessments	109
Phone Pretexting: Let’s Dial into Your Target	110
Reconnaissance	110
Preparing Pretext	111
Execution	113
Spear Phishing: Time to Reel In Your Target	115
Infrastructure	116
Formulating Your Email or Instant Message	128
Execution	130

Can You Hear Us from the Inside?	131
Remote Reconnaissance	131
Badge Cloning	132
Onsite Reconnaissance	139
Social Engineering in Person	140
Physical Building Bypasses	140
Test Your Skills	144
6 Advanced Post-Exploitation Techniques	147
An Overview of Post-Exploitation	148
Evaluating Detection and Response Capabilities and Understanding the Impact of a Breach	148
Improving Blue Team Capabilities	149
How to Maintain Access, Use Persistence Mechanisms, and Create Backdoors	152
Types of Backdoors	153
Reverse Shells vs. Bind Shells	160
Patching Binaries to Include a Backdoor	165
Avoiding Detection Using Code Obfuscation	167
System-Level Backdoors	170
Kernel-Level Backdoors	174
Hardware and Firmware Backdoors	178
Scheduled Jobs and Tasks	179
New Users	180
Command and Control (C2) and Covert Channels	180
How DNS Tunneling Works	181
HTTP/HTTPS Tunneling	183
Nontraditional Protocols	183
Additional C2 Techniques and Tools	183
How to Perform Lateral Movement	186
Post-Exploitation Scanning	186
Legitimate Utilities and Living Off the Land	189

PowerShell for Post-Exploitation Tasks	189
PowerSploit and Empire	190
BloodHound	193
Windows Management Instrumentation for Post-Exploitation Tasks	193
Sysinternals and PsExec	194
Windows Remote Management (WinRM) for Post-Exploitation Tasks	195
Living-Off-the-Land Bins, Libraries, and Useful Scripts	195
Post-Exploitation Privilege Escalation	197
How to Cover Your Tracks	198
Steganography	198
Test Your Skills	202
7 Active Directory and Linux Environments	207
Active Directory Fundamentals	208
Centralized Data Repository and High Availability	208
Security	209
Active Directory Architecture	211
AD Trust Relationships	213
Key Active Directory Protocols and Services	216
Microsoft Entra ID and Azure	223
Active Directory Attack Techniques	229
How to Set Up the Lab	229
AD Recon and Enumeration	231
Password Spray and ASREPRoast	239
Useful LDAP Queries for Windows Active Directory Assessments	242
BloodHound Is Awesome	246
Using Responder for Credential Harvesting in Active Directory Environments	247
Using Secretsdump to Extract Credentials	253
Using Lsassy to Extract Credentials from LSASS	256
Using DonPAPI to Target the Windows Data Protection Application Programming Interface (DPAPI)	257

ADCS Reconnaissance and Enumeration with Certipy and BloodHound	259
Exploitation Using Certipy	261
Abusing Discretionary Access Control Lists (DACLS) and Access Control Entries (ACEs)	262
The Golden SAML Attack	263
Microsoft Entra Connect (Azure AD Connect) Sync Account Takeover	265
Microsoft Entra Privileged Identity Management (PIM) Abuse	266
Advanced Linux Environment Attack Techniques	267
A Quick Refresher on How Buffer Overflows Work	267
How to Exploit Buffer Overflows	267
Examples of Vulnerable Code	270
Bypassing No-Execute (NX) Stack with Return-Oriented Programming (ROP)	273
Bypassing NX with ROP	275
Test Your Skills	277
Project 7-1	280
8 The Future of Red Teaming Beyond the AI Revolution	281
Understanding the Current State of AI in Red Teaming	282
Creating AI-Powered Offensive Security Tools	283
LangChain, Prompt Templates, LangSmith, LangGraph, and LlamaIndex	289
Examining Fine-Tuned Uncensored AI Models	297
Understanding Retrieval-Augmented Generation (RAG) for Red Teaming	299
Vector Embeddings	300
Vector Database Storage	302
User Query Processing	305
Semantic Search and Document Retrieval	305
Context Preparation for LLMs or SLMs	308
Response Generation and Post-Processing and Refinement	311
Red Teaming AI and Autonomous Systems	313
Keeping Pace with Rapidly Evolving Technologies	322
Test Your Skills	324

9 Introduction to Bug Bounty and Effective Reconnaissance 327

- Understanding Bug Bounty Programs 328
 - Types of Bug Bounty Programs 329
 - Attack Surface Management vs. Bug Bounties 331
 - Vulnerability Disclosure Programs (VDPs) vs. Bug Bounties 332
 - Getting Started: Preparing Yourself as an Ethical Hacker in Bug Bounty Programs 333
 - Understanding the Scope and Rules of Engagement 334
- Exploring Effective Reconnaissance 336
 - Active vs. Passive Reconnaissance 336
 - Understanding OSINT 337
 - Leveraging DNS for Recon 339
 - Identifying Technical and Administrative Contacts 342
 - Identifying Cloud vs. Self-Hosted Assets 344
 - Social Media Scraping 348
 - Cryptographic Flaws 348
 - Obtaining Sensitive Information About External and Internal Hosts Using Certificate Transparency 353
 - Leveraging Password Dumps 355
 - Recon from File Metadata 358
 - Strategic Search Engine Analysis/Enumeration 360
 - Website Archive/Caching 361
 - Public Source-Code Repositories, Secrets, and Other Sensitive Information 362
 - Performing Reconnaissance with Recon-ng 363
 - Shodan 369
 - Amass, Maltego, and Other OSINT Tools 371
 - Using Generative AI and the Gorilla Large Language Model to Interact with Tools Like Amass 372
 - Using Open-Interpreter to Interact with Recon Tools 373
- Performing Active Reconnaissance 377
 - Creating Your Own Scanner Using Python 380

Exploring the Different Types of Enumeration	382
Using BloodHound for a Bug Bounty	394
Packet Inspection and Eavesdropping	395
Understanding the Art of Performing Vulnerability Scans	396
Understanding the Types of Vulnerability Scans	396
Challenges to Consider When Running a Vulnerability Scan	398
Performing Web Application and API Recon	401
Directory and File Brute-Forcing	401
API Recon	404
Communicating Your Findings and Creating Effective Bug Bounty Reports	407
Test Your Skills	410
Exercises	412
10 Hacking Modern Web Applications and APIs	415
Overview of Web Application-Based Attacks, the OWASP Top 10 for Web Applications, and OWASP Top 10 for LLM Applications	416
The HTTP Protocol	416
Understanding Web Sessions	425
OWASP Top 10 for Web Applications	427
OWASP Top 10 for LLM Applications	428
Building Your Own Web Application Lab	429
Understanding Business Logic Flaws	430
Understanding Injection-Based Vulnerabilities	432
Hacking Databases and Exploiting SQL Injection Vulnerabilities	433
Command Injection Vulnerabilities	448
Lightweight Directory Access Protocol (LDAP) Injection Vulnerabilities	449
Exploiting Authentication-Based Vulnerabilities	451
Brute-Forcing Credentials	451
Understanding Session Hijacking	452
Understanding Redirect Attacks	458
Taking Advantage of Default Credentials	458

- Exploiting Kerberos Vulnerabilities 459
- Exploiting Authorization-Based Vulnerabilities 461
 - Understanding Parameter Pollution 461
 - Exploiting Insecure Direct Object Reference (IDOR) Vulnerabilities 462
- Understanding Cross-Site Scripting (XSS) Vulnerabilities 463
 - Reflected XSS Attacks 465
 - Stored XSS Attacks 466
 - XSS Evasion Techniques 470
 - XSS Mitigations 471
- Understanding Cross-Site Request Forgery and Server-Side Request Forgery Attacks 472
 - Server-Side Request Forgery Attacks 475
 - Exploiting an SSRF Vulnerability Using WebSploit Labs 477
- Understanding Clickjacking 480
- Exploiting Security Misconfigurations 480
 - Exploiting Directory Traversal Vulnerabilities 480
 - Understanding Cookie Manipulation Attacks 482
- Exploiting File Inclusion Vulnerabilities 482
 - Local File Inclusion Vulnerabilities 482
 - Remote File Inclusion Vulnerabilities 483
- Exploiting Insecure Code Practices 484
 - Comments in Source Code 484
 - Lack of Error Handling and Overly Verbose Error Handling 484
 - Hard-Coded Credentials 484
 - Race Conditions 485
 - Unprotected APIs 485
 - Hidden Elements 488
 - Lack of Code Signing 488
- Using Additional Web Application Hacking Tools 488
- Test Your Skills 493

11 Automating a Bug Hunt and Leveraging the Power of AI 497

- Traditional Bug Hunting Methods 498
 - Limitations of Manual Bug Hunting 499
- AI-Powered Automation in Bug Hunting 500
 - AI Capabilities of Bug Bounty Platforms 508
 - A Comprehensive View of an Organization’s External Risk Exposure 509
 - Vulnerability Prioritization Using AI 510
 - AI-Created Scanner Templates 512
- AI Model Training, Fine-Tuning, and RAG for Bug Bounties 516
 - Deploying AI Models 516
 - Fine-Tuning AI Models 517
 - Using RAG and AI Agents for Bug Bounty Hunting 519
 - Tool Calling 520
- Challenges of Using AI for Bug Bounty Hunting 521
 - Censored Models and Guardrails 521
 - Hallucinations or Confabulations 522
- Test Your Skills 523

Answers to Multiple-Choice Questions 527

Index 541

Register your copy of *Redefining Hacking: A Comprehensive Guide to Red Teaming and Bug Bounty Hunting in an AI-Driven World* at informit.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to informit.com/register and log in or create an account. Enter the product ISBN (**9780138363611**) and click Submit. Once the process is complete, you will find any available bonus content under “Registered Products.”

Acknowledgments

Omar Santos:

I would like to thank the technical editors for their time and technical expertise.

I would like to thank the Pearson team, especially James Manly and Christopher Cleveland, for their patience, guidance, and support.

Savannah Lazzara:

I would like to thank the technical editors for their amazing work making this book come to life, as well as my mentors, Travis Weathers, Matthew Eidelberg, and Ben Rollin, who have helped me develop my skill set over the years. Additionally, I can't thank the Pearson team enough, specifically James Manly and Christopher Cleveland for supporting us.

Wesley Thurner:

I would like to thank the technical editors for their expertise, James Manly, Christopher Cleveland, and the Pearson team for bringing this to fruition, their support and understanding.

This page intentionally left blank

About the Authors

Omar Santos is a Distinguished Engineer at Cisco who pioneers advancements in artificial intelligence security, cybersecurity research, ethical hacking, incident response, and vulnerability disclosure. As co-chair of the Coalition for Secure AI (CoSAI) and board member of the OASIS Open standards organization, he shapes the future of secure technology adoption across industries.

Omar drives innovation through multiple leadership roles, including as founder of OpenEoX and co-chair of the Forum of Incident Response and Security Teams (FIRST) PSIRT Special Interest Group. His commitment to cybersecurity education and community building is evident in his role as the co-founder and one of the leaders of the DEF CON Red Team Village and the chair of the Common Security Advisory Framework (CSAF) technical committee.

Omar has published over 20 books, created over 20 video courses, and contributed more than 40 academic research papers to the field. Omar's work in cybersecurity is also recognized through multiple granted patents. You can follow Omar in X at <https://x.com/santosomar> and LinkedIn at <https://www.linkedin.com/in/santosomar>.

Savannah Lazzara is a Security Engineer specializing in red teaming at a tech company. Savannah has multiple years of experience in security consulting working with many Fortune 500 corporations and has experience in carrying out security assessments, which include network assessments, social engineering exercises, physical facility penetration tests, and wireless assessments. Savannah also has experience in performing adversary simulation assessments, which include remote red team simulations, insider threat assessments, and onsite red team assessments. Savannah's area of expertise is focused on social engineering and physical security.

Savannah is the co-lead of Red Team Village. Additionally, Savannah has spoken at various cybersecurity conferences, including Source Zero Con, BSides, and others. Savannah has also spoken on multiple podcast shows including *The Hacker Factory* and *Hackerz and Haecksen*.

Wesley Thurner is a Principal Security Engineer on the Intuit Red Team, designing and leading transformative security initiatives across the organization. He accomplishes this through multiple internal roles. Wesley leads the internal Purple Team, is an Ambassador for cybersecurity education throughout the workforce, and leads working groups to define security policies.

Prior to Intuit, Wesley served as an exploitation operator in the US Department of Defense's most elite computer network exploitation (CNE) unit. There he led and developed multiple teams across a variety of roles in the US Air Force's premier selectively manned cyber attack squadron. Wesley is also a co-lead for the Red Team Village, a community-driven village bridging the gap between penetration testers and offensive operations. He has spoken at various cybersecurity conferences and authored numerous training courses as well as white papers.

This page intentionally left blank

About the Technical Reviewers

Matthew Eidelberg graduated from Seneca@York University with a Bachelor of Technology in Informatics and Security. He has more than ten years of experience in the cybersecurity industry, with a primary focus on adversary emulation. His specialty lies in research and development related to exploits and malware—specifically the creation of new techniques and tools for bypassing and evading endpoint security measures. Matthew joined Black Hills Information Security (BHIS) in August 2023. Prior to that, he spent eight years at Optiv, where he led the Adversary Simulation Services team.

Petar Radanliev is a post-doctoral research associate at the Department of Computer Science at the University of Oxford. He obtained his PhD at the University of Wales in 2014. He continued with post-doctoral research at Imperial College London, the University of Cambridge, Massachusetts Institute of Technology, and the Department of Engineering Science at the University of Oxford before moving to the Department of Computer Science. His current research focuses on artificial intelligence, cybersecurity, quantum computing, and blockchain technology. Before joining academia, Petar spent ten years as a cybersecurity manager for RBS (the largest bank in the world at the time) and five years as a lead penetration tester for the Ministry for Defence.

This page intentionally left blank

Introduction

The art and science of red teaming and security testing are at a fascinating crossroads with emerging technologies like artificial intelligence (AI). Gone are the days when cybersecurity was merely about finding buffer overflows or SQL injection vulnerabilities. Today's security professionals face an entirely new frontier: systems that reason, adapt, and learn.

The tools and techniques that served us well for decades are being transformed. AI-powered defense systems now detect and respond to threats in milliseconds, while traditional machine learning and generative AI models analyze patterns of behavior that would take human analysts months to uncover. Yet, paradoxically, this same technological revolution has created new vulnerabilities and attack surfaces that few could have imagined just years ago.

This book will help you gain relevant knowledge at a crucial moment in the evolution of offensive security testing. Many organizations increasingly rely on AI-driven systems to protect their assets. Consequently, the traditional approaches to red teaming, penetration testing, and bug bounty hunting must evolve. Red teamers now need to understand the skills needed for today's environment. Bug bounty hunters must adapt their tools and methodologies to probe systems that can predict and react to their every move.

Whether you're a seasoned red teamer, an aspiring bug bounty hunter, or an individual looking to understand the future of offensive security, this book will serve as your guide through this transformation. We'll explore the cutting-edge techniques being developed at the intersection of AI and security testing, from AI-powered reconnaissance to principles of exploit development. This book is your roadmap to navigating this new terrain, combining time-tested methodologies with emerging technologies to create a comprehensive approach to modern security testing.

Welcome to the future of ethical hacking!

Goals/Objectives/Approach of the Book

This book provides a comprehensive framework for transitioning from traditional security testing methodologies to modern and AI-enhanced methodologies, helping you to remain effective in an evolving threat landscape. This book delivers actionable, hands-on knowledge that you can immediately apply in real-world red team engagements and bug bounty hunting, with a focus on both current and emerging technologies.

This book will equip you to execute sophisticated red team operations that account for AI-powered defense systems and leverage AI and automation effectively in bug bounty hunting while maintaining the crucial human element.

You will learn how to build and maintain modern red team infrastructure and leverage custom tools and frameworks that combine traditional and new techniques.

Each chapter begins with core concepts and then provides step-by-step explanations and real-world case scenarios.

Targeted Reading Audience

This book is designed for beginners, experienced cybersecurity professionals transitioning to red team operations, and active bug bounty hunters looking to enhance their knowledge.

Additionally, it is designed for red teamers seeking to integrate AI into their methodology and any organizations building or improving their red team capabilities.

The only prerequisites are a basic understanding of networking, operating systems, and web technologies and a basic understanding of cybersecurity concepts and terminology.

Book Organization

Chapter 1: The Evolution of Penetration Testing, Red Teaming, and Bug Bounties

This chapter covers the transformative journey of security testing from its origins to today's AI-augmented practices. We explore how simple vulnerability assessments evolved into sophisticated red team operations, examining the pivotal developments that shaped modern security testing methodologies, tools, and frameworks.

The chapter highlights how the industry's landscape has been revolutionized by major shifts: the evolution of penetration testing frameworks, the emergence of professional red teaming incorporating physical security and social engineering, and the rise of bug bounty platforms that democratized security testing.

Chapter 2: Introduction to Red Teaming

This chapter goes over the fundamental principles and methodologies that distinguish red teaming from traditional penetration testing. You will learn how red teams operate as dedicated adversary emulation units, understanding the key components of engagement planning, threat actor profiling, and the establishment of clear objectives and rules of engagement. The chapter introduces essential frameworks and discusses how they guide modern red team operations.

Chapter 3: Red Team Infrastructure

This chapter covers the critical components of building and maintaining a resilient red team infrastructure. From designing covert Command and Control (C2) frameworks to implementing secure communication channels, the chapter covers essential concepts, including domain categorization, traffic redirection, and operational security measures.

Chapter 4: Modern Red Team Methodology and Tools

This chapter examines cutting-edge methodologies and tools that define modern red team operations. The chapter investigates how traditional attack frameworks are being enhanced by modern tools, exploring advanced techniques in initial access, lateral movement, and persistence. The chapter walks through each phase of a red team operation.

Chapter 5: Social Engineering and Physical Assessments

This chapter explores the human element of red teaming through advanced social engineering techniques and physical security assessments. The chapter examines how social engineering has evolved in the age of AI-powered behavioral analytics and security awareness training, covering sophisticated pretexting, phishing campaigns, and psychological manipulation techniques that can bypass modern defenses. Special attention is given to conducting safe and ethical social engineering exercises that provide valuable insights without causing harm.

Moving into physical security, the chapter details methodologies for conducting comprehensive physical penetration tests, including facility access assessments, physical security control bypass, and the convergence of cyber and physical attack vectors. You will learn about modern physical security tools, RFID cloning techniques, and tactics for blending cyber and physical methodologies.

Chapter 6: Advanced Post-Exploitation Techniques

This chapter details sophisticated post-exploitation methodologies that enable red teams to maintain persistent access while evading detection. The chapter covers advanced techniques for memory manipulation, process injection, and living-off-the-land techniques, with particular focus on bypassing modern endpoint detection and response (EDR) solutions and AI-powered security tools. You will learn how to develop and implement custom implants that can operate undetected in highly monitored environments. We discuss the importance of operational security during post-exploitation activities and provide strategies for cleaning up artifacts to maintain long-term covert presence.

Chapter 7: Active Directory and Linux Environments

This chapter explores advanced tactics for compromising modern Active Directory (AD) environments and Linux infrastructures. For Active Directory, the chapter covers sophisticated attack techniques including domain privilege escalation, credential abuse, and trust relationship exploitation, with special focus on bypassing modern security controls.

BloodHound is an open-source tool designed to map out relationships and permissions in Active Directory. It uses graph theory to visualize attack paths that could be exploited by attackers to escalate privileges or move laterally within a network. By analyzing these relationships, BloodHound helps security professionals understand how an attacker might compromise an environment and provides insights into misconfigurations or weaknesses that need remediation. SharpHound is the data collection component of BloodHound. It gathers information from AD environments by querying domain controllers and domain-joined systems. The collected data is then fed into BloodHound for analysis. In this chapter, you will learn details and review examples on how to use these tools.

The Linux section covers advanced exploitation techniques specific to modern Linux environments, including container escape, kernel exploitation, and privilege escalation in hardened environments. The chapter details methodologies for exploiting Linux systems leveraging common misconfigurations and vulnerabilities.

Chapter 8: The Future of Red Teaming Beyond the AI Revolution

This chapter covers how artificial intelligence and emerging technologies are revolutionizing red team operations. You will learn how AI is being leveraged to enhance attack capabilities, including autonomous reconnaissance systems, exploitation at scale, and advanced social engineering attacks using AI.

Chapter 9: Introduction to Bug Bounty and Effective Reconnaissance

This chapter provides a comprehensive introduction to professional bug bounty hunting in today's competitive landscape. The chapter examines the evolution of bug bounty programs, platform dynamics, and strategies for building a successful career as a bug hunter.

Chapter 10: Hacking Modern Web Applications and APIs

This chapter covers many techniques for testing modern web applications and APIs. We explore sophisticated attack methodologies against contemporary application architectures, including microservices, serverless functions, and GraphQL APIs.

Moving beyond traditional vulnerabilities, the chapter investigates emerging attack vectors in modern frameworks, cloud configurations, and authentication mechanisms.

Chapter 11: Automating a Bug Hunt and Leveraging the Power of AI

This chapter covers how AI and automation are revolutionizing the bug hunting process. The chapter covers practical tactics for building intelligent scanning systems, developing custom tools that leverage AI for vulnerability detection, and creating automated pipelines that can scale across multiple targets. You will learn how to combine traditional security tools with AI capabilities to enhance security assessments. Throughout the book you will learn about retrieval-augmented generation (RAG) and tools such as LangChain, LlamaIndex, Ollama, and many other AI frameworks.

Credits

Cover - Thinkhubstudio/Shutterstock

Figures 5.1, 5.2, 5.9 & 5.10 - OpenAI Inc

Figures 3.2, 3.3, 3.4, 3.16, 3.17, 3.21 - Microsoft Corporation

Figures 4.3, 4.4, 4.5, 4.6, 4.7 - Bright Data Ltd

Figure 4.10 – Caphyon

Figures 4.11, 4.12, 4.13, 4.15 - Microsoft Corporation

Figure 5.3 – SpoofCard

Figure 5.5 – Palo Alto Networks

Figures 5.6, 5.7, 5.17, 5.18 - Microsoft Corporation

Figure 5.8 – Breakpoint Software Development

Figure 5.13, 5.14, 5.15 & 8.8 - GitHub, Inc

Figure 6.5 – Jens Duttke

Figure 6.9 – GitHub, Inc

Figures 7.8, 7.9, 7.10, 7.11 - Google LLC

Figure 7.14 – Mayfly

Figures 8.1, 8.3, 9.2 & 9.4 - Google LLC

Figures 10.3 & 10.16 - Wireshark Foundation

Figures 10.7, 10.9, 10.10, 10.11, 10.12 - Apple Inc

Figures 10.20 & 10.21 - Mozilla Foundation

Figure 10.25 - PortSwigger Ltd

Figure 10.26 - ZAP Dev Team

Figures 11.5 – 11.12 - Apple Inc

Automating a Bug Hunt and Leveraging the Power of AI

Chapter Objectives

After reading this chapter and completing the exercises, you will be able to do the following:

- Use traditional bug hunting methods
- Employ AI-powered automation in bug hunting
- Understand AI model training, fine-tuning, and retrieval-augmented generation (RAG)
- Understand the challenges of using AI for bug bounty hunting

In Chapter 8, “The Future of Red Teaming Beyond the AI Revolution,” you learned about the current state of AI in red teaming, examining AI-powered offensive tools and techniques, fine-tuned uncensored AI models, and the application of retrieval-augmented generation (RAG) for red teaming purposes.

In this chapter, we will explore how to leverage AI for bug bounty hunting. You will learn about the methodologies and tools that can enhance your effectiveness as a bug bounty hunter, integrating AI to identify vulnerabilities more efficiently and accurately.

Traditional Bug Hunting Methods

Given the vast amount of information available online about bug bounty hunting, beginners might feel overwhelmed—which is entirely normal. To navigate this information overload, focus on a few high-quality resources and immediately apply what you’ve learned through hands-on practice.

Tips for Beginners

In Chapter 9, “Introduction to Bug Bounty and Effective Reconnaissance,” you learned how to get started with bug bounties. You should begin with the basics. For example, in web applications, you can start by understanding HTTP requests and the basics of how websites function. Choose a common vulnerability to specialize in, such as cross-site scripting (XSS), and concentrate your efforts there. PortSwigger offers an excellent Cross-Site Scripting (XSS) Cheat Sheet that lists XSS vectors across different frameworks and methods to bypass web application firewalls (WAFs).

Another area to explore is business logic vulnerabilities. Finding them is often straightforward once you grasp the application’s overall workflow. They require you to understand how the application is supposed to function and then identify edge cases that could have security or business implications. For instance, you might discover a way to bypass a paywall and access premium content without payment.

Information disclosure is another impactful type of vulnerability that can occur anywhere within an application. Your goal is to find instances where sensitive information—like user data, internal company details, old data backups, API keys, source code, or error messages—is inadvertently exposed. Identifying information that should remain confidential is crucial in this category.

As a beginner, you are advised to focus on one type of vulnerability and attempt to find it across multiple programs or targets. With determination and persistence, you’re likely to achieve success because bugs are always present in applications.

How do you shift from theory to practice? The field of bug bounty continues to evolve rapidly, prioritizing practical experience over theory. I (Omar) always say that cybersecurity, and especially ethical hacking, is like math. The more you practice, the better you will become.

Note

Finding bugs is essentially about identifying features that don’t work as intended. To do this effectively, you need to understand how the application should function and recognize deviations, especially from a security perspective. This task requires deep familiarity with the application and knowledge of common vulnerability areas. Be cautious of targeting only “low-hanging fruits”—simple bugs that are easy to find but are also the focus of many other hunters. Targeting only simple bugs increases the chance of submitting duplicate reports and can result in a low return on investment for your time and effort. Instead, thoroughly analyze applications to uncover more significant and less obvious vulnerabilities.

Limitations of Manual Bug Hunting

Manual bug hunting has been a fundamental aspect of ethical hacking for many years. Skilled cybersecurity professionals meticulously examine applications to uncover flaws that automated tools might miss. However, as technology becomes more complex and cyber threats more sophisticated, the limitations of manual bug hunting are becoming more pronounced.

Manual bug hunting is very time-consuming due to the extensive effort required for in-depth analysis. You must meticulously test different system configurations and system behaviors, which can be extremely labor-intensive. Modern applications often involve complex architectures, including microservices, APIs, and third-party integrations, making it impractical for individuals to scrutinize every component manually. This extensive time investment can lead to delayed vulnerability discovery, resulting in a slow threat response. The lengthy process of manual analysis can postpone the identification of critical security issues, increasing the risk window and providing attackers with more opportunities to exploit vulnerabilities.

Human error is an inevitable factor in manual bug hunting. You may overlook vulnerabilities due to cognitive limitations such as making assumptions or focusing on familiar attack vectors while neglecting others. Fatigue and attention lapses can occur during extended periods of manual testing, leading to decreased concentration and missed flaws. Knowledge gaps also contribute to oversight because no individual can have exhaustive knowledge of all potential vulnerabilities across diverse technologies.

Many individuals currently automate a lot of the bug bounty activities to learn about new bug bounties in platforms like HackerOne, Bugcrowd, Intigriti, and others. For example, you may scan any HackerOne program (bug bounty) with Nuclei, as shown in Example 11-1. Nuclei is an open-source vulnerability scanner developed by ProjectDiscovery, known for its speed, efficiency, and customizability. It uses a template-based approach, with YAML files defining the methods for detecting vulnerabilities across various targets, including web applications, cloud infrastructure, and networks.

EXAMPLE 11-1 Using Nuclei to Scan Hosts in Any HackerOne Bug Bounty

```
websploit$ python3 h1_2_nuclei.py -handle security
[i] Checking scope for: security
[i] Parsing scope items
[i] Wildcards in scope:      1
[i] Hosts in scope:         19
[i] Hosts out of scope:     3
[i] Checking subdomains with chaos
[i] Hosts in scope:         8
[i] Hosts out of scope:     4
[i] Removing out of scope items
[i] Unique hosts in scope:  9
[i] Saving hosts to: targets/security/chaos_security_250808.txt
[i] Resolving subdomains with httpx
[i] Output saved to: targets/security/httpx_security_250808.txt
[i] Number of live targets: 9
```

```
[i] Scanning targets with Nuclei  
[i] Output saved to: targets/security/nuclei_security_250808.txt  
[i] Vulnerabilities found: 8
```

The tools shown in Example 11-1 can be obtained from https://github.com/vavkamil/h1_2_nuclei.

To overcome the limitations of manual bug hunting, you should use automation that goes beyond vulnerability scanners and traditional hacking tools. Leveraging AI and machine learning can significantly help. You can adopt a hybrid testing approach that combines manual testing with automated tools to provide a more comprehensive security assessment. In the following sections, we will explore the benefits and limitations of using AI in bug hunting.

AI-Powered Automation in Bug Hunting

You’ve probably heard the phrase “the sky’s the limit.” When it comes to using AI for bug bounty hunting, that statement couldn’t be more accurate. The possibilities are amazing, and AI has the potential to revolutionize the way vulnerabilities are detected and exploited.

Figure 11-1 illustrates a simplified hierarchy of how bug bounty workflow entities are related.

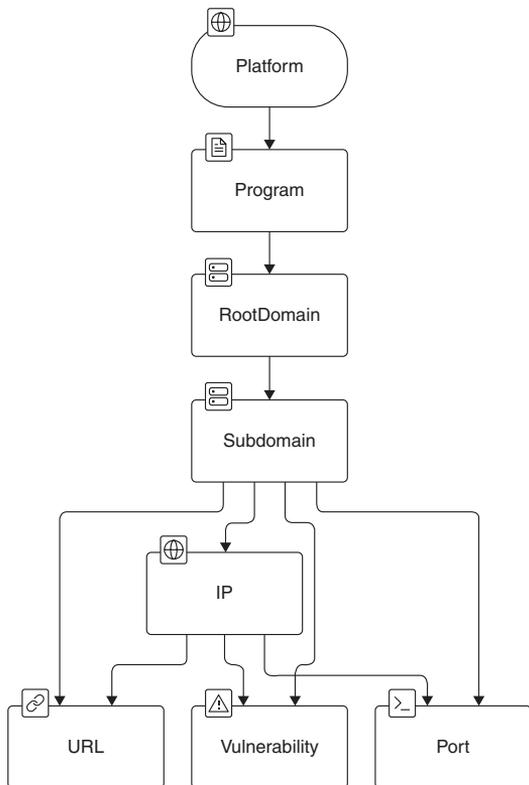


Figure 11-1
The Bug Bounty Hunter Workflow

At the top of the hierarchy are the bug bounty platforms like Bugcrowd, HackerOne, Intigriti, and others. Next, the program is the specific bug bounty program offered by a company or organization on the platform. For example, a company may set up a program with Bugcrowd to allow researchers to test their web applications for vulnerabilities.

The root domain represents the main domain or scope of the program. For example, if the company's main website is websploit.org, this would be the root domain in scope for bug testing. Programs often list which root domains are in scope for testing.

Beneath the root domain are subdomains, which are more specific areas of the website (such as api.websploit.org or admin.websploit.org). Bug bounty programs may specify which subdomains are part of the scope because certain subdomains may expose different services or applications.

The IP address represents the network layer, where subdomains are resolved to one or more IP addresses. This allows ethical hackers to probe network-level configurations, services, and open ports that could be vulnerable.

The URL represents specific web pages or API endpoints within the subdomain. These URLs may have parameters, authentication mechanisms, or other inputs that hackers target for testing vulnerabilities like XSS or SQL injection.

Ports represent the specific network ports on the IP address that could be open and potentially expose services. For example, port 80 (HTTP) and port 443 (HTTPS) are often tested.

Let's assume that now you want to automate the discovery, scanning, enumeration, and processing of these tasks using AI agents. You can use platforms like LangGraph and LangGraph Studio to create agents that will help accelerate these tasks with low-code requirements. Let's take a look at the diagram in Figure 11-2.

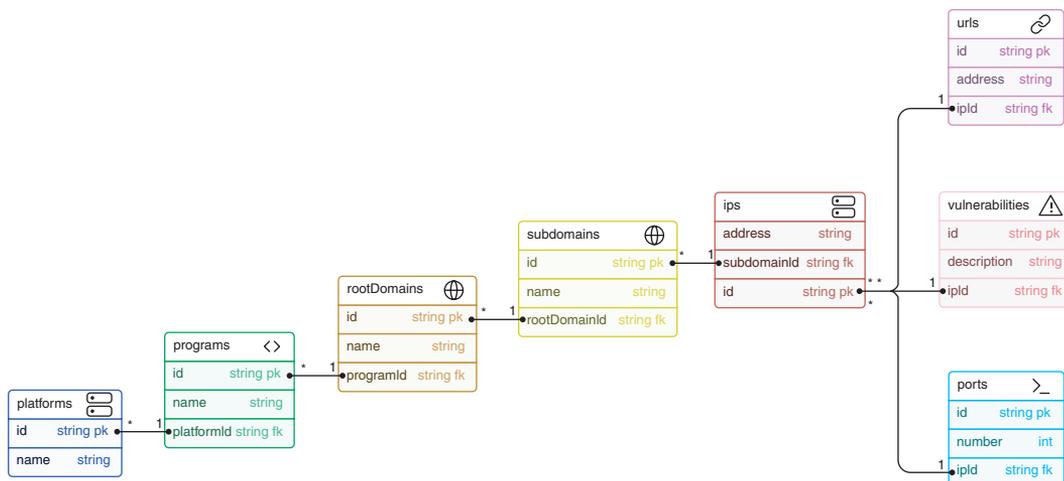


Figure 11-2

High-Level Building Blocks of an Automated Recon System

Figure 11-2 represents the architecture of the building blocks for an automated system to automate reconnaissance and vulnerability in bug bounty programs.

The *platforms* table represents the various bug bounty platforms such as Bugcrowd, HackerOne, and similar platforms. The following are the *platforms* fields:

- **id**: A unique identifier for each platform (primary key)
- **name**: The name of the platform (such as Bugcrowd, HackerOne)

The *programs* table represents individual bug bounty programs run by different organizations on the platforms. Each program specifies the scope and rules for identifying and reporting vulnerabilities. These are specific bug bounties listed under the *platforms*. The following are the *programs* fields:

- **id**: A unique identifier for each program (primary key)
- **name**: The name of the bug bounty program
- **platformId**: A foreign key that links the program to a specific platform

The *root domains* are the primary domains that are in scope for a bug bounty program. Each program may cover one or multiple root domains that you are allowed to test. The following are the *root domains* fields:

- **id**: A unique identifier for each root domain (primary key)
- **name**: The name of the root domain (such as websploit.org)
- **programId**: A foreign key linking the root domain to a specific bug bounty program

These are the *subdomains* under the root domains that are also in scope for the bug bounty. A root domain like websploit.org may have subdomains such as api.websploit.org, which are part of the same scope for vulnerability testing. The following are the *subdomains* fields:

- **id**: A unique identifier for each subdomain (primary key)
- **name**: The name of the subdomain
- **rootDomainId**: A foreign key linking the subdomain to a specific root domain

The *IPs* table stores IP addresses associated with subdomains. These IP addresses are the targets that may be probed for vulnerabilities such as open ports, services, and misconfigurations. The following are the fields in the *IPs* table:

- **id**: A unique identifier for each IP address of the corresponding subdomain
- **address**: The IP address of the subdomain
- **subdomainId**: A foreign key linking the IP address to a specific subdomain

The *ports* table contains information about open ports on specific IP addresses. Identifying open ports is critical for understanding the services running on the target, which can reveal vulnerabilities such as misconfigurations or exposure of sensitive services. The following are the fields for *ports*:

- **id:** A unique identifier for each port
- **number:** The port number (such as 80, 443, 22)
- **ipId:** A foreign key linking the port to a specific IP

The *URLs* table tracks specific web addresses within the subdomains that may need further inspection for vulnerabilities such as XSS, CSRF, or SQL injection. The following are the URLs fields:

- **id:** A unique identifier for each URL
- **address:** The specific URL address (such as `https://api.websploit.org/login`)
- **ipId:** A foreign key linking the URL to a specific IP address

The *vulnerabilities* table stores information about specific vulnerabilities found during the bug hunting process. This is one of the core elements of this automated system because it represents the outcomes of ethical hacking and bug hunting efforts. The following are the URL fields:

- **id:** A unique identifier for each vulnerability (primary key)
- **description:** A textual description of the vulnerability (such as “SQL Injection in the login page”).
- **ipId:** A foreign key linking the vulnerability to a specific IP

How can AI help with bug hunting in this architecture? AI can correlate information between different components, such as IPs, URLs, vulnerabilities flagged, and open ports, to paint a complete picture of potential attack surfaces. It can even analyze historical data from other scans to similar applications. AI could predict which services or ports are more likely to have security flaws based on the type of system, configuration, or past exploitation data.

AI models can be trained to identify common web vulnerabilities, such as XSS or SQL injection, by scanning and analyzing the URLs, parameters, and request-response patterns. AI can automate the process of identifying malicious input or output behavior, improving detection accuracy and reducing false positives.

AI can accelerate subdomain enumeration by using predictive models that recognize patterns in domain naming conventions, automatically suggesting potential subdomains or previously undiscovered targets. AI-driven DNS reconnaissance tools can quickly uncover new subdomains linked to root domains by cross-referencing large datasets from DNS history and certificates.

What if we give the explanation of Figure 11-2 to tools such as Cursor, GitHub Copilot, or OpenAI's models such as o1? I used OpenAI's o1, and it created Python code that uses SQLAlchemy, an object-relational mapping (ORM) library. Example 11-2 shows the Python code. Additionally, the code includes functions to process and interact with the data, such as adding records and querying vulnerabilities.

EXAMPLE 11-2 AI-Generated Code to Implement the Database to Support the Use Case of Figure 11-2

```
from sqlalchemy import (
    create_engine, Column, Integer, String, ForeignKey, Text, UniqueConstraint
)
from sqlalchemy.orm import sessionmaker, relationship
from sqlalchemy.ext.declarative import declarative_base

# Define the database
engine = create_engine('sqlite:///bug_bounty.db')
Base = declarative_base()

# Define the platforms
class Platform(Base):
    __tablename__ = 'platforms'
    id = Column(Integer, primary_key=True)
    name = Column(String, unique=True)
    programs = relationship('Program', back_populates='platform')

class Program(Base):
    __tablename__ = 'programs'
    id = Column(Integer, primary_key=True)
    name = Column(String)
    platform_id = Column(Integer, ForeignKey('platforms.id'))

    platform = relationship('Platform', back_populates='programs')
    root_domains = relationship('RootDomain', back_populates='program')

class RootDomain(Base):
    __tablename__ = 'root_domains'
    id = Column(Integer, primary_key=True)
    name = Column(String)
    program_id = Column(Integer, ForeignKey('programs.id'))

    program = relationship('Program', back_populates='root_domains')
    subdomains = relationship('Subdomain', back_populates='root_domain')

class Subdomain(Base):
    __tablename__ = 'subdomains'
    id = Column(Integer, primary_key=True)
    name = Column(String)
    root_domain_id = Column(Integer, ForeignKey('root_domains.id'))

    root_domain = relationship('RootDomain', back_populates='subdomains')
    ips = relationship('IP', back_populates='subdomain')
```

```
class IP(Base):
    __tablename__ = 'ips'
    id = Column(Integer, primary_key=True)
    address = Column(String)
    subdomain_id = Column(Integer, ForeignKey('subdomains.id'))

    subdomain = relationship('Subdomain', back_populates='ips')
    ports = relationship('Port', back_populates='ip')
    urls = relationship('URL', back_populates='ip')
    vulnerabilities = relationship('Vulnerability', back_populates='ip')

class Port(Base):
    __tablename__ = 'ports'
    id = Column(Integer, primary_key=True)
    number = Column(Integer)
    ip_id = Column(Integer, ForeignKey('ips.id'))

    ip = relationship('IP', back_populates='ports')

class URL(Base):
    __tablename__ = 'urls'
    id = Column(Integer, primary_key=True)
    address = Column(String)
    ip_id = Column(Integer, ForeignKey('ips.id'))

    ip = relationship('IP', back_populates='urls')

class Vulnerability(Base):
    __tablename__ = 'vulnerabilities'
    id = Column(Integer, primary_key=True)
    description = Column(Text)
    ip_id = Column(Integer, ForeignKey('ips.id'))

    ip = relationship('IP', back_populates='vulnerabilities')

# Create all tables
Base.metadata.create_all(engine)

# Create a session
Session = sessionmaker(bind=engine)
session = Session()

# Functions to process data
def add_platform(name):
    platform = Platform(name=name)
```

```
session.add(platform)
session.commit()
print(f"Added platform: {name}")

def add_program(name, platform_name):
    platform = session.query(Platform).filter_by(name=platform_name).first()
    if not platform:
        print(f"Platform {platform_name} not found.")
        return
    program = Program(name=name, platform=platform)
    session.add(program)
    session.commit()
    print(f"Added program: {name} under platform {platform_name}")

def add_root_domain(name, program_name):
    program = session.query(Program).filter_by(name=program_name).first()
    if not program:
        print(f"Program {program_name} not found.")
        return
    root_domain = RootDomain(name=name, program=program)
    session.add(root_domain)
    session.commit()
    print(f"Added root domain: {name} under program {program_name}")

def add_subdomain(name, root_domain_name):
    root_domain = session.query(RootDomain).filter_by(name=root_domain_name).
first()
    if not root_domain:
        print(f"Root domain {root_domain_name} not found.")
        return
    subdomain = Subdomain(name=name, root_domain=root_domain)
    session.add(subdomain)
    session.commit()
    print(f"Added subdomain: {name} under root domain {root_domain_name}")

def add_ip(address, subdomain_name):
    subdomain = session.query(Subdomain).filter_by(name=subdomain_name).first()
    if not subdomain:
        print(f"Subdomain {subdomain_name} not found.")
        return
    ip = IP(address=address, subdomain=subdomain)
    session.add(ip)
    session.commit()
    print(f"Added IP: {address} under subdomain {subdomain_name}")
```

```
def add_port(number, ip_address):
    ip = session.query(IP).filter_by(address=ip_address).first()
    if not ip:
        print(f"IP {ip_address} not found.")
        return
    port = Port(number=number, ip=ip)
    session.add(port)
    session.commit()
    print(f"Added port: {number} to IP {ip_address}")

def add_url(address, ip_address):
    ip = session.query(IP).filter_by(address=ip_address).first()
    if not ip:
        print(f"IP {ip_address} not found.")
        return
    url = URL(address=address, ip=ip)
    session.add(url)
    session.commit()
    print(f"Added URL: {address} to IP {ip_address}")

def add_vulnerability(description, ip_address):
    ip = session.query(IP).filter_by(address=ip_address).first()
    if not ip:
        print(f"IP {ip_address} not found.")
        return
    vulnerability = Vulnerability(description=description, ip=ip)
    session.add(vulnerability)
    session.commit()
    print(f"Added vulnerability to IP {ip_address}: {description}")

def get_vulnerabilities_by_program(program_name):
    program = session.query(Program).filter_by(name=program_name).first()
    if not program:
        print(f"Program {program_name} not found.")
        return []
    vulnerabilities = session.query(Vulnerability).join(IP).join(Subdomain).
    join(RootDomain).filter(
        RootDomain.program == program
    ).all()
    return vulnerabilities

# Example usage
if __name__ == '__main__':
    # Add data
    add_platform('HackerOne')
    add_program('Acme Corp Bug Bounty', 'HackerOne')
```

```
add_root_domain('acme.com', 'Acme Corp Bug Bounty')
add_subdomain('api.acme.com', 'acme.com')
add_ip('192.168.1.1', 'api.acme.com')
add_port(80, '192.168.1.1')
add_port(443, '192.168.1.1')
add_url('https://api.acme.com/login', '192.168.1.1')
add_vulnerability('SQL Injection in login page', '192.168.1.1')

# Retrieve vulnerabilities for a program
vulns = get_vulnerabilities_by_program('Acme Corp Bug Bounty')
for vuln in vulns:
    print(f"Vulnerability ID: {vuln.id}, Description: {vuln.description}")
```

The code in Example 11-2 provides a foundational structure for your automated system. You can extend it by adding more processing functions, integrating actual reconnaissance tools, and implementing vulnerability scanning methods.

Tip

For a real-world application, consider adding error handling, input validation, and logging. You might also want to integrate external libraries and tools for tasks like DNS enumeration, port scanning, and vulnerability assessment to automate the data population.

AI Capabilities of Bug Bounty Platforms

Bug bounty platforms like Bugcrowd are all leveraging AI to enhance their services, streamline operations, and provide more security solutions to their clients. AI's integration into these platforms enables continuous monitoring, automated vulnerability detection, and intelligent threat analysis, among other capabilities. Bugcrowd has an AI-driven service called Continuous Attack Surface Penetration Testing (CASPT).

CASPT uses AI models to continuously scan and identify new or modified assets within an organization's digital footprint. This effort includes detecting new applications, services, or infrastructure components that may have been added or altered since the last assessment.

AI helps establish a baseline of the organization's attack surface by analyzing historical data. It then monitors for deviations or changes, ensuring that any new vulnerabilities introduced by modifications are promptly identified. Based on the continuous monitoring of assets, AI determines optimal times to initiate penetration tests. By doing so, it guarantees that testing is both timely and relevant, focusing on areas with the highest risk or recent changes.

Bugcrowd's acquisition of a company named Informer helped it in advanced external attack surface management (EASM) capabilities. AI algorithms merge detailed asset information from Informer with Bugcrowd's existing vulnerability databases. You can initiate new penetration tests directly from the EASM dashboards, with AI managing the integration and ensuring that tests are aligned with the current state of the attack surface.

A Comprehensive View of an Organization's External Risk Exposure

An organization's attack surface extends far beyond what is readily visible. The modern typical ecosystem consists of web domains, subdomains, IP addresses, cloud services, APIs of hundreds of applications (including AI applications), and more—each potentially serving as an entry point for an attacker. The first step in automated external attack surface management is to gain visibility into both known and unknown assets that an adversary might exploit. AI-powered tools can scan, map, and inventory these assets, giving you an "attacker's perspective."

Figure 11-3 illustrates a process flow or a framework for external attack surface discovery and exploitation. The figure is structured with concentric circles, which could symbolize different layers or steps in the process. Each layer is associated with one of four core concepts: discovering assets, monitoring changes, getting actionable insights, and amplifying security testing.

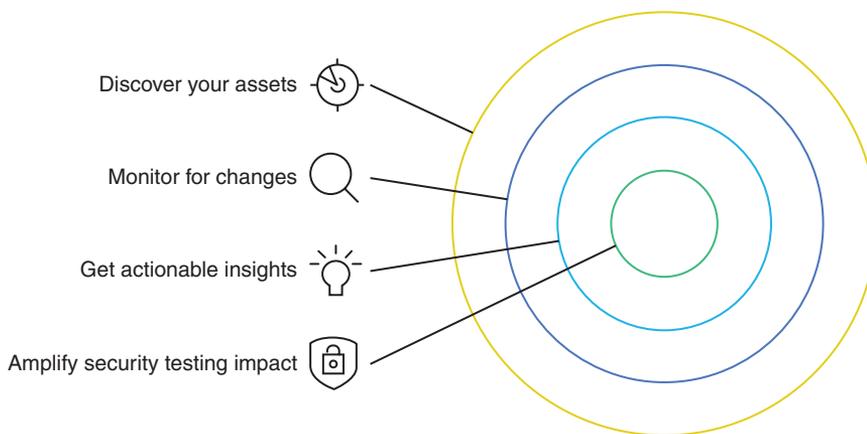


Figure 11-3

A Framework for External Attack Surface Discovery and Exploitation

The first goal is to find and identify all external assets, known and unknown, that could be part of an organization's footprint. Then you continuously watch for any alterations or updates in the external infrastructure or applications, ensuring that no new exposures go unnoticed. Subsequently, you turn the data gathered through monitoring into meaningful insights, which you can use to formulate your attack plan. You could integrate this process with red teaming, penetration testing, and crowdsourced bug bounties.

Vulnerability Prioritization Using AI

AI can assist in prioritizing vulnerabilities based on potential exploitability and risk. By analyzing factors such as Common Vulnerability Scoring System (CVSS) scores, the affected services, historical exploit data, and business impact, AI can help security teams focus on the most critical vulnerabilities first.

Let's start by defining CVSS, the Exploit Prediction Scoring System (EPSS), and CISA's Known Exploited Vulnerabilities (KEV) Catalog, and then examine how AI can assist in vulnerability prioritization using these systems.

A Quick Refresher About CVSS, EPSS, and CISA's KEV

CVSS is an open framework for communicating the characteristics and severity of software vulnerabilities. It provides a way to capture the principal technical characteristics of a vulnerability and produce a numerical score reflecting its severity (0–10). The CVSS score can then be translated into a qualitative representation (such as low, medium, high, and critical) to help organizations properly assess and prioritize their vulnerability management processes. You can access the latest CVSS specification at <https://www.first.org/cvss>.

CVSS consists of four metric groups: Base, Threat, Environmental, and Supplemental. Base metrics represent the intrinsic qualities of a vulnerability. The Threat metric group represents the traits of a vulnerability connected to potential threats, which can evolve over time, though they may remain consistent across different user environments.

The Environmental metric group captures the features of a vulnerability that are specific and relevant to an individual consumer's environment. This includes factors such as the existence of security controls that could reduce or eliminate the impact of a successful attack, as well as the criticality of the affected system within the overall technology infrastructure. The Supplemental metric group consists of metrics that offer additional context and measure external attributes of a vulnerability. The answers to these metrics are defined by the CVSS consumer, enabling them to use a risk analysis system tailored to their environment to assign locally relevant severity to the metrics and values.

EPSS is a data-driven effort for estimating the probability that a software vulnerability will be exploited in the wild. Unlike CVSS, which focuses on the technical severity of a vulnerability, EPSS aims to predict the likelihood of a vulnerability being exploited based on real-world data and machine learning techniques. You can access the latest EPSS specification at <https://www.first.org/epss>.

The following are a few key features of EPSS:

- Provides a probability score between 0 and 1 (0% to 100%)
- Uses current threat information and real-world exploit data
- Updates scores daily based on new information
- Helps prioritize vulnerabilities based on exploitation likelihood rather than just technical severity

The KEV Catalog is maintained by the Cybersecurity and Infrastructure Security Agency (CISA) and serves as an authoritative source of vulnerabilities that have been exploited in the wild. It is designed to help organizations prioritize remediation efforts for vulnerabilities that pose significant risks. You can access CISA's KEV at <https://www.cisa.gov/kev>.

KEV focuses on vulnerabilities with known exploit code or active exploitation and provides remediation due dates for each vulnerability. This list is updated regularly as new exploited vulnerabilities are identified.

How AI Can Help Bug Hunters and Organizations Prioritize Vulnerabilities Using CVSS, EPSS, and KEV

AI can help integrate data from CVSS, EPSS, and KEV along with other relevant sources to create a comprehensive view of vulnerability risks. Machine learning algorithms can analyze this data to identify patterns and correlations that humans might miss, providing a more nuanced understanding of vulnerability priorities. Figure 11-4 shows how you could build AI systems to analyze and prioritize vulnerabilities on a large scale using data from CVSS, EPSS, and KEV.

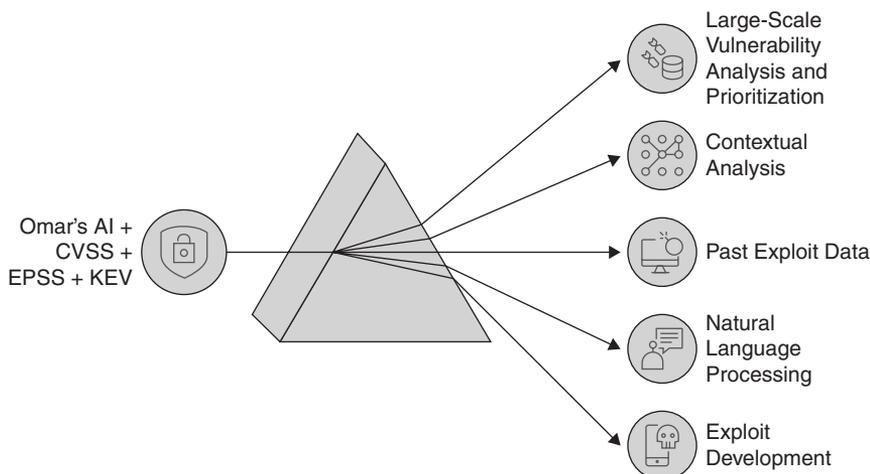


Figure 11-4

AI Analysis of CVSS, EPSS, and KEV

By leveraging historical data from CVSS, EPSS, and KEV, AI can develop predictive models to estimate the likelihood of future exploits. AI can consider the specific context of an organization or bug hunter's environment when prioritizing vulnerabilities. By analyzing factors such as the organization's network architecture, deployed assets, and security controls, you can use AI to provide tailored prioritization recommendations to your client that go beyond generic CVSS or EPSS scores.

AI-enabled systems can continuously monitor threat intelligence feeds, including updates to the KEV Catalog, and automatically adjust vulnerability discovery priorities based on new information.

You can leverage GenAI with NLP capabilities to analyze vulnerability descriptions, security advisories, and exploit discussions in forums to extract additional context and severity indicators. This information can be used to refine prioritization beyond what CVSS, EPSS, and KEV provide numerically.

By analyzing the characteristics of vulnerabilities that have progressed from CVSS scoring to EPSS high-probability to KEV listing, you can potentially predict which newly discovered vulnerabilities are most likely to follow a similar path. This analysis can help prioritize vulnerabilities that may not yet have high EPSS scores or KEV entries but have characteristics suggesting they may soon become critical.

Based on the vulnerability type and target system, AI could generate basic exploit templates or code skeletons, which an ethical hacker could then refine and customize. AI can parse and analyze large volumes of technical documentation, APIs, and code comments to identify potential weak points or unintended functionality that could be exploited.

AI-Assisted Exploit Development Using CISA's KEV Catalog and Other Information

AI systems could potentially adapt known exploit techniques to new environments or slightly different versions of software, helping to test the extent of a vulnerability. AI assistants could work alongside human researchers, offering suggestions, checking for errors, and providing relevant information during the exploit development process.

AI-Created Scanner Templates

The ProjectDiscovery Nuclei Scanner is a fast and extensible open-source vulnerability scanner designed to identify and mitigate security vulnerabilities across modern applications, infrastructure, cloud platforms, and networks.

Nuclei uses a template-driven approach, where each template is a YAML file that defines the steps needed to detect specific vulnerabilities. This approach allows for highly customizable and targeted scans. It also supports a massive library of community-curated templates, which are continually updated to include the latest vulnerabilities in many systems.

The ProjectDiscovery Cloud Platform (PDCP) is a cloud-based security solution targeted at delivering ongoing visibility into your external attack surface by identifying exploitable vulnerabilities and misconfigurations. It is designed to address multiple use cases and can scale to support the essential workflows that bug bounty hunters and security teams require to find vulnerabilities. You can access PDCP at <https://cloud.projectdiscovery.io>.

Figure 11-5 shows the PDCP dashboard.

As previously mentioned, Nuclei supports a massive library of community-curated templates. Figure 11-6 shows an example of a community template.

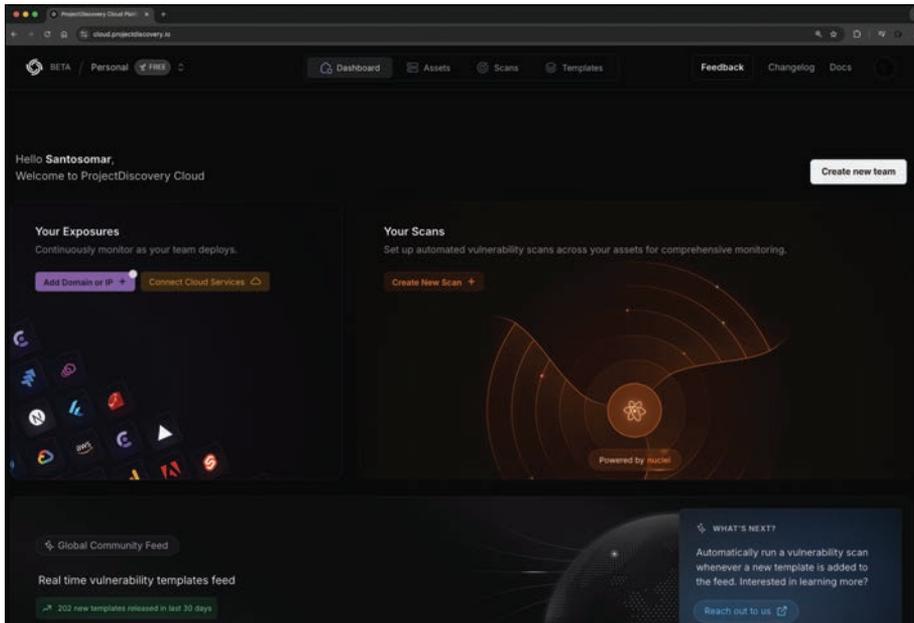


Figure 11-5
The PDCP Dashboard

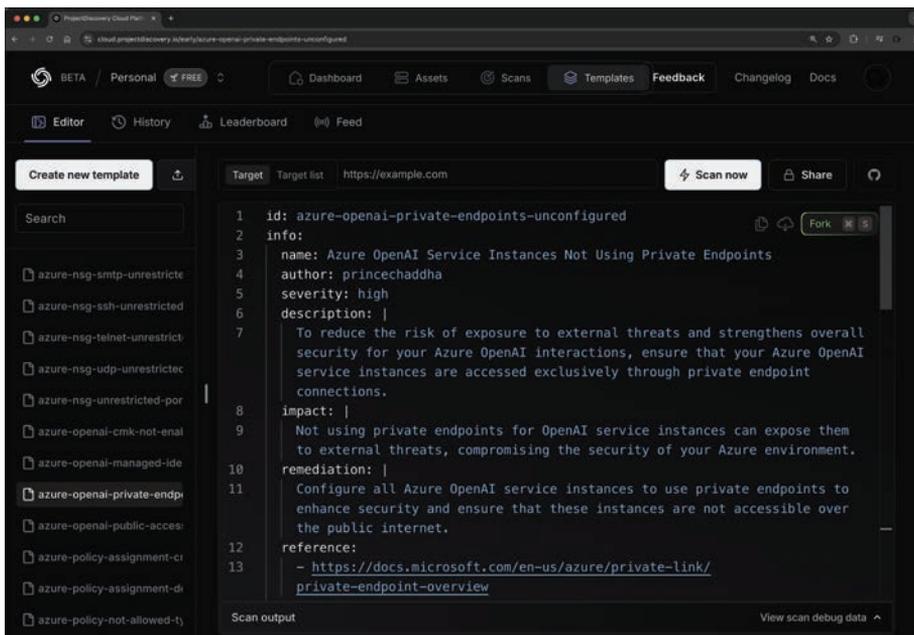


Figure 11-6
Community Template for Azure OpenAI Service Instances Not Using Private Endpoints

The Nuclei template in Figure 11-6 is designed to identify Azure OpenAI Service instances that are not configured to use private endpoints. It highlights that failing to use private endpoints can leave OpenAI service instances exposed to external attacks.

Figure 11-7 shows the interface of ProjectDiscovery's Cloud Platform, where Nuclei templates can be created using AI. In this example, the interface displays a vulnerability scan template related to an insecure direct object reference (IDOR) vulnerability.

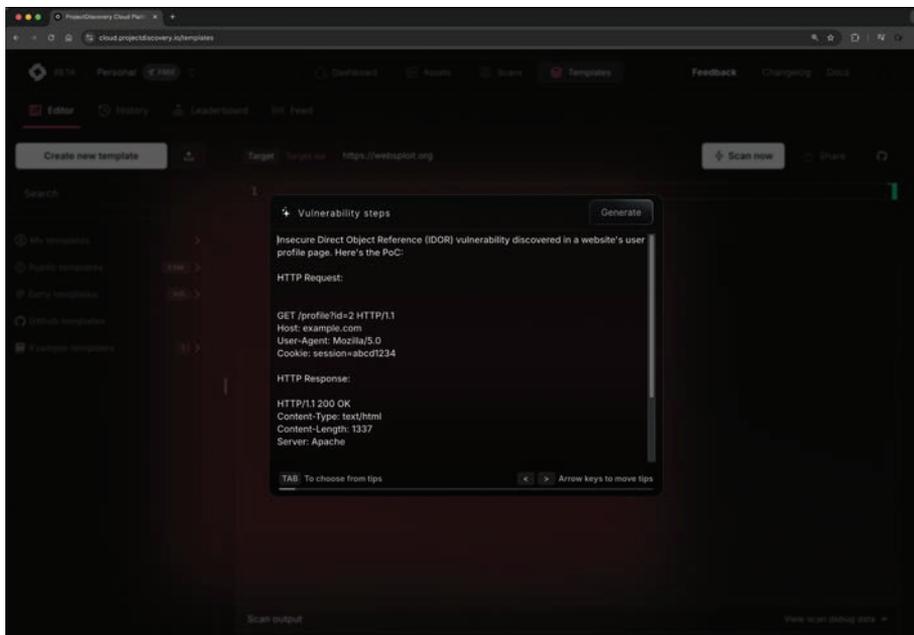


Figure 11-7
AI-Generated Nuclei Vulnerability Scanner Template

Example 11-3 shows the full template generated by the AI feature illustrated in Figure 11-7.

EXAMPLE 11-3 The AI-Generated Nuclei Scanner Template

```
id: userprofile-idor
```

```
info:
```

```
  name: Insecure Direct Object Reference (IDOR) in User Profile Page
  author: ProjectDiscoveryAI
  severity: high
  description: |
```

The application exposes sensitive information of a user (ID: 2) who is not the authenticated user (session: abcd1234), leading to an IDOR vulnerability.

```
http:
- raw:
  - |
    GET /profile?id=2 HTTP/1.1
    Host: {{Hostname}}
    User-Agent: Mozilla/5.0
    Cookie: session=abcd1234

matchers-condition: and
matchers:
- type: word
  words:
    - "Welcome, otheruser"
- type: status
  status:
    - 200
```

The AI model used in Example 11-3 generated a proof-of-concept (PoC) template that provides an HTTP request to potentially find IDORs. The generated request targets a profile page (**/profile?id=2**), indicating that by manipulating the id parameter in the URL, a user could access another user's profile.

After the template is generated, you can use it to scan a target application (websploit.org in the example shown in Figure 11-8).

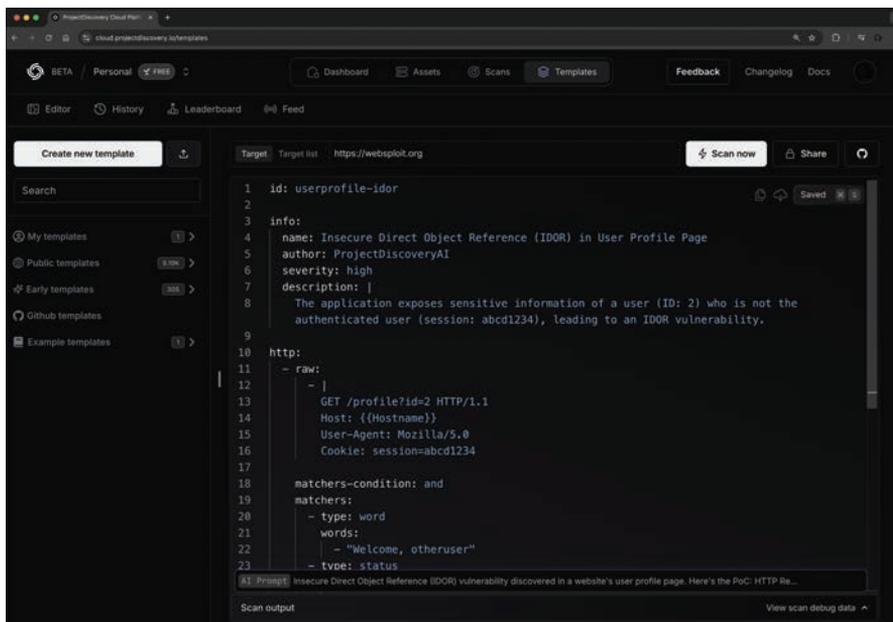


Figure 11-8
Scanning websploit.org Using the AI-Generated Template

AI Model Training, Fine-Tuning, and RAG for Bug Bounties

In Chapter 8, you learned about AI model training, fine-tuning, and RAG. As a security researcher participating in bug bounties, you can fine-tune and use RAG to enhance your ability to identify and report vulnerabilities.

Deploying AI Models

You can easily deploy AI models in cloud platforms such as Azure AI Studio, AWS Bedrock, and Google Cloud Vertex AI. Figure 11-9 shows a model deployment screen within Azure AI Studio, specifically for deploying the Meta-Llama model. The project resource is named omar-demo-project-23, which is the current environment in which the model will be deployed. The custom name for the endpoint is omar-ai-abc123. After deployment, the system will automatically generate an endpoint URL. The deployment is named meta-llama-3-1-8b-1 because it is based on the Meta-Llama 3.1 model using the 8-billion parameter version. Inferencing data collection is disabled, meaning the deployment will not collect data during inference runs.

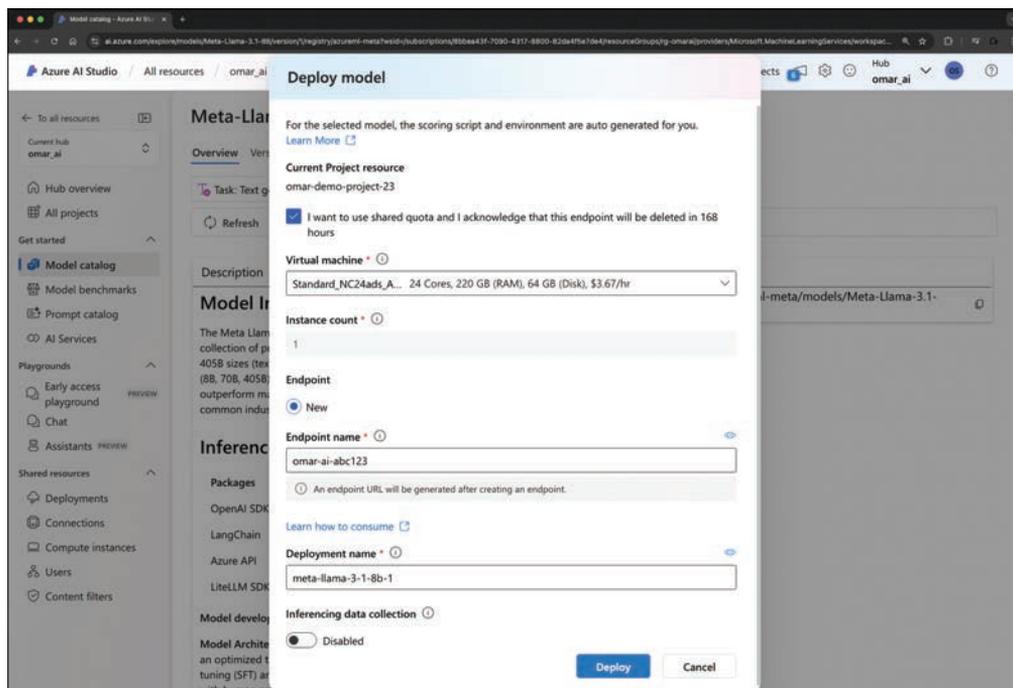


Figure 11-9
Azure AI Studio Model Deployment Screen

Fine-Tuning AI Models

Fine-tuning involves adjusting a pretrained model using a specific dataset to enhance its performance for particular tasks. This process is invaluable for bug hunters because it allows you to tailor AI models to recognize specific types of vulnerabilities or security issues unique to certain software or systems. By fine-tuning models with data from known vulnerabilities, you can improve the accuracy and relevance of your automated tools and potentially AI agents.

Figure 11-10 shows the interface of Google Cloud's Vertex AI and its Colab Enterprise environment. This notebook is used for fine-tuning the Meta-Llama 3.1 8B-Instruct model.

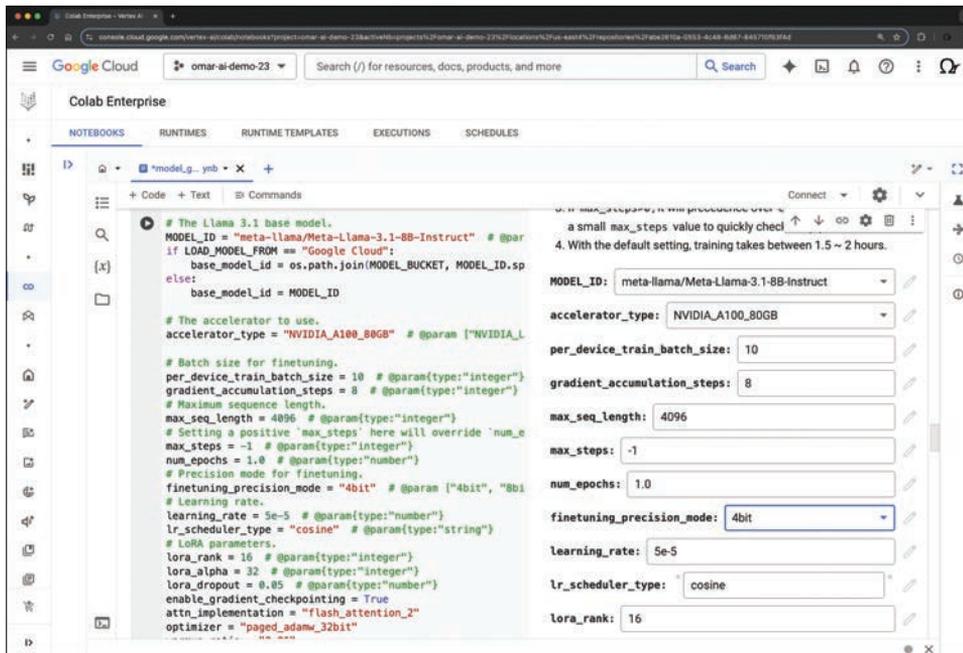


Figure 11-10
Fine-Tuning AI Models in Google Vertex AI

In Figure 11-10, the model is being run on an NVIDIA A100 GPU, as specified in the **accelerator_type** parameter.

The following are additional details about the fine-tuning parameters shown:

- **Batch Size:** The **per_device_train_batch_size** is set to 10, meaning each GPU will process one batch at a time.
- **Gradient Accumulation Steps:** This parameter is set to 8, meaning gradients will accumulate over 8 batches before updating the model weights, helping manage memory constraints.
- **Max Sequence Length:** The model will process sequences of up to 4096 tokens at a time.

- **Max Steps:** This parameter is set to -1 , which suggests there is no limit on the number of steps for training.
- **Epochs:** Training will run for 1.0 epoch, meaning the entire dataset will be passed through the model once.
- **Precision:** Fine-tuning is performed using 4-bit precision (`fine-tuning_precision_mode`), which helps reduce the memory usage while still being performant.
- **Learning Rate:** The learning rate is set to $5e-5$ (0.00005), which controls how much the model adjusts weights with each training step.
- **Learning Rate Scheduler:** The learning rate will follow a “cosine” decay schedule, gradually reducing as training progresses.

LoRA and QLoRA

Low-rank adaptation (LoRA) is a machine learning technique designed to efficiently fine-tune language models by introducing trainable low-rank matrices into each layer of a pretrained model while keeping the original weights frozen. This approach significantly reduces the number of trainable parameters, making the process faster and more memory-efficient without compromising performance. *QLoRA* builds on LoRA by incorporating 4-bit quantization, which further reduces memory usage and enables the fine-tuning of extremely large models, such as those with 65 billion parameters, on hardware with limited resources like a single 48GB GPU. This method maintains high performance levels comparable to full 16-bit fine-tuning by combining quantization with low-rank adapters, thus democratizing access to efficient model tuning.

The following are the LoRA parameters used in the example shown in Figure 11-10:

- **LoRA Rank:** Set to 16, this refers to the rank in the low-rank adaptation technique, commonly used to fine-tune large models efficiently by focusing only on specific layers.
- **LoRA Alpha:** Set to 32, this controls the scaling factor for LoRA, which influences the weight adjustment during fine-tuning.
- **LoRA Dropout:** Set to 0.05, meaning 5 percent of the neurons are randomly dropped during training to prevent overfitting.

There are many other fine-tuning settings, but the following are displayed in Figure 11-10:

- **Gradient Checkpointing:** Uses Enabled (True), which helps save memory by recomputing some layers during the backward pass.
- **Attention Implementation:** Uses `flash_attention_2`, a more memory-efficient attention mechanism that allows for faster processing of large sequences.

Using RAG and AI Agents for Bug Bounty Hunting

RAG combines the capabilities of language models with data retrieval from vector databases and other sources, providing contextually relevant information that enhances the model inference and reduces the likelihood of hallucinations or confabulations. For bug bounty hunters, RAG can be used to access up-to-date information about potential vulnerabilities or exploits as they emerge.

Many different tools such as LlamaIndex, LangChain, LangGraph, and cloud services such as Google Vertex AI can help you build RAG implementations and AI agents. Figure 11-11 shows the Google Vertex AI Agent Console, which is used for building and managing AI agents. A custom AI agent named Omar's Bug Bounty Hunter Agent is being configured.

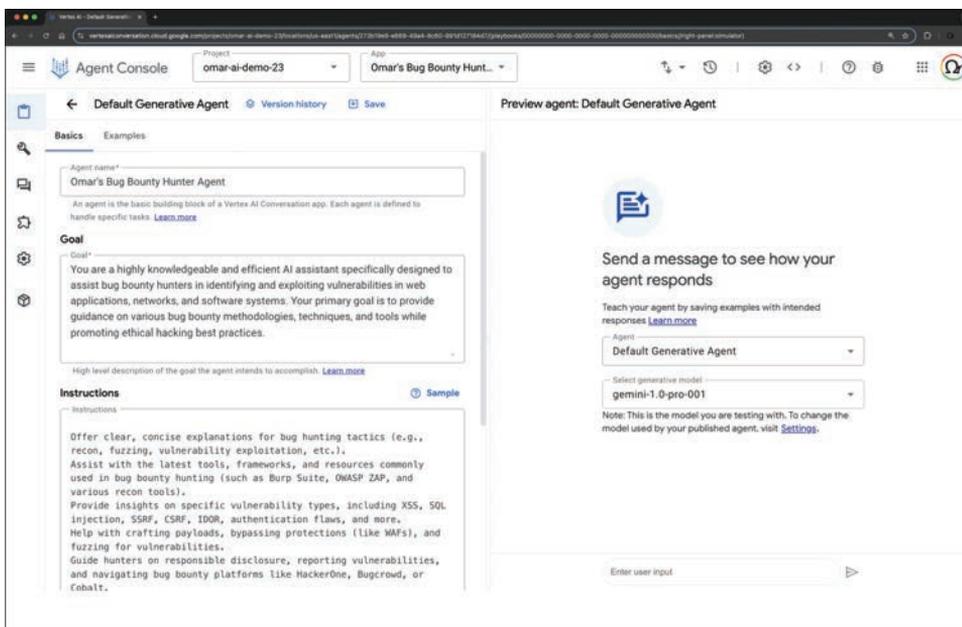


Figure 11-11
Creating AI Agents in Google Vertex AI

The configured “Goal” provides a description of the agent’s purpose, which is to assist bug bounty hunters in identifying and exploiting vulnerabilities while promoting ethical hacking practices. The “Instructions” section provides guidelines for the agent’s behavior. These instructions include offering explanations for bug hunting tactics, assisting with tools and frameworks, providing insights on various vulnerability types, and guiding responsible disclosure.

You can also configure data stores (using the built-in vector database) to use your data and very easily create a RAG deployment. These Vertex AI agents feature special state handlers known as *data store handlers*. These handlers allow the data store agent to engage in conversations with end users about the stored content.

Tool Calling

You can also link your agent to external tools (*tool calling*) and enable it to perform tasks more effectively. Figure 11-12 shows an example of tool calling in Google Vertex AI.

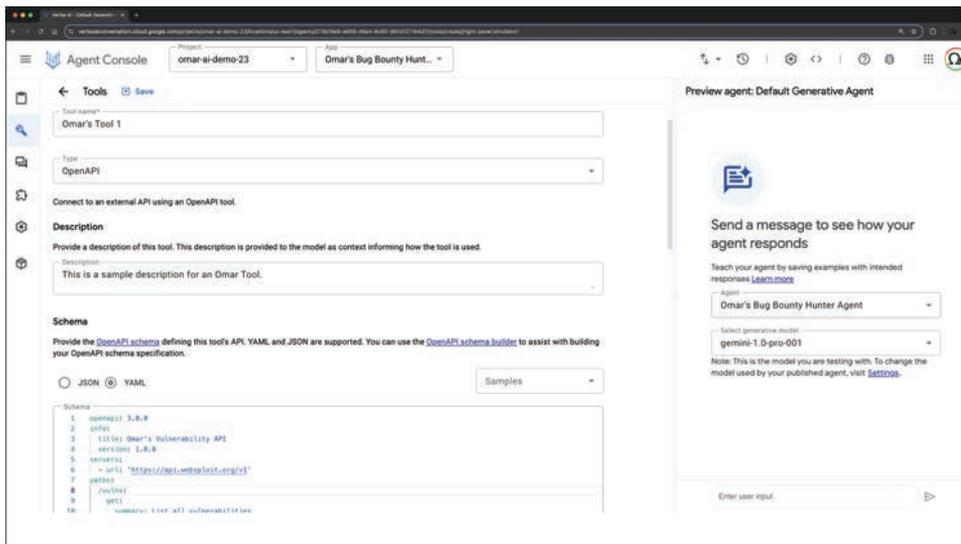


Figure 11-12
Tool Calling in Google Vertex AI

Use Case: Improving Reporting and Documentation

RAG and AI-driven tools can significantly enhance the quality, consistency, and efficiency of vulnerability reporting and documentation in the context of bug bounty hunting, red teaming, and traditional penetration testing. These advanced systems can take basic input about a discovered vulnerability and generate a structured report template, including standard sections like executive summary, technical details, impact assessment, and remediation recommendations. Automation can help you ensure that no crucial elements are overlooked in the reporting process.

You learned that when it comes to vulnerability severity assessment, AI can analyze the vulnerability details against databases of known issues to suggest a severity rating based on frameworks like CVSS and EPSS. The AI can then justify the severity score, explaining which factors contributed most significantly to the rating. This approach provides a clear, objective basis for prioritizing vulnerabilities.

One of the most valuable features of AI in this context is its ability to expand brief notes into clear, step-by-step reproduction instructions. It can suggest additional edge cases or variations

to test, ensuring comprehensive coverage of the vulnerability. For common vulnerability types, AI could even generate or suggest basic POC code snippets using tools like Cursor or GitHub Copilot.

Based on the vulnerability type and affected systems, AI can offer specific remediation recommendations using your data that could include best practices and other insights. It can prioritize these recommendations and estimate the effort required for each, helping your client design their mitigation plans effectively. Of course, AI can also improve the clarity and professionalism of your writing and fix grammatical issues. It can adjust the technical level of the content for different audiences, making the report accessible to both technical and nontechnical stakeholders.

AI can help you tailor the report to address stakeholders' specific concerns and technical stack. This customization makes the report more relevant and actionable for the client. AI can also assist in managing different versions of the report, highlighting changes and allowing for collaborative editing with comments and suggestions. You can leverage AI in this comprehensive manner to produce more thorough, professional, and actionable reports in less time.

Challenges of Using AI for Bug Bounty Hunting

Not everything is perfect with AI. There are several challenges when you integrate AI into bug bounty hunting, especially with AI models that have guardrails or are censored.

Censored Models and Guardrails

Censorship and *guardrails* refer to the restrictions imposed on AI models to prevent them from generating harmful, unethical, or legally problematic content. While these guardrails are essential for ensuring ethical AI usage, they can impede bug bounty hunting efforts in several ways:

- **Limited Payload Generation:** AI models with strong guardrails may refuse to create or assist in crafting potentially malicious payloads, even when these are essential for security testing. For example, if you're trying to identify SQL injection or cross-site scripting (XSS) vulnerabilities, an AI model might block requests to generate queries or scripts that it deems dangerous, even if they are being used ethically within a penetration testing environment.
- **Restricted Vulnerability Testing:** Many censored models are programmed to avoid generating code or commands that could be used in hacking or exploitation scenarios. In bug bounty hunting, ethical hackers need to test for security flaws like buffer overflows, file inclusions, or even zero-day vulnerabilities. AI models with tight restrictions may prevent access to code snippets, payloads, or scripts needed to identify these flaws. This is why models like White Rabbit Neo are becoming very popular among the cybersecurity community.
- **False Positives/Negatives:** Guardrails and censoring may lead to false negatives, where the model fails to help detect a valid vulnerability in scanners, fuzzers, and other systems because it is avoiding potentially "dangerous" outputs.

Hallucinations or Confabulations

Hallucinations or *confabulations* in the context of AI refer to instances where the AI model generates information that is factually incorrect, made up, or not grounded in reality. This is an important concept to understand when working with AI systems, especially in critical areas like security research and bug bounty hunting. Let me elaborate. Hallucinations occur when an AI model produces content that seems plausible but is actually false or nonsensical. This can happen because AI models are trained on vast amounts of data and learn to generate human-like text, but don't have true understanding or real-world knowledge in the way humans do. AI models make statistical associations between words and concepts, which can sometimes lead to incorrect outputs. In bug bounty hunting, hallucinations can present serious challenges, especially when accuracy and precision are critical. This is why "grounding" by using techniques such as RAG is really important.

The following are a few examples of hallucinations or confabulations in ethical hacking scenarios:

- **Incorrect Vulnerability Reporting:** AI models prone to hallucinations may identify vulnerabilities that do not exist. This could lead to bug bounty hunters wasting significant time chasing false vulnerabilities that were never present in the system.
- **Faulty Code or Exploits:** If an AI hallucinates while generating exploitation scripts or proofs of concept, it may provide incorrect or incomplete code, leading to failed attempts to reproduce vulnerabilities or broken exploitation logic. For example, an AI might hallucinate a new method to exploit a flaw, but the code it generates is nonfunctional or dangerous to the system itself.
- **Misleading Contextual Understanding:** AI hallucinations can also affect how models interpret security-related context. For example, when analyzing complex security protocols or responses from a web application firewall (WAF) or a cloud security application, an AI model might generate completely fabricated conclusions or explanations, misleading a bug hunter into wrong assumptions about the system's behavior.
- **Incomplete Testing Flows:** AI models, due to guardrails or insufficient training on adversarial tasks, may fail to complete multistep attack scenarios that are essential for bug bounty hunting. These multistep attacks often involve intricate chains of vulnerabilities that require creative, out-of-the-box thinking, which current AI models might not be equipped for.
- **Inability to Handle Zero-Days:** AI models may struggle to discover zero-day vulnerabilities, especially since they rely on training data. Zero-days are unknown vulnerabilities, meaning no data exists to train models on these flaws. This is a significant limitation when trying to use AI in environments where cutting-edge or undiscovered vulnerabilities are the primary focus.

Compliance with Bug Bounty Guidelines

Many bug bounty programs have specific rules about how vulnerabilities should be tested or reported. If AI-generated exploits or tests do not align with these guidelines, the hunter risks violating the program's terms, which can lead to disqualification or legal ramifications.

Test Your Skills

Multiple-Choice Questions

1. Which of the following is NOT a primary goal in the framework for external attack surface discovery and exploitation?
 - A. Discover
 - B. Monitor
 - C. Actionable insights
 - D. Patch management
2. What is the purpose of the Common Vulnerability Scoring System (CVSS)?
 - A. To predict the likelihood of a vulnerability being exploited
 - B. To communicate the characteristics and severity of software vulnerabilities
 - C. To list vulnerabilities that have been exploited in the wild
 - D. To generate exploit code for known vulnerabilities
3. What does EPSS stand for in the context of vulnerability assessment?
 - A. Exploit Prediction Scoring System
 - B. External Penetration Security System
 - C. Enhanced Protection Scanning Service
 - D. Exploit Prevention Scoring Standard
4. Which organization maintains the Known Exploited Vulnerabilities (KEV) Catalog?
 - A. NIST
 - B. OWASP
 - C. CISA
 - D. ISO
5. What is the primary purpose of the ProjectDiscovery Nuclei Scanner?
 - A. To generate AI models for vulnerability detection
 - B. To create vulnerability reports
 - C. To identify and mitigate security vulnerabilities across various platforms
 - D. To manage bug bounty programs

6. What type of file format does Nuclei use for its templates?
 - A. STIX
 - B. XML
 - C. YAML
 - D. TXT
7. What is LoRA in the context of AI model fine-tuning?
 - A. Low-range analysis
 - B. Low-rank adaptation
 - C. Linear optimization for rapid advancement
 - D. Logarithmic rational approximation
8. What does RAG stand for in the context of AI and bug bounty hunting?
 - A. Rapid attack generation
 - B. Retrieval-augmented generation
 - C. Risk assessment guide
 - D. Responsive AI governance
9. What is a potential challenge of using AI models with strong guardrails in bug bounty hunting?
 - A. Increased accuracy in vulnerability detection
 - B. Faster processing of large datasets
 - C. Limited ability to generate potentially malicious payloads for testing
 - D. Improved compliance with bug bounty program guidelines
10. What is a confabulation in the context of AI-assisted bug bounty hunting?
 - A. A visual glitch in the AI interface
 - B. A sudden increase in processing speed
 - C. Generation of factually incorrect or made-up information
 - D. A temporary loss of connection to the AI service
11. Which of the following is NOT a metric group in the Common Vulnerability Scoring System (CVSS) version 4.x?
 - A. Base
 - B. Threat

- C. Environmental
 - D. Exploitation
12. What is the primary advantage of using QLoRA over traditional LoRA?
- A. It allows for fine-tuning of larger models on limited hardware.
 - B. It increases the number of trainable parameters.
 - C. It eliminates the need for GPU acceleration.
 - D. It provides more accurate results in all scenarios.
13. What is the primary purpose of tool calling in AI agents for bug bounty hunting?
- A. To generate more accurate vulnerability reports
 - B. To enable the AI to perform tasks using external tools
 - C. To improve the AI's natural language processing capabilities
 - D. To create new exploits automatically using an IDE
14. What is the purpose of gradient accumulation in the context of AI model fine-tuning?
- A. To increase the learning rate
 - B. To reduce memory constraints
 - C. To improve model accuracy
 - D. To speed up the training process

Index

Numerics

0365-realminfo, 76

A

access, web shell, 154, 155

accounts, 151–152. *See also* authentication
ASREProasting, 240–242
creating on compromised systems, 181
developer, 167
SSO (single sign-on), 210

ACE (access control entry), 262–263

active reconnaissance, 336–337, 377, 378–380

AD (Active Directory), 207, 208. *See also*
BloodHound; LDAP (Lightweight directory
Access Protocol)

ACE (access control entry), 262–263

attack techniques, 229

centralized data repository, 208, 211

DACLs (discretionary access control lists),
262–263

DC (domain controller), 208, 234–235

DNS (Domain Name Service), 221–223

domains, 212–213

forest, 211–213, 215–216

functional level, 214

root domain, 214–215

high availability, 209

Kerberoasting, 243, 244–246

Kerberos, 219

lab setup, 229–230

LDAP (Lightweight directory Access
Protocol), 216–218

MFA (multifactor authentication), 209

nmap scan report, 233–234

object, DN (distinguished name), 216

OUs (organizational units), 213

policies, 216–217

reconnaissance, 231

replication, 208

inbound, 209

outbound, 208–209

schema, 211, 213

security, audits, 209

SSO (single sign-on), 210

two-way transitive trusts, 213–214

user enumeration, 235–236, 237–239

using cme, 235–236

using enumdomgroups, 237–238

using net rpc, 238

using rpcclient, 237

Zero Trust, 209

ADCS (Active Directory Certificate Services),
100, 259

- components, 259
- enumeration, 259–260
- finding vulnerabilities with Certipy
 - options, 100–101
 - templates, 101–103
- vulnerabilities, 100
- ADFS (Active Directory Federation Services),
 - Golden SAML Attack, 263–265
- ADSI (Active Directory Service Interfaces), 211
- Advanced Installer, 92
- advanced operators, Google, 360
- adversarial emulation, 9–10, 24
- agent
 - AI, 519
 - C2, 3
- AI, 14. *See also* LangSmith; model, AI; RAG (retrieval-augmented generation)
 - assisted bug bounty hunting, 510
 - AI agents, 519
 - censored models and guardrails, 521
 - integrating data from CVSS, EPSS, and KEV, 511–512
 - RAG, 519
 - tool calling, 520
 - generated Nuclei scanner templates, 512, 513–515
 - hallucination, 372–373, 522
 - malicious model, 3
 - model
 - ChatGPT, 3, 112, 128
 - deploying, 516
 - embeddings, 300–302
 - fine-tuning, 517–518
 - OpenAI, 302
 - poisoning, 2–3
 - WhiteRabbitNeo, 297–298
 - penetration testing tools, 13
 - powered OSINT, 284–289
 - powered SSL/TLS analysis, 291–294
 - prompt engineering, 4
 - public platform, 2
 - red teaming, 282–283, 311–314
 - garak, 319–321
 - keeping pace with evolving technologies, 322–323
 - Prompt Fuzzer, 317–319
 - TAP (Tree of Attacks with Pruning), 315
 - tools for testing AI applications and model security, 316, 322
 - system prompts, 290
 - training
 - data source, 4
 - tainting the data, 1–4
 - using confidential data, 282–283
 - vulnerabilities, 2–3
 - XDR platforms, 150
- alarms, 188
- alert, unauthorized access, 151
- algorithm/algorithmic
 - BM25, 306–307
 - red teaming, 313
- Amass, 371–372
- Anderson, James P., “Computer Security Technology Planning Study”, 7, 8
- ANN (Approximate Nearest Neighbor), 306
- Apache redirector
 - adding in the IP of your C2 server in /etc/hosts, 40–42
 - config, 39–40
 - installing, 38–39
 - installing certificates for your domain, 39
- APIBench, 373
- API/s (application programming interface/s)
 - OWASP Enterprise Security, 447
 - recon, 404–407
 - REST (representational state transfer), 423
 - unprotected, 485–487
- application-level backdoors. *See also* web application/s
 - bind shell
 - comparison with reverse shell, 164
 - how they work, 162–164
 - modified applications and hidden functionalities in legitimate software, 157
 - patching binaries to include a backdoor, 166–167

- reverse shell, 157–161
 - comparison with bind shell, 164
 - how they work, 160–161
 - supply chain, 167
 - web shell, 153
 - access, 154, 155
 - ASP, 154
 - code and command obfuscation, 156
 - header, 156
 - JavaScript, 155
 - PHP, 154
 - storage, 153
 - whoami command, 154
 - APTs (advanced persistent threats), 148
 - ARPANET, 9
 - Arsenal Kit, 35
 - ASLR (Address Space Layout Randomization), 272–273
 - ASM (attack surface management), 331
 - comparison with bug bounty programs, 332
 - tools, 331–332
 - ASN (autonomous system number), 65–66
 - ASP web shell, 154
 - ASREPROasting, 240–242
 - attack. *See also* social engineering
 - adversarial emulation, 9–10
 - Kerberoasting, 99–100
 - scenarios, 148
 - using a C2 server, 185–186
 - watering hole, 2
 - auditing, 151, 209
 - authentication
 - AD (Active Directory)
 - Kerberos, 219
 - NTLM (NT LAN Manager), 218–219
 - identifying type of, 84–87
 - Kerberoasting, 99–100, 228
 - Kerberos
 - Golden Ticket attack, 459
 - unconstrained delegation, 459–460
 - vulnerabilities, 228
 - LDAP (Lightweight directory Access Protocol), 216–218
 - multifactor, 19, 209
 - NTLM (NT LAN Manager), 218, 219
 - vulnerability scans, 397–398
 - automation. *See also* AI
 - bug hunting
 - AI-powered, 500–508
 - Nuclei, 499–500. *See also* Nuclei
 - reconnaissance, 501–502
 - avatar generator, 2
 - awk command, 196–197
 - Azure AD Connect Sync Account Takeover, 265–266
 - Azure CDN, spear phishing infrastructure, 117–118
- ## B
- backdoors, 35–36, 152, 153
 - application-level
 - bind shell, 162–164
 - modified applications and hidden functionalities in legitimate software, 157
 - patching binaries to include, 165–167
 - reverse shell, 157–161, 164
 - web shell, 153–157
 - hardware and firmware, 179–180
 - kernel-level
 - bootkit, 179
 - rootkit, 174–178
 - persistence, 180
 - scheduled jobs and tasks, 180
 - supply chain, 167
 - system-level, 171–174
 - badge/s
 - card credentials, 134
 - cloning, 135–138
 - creating, 138–139
 - facility code, 134
 - high-frequency cards, 133–134
 - iClass legacy cards, 134–135
 - LF cards, 133
 - magstripe, 132

- parity bits, 134
- writing to a, 137
- BadUSB device, 180
- bash script, 0365-realminfo, 76
- behavioral analysis, 149
- binaries
 - GTFOBins, 196–197
 - Living-Off-the-Land, 197–198
 - trojanized, 171–172
- bind shell, 162–164
- black hat, 3
- BlackHills Information Security, 262
- blacklist, 120
- blind SQL injection, 442–444
- BloodHound, 194, 246
 - data collection using the Python ingestor, 246–247
 - queries, 247
 - using for bug bounty, 394–395
- blue team/ing
 - behavioral analysis, 149
 - improving capabilities, 149–152
 - playbooks, 150
 - red team collaboration, 20, 152
 - threat-hunting hypotheses, 150
 - war games, 9
- BM25, 306–307
- BOF (Beacon Object File), 97
- Boolean-based SQL injection, 442–444
- bootkit, 179
- boot-start driver, 177–178
- breach
 - assessing impact, 148
 - presumption of, 149
 - through social engineering, 109–110
- Bright Data, 67
- browser
 - “Bugs Bounty” initiative, 10
 - cookies, 4
 - developer tools, 406
- brute-force attack, 451
 - credential stuffing, 452
 - password spraying, 452
 - specialized tools, 452
- buffer overflow, 267–269, 271–272
- bug bounty hunting, 2, 6, 328
 - AI-assisted, 510
 - AI agents, 519
 - hallucination, 522
 - integrating data from CVSS, EPSS, and KEV, 511–512
 - tool calling, 520
 - using RAG, 519
 - versus ASM solutions, 332
 - automation
 - AI-powered, 500–508
 - Nuclei, 499–500
 - benefits, 328
 - challenges, 329
 - collaborative approach, 10
 - communicating your findings, 407–409
 - evolution, 6–7
 - future of, 12–13
 - getting started, 333
 - limitations of manual bug hunting, 499–500
 - ML models, 14
 - Mozilla, 11
 - Netscape, 10
 - OSINT (open-source intelligence), 337–338
 - platforms, 330–331, 502, 508–509
 - private, 330, 331
 - public, 331
 - reports, 334
 - rewards, 11
 - ROEs (rules of engagement), 334–336
 - scope, 334–336
 - traditional, 498
 - types of, 329
 - using BloodHound, 394–395
 - using RAG, 519
 - versus VDPs (vulnerability disclosure programs), 332–333
 - VRTX (Versatile Real-Time Executive) operating system, 10
 - workflow, 500–501
- Bugcrowd, 11
- “Bugs Bounty” initiative, 10
- building your own web application, 429–430

- Burp Suite, 405, 488–490
 - BurpGPT, 282, 283
 - business logic flaws, 430, 498
 - CWE (Common Weakness Enumeration) ID 840, 430–431
 - exploiting, 431
 - OWASP recommendations, 430
 - bypass tools, 140
 - DDT (Double Door Tool), 141
 - J Tool, 143
 - Quick Jim, 140–141
 - REX Sensor Bypass Gun, 142–143
 - UDT (Under the Door Tool), 141–142
- C**
-
- C2 (Command and Control), 33, 181, 185. *See also* Mythic
 - beacon, 4
 - Cobalt Strike, 35
 - communication channels, 36
 - covert channels, 181. *See also* covert channels
 - DNS tunneling, 182–183
 - Dropbox, 181–182
 - Google Drive and Google Sheets, 182
 - HTTP/HTTPS tunneling, 183–184
 - using nontraditional protocols, 184
 - X, 182
 - frameworks
 - C2 Matrix, 186
 - Cobalt Strike, 35
 - Mythic, 35–36, 165
 - Nighthawk, 35
 - TrevorC2, 186
 - Mythic, 35–36, 165
 - tools, 185–186
 - capturing
 - access control card data, 135–138
 - credentials, 4
 - HTTP packets, 420–421
 - NTLMv2 hash, 251
 - card credentials, 134
 - CBEST (Bank of England's Cyber Assessment Framework), 21–22
 - CDN (content delivery network), 33–34, 37
 - with reverse proxy using rewrites, 37
 - spear phishing infrastructure, 117–118
 - traffic, 34
 - censorship, AI model, 521
 - certificate/s
 - information gathering using CertSPY, 354–355
 - installing
 - on your C2 server, 41–43
 - for your domain, 39
 - mapping, 262
 - obtaining information from, 349–350
 - revocation, 348
 - SSL, 348
 - transparency, 66, 283–284
 - certification, red teamer, 322
 - Certipy
 - ADCS enumeration, 259–260
 - ADCS exploitation, 261–262
 - ESC1 request certificate for administrator, 103
 - options, 100–101
 - templates, 101–103
 - CertSPY, obtaining certificate information, 354–355
 - CGC (control group coordinator), CBEST (Bank of England's Cyber Assessment Framework), 22
 - ChatGPT, 3
 - creating pretext script, 112
 - spear phishing text formulation, 128
 - chipset, smart card, 134
 - Chrome, cookies, 4
 - CISO (chief information security officer), 149
 - clickjacking, 480
 - cloning badges, 135–138
 - cloud, 33
 - assets, identifying, 344–348
 - spear phishing infrastructure, 116
 - Cloud Checker script, 345–348
 - cloudproxy, 67

- cme, 231, 232–233, 235–236, 239
- Cobalt Strike, 35
- code. *See also* command/s
 - compression, 156
 - creating a malicious daemon in Rust, 172–174
 - Empire post-exploitation tool, 192–194
 - error handling, 484
 - gadgets, 273, 274–275
 - hard-coded credentials, 484
 - Linux kernel module rootkit, 174–175
 - Metasploit RDP post-exploitation module, 189
 - obfuscation, 156, 167–171
 - Python, 4
 - signing, 488
 - source
 - comments, 484
 - repositories, 362–363
 - with reverse shell, 160
 - starting a web service to expose PowerShell scripts to compromised hosts, 192
 - vulnerable, 270–272
- collaboration
 - bug bounty programs, 10
 - red team/blue team, 20, 152
 - red team/engineering team, 27
- COM hijacking, 96
- command/s. *See also* tool/s
 - awk, 196–197
 - cme, 239
 - curl, 424
 - Evilginx3, 124–126
 - find, 196
 - injection, 448–449
 - Invoke-MFASweep, 87
 - Linux
 - dig, 340
 - lsmod, 176
 - sudo update-ca-certificates, 68
 - marketplace refresh, 366
 - Meterpreter, 164–165
 - NameSpi, 79–81
 - Netcat, 162–163
 - nslookup, 234–235
 - obfuscation, 156
 - payload, 180
 - pip3 install dnspython, 341
 - PowerShell, 189–190
 - PsExec, 195–196
 - ScareCrow, 91
 - steghide, 199–202
 - tar, 196
 - whoami, 154
 - whois, 342–344, 345
- comments in source code, 484
- communications
 - C2, 36
 - red team, 34, 65
- community profiles, Mythic, 36
- competitor analysis, 4
- confabulation. *See* hallucination
- config, Apache redirector, 39–40
- context preparation, 308–310
- continuous improvement, 25
- continuous testing, 12–13
- control flow obfuscation, 168
- controls, 5
 - physical security, 110
 - technical, 110
- cookie/s
 - extracting, 4
 - HTTP session, 427
 - HTTPOnly flag, 457
 - manipulation attacks, 482
 - nonpersistent, 457
 - stealing, 468–469
- covering your tracks, 199–202
- covert channels, 181
 - DNS tunneling, 182–183
 - Dropbox, 181–182
 - Google Drive and Google Sheets, 182
 - HTTP/HTTPS tunneling, 183–184
 - using nontraditional protocols, 184
 - X, 182
- CrackMapExec. *See* cme

- creating
 - badges, 138–139
 - C2 (Command and Control), 185–186
 - Microsoft Entra domain, 224–228
 - credential/s. *See also* AD (Active Directory); authentication
 - brute-forcing, 451
 - capturing, 4
 - card, 134
 - default, 458–459
 - extracting
 - using Lsassy, 256–257
 - using Secretsdump, 253–256
 - hard-coded, 484
 - identifying associated email addresses, 81
 - identifying authentication type, 84–87
 - obtaining, 97, 148
 - storage, 151
 - stuffing, 452
 - writing to a badge, 137
 - cross-site scripting, 463–464
 - DOM XSS, 469–470
 - mitigations, 471–472
 - redirecting users to malicious sites, 467–468
 - reflected XSS attacks, 465–466
 - stealing cookies, 468–469
 - stored XSS attacks, 466–467
 - XSS evasion techniques, 471
 - CSIRT (computer security incident response team), 149
 - CSRA (Certificate Services Remote Administration), 260
 - CSRF (cross-site request forgery), 472–475
 - CS-Situational-Awareness-BOF, 97
 - curl command, 424
 - CVSS (Common Vulnerability Scoring System), 510, 511–512
 - CWE (Common Weakness Enumeration) ID 840, 430–431
 - cybersecurity, 5–6. *See also* bug bounty hunting; penetration testing; red teaming
 - hybrid engagements, 12
 - purple team engagements, 12–13
- ## D
-
- DACLs (discretionary access control lists), abusing, 262–263
 - daemon, malicious, 172–174
 - dark web, 81
 - data collection, BloodHound, 246
 - data source, training, 4
 - database
 - Google Hacking, 360–361
 - SQL, fingerprinting, 440–441
 - vector, 302–305
 - DC (domain controller)
 - AD (Active Directory), 208
 - enumeration, 234–235
 - high availability, 209
 - inbound replication, 209
 - outbound replication, 208–209
 - DDT (Double Door Tool), 141
 - default passwords, 458–459
 - dense vector search, 306
 - dig command, 340
 - DirBuster, 491–492
 - directory
 - enumeration, 401–402
 - force-testing, 401
 - traversal, 480–481
 - DN (distinguished name), 216
 - DNS (Domain Name Service), 34
 - AD implementation, 222–223
 - exfiltration, 183
 - gathering domain information, 97–99
 - lookups, 339–340
 - querying user realm, 76–78
 - subdomain discovery, 71–73
 - using dnsx, 74–75
 - using httpx, 75–76
 - using subfinder, 73–74
 - tunneling, 182–183
 - dns2tcp, 183
 - dnsx, targeting h4cker.org, 74–75
 - document ranking, 307
 - documentation, phone pretexting, 115
 - DoH (DNS over HTTPS), 34, 36

DOM (Document Object Model) XSS, 469–470
 domain/s. *See also* AD (Active Directory);
 certificate/s; DNS (Domain Name Service
 AD (Active Directory), 212–213, 222–223
 federated, 76, 78
 managed, 76
 Microsoft Entra, creating, 224–228
 MITRE ATT&CK Framework, 20
 spear phishing infrastructure, 116–118
 DonPAPI, extracting DPAPI sensitive information,
 257–259
 Doppelgänger, 135–136
 Dropbox, as covert channel, 181–182
 DuckyScript, 180

E

E5 licenses, 130
 eavesdropping, 395–396
 email
 associating with leaked credentials, 81
 lure URL, testing, 130
 embeddings, 300. *See also* vector embeddings
 Empire, 192–194
 emulating, industry-specific threats, 19
 encryption
 vector database, 305
 web session, 427
 engagements
 lack of findings, 27–28
 red team, 28
 adversarial emulation, 24
 full-scope, 22–23
 objective-based, 23–24
 purple team, 24–25
 tabletop, 25–26
 engineering team, red team collaboration, 27
 enumeration, 339. *See also* reconnaissance
 ADCS (Active Directory Certificate Services),
 259–260
 DC (domain controller), 234–235
 directory, 401–402
 group, 384–385
 host, 382
 network share, 385–386
 passive, KnockKnock.py, 82–84
 search engine, 360–361
 SMB, 386–392
 subdomain
 identifying headers, 76
 using dnsx, 74–75
 using httpx, 75–76
 using subfinder, 73–74
 user, 83–85, 382–384
 using cme, 235–236
 using rpcclient, 235–236
 web page/web application, 392–394
 EPSS (Exploit Prediction Scoring System), 510,
 511–512
 ethical hacking, 10, 11
 external eavesdropping, 396
 getting started in, 333
 Linux distributions for, 429–430
 Evilginx3, 119
 commands, 124–126
 configuring the server, 128
 OpenGraph settings, 123
 phishlet, 120
 auth_urls section, 121
 force_post section, 121–122
 intercept section, 122
 lures options, 126–128
 pulling and building the package, 120
 redirector, 123
 virtual instance settings commands, 124
 YAML, 120–121
 execution, spear phishing, 130–131
 exercises, tabletop, 25–26, 151
 exfiltration, 19, 183
 Exif tool, 358–360
 exploit/ing, 5
 ADCS, 261–262
 asking an AI model to create an, 297
 buffer overflow, 269
 business logic flaws, 431
 Kerberos, 459

- ROP (return-oriented programming), 273, 274–275
- SQL injection vulnerabilities, 439–440
- external attack surface, 509
- extracting credentials
 - using Lsassy, 256–257
 - using Secretsdump, 253–256

F

- facility code, 134
- fake organization, creating, 2
- federation, 76, 78
- feedback
 - purple team engagement, 25
 - threat intelligence team, 26
- ffuf, 402–404
- find command, 196
- fine-tuning AI models, 517–518
- fingerprinting
 - database, 440–441
 - web application development framework, 456
- firmware
 - backdoors, 179–180
 - vulnerabilities, 180
- flag/s
 - HTTPonly, 457
 - KnockKnock.py, 82–83
 - subfinder, 73
- Flare, 81
- Flipper Zero, as BadUSB device, 180
- forest, AD (Active Directory), 211–213, 215–216
 - functional level, 214
 - root domain, 214–215
- forgery
 - cross-site request, 472–475
 - server-side request, 421–476, 477–480
- Foster, Garrett, 104
- frameworks. *See also* TTPs (tactics, techniques, and procedures)
 - C2 (Command and Control)
 - C2 Matrix, 186

- Cobalt Strike, 35
- Mythic, 35–36, 165
- Nighthawk, 35
- TrevorC2, 186
- CVSS (Common Vulnerability Scoring System), 510
- Evilginx3, 119–120
- LangChain, 289
 - prompt template, 289–290
 - system prompts, 290
- MITRE ATT&CK, 150
- OSINT, 338
- penetration testing, 7
- post-exploitation, Empire, 192–194
- red team, 20, 27
 - CBEST (Cyber Assessment Framework), 21–22
 - MITRE ATT&CK, 20, 35, 54–56
 - TIBER-EU (Threat Intelligence-Based Ethical Red Teaming), 21
 - UKC (Unified Kill Chain), 20–21
- web application development, fingerprinting, 456
- web development, 427
- Zero Trust, 209
- full-scope engagement, 22–23
- functions
 - safe, 271–272
 - strncpy(), 270, 271, 271–272
 - system(), 273, 275–277
- fuzzing, 157

G

- gadgets, 273, 274–275
- garak, probes, 319–321
- GetNPUUsers.py, 241–242, 245
- GitHub, 2, 338
- GOAD, Active Directory lab setup, 229–230
- gobuster, 401–402
- Golden SAML Attack, 263–265
- “golden thread” principle, 21–22
- Google Drive, as covert channel, 182

Google hacking
 advanced operators, 360
 GHDB (Google Hacking Database), 360–361
 Gorilla, 372–373
 group enumeration, 384–385
 Group Policy, 222. *See also* policy/ies
 GTFOBins, 196–197
 guardrails, 521
 GUI, Mythic, 43–44

H

h4cker.org
 DNS lookups, 339–340
 DNS resolution details, 340
 targeting, 73–74
 h8mail, 355–358
 HackerOne, 11, 78, 329
 hacking
 ethical, 10, 11
 Google
 advanced operators, 360
 GHDB (Google Hacking Database),
 360–361
 hallucination, 372–373, 522
 hard-coded credentials, 484
 hardware, backdoors, 179–180
 hash, cracking, 245–246
 hashcat, 245–246, 251–252
 haveibeenpwned.com, 81
 header
 boot-start driver, 177–178
 web shell, 156
 help desk, impersonation, 112–113
 help option, Mythic, 52–53
 Hexed.it, 166–167
 HID 2K iClass card, 137
 high availability, AD (Active Directory), 209
 high-frequency smart cards, 133–134
 horizontal privilege escalation, 199
 host enumeration, 382
 HTTP, 36, 416–417
 client/server exchange, 423–424
 messages, 419–420, 422
 multiplexing, 424

packet capture using Wireshark, 420–421
 parameter pollution, 461–462
 proxy, 417
 request/response, 418–419, 421–422
 session, 425–426
 cookies, 427
 encryption, 427
 state, 426
 token, 426
 URL, 422–423
 httpx, 75–76
 “human-in-the-loop,” RAG system, 311
 Hunter & Ready, VRTX (Versatile Real-Time
 Executive) operating system, 10
 Hunter.io, 79
 hybrid search retrieval, 306–307

I

IAM (identity and access management), 148. *See also* AD (Active Directory)
 Microsoft Entra, 223
 features, 223–224
 portal, 224
 iClass legacy cards, 134–135
 IDaaS (Identity as a Service), 223
 identifying, vulnerabilities, 18
 IDOR (insecure direct object reference), 462–463
 Impacket, 231
 impersonation, help desk, 112–113
 inbound replication, 209
 incident response, 13–14
 indexing, 306
 industry-specific threats, emulating, 19
 infection, TDL3 rootkit, 177–179
 information density optimization, 309
 information retrieval
 BM25, 306–307
 dense vector search, 306
 document ranking, 307
 hybrid search, 306–307
 indexing, 306
 MMR (Maximal Marginal Relevance),
 307–308
 semantic search, 305

- infrastructure
 - private branch exchange, 114–115
 - red team, 33–34
 - C2 communication channels, 36
 - redirectors, 37. *See also* redirectors
 - spear phishing
 - cloud, 116
 - domains, 116–118
 - Evilginx3, 119–128. *See also* Evilginx3
 - security settings, 116
 - virtual instance, 116
 - initial access, 81
 - identifying authentication type, 84–87
 - MDM discovery, 87–88
 - password spraying, 82, 84
 - payload preparation, 88–89
 - Mangle, 91–95
 - ScareCrow, 89–91
 - user enumeration, 83–85
 - username validation, 82–83
 - injection, 432
 - command, 448–449
 - LDAP, 449–450
 - SQL, 436–438
 - blind, 442–444
 - categories, 438–440
 - exploiting vulnerabilities, 439–440
 - fingerprinting a database, 440–441
 - mitigation, 447
 - out-of-band exploitation, 444–445
 - stacked queries, 445–446
 - stored procedure, 446–447
 - time-based, 445
 - time-delay, 446
 - UNION exploitation technique, 442
 - inlining, 168–171
 - Instagram, targeting, 78–79
 - installing
 - Apache, 38–39
 - certificates on your C2 server, 41–43
 - CertSPY, 354
 - dnspython library, 341
 - Gorilla, 373
 - Mythic, 43
 - Proxy Manager, 67
 - proxychains4, 53
 - Internet, 9
 - Internet Archive Wayback Machine, 361–362
 - Invoke-MFASweep cmdlet, 87
 - IOCs (indicators of compromise), 150
 - iptables, 37
-
- ## J
-
- J Tool, 143
 - JavaScript web shell, 155
 - JITT (just-in-time training), 2
 - jitter, 34
 - job listings, leveraging, 348
 - Juice Shop, 437
 - Jupyter notebooks, 4
-
- ## K
-
- Keras lambda layer, 3
 - Kerberoasting, 99–100, 228, 243, 244–246
 - Kerberos, 219
 - Golden Ticket attack, 459
 - unconstrained delegation, 459–460
 - vulnerabilities, 228
 - kernel-level backdoors
 - bootkit, 179
 - rootkit
 - TDL3, 176–179
 - types of, 174
 - KEV Catalog, 511–512
 - kickoff call, red team, 65
 - KnockKnock.py, 82–84
 - error codes, 82
 - flags, 82–83
 - user enumeration, 83–85
-
- ## L
-
- lab setup, GOAD, 229–230
 - lambda layer, 3
 - LangChain, 289

AI-powered SSL/TLS analysis, 291–294
 Expression Language, 296
 hybrid search, 307
 LangGraph, 296
 LangSmith, 295
 playground, 295
 prompt template, 296
 prompt template, 289–290
 system prompts, 290

lateral movement, 97, 187
 Kerberoasting, 99–100
 post-exploitation activities, 151
 scanning, 187–189
 SCCM abuse, 104–105
 through ADCS (Active Directory Certificate Services), 101–103

launching, Responder, 248–249

layout, report, 105–106

LDAP (Lightweight directory Access Protocol), 216–218
 injection vulnerabilities, 449–450
 potential exploits, 229
 queries, 243–244
 query operators, 242–243
 retrieving all user accounts in a domain, 243

ldapsearch, 244

leaked data, intelligence gathering, 81. *See also* credential/s

LF cards, 133

LFI (local file inclusion) vulnerability, 482–483

licence, E5, 130

LinkedIn profile, intelligence gathering, 81

Linux, 207
 buffer overflow, 267–269
 dig command, 340
 distributions for ethical hacking, 429–430
 GTFOBins, 196–197
 kernel module rootkit
 code, 174–175
 TDL3, 176–179
 Ismod command, 176
 sudo update-ca-certificates command, 68

LlamaIndex, 296–297, 307

LLMNR/NBT-NS spoofing attack, 250

LLMs (large language models), 282. *See also* LangChain; LlamaIndex; model, AI
 context preparation, 308–310
 Gorilla, 372–373
 local, 283
 OWASP Top 10, 428–429

.lnk shortcut, modification, 96

LOA (letter of authorization), 131

local LLMs, 283

logs, CT (certificate transparency), AI-powered OSINT analysis, 284–289

LOLBAS (Living-Off-the-Land Binaries, Scripts, and Libraries), 197–198

LoRA (low-rank adaptation), 518

LOTL (living off the land), 152, 197–198

LSA (Local Security Authority), 255

Lsassy, extracting credentials, 256–257

Ismod command, 176

lure URL, 120, 130

M

magstripe, 132

malicious model, AI, 3

Maltego, 372

malware, 7, 152
 DNS tunneling, 182–183
 rootkit, 174
 TDL3, 177–179
 types of, 174

managed domain, 76

Mangle, 91–95

marketplace refresh command, 366

MassDNS, 342

MDM (Mobile Device Management), 87–88

memory
 ASLR (Address Space Layout Randomization), 272–273
 NX Bit, bypassing, 275–277
 RIP (return instruction pointer), 273–274

messages
 database error, 440–441
 HTTP, 419–420, 422
 metadata, 78–79, 358–360

- Metasploit, 188
 - Meterpreter module, commands and cmdlets, 164–165
 - RDP post-exploitation module, 189
- Metasploit Unleashed, 165
- Meterpreter, 164–165. *See also* bind shell; reverse shell
- methodology
 - communication, 65
 - penetration testing, 7
 - planning, 64
 - reconnaissance, 65–66
- MFA (multifactor authentication), 19, 209
- MFASweep, 84–87
- Michael, Duane, 104
- Microsoft Entra
 - creating a domain, security settings, 224–227
 - features, 223–224
 - PIM (Privileged Identity Management) abuse, 266–267
 - portal, 224
 - SyncJacking, 265–266
 - vulnerabilities, 227–229
- Microsoft Teams, username validation, 82
- military
 - origins of red teaming, 9
 - Sigma War Games, 9
 - Tiger Team, 7
- misconfiguration, 151–152
- Misconfiguration-Manager, 104
- MITRE
 - ATT&CK Framework, 20, 35, 54–56, 150
 - CWE (Common Weakness Enumeration) ID 840, 430–431
- ML (machine learning), 14
 - LoRA (low-rank adaptation), 518
 - malicious model, 3
 - models, 14
 - penetration testing tools, 13
 - public platform, 2
- MMR (Maximal Marginal Relevance), 307–308
- model, AI. *See also* training
 - censored, 521
 - ChatGPT, 3
 - creating pretext script, 112
 - spear phishing text formulation, 128
 - deploying, 516
 - embeddings, 300–302
 - fine-tuning, 517–518
 - OpenAI, 302
 - poisoning, 2–3
 - WhiteRabbitNeo, 297–298
- module
 - Empire, 192–194
 - Metasploit, 189
 - Meterpreter, 164–165
 - PowerSploit, 190–192
 - Python, 231
 - Recon-ng, 365–369
- Mozilla, bug bounty program, 11
- Mudge, Ralph, 35
- Mythic, 35–36, 165
 - accessing the GUI, 43–44
 - Apache redirector
 - adding in the IP of your C2 server in /etc/hosts, 40–42
 - config, 39–40
 - installing Apache, 38–39
 - installing certificates for your domain, 39
 - checking the connection from end to end, 47
 - editing the C2 config, 45–47
 - help option, 52–53
 - installing, 43
 - installing certificates on your C2 server, 41–43
 - installing proxychains4, 53
 - installing the C2 profile and agent, 43
 - login, 44
 - MITRE ATT&CK mappings, 54–56
 - operator box, 54
 - payloads, 47–52
 - reporting, 54, 57–59
 - socks, 53

SOCKS proxy, 53–54
 starting, 43
 updating the C2 profile config, 45–46

N

NameSpi, 79–81

Netcat, 161

 commands, 163

 nc -lvp command, 162–163

Netscape, “Bugs Bounty” initiative, 10

network

 segmentation, 187–188

 share enumeration, 385–386

 vulnerability scanning

 bandwidth limitations, 400

 choosing a time to run, 397–398

 determining protocols in use, 399

 fragile systems/nontraditional assets, 400

 query throttling, 400

Nighthawk, 35

Nikon COOLPIX P1000, 139

Nikto, 393–394

NLP (natural language processing), 282

Nmap, 340

 scan report, 233–234

 Scripting Engine, 378–379, 385, 393

 SYN scan, 378

nonpersistent cookies, 457

NSA (National Security Agency), 9

nslookup command, 234–235

NTLM (NT LAN Manager), 218

 potential exploits, 229

 version 2, 219

 vulnerabilities, 227–228

ntlmrelayx tool, 252

Nuclei, 499–500, 512, 513–515

NX Bit, 273, 275–277

O

obfuscation

 code, 167–171

 control flow, 168

 shellcode, 92

object

 DN (distinguished name), 216

 Group Policy, 222

objective-based engagement, 23–24

obtaining credentials, 97

OneDrive, username validation, 82

onsite reconnaissance, 139

Open Interpreter, 373–374

 contents of the created file, 375–377

 passive reconnaissance, 374

 plan of action, 375

OpenAI, 302

OpenGraph, 123

open-source software, backdoor, 157

 embedding, 167

 reverse shell, 157–160

OpenSSL, obtaining certificate information,
 349–350

operating system. *See also* Linux; Windows

 ASLR (Address Space Layout

 Randomization), 272–273

 No-Execute bit, 273

 NX Bit, 273

operator box, 54

operator/s

 Google search engine, 360

 LDAP query, 242–243

 UNION, 442

organizations. *See also* AD (Active Directory);

 bug bounty hunting

 external attack surface, 509

 hybrid engagements, 12

 identifying the SSO provider, 76

 internal security, 11

 username collection, 79–81

OSINT (open-source intelligence), 2, 111, 338.

See also reconnaissance

 AI-powered, 283, 284–289

 bug bounty program, 337–338

 tools

 Amass, 371–372

 Maltego, 372

OUs (organizational units), AD (Active Directory),
 213

outbound replication, 208–209

out-of-band exploitation, 444–445
 OWASP (Open Web Application Security Project)
 Cheat Sheet, 426
 Clickjacking Defense Cheat Sheet, 480
 Enterprise Security API, 447
 Error Handling Cheat Sheet, 484
 Juice Shop, 437
 recommendations for business logic security, 430
 testing guide, 7
 Top 10 for LLM Applications, 428–429
 Top 10 for Web Applications, 427–428
 WebGoat, 437
 ZAP (Zed Attack Proxy), 462, 490–492

P

packet
 capture, HTTP, 420–421
 inspection, 395–396
 PAM (Privileged Access Management), 151–152
 parameter pollution, 461–462
 parity bits, 134
 participants
 CBEST (Bank of England’s Cyber Assessment Framework), 21–22
 TIBER-EU (Threat Intelligence-Based Ethical Red Teaming), 21
 passive enumeration, KnockKnock.py, 82–84
 passive reconnaissance, 337
 Gorilla, 373
 Open-Interpreter, 374
 pass-the-hash, 187
 password/s
 cracking, 245–246, 251–252
 default, 458–459
 dump, 355–358
 policies, 236–237, 239–240, 242
 spraying, 82, 84, 239, 240, 452
 synchronization, 228
 wordlists, 157
 payload/s, 180
 fuzzing, 157
 Mythic, 47–52
 preparation, 88–89
 Mangle, 91–95
 ScareCrow, 89–91
 PC (program counter), 274
 PCA (principal component analysis), 300
 PDCP (Project Discovery Cloud Platform), 512–513
 penetration testing, 2, 6
 AI tools, 13
 frameworks, 7
 future of, 12–13
 methodology, 7
 versus red teaming, 18, 27
 standards, 7
 Tiger Team, 7
 tools, 9
 persistence, 96, 152
 backdoor, 180
 post-exploitation activities, 148
 personnel, selecting a target, 110–111
 phishing, 2–3, 13–14, 151–152
 phishlet, 120
 auth_urls section, 121
 force_post section, 121–122
 intercept section, 122
 listing, 125–126
 lures options, 126–128
 template, 120
 phone pretexting
 documentation, 115
 documenting and organizing your data, 111
 preparing pretext, 111–112
 execution, 113–115
 impersonating the help desk, 112–113
 SpoofCard, 114–115
 reconnaissance, 110–111
 selecting a target, 110–111
 PHP web shell, 154

- physical security, 110, 131. *See also* badge/s
 - bypass tools, 140
 - DDT (Double Door Tool), 141
 - J Tool, 143
 - Quick Jim, 140–141
 - REX Sensor Bypass Gun, 142–143
 - UDT (Under the Door Tool), 141–142
- PIM (Privileged Identity Management) abuse, 266–267
- pip3 install dnspython, 341
- planning, red team, 64
- platform/s
 - bug bounty, 330–331, 332–333, 508–509
 - Burp Suite, 405
 - XDR (extended detection and response), 149–150
- playbook, blue team, 150
- POCs (points of contact), 131
- policy/ies
 - AD (Active Directory), 209, 216–217
 - PAM (Privileged Access Management), 151–152
 - password, 236–237, 239–240, 242
 - scanning, 400
- port scan, 377
 - Nmap SYN scan, 378
 - TCP, 380–381
- portal, Microsoft Entra, 224
- post exploitation, 105, 147–148, 157–160, 194.
 - See also* C2 (Command and Control)
 - APTs (advanced persistent threats), 148
 - backdoors, 152, 153. *See also* backdoors
 - application-level backdoors, 153–171
 - hardware and firmware, 179–180
 - kernel-level, 174–179
 - patching binaries to include, 165–167
 - scheduled jobs and tasks, 180
 - system-level, 171–174
 - code obfuscation, 167–168
 - control flow obfuscation, 168
 - inlining and unrolling, 168–171
 - covering your tracks, 199–202
 - creating accounts on compromised systems, 181
 - evaluating detection and response capabilities, 148
 - improving blue team capabilities, 149
 - advanced monitoring tools, 149
 - playbooks, 150
 - tabletop exercises, 151
 - threat-hunting hypotheses, 150
 - lateral movement, 151, 187, 187–189
 - password/s, 187
 - persistence, 148, 152
 - privilege escalation, 198–199
 - tools
 - Empire, 192–194
 - PowerShell commands, 189–190
 - PsExec, 195–196
 - Sysinternals, 195
 - WinRM (Windows Remote Management), 196
 - WMI (Windows Management Instrumentation), 194–195
- PowerShell, commands and cmdlets, 189–190
- PowerSploit, 190–192
- predictive threat modeling, 13–14
- presumption of breach, 149
- pretext script, 112
- private branch exchange, standing up, 114–115
- private bug bounty programs, 329, 330, 331
- privilege escalation, 148
 - horizontal, 199
 - post-exploitation, 198–199
- probes, garak, 319–321
- prompt. *See also* information retrieval
 - engineering, 4
 - injection, 319, 321–322
 - template, 289–290, 296
- Prompt Fuzzer, 317–319
- protocol. *See also* HTTP
 - stateful, 417
 - stateless, 417
- proxy, HTTP, 417
- Proxy Manager, 67, 68
 - installing, 67
 - setup, 68–71
 - testing, 71

proxychains, 255
 proxychains4, installing, 53
 PsExec, 195–196
 PTES (Penetration Testing Execution Standard), 7
 public bug bounty programs, 329, 331
 public source-code repositories, 362–363
 purple team engagement, 12–13, 24–25
 pymeta.py, targeting Instagram, 78–79
 Python, 3, 4

- BloodHound.py, 246
- creating your own scanner, 380–381
- GetNPUsers.py, 241–242, 245
- modules, 231
- scapy, 380–381
- script, 4, 341, 345–348

Q

query/ies. *See also* information retrieval; prompt

- BloodHound, 247
- embeddings, 300–302
- hackerone.com, 78
- LDAP, 243–244
- operators, LDAP, 242–243
- Recon-ng, 368
- SQL, 433–435
- throttling, 400
- user realm, 76–78

 Quick Jim, 140–141

R

race conditions, 485
 RAG (retrieval-augmented generation), 296, 299–300. *See also* information retrieval

- for bug bounty hunting, 519
- context preparation, 308–310
- document ranking, 307
- “human-in-the-loop”, 311
- hybrid search retrieval, 306–307
- MMR (Maximal Marginal Relevance), 307–308
- in red teaming, 300

response generation, post-processing and refinement, 311–313
 semantic search, 304
 user query processing, 305
 vector embeddings, 300–302
 vector storage database, 302–305
 RAG-Fusion, 308
 ransomware, 19
 RAPTOR, 308
 RDP (Remote Desktop Protocol), 189
 reconnaissance, 336. *See also* enumeration

- active, 336–337, 377
 - Nmap SYN scan, 378
 - port scan, 377
- AD (Active Directory), 231. *See also* AD (Active Directory), user enumeration
- advanced operators, 360
- API, 404–407
- ASN (autonomous system number), 65–66
- automation, 501–502
- certificate transparency, 66
- DNS (Domain Name Service)
 - lookups, 339–340
 - obtaining domain information using dig, 340
 - querying user realm, 76–78
 - subdomain discovery, 71–73
 - targeting h4cker.org with dnsx, 74–75
 - targeting h4cker.org with httpx, 75–76
 - targeting h4cker.org with subfinder, 73–74
- identifying cloud vs. self-hosted assets, 344–348
- identifying technical and administrative contacts, 342–344
- leaked data, 81
- leveraging job listings, 348
- metadata, 78–79, 358–360
- methods, 337
- obtaining certificate information, 349–350, 354–355
- onsite, 139

- passive, 337
 - Gorilla, 373
 - Open-Interpreter, 374
 - password dumps, 355–358
 - phone pretexting, 110–111
 - proxying traffic, 66–71
 - public source-code repositories, 362–363
 - revealing weak cryptographic implementations, 350–353
 - satellite images of facility, 131–132
 - social media, 132
 - software discovery, 81
 - username collection, 79–81
 - web application, 401–404
- Recon-ng, 369
- help command, 364
 - modules, 365–369
 - setting the source domain and running the query, 368
 - starting, 363–364
- red team/ing, 6. *See also* C2 (Command and Control)
- adversarial emulation, 9–10, 18–19
 - algorithmic, 311–314
 - keeping pace with evolving technologies, 322–323
 - Prompt Fuzzer, 317–319
 - TAP (Tree of Attacks with Pruning), 315
 - tools for testing AI applications and model security, 316, 322
- blue team collaboration, 20, 152
- challenges, 26
- lack of findings, 27–28
 - leveling up other teams, 26–27
- communications, 34, 65
- current state of AI in, 282–283
- engagements, 9–10, 28
- adversarial emulation, 24
 - full-scope, 22–23
 - objective-based, 23–24
 - purple team, 24–25
 - tabletop, 25–26
- evolution, 7
- frameworks, 20, 27
- CBEST (Bank of England’s Cyber Assessment Framework), 21–22
 - MITRE ATT&CK, 20, 35
 - TIBER-EU (Threat Intelligence-Based Ethical Red Teaming), 21
 - UKC (Unified Kill Chain), 20–21
- future of, 12–13
- importance of, 18
- challenging assumptions, 18–19
 - emulating industry-specific threats, 19
 - identifying critical vulnerabilities, 18
- infrastructure, 33–34, 37. *See also* C2 (Command and Control); redirectors
- initial access, 81
- identifying authentication type, 84–87
 - password spraying, 82–83
 - user enumeration, 83–85
 - username validation, 82–83
- kickoff call, 65
- lateral movement, 97, 187
- Kerberoasting, 99–100
 - SCCM abuse, 104–105
 - through ADCS (Active Directory Certificate Services), 101–103
- military origins, 9
- versus penetration testing, 18, 27
- persistence, 96
- phone pretexting
- documentation, 115
 - documenting and organizing your data, 111
 - execution, 113–115
 - impersonating the help desk, 112–113
 - preparing pretext, 111–112
 - reconnaissance, 110–111
 - selecting a target, 110–111
 - SpoofCard, 114–115
- planning, 64
- post-exploitation activities, 105, 148
- APTs (advanced persistent threats), 148
 - backdoors. *See* backdoors
 - code obfuscation, 167–171

- creating accounts on compromised systems, 181
- evaluating detection and response capabilities, 148
- improving blue team capabilities, 149–152
- persistence, 148
- RAG (retrieval-augmented generation), 300
- reconnaissance, 65
 - ASN, 65–66
 - certificate transparency, 66
 - DNS, 71–78
 - leaked data, 81
 - metadata, 78–79
 - proxying traffic, 66–71
 - targeting h4cker.org with subfinder, 73–74
 - username collection, 79–81
- reporting, 105–106
- ROEs (rules of engagement), 64
- social engineering, 140
- spear phishing, 115. *See also* spear phishing
 - Azure CDN, 117–118
 - domains, 116–118
 - E5 licenses, 130
 - Evilginx3, 119–128. *See also* Evilginx3
 - execution, 130–131
 - security settings, 116
 - text formulation, 128
 - virtual instance, 116
- tainting training data, 1–4
- TTPs (tactics, techniques, and procedures), 17–18
- redirect attacks, 458
- redirector, 37, 120
 - Apache
 - adding in the IP of your C2 server in /etc/hosts, 40–42
 - config, 39–40
 - installing, 38–39
 - installing certificates for your domain, 39
 - Evilginx3, 123
- reflected XSS attacks, 465–466
- remote reconnaissance
 - satellite images of facility, 131–132
 - social media, 132
- replication, 208
 - inbound, 209
 - outbound, 208–209
- report/s
 - audit, 151
 - bug bounty, 334, 407–409
 - layout, 105–106
 - Mythic, 54, 57–59
 - nmap scan, 233–234
 - technical review, 106
- request/response model, HTTP, 418–419
- resources, controls, 5
- Responder, 247–248, 253
 - capturing the NTLMv2 hash, 251
 - launching, 248–249
 - network activity, 250
- REST (representational state transfer) API, 423
- reverse proxy, 37
- reverse shell, 157–160
 - comparison with bind shell, 164
 - how it works, 160–161
- rewards, bug bounty programs, 11
- REX Sensor Bypass Gun, 142–143
- RFI (remote file inclusion) vulnerabilities, 483
- RID (relative identifier), 237
- RIP (return instruction pointer), 273–274
- RMM tools, 189
- ROEs (rules of engagement), 9
 - bug bounty programs, 334–336
 - red team, 64
- root domain, AD forest, 214–215
- rootkit
 - code, 174–175
 - TDL3, 176–179
 - types of, 174
- ROP (return-oriented programming), 273, 274–275
 - bypassing NX, 273, 275–277
 - gadgets, 273, 274–275
- rpcclient, 235–236
- Rust, creating a malicious daemon in, 172–173

S

- safe functions, 271–272
- Salt Typhoon, 152
- scapy, 380–381
- ScareCrow, 89–91
- SCCM abuse, 104–105
- scheduled jobs and tasks, as backdoor, 180
- schema, AD (Active Directory), 211, 213
- scope, bug bounty programs, 334–336
- script/s
 - AI-powered OSINT, 284–289
 - bash, 0365-realminfo, 76
 - Cloud Checker, 345–348
 - GetNPUsers.py, 241–242
 - GetUserSPNs.py, 245
 - Living-Off-the-Land, 197–198
 - NSE (Nmap Scripting Engine), 379, 385, 393
 - PowerSploit, 190–192
 - pretext, 112
 - Python, 4, 341, 345–348
 - secretsdump, 255–256
 - web shell, 153
 - access, 154, 155
 - ASP, 154
 - code and command obfuscation, 156
 - JavaScript, 155
 - PHP, 154
 - storage, 153
 - whoami command, 154
 - WMI (Windows Management Instrumentation), 194–195
- SecLists, 157
- second generation smart cards, 133
- Secretsdump, extracting credentials, 253–256
- security, 5–6. *See also* authentication; IAM (identity and access management)
 - ASLR (Address Space Layout Randomization), 272–273
 - awareness training, 26–27
 - NX Bit, bypassing, 275–277
 - physical, 110, 131
 - posture, 18–19, 64
 - Zero Trust, 209
- semantic search, 304, 305
- sensitive data patterns, 157
- server
 - command and control, 34–35
 - Evilginx3, configuring, 128
 - LDAP, 218
 - web, 416
- service/s
 - Identity as a, 223
 - malicious, 172–173
- session, 120, 417
 - hijacking, 4, 452–458
 - HTTP, 425–426
 - cookies, 427
 - encryption, 427
 - state, 426
 - token, 426
 - management frameworks, 427
- setup, Proxy Manager, 68–71
- shellcode obfuscation tools, 92
- Shodan, 369–371
- Sigma War Games, 9
- single point of failure, 210
- smart cards. *See also* badge/s
 - chipset, 134
 - facility code, 134
 - high-frequency, 133–134
 - iClass legacy, 134–135
 - low-frequency RFID, 133
 - second generation, 133
- SMB
 - enumeration, 386–392
 - signing, 252, 253
- SOAP (Simple Object Access Protocol), 423
- SOC (security operations center), 19
- socket, 37
- social engineering, 7, 9–10, 13–14
 - breaches, 109–110
 - phishing, 151–152
 - phone pretexting
 - documentation, 115
 - documenting and organizing your data, 111
 - execution, 113–115

- impersonating the help desk, 112–113
 - preparing pretext, 111–112
 - reconnaissance, 110–111
 - selecting a target, 110–111
 - SpoofCard, 114–115
- spear phishing, 115. *See also* spear phishing
 - E5 licenses, 130
 - Evilginx3, 119–128
 - execution, 130–131
 - infrastructure, 116–128
 - testing the lure URL, 130
 - text formulation, 128
- tailgating, 140
- social media
 - OSIN (open-source intelligence), 132
 - scraping, 348
- sock puppet account, 111
- socks, 35, 53
- SOCKS proxy, Mythic, 53–54
- software, discovery, 81
- source code, 157–160
 - comments, 484
 - repositories, 362–363
 - with reverse shell, 160
- spear phishing, 115. *See also* Evilginx3
 - execution, 130–131
 - infrastructure
 - Azure CDN, 117–118
 - cloud, 116
 - domains, 116–118
 - security settings, 116
 - virtual instance, 116
 - lure URL, testing, 130
 - text formulation, 128
- SPN (service principal name), 99–100
- SpoofCard, 114–115
- SprayHound, 240
- SQL
 - fingerprinting a database, 440–441
 - injection attack, 436–438
 - categories, 438–440
 - exploiting vulnerabilities, 439–440
 - mitigation, 447
 - out-of-band exploitation, 444–445
 - stored procedure, 446–447
 - time-based, 445
 - time-delay, 446
 - UNION exploitation technique, 442
 - using Booleans, 442–444
 - query, 433–435, 445–446
 - statements, 433, 435–436
- SRI (Subresource Integrity), 488
- SSO (single sign-on), AD (Active Directory), 210
- SSRF (server-side request forgery), 477–480
- stacked queries, 445–446
- STACKTITAN, 114–115
- standards, penetration testing, 7
- stateful protocol, 417
- stateless protocol, 417
- statement, SQL, 433, 435–436
- Stealth Wiegand Data Interpreter (WDI), 134
- steganography, 199–202
- steghide command, 199–202
- storage
 - credentials, 151
 - session ID, 360
 - vector embeddings, 302–305
 - web shell, 153
- stored procedure, 446–447
- stored XSS attacks, 466–467
- strcpy(), 270, 271
- subdomain discovery
 - identifying headers, 76
 - using dnsx, 74–75
 - using httpx, 75–76
 - using subfinder, 73–74
- subfinder
 - flags, 73
 - targeting h4cker.org, 73–74
- sudo update-ca-certificates command, 68
- supply chain backdoor, 167
- SyncJacking, 265–266
- Sysinternals, 195
- system(), 275–277
- system binary, trojanized, 171–172
- system prompts, 290

T

- T5577 card, writing to, 137–138
- tailgating, 140
- tainting training data, 1–4
- TAP (Tree of Attacks with Pruning), 315
- tar command, 196
- TCP port scan, 380–381
- TDL3 rootkit, 176–179
- Teamfiltration, 84
- teams, leveling up, 26–27
- technical review, report, 106
- template
 - bug bounty program scope, 336
 - Nuclei, 512
 - phishlet, 120
 - prompt, 289–290, 296
- TensorFlow Embedding Projector, 301
- testing
 - AI
 - garak, 319–321
 - Prompt Fuzzer, 317–319
 - continous, 12–13
 - directory force-, 401
 - lure URL, 130
- text editors, Vim, 197
- TF-IDF (Term Frequency-Inverse Document Frequency), 300
- Thompson, Chris, 104
- threat/s
 - industry-specific, 19
 - intelligence, 21, 24, 26, 150, 428–429
- TIBER-EU (Threat Intelligence-Based Ethical Red Teaming), 21
- Tiger Team, 7
- time-based SQL injection, 445
- token, HTTP session, 426
- tool/s. *See also* building bypasses; frameworks
 - Advanced Installer, 92
 - AI-powered, 282, 283, 289–295
 - Arsenal Kit, 35
 - ASM (attack surface management), 331–332
 - BloodHound
 - data collection using the Python ingestor, 246–247
 - queries, 247
 - Bright Data, 67
 - browser developer, 406
 - BurpGPT, 282
 - calling, 520
 - Certipy
 - ADCS enumeration, 259–260
 - ADCS exploitation, 261–262
 - ESC1 request certificate for administrator, 103
 - options, 100–101
 - vulnerabilities, template, 101–103
 - CertSPY, 354–355
 - cloudproxy, 67
 - cme, 231, 232–233, 235–236
 - code obfuscation, 170
 - dauthi, MDM discovery, 87–88
 - dnsx, targeting h4cker.org, 74–75
 - DonPAPI, extracting DPAPI sensitive information, 257–259
 - Doppelgänger, 135–136
 - Evilginx3, 119. *See also* phishlet
 - commands, 124–126
 - configuring the server, 128
 - OpenGraph settings, 123
 - phishlet, 120
 - pulling and building the package, 120
 - redirector, 123
 - virtual instance settings commands, 124
 - YAML, 120–121
 - Exif, 358–360
 - ffuf, 402–404
 - Flare, 81
 - gobuster, 401–402
 - h8mail, 355–358
 - hashcat, password cracking, 251–252
 - httpx, 75–76
 - Hunter.io, 79
 - Impacket, 231
 - KnockKnock.py, 82
 - error codes, 82
 - flags, 82–83
 - user enumeration, 83–85
 - ldapsearch, 244
 - legitimate, 189
 - Lsassy, extracting credentials, 256–257
 - Mangle, 91–95

- MassDNS, 342
 - MFASweep, 84–87
 - NameSpi, 79–81
 - Nikto, 393–394
 - Nmap, 340
 - scan report, 233–234
 - Scripting Engine, 379
 - SYN scan, 378
 - ntlmrelayx, 252
 - Open Interpreter, 373–374
 - passive reconnaissance, 374
 - plan of action, 375
 - OSINT (open-source intelligence)
 - Amass, 371–372
 - Maltego, 372
 - OWASP ZAP (Zed Attack Proxy), 462
 - PAM (Privileged Access Management), 151
 - penetration testing, 9, 13
 - post-exploitation
 - BloodHound, 194
 - Empire, 192–194
 - PsExec, 195–196
 - WinRM (Windows Remote Management), 196
 - Prompt Fuzzer, 317–319
 - Proxy Manager, 67, 68
 - installing, 67
 - setup, 68–71
 - testing, 71
 - proxychains, 255
 - pymeta.py, targeting Instagram, 78–79
 - Recon-ng, 363–369
 - help command, 364
 - modules, 365–369
 - setting the source domain and running the query, 368
 - starting, 363–364
 - Responder, 247–248, 253
 - capturing the NTLMv2 hash, 251
 - launching, 248–249
 - network activity, 250
 - RMM, 189
 - rpcclient, 235–236
 - ScareCrow, 89–91
 - Secretsdump, extracting credentials, 253–256
 - shellcode obfuscation, 92
 - Shodan, 369–371
 - SpoofCard, 114–115
 - SprayHound, 240
 - steghide, 199–202
 - Teamfiltration, 84
 - TensorFlow Embedding Projector, 301
 - testssl.sh, 350–353
 - TruffleHog, 362–363
 - Try-SQL Editor, 435
 - Vibe, 97–99
 - WhatBreach, 357–358
 - WMI (Windows Management Instrumentation), 194–195
 - ZAP, 404
 - traditional bug hunting, 498
 - traffic
 - CDN (content delivery network), 34
 - DNS, 182–183
 - proxying, 66–71. *See also* Proxy Manager
 - training, 3–4
 - data source, 4
 - security awareness, 26–27
 - tainting the data, 1–4
 - using confidential data, 282–283
 - trojanized system binary, 171–172
 - TruffleHog, 362–363
 - TrustedSec BOF, 97
 - Try-SQL Editor, 435
 - TTPs (tactics, techniques, and procedures), 7, 12–13, 17–18, 35
 - TTXs (tabletop exercises), 25–26, 151
 - tunneling
 - DNS, 182–183
 - HTTP/HTTPS, 182–183
 - two-way transitive trusts, AD (Active Directory), 213–214
-
- U**
- UDT (Under the Door Tool), 141–142
 - UKC (Unified Kill Chain), 20–21
 - unauthenticated scans, 397

unauthorized access, alerting, 151
 unconstrained Kerberos delegation, 459–460
 unrolling, 168–171
 URL, 422–423

- lure, 120, 130
- query string manipulation, 439–440

 user enumeration, 382–384

- AD (Active Directory)
 - using cme, 235–236
 - using enumdomgroups, 237–238
 - using net rpc, 238
 - using rpcclient, 237
- using KnockKnock.py, 83–85

 user query processing, 305
 username

- collection, 79–81
- validation, 82

 utilities, legitimate, 189. *See also* tool/s

V

VDPs (vulnerability disclosure programs), versus

- bug bounties, 332–333

 vector database, encryption, 305
 vector embeddings, 300–305
 Vibe, 97–99
 Vim, 197
 virtual instance, spear phishing infrastructure, 116
 Volt Typhoon, 152
 VRTX (Versatile Real-Time Executive) operating system, 10
 vulnerability/ies. *See also* penetration testing

- ADCS (Active Directory Certificate Services), 100, 101–103, 262
- AI, 2–3
- ASREPROasting, 241
- code, 270–272
- firmware, 180
- identifying, 18
- Kerberos, 228
- Microsoft Entra, 227–229
- NTLM (NT LAN Manager), 219, 227–228
- scans, 282, 396

authenticated, 397–398
 bandwidth limitations, 400
 choosing a time to run, 398–399
 determining protocols in use, 399
 fragile systems/nontraditional assets, 400
 network topology, 399
 Nuclei, 512, 513–515
 query throttling, 400
 unauthenticated, 397
 severity assessment, 510–512, 520
 single point of failure, 210
 web application, 95–96

- authentication-based, 451–460
- business logic flaws, 430, 431
- command injection, 448–449
- directory traversal, 480–481
- IDOR (insecure direct object reference), 462–463
- Kerberos, 459–460
- LDAP injection, 449–450
- LFI (local file inclusion), 482–483
- parameter pollution, 461–462
- RFI (remote file inclusion), 483
- SQL injection, 433–447
- unprotected APIs, 485–487
- XSS (cross-site scripting), 463–470

W

war games, 9
 watering hole attack, 2
 Wayback Machine, 361–362
 web application, 416, 459–460. *See also* injection

- building your own, 429–430
- clickjacking, 480
- command injection, 448–449
- cookie manipulation attacks, 482
- cross-site scripting, 463–464
 - DOM XSS, 469–470
 - mitigations, 471–472
- redirecting users to malicious sites, 467–468
- reflected XSS attacks, 465–466
- stealing cookies, 468–469

- stored XSS attacks, 466–467
 - XSS evasion techniques, 471
 - CSRF (cross-site request forgery), 472–475
 - development framework, fingerprinting, 456
 - directory traversal, 480–481
 - hacking tools, 488–492
 - HTTP, 416–417
 - client/server exchange, 423–424
 - messages, 419–420, 422
 - multiplexing, 424
 - packet capture using Wireshark, 420–421
 - parameter pollution, 461–462
 - proxy, 417
 - request/response, 418–419, 421–422
 - session, 425–427
 - URL, 422–423
 - web session, 425–427
 - IDOR (insecure direct object reference), 462–463
 - LFI (local file inclusion) vulnerability, 482–483
 - OWASP Top 10 for, 427–428
 - parameter tampering, 488
 - reconnaissance, 401–404
 - redirect attacks, 458
 - RFI (remote file inclusion) vulnerabilities, 483
 - session hijacking, 452–458
 - SSRF (server-side request forgery), 475–476, 477–480
 - unprotected APIs, 485–487
 - vulnerabilities, 95–96
 - web page enumeration, 392–394
 - web proxy, 488–490
 - web shell, 153, 156
 - access, 154, 155
 - ASP, 154
 - code and command obfuscation, 156
 - header, 156
 - JavaScript, 155
 - PHP, 154
 - storage, 153
 - whoami command, 154
 - WebGoat, 437
 - website archive/caching, 361–362
 - WebSploit Labs, 333, 477–480
 - WhatBreach, 357–358
 - White Cell, 9
 - WhiteRabbitNeo, 297–298
 - whoami command, 154
 - whois command, 342–344, 345
 - Wiegotcha, 135
 - Williams, Walter C., “Program Management in Design and Development”, 7
 - Windows
 - extracting DPAPI sensitive information, 257–259
 - LSA (Local Security Authority), 255
 - scheduled jobs and tasks, as backdoor, 180
 - Sysinternals, 195
 - WinRM (Windows Remote Management), 196
 - Wireshark, 420–421
 - WMI (Windows Management Instrumentation), 194–195
 - wordlists, 157
-
- ## X
-
- X, as covert channel, 182
 - XDR (extended detection and response) systems, 149–150
 - XSS (cross-site scripting). *See* cross-site scripting
-
- ## Y-Z
-
- YAML, Evilginx3, 120–121
 - ZAP (Zed Attack Proxy), 404, 462, 490–492
 - Zero Trust, 209