

Network Security, Firewalls, and VPNs

THIRD EDITION

J. Michael Stewart | Denise Kinsey

Contents

1. **Cover**
2. **Title Page**
3. **Copyright Page**
4. **Dedication**
5. **Brief Contents**
6. **Contents**
7. **Preface**
8. **Acknowledgments**
9. **About the Authors**
10. **CHAPTER 1 Fundamentals of Network Security**
 1. **What Is Network Security?**
 1. **What Is Trust?**
 2. **Who—or What—Is Trustworthy?**
 3. **What Are Security Objectives?**
 2. **What Are You Trying to Protect?**
 1. **Seven Domains of a Typical IT Infrastructure**
 3. **Goals of Network Security**
 4. **How Can You Measure the Success of Network Security?**
 5. **Why Are Written Network Security Policies Important?**
 1. **Planning for the Worst**
 6. **Who Is Responsible for Network Security?**
 7. **Enhancing the Security of Wired Versus Wireless LAN Infrastructures**
 8. **Internal and External Network Issues**
 9. **Common Network Security Components Used to Mitigate Threats**
 1. **Hosts and Nodes**
 2. **Firewalls**
 3. **Virtual Private Networks**
 4. **Proxy Servers**
 5. **Network Address Translation**
 6. **The Domain Name System**
 7. **Directory Services**
 8. **Intrusion Detection Systems and Intrusion Prevention Systems**
 9. **Network Access Control**

- 10. TCP/IP Basics
 - 1. OSI Reference Model
 - 2. Sub-Protocols
 - 3. Headers and Payloads
 - 4. Filtering on Addresses
- 11. CHAPTER SUMMARY
- 12. KEY CONCEPTS AND TERMS
- 13. CHAPTER 1 ASSESSMENT

11. CHAPTER 2 Network Security Threats

- 1. Hackers and Their Motivation
- 2. Favorite Targets of Hackers
- 3. Threats from Internal Personnel and External Entities
- 4. The Hacking Process
 - 1. Reconnaissance
 - 2. Scanning
 - 3. Enumeration
 - 4. Attacking
 - 5. Post-Attack Activities
- 5. Common IT Infrastructure Threats
 - 1. Hardware Failures and Other Physical Threats
 - 2. Natural Disasters
 - 3. Accidents
- 6. Malicious Code (Malware)
 - 1. Advanced Persistent Threat
- 7. Fast Growth and Overuse
- 8. Wireless Versus Wired
- 9. Eavesdropping
- 10. Hijack and Replay Attacks
- 11. Insertion Attacks
- 12. Fragmentation Attacks
- 13. Buffer Overflows
- 14. Session Hijacking, Spoofing, and Man-in-the-Middle Attacks
 - 1. Session Hijacking
 - 2. Spoofing Attacks
 - 3. Man-in-the-Middle Attacks
- 15. Covert Channels
- 16. Network and Resource Availability Threats
 - 1. Denial of Service (DoS)
 - 2. Distributed Denial of Service (DDoS)
- 17. Hacker Tools
- 18. Social Engineering
- 19. CHAPTER SUMMARY
- 20. KEY CONCEPTS AND TERMS
- 21. CHAPTER 2 ASSESSMENT

12. CHAPTER 3 Common Network Topologies and Infrastructures

- 1. What Is a Network Topology?

2. Types of Network Devices
3. What Differentiates Logical and Physical Topologies?
 1. Types of Physical Topologies
 2. Logical Topology
 3. Creating Logical Topologies
4. Differences Between Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6)
 1. IPSec and IPv6
5. Examples of Network Infrastructures and Related Security Concerns
 1. Workgroups
 2. SOHO Networks
 3. Client/Server Networks
 4. LAN Versus WAN
 5. Thin Clients and Terminal Services
 6. Remote Control, Remote Access, and VPN
 7. Boundary Networks
6. CHAPTER SUMMARY
7. KEY CONCEPTS AND TERMS
8. CHAPTER 3 ASSESSMENT

13. CHAPTER 4 Network Design Considerations

1. Network Design and Defense in Depth
 1. Achieving Defense in Depth through Layering
 2. Planning for Scalability
 3. Working with Senior Management
2. Protocols
3. Common Types of Addressing
 1. IPv6
4. Controlling Communication Pathways
 1. Router Configuration
 2. Encrypted Protocols
 3. Physical Access Management
 4. Filtering
5. Intrusion Detection Systems and Intrusion Prevention Systems
6. Hardening Systems
7. Equipment Selection
8. Authentication, Authorization, and Accounting
9. Communication Encryption
10. Hosts: Local-Only or Remote and Mobile
11. Redundancy
12. Endpoint Security
 1. Clients
 2. Servers
 3. Routers
 4. Switches
 5. Firewalls and Proxies
13. Risk Assessment and Management
14. What Are Zones of Risk?

- 15. CHAPTER SUMMARY
- 16. KEY CONCEPTS AND TERMS
- 17. CHAPTER 4 ASSESSMENT

14. CHAPTER 5 Firewall Fundamentals

- 1. What Is a Firewall?
 - 1. What Firewalls Cannot Do
- 2. Why Do You Need a Firewall?
- 3. How Firewalls Work and What Firewalls Do
- 4. Types of Firewalls
- 5. Individual and SOHO Firewall Options
 - 1. Managing the Firewall on an ISP Connection Device
 - 2. Converting a Home Router into a Firewall
- 6. Uses for Host Software Firewalls
 - 1. Examples of Software Firewall Products
 - 2. Using Windows 10's Host Software Firewall
 - 3. Using a Linux Host Software Firewall
- 7. Uses for Commercial Software Network Firewalls
- 8. Uses for Hardware/Appliance Firewalls
- 9. Next-Generation Firewalls
- 10. What Are Virtual Firewalls?
- 11. Dual-Homed and Triple-Homed Firewalls
- 12. Ingress and Egress Filtering
- 13. Types of Filtering
 - 1. Static Packet Filtering
 - 2. Stateful Inspection and Dynamic Packet Filtering
 - 3. Network Address Translation
 - 4. Application Proxy
 - 5. Circuit Proxy
 - 6. Content Filtering
- 14. Selecting the Right Firewall for Your Needs
- 15. The Difference Between Buying and Building a Firewall
- 16. CHAPTER SUMMARY
- 17. KEY CONCEPTS AND TERMS
- 18. CHAPTER 5 ASSESSMENT

15. CHAPTER 6 Firewall Implementation

- 1. Examining Your Network and Its Security Needs
 - 1. What to Protect and Why
 - 2. Preserving Privacy
- 2. Proper Firewall Implementation Procedure
- 3. Constructing, Configuring, and Managing a Firewall
- 4. pfSense
- 5. pfSense Requirements
- 6. Planning a Firewall Implementation with pfSense
 - 1. Firewalling a Small Organization: Packet Filtering or Application-Level Firewall, a Proxy Implementation
 - 2. Firewalling Medium and Large Organizations: Application-Level Firewall and Packet Filtering, a Hybrid System

- 3. Firewalling in a Subnet Architecture
- 7. Installing the pfSense Firewall
- 8. Configuring a Firewall with pfSense
- 9. Elements of Firewall Deployment
- 10. Testing and Troubleshooting
- 11. CHAPTER SUMMARY
- 12. KEY CONCEPTS AND TERMS
- 13. CHAPTER 6 ASSESSMENT

16. CHAPTER 7 Firewall Deployment Considerations

- 1. Common Security Strategies for Firewall Deployments
 - 1. Security Through Obscurity
 - 2. Least Privilege
 - 3. Simplicity
 - 4. Defense in Depth
 - 5. Diversity of Defense
 - 6. Chokepoint
 - 7. Weakest Link
 - 8. Fail-Safe
 - 9. Forced Universal Participation
- 2. Authentication, Authorization, and Accounting
- 3. Placement of Network Hardware Firewalls
- 4. Benefit and Purpose of Reverse Proxy
- 5. Use and Benefit of Port Forwarding
- 6. Considerations for Selecting a Bastion Host OS
- 7. Monitoring and Logging
- 8. Understanding and Interpreting Firewall Logs and Alerts
- 9. Intrusion Detection Systems and Intrusion Prevention Systems
- 10. Security Event and Information Management
- 11. Evaluating Needs and Solutions in Designing Security
- 12. What Happens When Security Gets in the Way of Doing Business?
- 13. CHAPTER SUMMARY
- 14. KEY CONCEPTS AND TERMS
- 15. CHAPTER 7 ASSESSMENT

17. CHAPTER 8 Configuring Firewalls

- 1. Firewall Rules
 - 1. Inbound and Outbound Communications
 - 2. Access Control Lists
- 2. Composing Firewall Rules
- 3. Ordering Firewall Rules
- 4. What Should You Allow and What Should You Block?
- 5. Essential Elements of a Firewall Policy
- 6. Limitations of Firewalls
- 7. Improving Performance
- 8. The Downside of Encryption with Firewalls
- 9. Firewall Enhancements
- 10. Management Interfaces
- 11. CHAPTER SUMMARY

12. KEY CONCEPTS AND TERMS

13. CHAPTER 8 ASSESSMENT

18. CHAPTER 9 VPN Fundamentals

1. What Is a Virtual Private Network?

2. What Are the Benefits of Deploying a VPN?

3. What Are the Limitations of a VPN?

4. What Are Effective VPN Policies?

5. VPN Deployment Models and Architecture

1. VPN Deployment Models

2. VPN Architectures

6. Tunnel Versus Transport Mode

7. The Relationship Between Encryption and VPNs

1. Symmetric Cryptography

2. Asymmetric Cryptography

3. Hashing

4. Establishing VPN Connections with Cryptography

5. Digital Certificates

8. What Is VPN Authentication?

9. What Is VPN Authorization?

10. CHAPTER SUMMARY

11. KEY CONCEPTS AND TERMS

12. CHAPTER 9 ASSESSMENT

19. CHAPTER 10 VPN Management

1. VPN Management Best Practices

1. Build in Redundancy

2. Choose the Right VPN Product for Your Environment

3. Develop a VPN Policy

4. Prohibit Split Tunneling

5. Ensure Client Security

6. Practice Vulnerability Management

7. Use Multifactor Authentication

8. Document Your Implementation Plan

9. Monitor VPN Availability

10. Perform Regular Reviews, Backups, and Updates

2. Developing a VPN Policy

3. Developing a VPN Deployment Plan

1. Bypass Deployment

2. Internally Connected Deployment

3. DMZ-Based Implementation

4. VPN Threats and Exploits

5. Commercial Versus Open-Source VPNs

6. Differences Between Personal and Enterprise VPNs

7. Balancing Anonymity and Privacy

8. Protecting VPN Security to Support Availability

9. The Importance of User Training

10. VPN Troubleshooting

11. CHAPTER SUMMARY

12. **KEY CONCEPTS AND TERMS**

13. **CHAPTER 10 ASSESSMENT**

20. **CHAPTER 11 VPN Technologies**

1. **Differences Between Software and Hardware Solutions**

1. **Software VPNs**

2. **Hardware VPNs**

2. **Differences Between Layer 2 and Layer 3 VPNs**

3. **Internet Protocol Security (IPSec)**

4. **Layer 2 Tunneling Protocol (L2TP)**

5. **Secure Sockets Layer (SSL)/Transport Layer Security (TLS)**

1. **SSL/TLS and VPNs**

6. **Secure Shell (SSH) Protocol**

7. **Establishing Performance and Stability for VPNs**

1. **Performance**

2. **Stability**

8. **Using VPNs with Network Address Translation (NAT)**

9. **Types of Virtualization**

1. **Desktop Virtualization**

2. **SSL VPN Virtualization**

10. **CHAPTER SUMMARY**

11. **KEY CONCEPTS AND TERMS**

12. **CHAPTER 11 ASSESSMENT**

21. **CHAPTER 12 VPN Implementation**

1. **Operating System–Based VPNs**

2. **VPN Appliances**

1. **Configuring a Typical VPN Appliance**

2. **Client-Side Configuration**

3. **Remote Desktop Protocol**

4. **Using Remote Control Tools**

5. **Using Remote Access**

1. **The Technology for Remote Use**

2. **Choosing Between IPSec and SSL/TLS Remote Access VPNs**

6. **Remote Desktop Services**

1. **RD RemoteApp**

2. **RD Web Access**

3. **RDS and Hosted Services**

7. **Microsoft DirectAccess**

8. **DMZ, Extranet, and Intranet VPN Solutions**

1. **Intranet VPNs**

2. **Extranet VPNs**

9. **Internet Café VPNs**

10. **Online Remote VPN Options**

1. **Security**

2. **Wake-on-LAN Support**

3. **File Sharing**

4. **Remote Printing**

5. **Mac and Mobile Device Support**

11. **The Tor Application**
12. **Planning a VPN Implementation**
 1. **Download and Verify Installation Files**
 2. **Installation Tips**
 3. **Configuring OpenVPN**
 4. **Running OpenVPN**
 5. **Testing and Troubleshooting**
 6. **OpenVPN Private Tunnel**
13. **VPN Implementation Best Practices**
14. **CHAPTER SUMMARY**
15. **KEY CONCEPTS AND TERMS**
16. **CHAPTER 12 ASSESSMENT**

22. CHAPTER 13 Firewall Security Management

1. **Best Practices for Firewall Management**
2. **Security Measures in Addition to a Firewall**
3. **Mitigating Firewall Threats and Exploits**
4. **Concerns Related to Tunneling Through or Across a Firewall**
5. **Testing Firewall Security**
6. **Important Tools for Managing and Monitoring a Firewall**
7. **Troubleshooting Firewalls**
8. **Detecting Firewall Threats and Responding to Incidents**
9. **CHAPTER SUMMARY**
10. **KEY CONCEPTS AND TERMS**
11. **CHAPTER 13 ASSESSMENT**

23. CHAPTER 14 Best Practices for Network Security Management

1. **Essentials of Network Security Management**
2. **Network Security Management Tools**
 1. **Security Plan and Security Policy**
 2. **Security Checklist**
3. **Limiting Network Access and Implementing Encryption**
4. **Physical Security**
5. **Techniques for Preventing and Deterring Incidents**
 1. **Imposing Compartmentalization**
 2. **Using Honeypots, Honeynets, and Padded Cells**
 3. **Antivirus, Monitoring, and Other Host Security Controls**
 4. **Backup and Recovery**
 5. **User Training, Security Awareness, and Ongoing Education**
6. **Incident Response**
 1. **Preparation**
 2. **Detection and Analysis**
 3. **Containment**
 4. **Eradication**
 5. **Recovery**
 6. **Post-Incident Follow-Up**
7. **Fail-Secure, Fail-Open, and Fail-Close Options**
8. **Network Security Assessments and Testing**

1. Security Assessment
2. Configuration Scans
3. Vulnerability Scanning
4. Penetration Testing
5. Post-Mortem Assessment Review
9. Change Management
10. Compliance Auditing
11. Network Security Troubleshooting
12. CHAPTER SUMMARY
13. KEY CONCEPTS AND TERMS
14. CHAPTER 14 ASSESSMENT

24. CHAPTER 15 Emerging Technology and Regulatory Considerations

1. What the Future Holds for Network Security, Firewalls, and VPNs
 1. Threats
 2. Firewall Capabilities
 3. Encryption
 4. Authentication
 5. Metrics
 6. Focus
 7. Securing the Cloud
 8. Securing Mobile Devices
 9. Internet of Things
 10. Mobile IP
 11. Bring Your Own Device (BYOD)
2. Resource Sites for Network Security, Firewalls, and VPNs
 1. Firewall Vendors
 2. Virtual Private Network Vendors
 3. Network Security Websites
 4. Network Security Magazine Websites
3. Tools for Network Security, Firewalls, and VPNs
 1. Commercial Off-the-Shelf (COTS) Software
 2. Open-Source Applications and Tools
4. The Vanishing Network Perimeter
 1. The Impact of Ubiquitous Wireless and Mobile Connectivity
 2. Making Wireless and Mobile Connectivity More Secure
5. Potential Uses of Security Technologies
6. Regulations and Compliance
7. Specialized Firewalls Available
8. Emerging Network Security Technologies
 1. Data Leakage Prevention
 2. Biometrics
 3. Virtualization Security
 4. IP Version 6
 5. VPNs, Firewalls, and Virtualization
9. CHAPTER SUMMARY
10. KEY CONCEPTS AND TERMS

11. **CHAPTER 15 ASSESSMENT**

25. **APPENDIX A Answer Key**

26. **APPENDIX B Standard Acronyms**

27. **Glossary of Key Terms**

28. **References**

29. **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. **xxii**
23. **xxiii**
24. **xxiv**
25. **1**
26. **2**
27. **3**
28. **4**
29. **5**
30. **6**
31. **7**
32. **8**
33. **9**
34. **10**
35. **11**
36. **12**

37. **13**
38. **14**
39. **15**
40. **16**
41. **17**
42. **18**
43. **19**
44. **20**
45. **21**
46. **22**
47. **23**
48. **24**
49. **25**
50. **26**
51. **27**
52. **28**
53. **29**
54. **30**
55. **31**
56. **32**
57. **33**
58. **34**
59. **35**
60. **36**
61. **37**
62. **38**
63. **39**
64. **40**
65. **41**
66. **42**
67. **43**
68. **44**
69. **45**
70. **46**
71. **47**
72. **48**

73. **49**
74. **50**
75. **51**
76. **52**
77. **53**
78. **54**
79. **55**
80. **56**
81. **57**
82. **58**
83. **59**
84. **60**
85. **61**
86. **62**
87. **63**
88. **64**
89. **65**
90. **66**
91. **67**
92. **68**
93. **69**
94. **70**
95. **71**
96. **72**
97. **73**
98. **74**
99. **75**
100. **76**
101. **77**
102. **78**
103. **79**
104. **80**
105. **81**
106. **82**
107. **83**
108. **84**

109. **85**
110. **86**
111. **87**
112. **88**
113. **89**
114. **90**
115. **91**
116. **92**
117. **93**
118. **94**
119. **95**
120. **96**
121. **97**
122. **98**
123. **99**
124. **100**
125. **101**
126. **102**
127. **103**
128. **104**
129. **105**
130. **106**
131. **107**
132. **108**
133. **109**
134. **110**
135. **111**
136. **112**
137. **113**
138. **114**
139. **115**
140. **116**
141. **117**
142. **118**
143. **119**
144. **120**

145. **121**
146. **122**
147. **123**
148. **124**
149. **125**
150. **126**
151. **127**
152. **128**
153. **129**
154. **130**
155. **131**
156. **132**
157. **133**
158. **134**
159. **135**
160. **136**
161. **137**
162. **138**
163. **139**
164. **140**
165. **141**
166. **142**
167. **143**
168. **144**
169. **145**
170. **146**
171. **147**
172. **148**
173. **149**
174. **150**
175. **151**
176. **152**
177. **153**
178. **154**
179. **155**
180. **156**

181.	157
182.	158
183.	159
184.	160
185.	161
186.	162
187.	163
188.	164
189.	165
190.	166
191.	167
192.	168
193.	169
194.	170
195.	171
196.	172
197.	173
198.	174
199.	175
200.	176
201.	177
202.	178
203.	179
204.	180
205.	181
206.	182
207.	183
208.	184
209.	185
210.	186
211.	187
212.	188
213.	189
214.	190
215.	191
216.	192

217. **193**
218. **194**
219. **195**
220. **196**
221. **197**
222. **198**
223. **199**
224. **200**
225. **201**
226. **202**
227. **203**
228. **204**
229. **205**
230. **206**
231. **207**
232. **208**
233. **209**
234. **210**
235. **211**
236. **212**
237. **213**
238. **214**
239. **215**
240. **216**
241. **217**
242. **218**
243. **219**
244. **220**
245. **221**
246. **222**
247. **223**
248. **224**
249. **225**
250. **226**
251. **227**
252. **228**

253. **229**
254. **230**
255. **231**
256. **232**
257. **233**
258. **234**
259. **235**
260. **236**
261. **237**
262. **238**
263. **239**
264. **240**
265. **241**
266. **242**
267. **243**
268. **244**
269. **245**
270. **246**
271. **247**
272. **248**
273. **249**
274. **250**
275. **251**
276. **252**
277. **253**
278. **254**
279. **255**
280. **256**
281. **257**
282. **258**
283. **259**
284. **260**
285. **261**
286. **262**
287. **263**
288. **264**

289. **265**
290. **266**
291. **267**
292. **268**
293. **269**
294. **270**
295. **271**
296. **272**
297. **273**
298. **274**
299. **275**
300. **276**
301. **277**
302. **278**
303. **279**
304. **280**
305. **281**
306. **282**
307. **283**
308. **284**
309. **285**
310. **286**
311. **287**
312. **288**
313. **289**
314. **290**
315. **291**
316. **292**
317. **293**
318. **294**
319. **295**
320. **296**
321. **297**
322. **298**
323. **299**
324. **300**

325. **301**
326. **302**
327. **303**
328. **304**
329. **305**
330. **306**
331. **307**
332. **308**
333. **309**
334. **310**
335. **311**
336. **312**
337. **313**
338. **314**
339. **315**
340. **316**
341. **317**
342. **318**
343. **319**
344. **320**
345. **321**
346. **322**
347. **323**
348. **324**
349. **325**
350. **326**
351. **327**
352. **328**
353. **329**
354. **330**
355. **331**
356. **332**
357. **333**
358. **334**
359. **335**
360. **336**

361. **337**
362. **338**
363. **339**
364. **340**
365. **341**
366. **342**
367. **343**
368. **344**
369. **345**
370. **346**
371. **347**
372. **348**
373. **349**
374. **350**
375. **351**
376. **352**
377. **353**
378. **354**
379. **355**
380. **356**
381. **357**
382. **358**
383. **359**
384. **360**
385. **361**
386. **362**
387. **363**
388. **364**
389. **365**
390. **366**
391. **367**
392. **368**
393. **369**
394. **370**
395. **371**
396. **372**

397. **373**
398. **374**
399. **375**
400. **376**
401. **377**
402. **378**
403. **379**
404. **380**
405. **381**
406. **382**
407. **383**
408. **384**
409. **385**
410. **386**
411. **387**
412. **388**
413. **389**
414. **390**
415. **391**
416. **392**
417. **393**
418. **394**
419. **395**
420. **396**
421. **397**
422. **398**
423. **399**
424. **400**
425. **401**
426. **402**
427. **403**
428. **404**
429. **405**
430. **406**
431. **407**
432. **408**

433. **409**
434. **410**
435. **411**
436. **412**
437. **413**
438. **414**
439. **415**
440. **416**
441. **417**
442. **418**
443. **419**
444. **420**
445. **421**
446. **422**
447. **423**
448. **424**
449. **425**
450. **426**
451. **427**
452. **428**
453. **429**
454. **430**
455. **431**
456. **432**
457. **433**
458. **434**
459. **435**
460. **436**
461. **437**
462. **438**
463. **439**
464. **440**
465. **441**
466. **442**
467. **443**
468. **444**

469. **445**
470. **446**
471. **447**
472. **448**
473. **449**
474. **450**
475. **451**
476. **452**
477. **453**
478. **454**
479. **455**
480. **456**
481. **457**
482. **458**
483. **459**
484. **460**
485. **461**
486. **462**
487. **463**
488. **464**
489. **465**
490. **466**
491. **467**
492. **468**
493. **469**
494. **470**
495. **471**
496. **472**
497. **473**
498. **474**
499. **475**
500. **476**
501. **477**
502. **478**
503. **479**
504. **480**

505. **481**

506. **482**

INFORMATION SYSTEMS SECURITY & ASSURANCE SERIES

ISSA

Network Security, Firewalls, and VPNs

THIRD EDITION

J. Michael Stewart | Denise Kinsey



JONES & BARTLETT
LEARNING



World Headquarters

Jones & Bartlett Learning

5 Wall Street

Burlington MA 01803

978-443-5000

info@jblearning.com

www.jblearning.com

Jones & Bartlett Learning books and products are available through most bookstores and online booksellers. To contact Jones & Bartlett Learning directly, call 800-832-0034, fax 978-443-8000, or visit our website, **www.jblearning.com**.

Substantial discounts on bulk quantities of Jones & Bartlett Learning publications are available to corporations, professional associations, and other qualified organizations. For details and specific discount information, contact the special sales department at Jones & Bartlett Learning via the above contact information or send an email to **specialsales@jblearning.com**.

Copyright © 2021 by Jones & Bartlett Learning, LLC, an Ascend Learning Company

All rights reserved. No part of the material protected by this copyright may be reproduced or utilized in any form, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the copyright owner.

The content, statements, views, and opinions herein are the sole expression of the respective authors and not that of Jones & Bartlett Learning, LLC. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement or recommendation by Jones & Bartlett Learning, LLC and such reference shall not be used for advertising or product endorsement purposes. All trademarks displayed are the trademarks of the parties noted herein. *Network Security, Firewalls, and VPNs, Third Edition* is an independent publication and has not been authorized, sponsored, or otherwise approved by the owners of the trademarks or service marks referenced in this product.

There may be images in this book that feature models; these models do not necessarily endorse, represent, or participate in the activities represented in the images. Any screenshots in this product are for educational and instructive purposes only. Any individuals and scenarios featured in the case studies throughout this product may be real or fictitious, but are used for instructional purposes only.

18365-8

Production Credits

VP, Product Management: Amanda Martin

Director of Product Management: Laura Pagluica

Product Manager: Ned Hinman

Content Strategist: Melissa Duffy

Development Editor: Kim Lindros

Technical Editor: Jeff Parker

Content Coordinator: Paula-Yuan Gregory

Project Manager: Kristen Rogers
Project Specialist: John Fuller
Senior Digital Project Specialist: Angela Dooley
Marketing Manager: Michael Sullivan
Product Fulfillment Manager: Wendy Kilborn
Composition: Exela Technologies
Cover Design: Timothy Dziewit
Text Design: Kristin E. Parker
Media Development Editor: Faith Brosnan
Rights Specialist: James Fortney
Cover Image (Title Page, Part Opener, Chapter Opener): ©
Digital_Art/Shutterstock
Printing and Binding: McNaughton & Gunn

Library of Congress Cataloging-in-Publication Data

Names: Stewart, James Michael, author. | Kinsey, Denise, author.
Title: Network security, firewalls, and VPNs / J. Michael Stewart,
Denise Kinsey.
Description: Third edition. | Burlington, MA : Jones & Bartlett
Learning, [2021] | Includes bibliographical references and index. |
Summary: "Network Security, Firewalls, and VPNs is written for a
course on Network Security and Perimeter Defense, with special
emphasis placed on Firewalls and VPNs - two key tools for
securing networks. It is currently divided into three parts - Part 1 :
Foundations of Network Security, Part 2 : Technical Overview of
Network Security, Firewalls, and VPNs, and Part 3 :
Implementation, Resources, and the Future"—Provided by
publisher.
Identifiers: LCCN 2020022142 | ISBN 9781284183658 (paperback)
Subjects: LCSH: Firewalls (Computer security) | Computer networks
—Security measures.
Classification: LCC TK5105.59 .S7585 2020 | DDC 005.8—dc23

LC record available at [**https://lccn.loc.gov/2020022142**](https://lccn.loc.gov/2020022142)

6048

Printed in the United States of America

24 23 22 21 20 10 9 8 7 6 5 4 3 2 1

This book is dedicated to all who desire to learn more about network security, firewalls, and VPNs. May you gain foundational knowledge with this resource.



© Digital_Art/Shutterstock

Brief Contents

Preface

Acknowledgments

About the Authors

PART 1

Foundations of Network Security

CHAPTER 1

Fundamentals of Network Security

CHAPTER 2

Network Security Threats

CHAPTER 3	<u>Common Network Topologies and Infrastructures</u>
CHAPTER 4	<u>Network Design Considerations</u>
PART 2	<u>Securing the Perimeter with Firewalls</u>
CHAPTER 5	<u>Firewall Fundamentals</u>
CHAPTER 6	<u>Firewall Implementation</u>
CHAPTER 7	<u>Firewall Deployment Considerations</u>
CHAPTER 8	<u>Configuring Firewalls</u>
PART 3	<u>Establishing Remote Access with VPNs</u>
CHAPTER 9	<u>VPN Fundamentals</u>
CHAPTER 10	<u>VPN Management</u>
CHAPTER 11	<u>VPN Technologies</u>
CHAPTER 12	<u>VPN Implementation</u>
PART 4	<u>Implementing Network Security</u>
CHAPTER 13	<u>Firewall Security Management</u>
CHAPTER 14	<u>Best Practices for Network Security Management</u>

CHAPTER 15

**Emerging Technology and Regulatory
Considerations**

APPENDIX A

Answer Key

APPENDIX B

Standard Acronyms

Glossary of Key Terms

References

Index



© Digital_Art/Shutterstock

Contents

Preface

Acknowledgments

About the Authors

PART 1

Foundations of Network Security

CHAPTER 1

Fundamentals of Network Security

What Is Network Security?

What Is Trust?

Who—or What—Is Trustworthy?

What Are Security Objectives?

What Are You Trying to Protect?

**Seven Domains of a Typical IT
Infrastructure**

Goals of Network Security

**How Can You Measure the Success of
Network Security?**

**Why Are Written Network Security
Policies Important?**

Planning for the Worst

**Who Is Responsible for Network
Security?**

**Enhancing the Security of Wired
Versus Wireless LAN
Infrastructures**

Internal and External Network Issues

**Common Network Security
Components Used to Mitigate
Threats**

Hosts and Nodes

Firewalls

Virtual Private Networks

Proxy Servers

Network Address Translation

The Domain Name System

Directory Services

**Intrusion Detection Systems and Intrusion
Prevention Systems**

Network Access Control

TCP/IP Basics

OSI Reference Model

Sub-Protocols

Headers and Payloads

Filtering on Addresses

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 1 ASSESSMENT

CHAPTER 2

Network Security Threats

Hackers and Their Motivation

Favorite Targets of Hackers

**Threats from Internal Personnel and
External Entities**

The Hacking Process

Reconnaissance

Scanning

Enumeration

Attacking

Post-Attack Activities

Common IT Infrastructure Threats

**Hardware Failures and Other Physical
Threats**

Natural Disasters

Accidents

Malicious Code (Malware)

Advanced Persistent Threat

Fast Growth and Overuse

Wireless Versus Wired

Eavesdropping

Hijack and Replay Attacks

Insertion Attacks

Fragmentation Attacks

Buffer Overflows

Session Hijacking, Spoofing, and Man-in-the-Middle Attacks

Session Hijacking

Spoofing Attacks

Man-in-the-Middle Attacks

Covert Channels

Network and Resource Availability Threats

Denial of Service (DoS)

Distributed Denial of Service (DDoS)

Hacker Tools

Social Engineering

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 2 ASSESSMENT

CHAPTER 3

Common Network Topologies and Infrastructures

What Is a Network Topology?

Types of Network Devices

**What Differentiates Logical and
Physical Topologies?**

Types of Physical Topologies

Logical Topology

Creating Logical Topologies

**Differences Between Internet Protocol
Version 4 (IPv4) and Internet
Protocol Version 6 (IPv6)**

IPSec and IPv6

**Examples of Network Infrastructures
and Related Security Concerns**

Workgroups

SOHO Networks

Client/Server Networks

LAN Versus WAN

Thin Clients and Terminal Services

Remote Control, Remote Access, and VPN

Boundary Networks

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 3 ASSESSMENT

CHAPTER 4

Network Design Considerations

Network Design and Defense in Depth

**Achieving Defense in Depth through
Layering**

Planning for Scalability
Working with Senior Management

Protocols

Common Types of Addressing

IPv6

Controlling Communication Pathways

Router Configuration

Encrypted Protocols

Physical Access Management

Filtering

**Intrusion Detection Systems and
Intrusion Prevention Systems**

Hardening Systems

Equipment Selection

**Authentication, Authorization, and
Accounting**

Communication Encryption

**Hosts: Local-Only or Remote and
Mobile**

Redundancy

Endpoint Security

Clients

Servers

Routers

Switches

Firewalls and Proxies

Risk Assessment and Management

What Are Zones of Risk?

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 4 ASSESSMENT

PART 2

Securing the Perimeter with Firewalls

CHAPTER 5

Firewall Fundamentals

What Is a Firewall?

What Firewalls Cannot Do

Why Do You Need a Firewall?

**How Firewalls Work and What
Firewalls Do**

Types of Firewalls

**Individual and SOHO Firewall
Options**

**Managing the Firewall on an ISP
Connection Device**

Converting a Home Router into a Firewall

Uses for Host Software Firewalls

Examples of Software Firewall Products
Using Windows 10's Host Software Firewall
Using a Linux Host Software Firewall

**Uses for Commercial Software
Network Firewalls**

**Uses for Hardware/Appliance
Firewalls**

Next-Generation Firewalls

What Are Virtual Firewalls?

**Dual-Homed and Triple-Homed
Firewalls**

Ingress and Egress Filtering

Types of Filtering

Static Packet Filtering

**Stateful Inspection and Dynamic Packet
Filtering**

Network Address Translation

Application Proxy

Circuit Proxy

Content Filtering

**Selecting the Right Firewall for Your
Needs**

**The Difference Between Buying and
Building a Firewall**

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 5 ASSESSMENT

CHAPTER 6

Firewall Implementation

**Examining Your Network and Its
Security Needs**

**What to Protect and Why
Preserving Privacy**

Proper Firewall Implementation Procedure

Constructing, Configuring, and Managing a Firewall

pfSense

pfSense Requirements

Planning a Firewall Implementation with pfSense

Firewalling a Small Organization: Packet Filtering or Application-Level Firewall, a Proxy Implementation

Firewalling Medium and Large Organizations: Application-Level Firewall and Packet Filtering, a Hybrid System

Firewalling in a Subnet Architecture

Installing the pfSense Firewall

Configuring a Firewall with pfSense

Elements of Firewall Deployment

Testing and Troubleshooting

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 6 ASSESSMENT

CHAPTER 7

Firewall Deployment Considerations

Common Security Strategies for Firewall Deployments

Security Through Obscurity

Least Privilege

Simplicity

Defense in Depth

Diversity of Defense

Chokepoint

Weakest Link

Fail-Safe

Forced Universal Participation

**Authentication, Authorization, and
Accounting**

**Placement of Network Hardware
Firewalls**

Benefit and Purpose of Reverse Proxy

Use and Benefit of Port Forwarding

**Considerations for Selecting a Bastion
Host OS**

Monitoring and Logging

**Understanding and Interpreting
Firewall Logs and Alerts**

**Intrusion Detection Systems and
Intrusion Prevention Systems**

**Security Event and Information
Management**

**Evaluating Needs and Solutions in
Designing Security**

**What Happens When Security Gets in
the Way of Doing Business?**

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 7 ASSESSMENT

CHAPTER 8

Configuring Firewalls

Firewall Rules

Inbound and Outbound Communications Access Control Lists

Composing Firewall Rules

Ordering Firewall Rules

What Should You Allow and What Should You Block?

Essential Elements of a Firewall Policy

Limitations of Firewalls

Improving Performance

The Downside of Encryption with Firewalls

Firewall Enhancements

Management Interfaces

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 8 ASSESSMENT

PART 3

Establishing Remote Access with VPNs

CHAPTER 9

VPN Fundamentals

What Is a Virtual Private Network?

What Are the Benefits of Deploying a VPN?

What Are the Limitations of a VPN?

What Are Effective VPN Policies?

VPN Deployment Models and Architecture

VPN Deployment Models

VPN Architectures

Tunnel Versus Transport Mode

The Relationship Between Encryption and VPNs

Symmetric Cryptography

Asymmetric Cryptography

Hashing

Establishing VPN Connections with Cryptography

Digital Certificates

What Is VPN Authentication?

What Is VPN Authorization?

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 9 ASSESSMENT

CHAPTER 10

VPN Management

VPN Management Best Practices

Build in Redundancy

Choose the Right VPN Product for Your Environment

Develop a VPN Policy

Prohibit Split Tunneling

Ensure Client Security

Practice Vulnerability Management

Use Multifactor Authentication

Document Your Implementation Plan

Monitor VPN Availability

Perform Regular Reviews, Backups, and Updates

Developing a VPN Policy

Developing a VPN Deployment Plan

Bypass Deployment

Internally Connected Deployment

DMZ-Based Implementation

VPN Threats and Exploits

Commercial Versus Open-Source VPNs

Differences Between Personal and Enterprise VPNs

Balancing Anonymity and Privacy

Protecting VPN Security to Support Availability

The Importance of User Training

VPN Troubleshooting

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 10 ASSESSMENT

CHAPTER 11

VPN Technologies

Differences Between Software and Hardware Solutions

Software VPNs

Hardware VPNs

Differences Between Layer 2 and Layer 3 VPNs

Internet Protocol Security (IPSec)

Layer 2 Tunneling Protocol (L2TP)

Secure Sockets Layer (SSL)/Transport Layer Security (TLS)

SSL/TLS and VPNs

Secure Shell (SSH) Protocol

Establishing Performance and Stability for VPNs

Performance

Stability

Using VPNs with Network Address Translation (NAT)

Types of Virtualization

Desktop Virtualization

SSL VPN Virtualization

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 11 ASSESSMENT

CHAPTER 12

VPN Implementation

Operating System–Based VPNs

VPN Appliances

Configuring a Typical VPN Appliance

Client-Side Configuration

Remote Desktop Protocol

Using Remote Control Tools

Using Remote Access

The Technology for Remote Use

Choosing Between IPSec and SSL/TLS

Remote Access VPNs

Remote Desktop Services

RD RemoteApp

RD Web Access

RDS and Hosted Services

Microsoft DirectAccess

DMZ, Extranet, and Intranet VPN

Solutions

Intranet VPNs

Extranet VPNs

Internet Café VPNs

Online Remote VPN Options

Security

Wake-on-LAN Support

File Sharing

Remote Printing

Mac and Mobile Device Support

The Tor Application

Planning a VPN Implementation

Download and Verify Installation Files

Installation Tips

Configuring OpenVPN

Running OpenVPN

Testing and Troubleshooting

OpenVPN Private Tunnel

VPN Implementation Best Practices

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 12 ASSESSMENT

PART 4

Implementing Network Security

CHAPTER 13

Firewall Security Management

**Best Practices for Firewall
Management**

**Security Measures in Addition to a
Firewall**

**Mitigating Firewall Threats and
Exploits**

**Concerns Related to Tunneling
Through or Across a Firewall**

Testing Firewall Security

**Important Tools for Managing and
Monitoring a Firewall**

Troubleshooting Firewalls

**Detecting Firewall Threats and
Responding to Incidents**

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 13 ASSESSMENT

CHAPTER 14

**Best Practices for Network Security
Management**

**Essentials of Network Security
Management**

Network Security Management Tools

**Security Plan and Security Policy
Security Checklist**

**Limiting Network Access and
Implementing Encryption**

Physical Security

**Techniques for Preventing and
Deterring Incidents**

Imposing Compartmentalization

**Using Honeypots, Honeynets, and Padded
Cells**

**Antivirus, Monitoring, and Other Host
Security Controls**

Backup and Recovery

**User Training, Security Awareness, and
Ongoing Education**

Incident Response

Preparation

Detection and Analysis

Containment

Eradication

Recovery

Post-Incident Follow-Up

**Fail-Secure, Fail-Open, and Fail-Close
Options**

**Network Security Assessments and
Testing**

Security Assessment

Configuration Scans

Vulnerability Scanning

Penetration Testing

Post-Mortem Assessment Review

Change Management

Compliance Auditing

Network Security Troubleshooting

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 14 ASSESSMENT

Emerging Technology and Regulatory Considerations

What the Future Holds for Network Security, Firewalls, and VPNs

Threats

Firewall Capabilities

Encryption

Authentication

Metrics

Focus

Securing the Cloud

Securing Mobile Devices

Internet of Things

Mobile IP

Bring Your Own Device (BYOD)

Resource Sites for Network Security, Firewalls, and VPNs

Firewall Vendors

Virtual Private Network Vendors

Network Security Websites

Network Security Magazine Websites

Tools for Network Security, Firewalls, and VPNs

Commercial Off-the-Shelf (COTS) Software

Open-Source Applications and Tools

The Vanishing Network Perimeter

The Impact of Ubiquitous Wireless and Mobile Connectivity

Making Wireless and Mobile Connectivity More Secure

Potential Uses of Security Technologies

Regulations and Compliance

Specialized Firewalls Available

**Emerging Network Security
Technologies**

Data Leakage Prevention

Biometrics

Virtualization Security

IP Version 6

VPNs, Firewalls, and Virtualization

CHAPTER SUMMARY

KEY CONCEPTS AND TERMS

CHAPTER 15 ASSESSMENT

APPENDIX A

Answer Key

APPENDIX B

Standard Acronyms

Glossary of Key Terms

References

Index



© Digital_Art/Shutterstock

Preface

Purpose of This Text

This text is part of the Information Systems Security & Assurance Series from Jones & Bartlett Learning (www.jblearning.com). Designed for courses and curricula in IT Security, Cybersecurity, Information Assurance, and Information Systems Security, this series features a comprehensive, consistent treatment of the most current research and trends in this critical subject area. These titles deliver fundamental information security principles packed with real-world applications and examples. Authored by Certified Information Systems Security Professionals (CISSPs), they deliver comprehensive

information on all aspects of information security. Reviewed word for word by leading technical experts in the field, these texts are not just current, but forward-thinking—putting you in the position to solve the cybersecurity challenges not just of today, but of tomorrow as well.

The first part of this text focuses on common network security topics and the business challenges and threats that you face as soon as you physically connect your organization's network to the public Internet. It presents key concepts and terms and reveals what hackers do when trying to access your network, thus providing you with the necessary foundation in network security for the discussions that follow.

Part 2 defines firewalls, providing you with an understanding of how to use these tools as security countermeasures to solve business challenges. It discusses how to select and deploy firewalls and the tools for managing and monitoring them. It focuses on the practical, giving concrete, step-by-step examples of how to implement a firewall.

Part 3 focuses on virtual private networks (VPNs) and reviews implementing a VPN, the technologies involved, and VPN-management best practices. It also discusses what challenges the future holds for information security professionals involved in network security. It covers the tools and resources available to the professional and scans the horizon of emerging technologies.

Learning Features

The writing style of this text is practical and conversational. Step-by-step examples of information security concepts and procedures are

presented throughout. Each chapter begins with a statement of learning objectives. Illustrations are used both to clarify the material and to vary the presentation. The text is sprinkled with Notes, Tips, FYIs, Warnings, and sidebars to alert the reader to additional and helpful information related to the subject under discussion.

Chapter Assessments appear at the end of each chapter, with solutions provided in the back of the text. Chapter summaries are included in the text to provide a rapid review or preview of the material and to help students understand the relative importance of the concepts presented.

Audience

The material is suitable for undergraduate or graduate students in computer science majors, information science majors, cybersecurity majors, students at a two-year technical college or community college who have a basic technical background, or readers who have a basic understanding of IT principles and want to expand their knowledge.

New to this Edition

This edition provides an improved format, where each major topic (network security, firewalls, VPNs) is examined from introduction through advanced topics before moving to a different topic and includes better integration among the topics. This edition also reflects the latest software versions and technology. The final chapters introduce encompassing best practices for network security and consider the future of technology, regulatory considerations, and people and process management.

Cloud Labs

This text is accompanied by Cybersecurity Cloud Labs. These hands-on virtual labs provide immersive mock IT infrastructures where students can learn and practice foundational cybersecurity skills as an extension of the lessons in this text. For more information or to purchase the labs, visit [**go.jblearning.com/stewart3e**](https://go.jblearning.com/stewart3e).



© Digital_Art/Shutterstock

Acknowledgments

Dr. Denise Kinsey would like to thank Rob and M.C. for your undying support and encouragement; the development team and everyone at JBL with special thanks to Ned, Melissa, and Kim. I appreciate all of your comments and ideas in building this edition.

Jones & Bartlett Learning thanks all of the people who reviewed the second edition of this text. Your feedback helped to shape this revision.

Georgia Brown
American Military University
ECPI

EC Council University

Western Governors University

Tidewater Community College

Shana Bumpas

John Tyler Community College

Colin T. Carlile

Our Lady of the Lake University

Casimer DeCusatis

Marist College

Charles DeSassure

Tarrant County College

M. Lamine Diop

Blue Ridge Community College

Perry Hood

Wiregrass Georgia Technical College

Richard McKinney

Nashville State Community College

Teresa Nash

Lewis-Clark State College

Nicholas Peirce

Thomas Nelson Community College

Richard S. Riddle Jr.

Danville Community College

Jianhua Yang

Columbus State University

Tatyana Zidarov

Hagerstown Community College



© Digital_Art/Shutterstock

About the Authors

James Michael Stewart has been working with computers and technology for more than 25 years. His work focuses on security, certification, and various operating systems. Recently, Michael has been teaching job-skill and certification courses such as CISSP, CEH, and Security+. He is the primary author of the *CISSP Study Guide, 4th Edition* and the *Security+ 2008 Review Guide*. In addition, Michael has written numerous texts on other security and Microsoft certification and administration topics. He has developed certification courseware and training materials as well as presented these materials in the classroom. Michael holds the following certifications: CISSP, ISSAP, SSCP, MCT, CEI, CEH, TICSA, CIW SA, Security+, MCSE+Security: Windows 2000, MCSA Windows Server 2003, MCDST, MCSE NT & W2K, MCP+I, Network+, iNet+. He

graduated in 1992 from the University of Texas at Austin with a bachelor's degree in philosophy.

Denise Kinsey, Ph.D has worked in computer networking and cybersecurity for over 20 years. She has designed, implemented, and managed IT and OT projects for government and the private sector. Denise is the author of several IT and cybersecurity texts meant to educate and empower readers to learn more about technology and implement secure systems for their employers and their homes. When not solving cybersecurity problems she loves to work with her students to improve the networks and cybersecurity of nonprofits and local businesses. Denise holds numerous certifications, including CISSP, C|CISO, Security+, and many others. Dr. Kinsey is an Associate Professor.



© Digital_Art/Shutterstock

Foundations of Network Security

CHAPTER 1

Fundamentals of Network Security

CHAPTER 2


Network Security Threats

CHAPTER 3

Common Network Topologies and Infrastructures

CHAPTER 4

Network Design Considerations



© Digital_Art/Shutterstock

Fundamentals of Network Security

COMPUTER NETWORK SECURITY is very complex. New threats from inside and outside networks appear constantly.

In response, the security community is constantly developing new products and procedures to defend against threats of the past and unknown threats of the future.

As companies merge, people lose their jobs, new equipment comes online, and **business tasks** change, people do not always do what you expect. Network security configurations that worked well yesterday might not work quite as well tomorrow. In an ever-changing business climate, whom should you trust? Has your trust been violated? How would you even know? Who is attempting to harm your network this time? And why?

Because of these complex issues, you need to understand the essentials of network security. This chapter will introduce you to the basic elements of network security. Once you have a firm grasp of these fundamentals, you will be well equipped to put effective security measures into practice on your organization's network.

CHAPTER 1 TOPICS

This chapter covers the following topics and concepts:

- What is meant by network security
- The items protected within the seven domains of a typical IT infrastructure
- The goals of network security
- How to assess the success of a network security implementation
- Why written network security policies are important
- Who is responsible for network security
- Controls to enhance wired and wireless local area network (LAN) infrastructure security
- Examples of internal and external network issues
- The common network security components used to mitigate threats throughout the IT infrastructure
- The basics of Transmission Control Protocol/Internet Protocol (TCP/IP)

CHAPTER 1 GOALS

When you complete this chapter, you will be able to:

- Describe the key concepts and terms associated with network security

- Describe the importance of a written security policy; explain how policies help minimize risk exposure and threats to a network infrastructure
- Define network security roles and responsibilities; identify who within an IT organization is accountable for these security implementations
- Identify examples of network security concerns or threats that require enhanced security countermeasures to properly minimize risk exposure and threats
- Describe the security requirements needed for wired versus wireless LAN infrastructures in order to provide an enhanced level of security
- Compare and contrast common network security components and devices, and their use throughout the IT infrastructure

WHAT IS NETWORK SECURITY?

Network security is the control of unwanted intrusion into, use of, or damage to communications on a computer network. This includes monitoring for abuses, looking for protocol errors, blocking nonapproved transmissions, and responding to problems promptly and accurately. Network security is also about supporting essential communication necessary to the organization's mission and goals, avoiding the unapproved use of resources, and ensuring the integrity of the information traversing the network.

Network security includes elements that prevent unauthorized access and actions while supporting desirable activities, such as accessing customer records or managing inventory. This is hard to do efficiently, cost effectively, and transparently. Efficient network security provides quick and easy access to **resources** for users. Cost-effective network security controls user access to resources and services, as long as the risks justify the control costs. Transparent network security supports the mission and goals of the organization through enforcement of the organization's network security policies, without unnecessary steps or delay of valid users performing legitimate tasks.

Computer networking technology is changing and improving faster today than ever before. Wireless connectivity is now a realistic option for most organizations and individuals. Malicious hackers are becoming more adept at using every means available to steal identities, intellectual property, and money.

Today, many organizations spend more time, money, and effort protecting their assets than they do during the initial installation of the network. **Threats**, both internal and external, can cause a catastrophic system failure or compromise. Such security breaches can even result in a company going out of business. Without network security, many businesses and individuals would not be able to work productively.

Network security must support workers in their jobs, while protecting against compromise, maintaining high performance, and keeping costs to a minimum. This can be incredibly challenging to achieve; but, with proper planning and execution, many organizations have successfully accomplished this task.

Network security begins with trust.

What Is Trust?

Trust is confidence in your expectation that others will act in your best interest. With computers and networks, trust is the confidence that other users will act in accordance with your organization's security rules and corporate mission and vision. You trust that they will not attempt to violate the stability, privacy, or integrity of the network and its resources. Trust is the belief that others are trustworthy.

Unfortunately, sometimes people violate your trust. They may do this by accident, oversight, or ignorance that the expectation even existed. In other situations, they violate trust deliberately. Because these people can be either internal personnel or external hackers, it's difficult to know whom to trust.

So, how can you answer the question, "Who is trustworthy?" You begin by realizing that trust is based on past experiences and behaviors. Trust is usually possible between people who already know each other. It's neither easy nor desirable to trust strangers. However, once you've defined a set of rules and everyone agrees to abide by those rules, you have established conditional trust. Over time, as people demonstrate that they are willing to abide by the rules and meet the expectations of conduct, you can consider them trustworthy.

Trust can also come from using a third-party method. If a trustworthy third party knows each of the entities and states that both are trustworthy, then you can extend conditional trust to each of the entities. Over time, someone's behavior shows whether or not the initial conditional trust was merited.

Ultimately, network security is based on trust. Organizations assume that their employees, as well as all of the computers and network devices, are trustworthy. But not all trust is necessarily the same. You can (and probably should) operate with different levels or layers of trust. Those with a higher level of trust can be assigned greater permissions and privileges. In terms of networking, **permission** is the authorization to access an asset, and **privilege** is the ability granted on the network. If someone or something violates your trust, you remove the violator's access to the secure environment. For example, organizations terminate an untrustworthy employee or replace a defective operating system.

A common example of a third-party trust system is the use of digital certificates that a public certificate authority issues. As shown in **FIGURE 1-1**, a user communicates with a web e-commerce **server**. The user does not initially know whether a web server is what it claims to be or if someone is “spoofing” its identity. Once the user examines the digital certificate issued to the web server from a certificate authority, the user can then trust that the identity of the website is valid. This occurs because both the user and the website have a common, trustworthy third party that each knows.

Certificate Authority



Certificate issued
to User



Certificate issued
to Web site



Web User



Initial unknown trust
relationship



Web E-commerce Server

FIGURE 1-1 An example of a third-party trust system.

Who—or What—Is Trustworthy?

Determining who or what is trustworthy is an ongoing activity of every organization, in both global corporations and a family's home network. In both cases, you offer trust to others on a conditional basis. This conditional trust changes over time, based on adherence to or violation of desired and prescribed behaviors.

If a program causes problems, it loses your trust and you remove it from the system. Similarly, if a user violates security, that person loses your trust and you might revoke access privileges. When a worker abides by the rules, your trust grows and privileges may increase if necessitated by additional job responsibilities or need. In the same way, when an Internet site does not cause harm, you deem it trustworthy and allow access to that site.

To review, trust is subjective, tentative, and changes over time. You can offer trust based on the reputation of a third party. You withhold trust when others violate the rules. Trust stems from actions in the past and can grow based on future behaviors.

In network security, trust is complex. Extending trust to employees without proper background investigation can be devastating. A network is only as secure as its weakest link. You need to scrutinize every aspect of a network, including software, hardware, configuration, communication patterns, content, and the users, to maintain network security. Otherwise, you will not be able to accomplish the security objectives of your organization.

What Are Security Objectives?

Security objectives are goals an organization strives to achieve through its security efforts. Typically, organizations recognize three primary security objectives:

- Confidentiality/privacy
- Integrity/nonrepudiation
- Availability/uptime

Confidentiality is the protection against unauthorized access, while providing authorized users access to resources without obstruction. Confidentiality ensures that data is not intentionally or unintentionally disclosed to anyone without a valid need to know. A job description defines the person's need to know. If a task does not require access to a specific resource, then that person does not have a need to access that resource.

Integrity is the protection against unauthorized changes, while allowing for authorized changes performed by authorized users. Integrity ensures that data remains consistent, both internally and externally. Consistent data does not change over time and remains in sync with the real world. Integrity also protects against accidents and hacker modification by **malicious code**, or software written with malicious intent.

Availability is the protection against **downtime**, loss of data, and blocked access, while providing consistent uptime, protecting data, and supporting authorized access to resources. Availability ensures that users can get their work done in a timely manner with access to the proper resources. Availability, for example, ensures that web resources are accessible, so sales are not lost.

Authentication is the proof or verification of a user's identity before granting access to a secured area. This can occur on a network or in the physical, real world. While the most common form of authentication is a password, passwords can be accidentally shared or stolen, especially if written down. Multifactor authentication is the method most network administrators prefer for secure login, as it requires multiple proofs of one's identity.

Authorization is controlling what users are allowed and not allowed to do. Authorization is dictated by the organization's security structure, which may focus on discretionary access control (DAC), mandatory access control (MAC), or role-based access control (RBAC). Authorization restricts access based on need to know and users' job descriptions. Authorization is also known as **access control**.

Nonrepudiation is the security principle that prevents a user from being able to deny having performed an action. For example, nonrepudiation prevents a sender from denying having sent a message. Public-key cryptography commonly provides nonrepudiation services.

Privacy protects the confidentiality, integrity, and availability of personally identifiable or sensitive data. Private data can include financial records or medical information. Privacy prevents the unauthorized watching and monitoring of users and employees.

Maintaining and protecting these security objectives can be a challenge. As with most difficult tasks, breaking security down into simpler or smaller components will help you to understand and ultimately accomplish this objective. To support security objectives, you need to know exactly what you are trying to protect.

WHAT ARE YOU TRYING TO PROTECT?

In terms of security, the things you want to protect are known as assets. An **asset** is anything used to conduct business. Any object, computer, program, piece of data, or other logical or physical component that employees need to accomplish a task is an asset.

Assets do not have to be expensive, complicated, or large. In fact, many assets are relatively inexpensive, commonplace, and variable in size. But no matter the characteristics, an asset needs protection. When assets are unavailable for whatever reason, employees cannot complete their work.

For most organizations, whether small or large, the assets of highest concern include business and personal data. If this information is lost, damaged, or stolen, serious complications result. Businesses can fail. Individuals can lose money. Identities can be stolen. Lives can be ruined.

What causes these problems? What violates network security? The answer includes accidents, ignorance, and hackers. Accidents happen, including hardware failures and natural disasters. Poor training equals ignorance. Workers with the best of intentions can damage systems if they do not know proper procedures or if they lack necessary skills. Overworked and rushed personnel overlook issues that can result in asset compromise or loss. Malicious hackers can launch attacks and **exploits** against the network, seeking to gain access or simply to cause damage.

Seven Domains of a Typical IT Infrastructure

No aspect of an IT infrastructure is without risk, nor is it immune to the scrutiny of a hacker. **Risk** is the likelihood that a threat will take advantage of a **vulnerability**. When designing and implementing network security, you need to analyze every one of the seven **domains** of a typical IT infrastructure (**FIGURE 1-2**) for potential vulnerabilities and weaknesses. Security measures must be detailed, focused, and exhaustive. You must consider every possible avenue of attack; assess risk; and if the risk is sufficient, apply a **countermeasure**. Failing to do so will leave an open pathway for a hacker. A hacker needs only one crack in your defenses to begin chipping away at the security of the entire network.




FIGURE 1-2 The seven domains of a typical IT infrastructure.

Each of the seven domains of a typical IT infrastructure has unique aspects that need security improvements. The following sections describe important foundational network security issues related to the seven domains.

User Domain

This domain refers to actual users, whether they are employees, consultants, contractors, or other third-party users. Any user who accesses and uses the organization's IT infrastructure must review and sign an acceptable use policy (AUP) prior to being granted access to the organization's IT resources and infrastructure. This domain should also focus on training, strong authentication, granular authorization, and detailed accounting. Additionally, many of the protections added to other domains provide additional protections for and against the User Domain.

Social engineering is a hacker attack against people. This type of attack attempts to fool people through clever wording, lying, misdirection, relationship manipulation, fear, psychological tricks, and confidence games. The results of social engineering attacks can include users giving away private information or performing actions on a computer that cause a reduction in network security.

Workstation Domain

This domain refers to the end user's desktop devices, such as a desktop computer, laptop, Voice over Internet Protocol (VoIP) telephone, or other endpoint device. Workstation devices typically require security countermeasures such as antivirus, anti-spyware, and vulnerability software patch management to maintain the integrity of the device.

The workstation, client, or standalone home system can be a target. These types of hosts are often less secure than local area network servers. This is not on purpose, because most security focuses on core components of the infrastructure, rather than those seemingly on the periphery. Many times, the security measures on a workstation are old, out-of-date, or improperly installed and configured. System hardening, communication protection, and positioning of work stations are critical to security.

LAN Domain

This domain refers to the physical and logical **local area network (LAN)** technologies (i.e., the wired or wireless topology built from devices and communication protocols) used to support workstation connectivity to the organization's network infrastructure. Protocols, addressing, topology, and communication encryption provide security for this domain.

The private LAN, from home offices to large corporations, is a common target for attack. A LAN often consists of dozens to thousands of hosts, which can make securing every host difficult. Once a hacker gains access to one system on the network, the rest of the LAN is vulnerable to attack. The compromise of a single host can lead to the compromise of the entire infrastructure.

LAN-to-WAN Domain

This domain refers to the organization's internetworking and interconnectivity point between the LAN and the WAN network infrastructures. Routers, firewalls, switches, proxies, demilitarized zones (DMZs), intrusion detection systems (IDSs), intrusion prevention systems (IPSs), and communication encryption are commonly used as security monitoring devices in this domain.

The **wide area network (WAN)** connections between LAN locations, especially those controlled by third-party entities, are targets. Transition interfaces between a private LAN and a WAN connection can be potential weak points. If a compromise is successful against a WAN connection, the malicious traffic inbound across the WAN/LAN interface is unlikely to be filtered.

Remote Access Domain

This domain refers to the authorized and authenticated remote access procedures for users to remotely access the organization's IT infrastructure, systems, and data. Remote access is always a popular target of hackers, as it removes the need for the hacker to be physically present to access and attack the LAN. Hackers anywhere in the world with an Internet or telephone connection can still reach out to attack any seemingly isolated target. Remote access is an invitation to hackers to try to breach your defenses.

Remote access solutions typically involve encrypted remote browser access or encrypted virtual private network (VPN) tunnels for secure remote communications.

WAN Domain

Organizations with remote locations require a WAN for interconnection. Organizations typically outsource WAN connectivity from service providers for end-to-end connectivity and bandwidth.

The WAN Domain includes networks, such as those using Asynchronous Transfer Mode (ATM) or Frame Relay, owned by a **telco** or a carrier network company that leases access to corporations. Often, the privacy of these WAN connections is electronic isolation rather than encryption. Hackers focusing on a specific target may attempt to breach the electronic isolation of the carrier network, rather than focusing on the endpoint LANs. This domain typically includes routers, circuits, switches, firewalls, and equivalent gear at remote locations, sometimes under a managed service offering by the service provider. Protocol selection, addressing schemes, and communication encryption are elements of securing this domain.

System/Application Domain

This domain refers to the hardware, operating system software, database software, client/server applications, and data that are typically housed in the organization's data center and/or computer rooms. Collections of servers hosting applications, virtualized systems, or databases are valuable targets. Sometimes the data hosted is the target. Sometimes the computing power of the servers is the resource the hacker wishes to seize. Network design, authentication, authorization, accounting, and node security are important security concerns for this domain.

Network administrators need to recognize that the potential for compromised security exists throughout an organization, which leads to the need for adequate network security. Starting from the knowledge that risk exists and threats loom, network security administrators can design and implement appropriate countermeasures and safeguards. The first step is to recognize that the potential for compromised security exists throughout an organization. The next step is to comprehend the goals of network security.

GOALS OF NETWORK SECURITY

Network security goals vary from organization to organization.

Often, however, there are a few common mandates:

- Ensure the confidentiality of resources
- Protect the integrity of data
- Maintain availability of the IT infrastructure
- Ensure the privacy of personally identifiable data
- Enforce access control
- Monitor the IT environment for violations of policy
- Support business tasks and the overall mission of the organization

In order to accomplish your organization's security goals, you need to write down those goals and develop a thorough plan to execute them.

Without a written plan, security will be haphazard at best and will likely fail to protect your assets. With a written plan, network security is on the path to success. Once you define them, these security goals will become your organization's roadmap for securing the entire IT infrastructure.

HOW CAN YOU MEASURE THE SUCCESS OF NETWORK SECURITY?

An organization measures the security of its network by how well its stated security goals are accomplished and its security standards are maintained. In essence, this becomes the organization's baseline definition for information systems security. For example, if private information on the network does not leak to outsiders, then your efforts to maintain confidentiality were successful. Or, if employees are able to complete their work on time and on budget, then your efforts to provide system integrity protection were successful.

If violations take place that compromise your assets or prevent the accomplishment of a security goal, however, then network security was less than successful. With well-designed and executed security, accidents, mistakes, and even intentional harmful exploits will challenge your best efforts. The perfect security components do not exist. All have weaknesses, limitations, backdoors, workarounds, programming bugs, or some other exploitable element.

Fortunately, though, successful security does not rely on the installation of a single defensive component. Instead, good network security relies on an interweaving of multiple effective security components. You don't have only one lock on your house. By combining multiple protections, defenses, and detection systems, you can rebuff many common, easy hacker exploits.

Network security success is not about preventing all possible attacks or compromises. Instead, you work to continually improve the state of security so that in the future, the network is better protected than it was in the past. As hackers create new exploits, security professionals learn about them, adapt their methods and systems, and establish new defenses. Successful network security is all about constant vigilance, not about creating an end product. Security is an ongoing effort that constantly changes to meet the challenge of new threats, requiring that network security professionals continually learn and adapt.

WHY ARE WRITTEN NETWORK SECURITY POLICIES IMPORTANT?

With a written **security policy**, all security professionals strive to accomplish the same end: a successful, secure work environment. Written plans and procedures focus security efforts and resources on the most important tasks to support your organization's overall security objectives. Without solid and defined goals, your security efforts would be chaotic and hard to manage.

The plan provides a common reference against which security tasks are compared. It serves as a measuring tool to judge whether security efforts are helping rather than hurting your organization's security objectives. By following the written plan, you can track progress so that you install and configure all the necessary components. A written plan validates what you do, defines what you still need to do, and guides you on how to repair the infrastructure when necessary. Without a written security policy, workers won't have a reliable guide on what to do, and judging security success will be impossible. Without a written policy, you have no security.

Planning for the Worst

Things invariably go wrong. Users make mistakes. Malicious code finds its way into your network. Hackers discover vulnerabilities to exploit. In anticipating problems that threaten security, you must plan for the worst.

This type of planning has many names, including contingency planning, worst-case scenario planning, business continuity planning, disaster recovery planning, and continuation-of-operations planning. The name is not important. What is crucial is that you construct a team and do the planning.

When problems occur, shift into response mode: respond, contain, and repair. Respond to all failures or security breaches to minimize damage, cost, and downtime. Contain threats to prevent them from spreading or affecting other areas of the infrastructure. Repair damage promptly to return systems to normal status quickly and efficiently. Remember, the goals of security are confidentiality, integrity, and availability. Keep these foremost in mind as you plan for the worst.

The key purpose of planning for problems is to be properly prepared to protect your infrastructure. With a little luck, a major catastrophe will not occur. But it is better to prepare and not need the response plan than to allow problems to cause your business to fail.

WHO IS RESPONSIBLE FOR NETWORK SECURITY?

Network security is the responsibility of everyone who uses the network within an organization. No one has the luxury of ignoring security rules. This applies to global corporations as well as home networks. Every person is responsible for understanding how to support and maintain network security. The weakest link rule applies here: If only one person fails to fulfill this responsibility, security for all will suffer.

Senior management has the ultimate and final responsibility for security. This is for good reason; senior management is the most accountable for the protection of the organization's assets. Without the approval and support of senior management, no security effort can succeed. Senior management must ensure the creation of a written security policy that all personnel understand and follow.

Senior management also assigns the responsibility for designing, writing, and executing the security plan to the IT staff. Ideally, the result of these efforts is a secure network infrastructure. The security staff, in turn, must effectively manage all assets, system vulnerabilities, imminent threats, and pertinent defenses. Their task is to design, execute, and maintain security throughout the organization.

In their role as overseers of groups of personnel, managers and supervisors must ensure that employees have all the tools and resources to accomplish their work. Managers must also ensure that workers are properly trained; possess necessary skills; and know

procedures, policies, boundaries, and restrictions. Employees can mount a legitimate legal case against an organization that requires them to perform work for which they are not properly trained.

Network administrators manage all of the organization's computer resources. Resources include file servers, network access, databases, printer pools, and applications. The network administrator's job is to ensure that resources are functional and available for users, while enforcing confidentiality and network integrity.

An organization's workers are the network users and operators. They ultimately do the work the business needs to accomplish. Users create products, provide services, perform tasks, input data, respond to queries, and much more. Job descriptions may apply to a single user or a group of users. Each job description defines a user's tasks. Users must perform these tasks within the limitations of network security.

Auditors watch for problems and violations. Auditors investigate the network, looking for anything not in compliance with the written security policy. Auditors watch the activity of systems and users to look for violations, trends toward **bottlenecks**, and attempts to perform violations. The information uncovered by auditors can help improve the security policy, adjust security configurations, or guide investigators toward apprehending security violators.

All of these **roles** exist within every organization. Sometimes different individuals perform these roles. In other situations, a single person performs all of these roles. In either case, these roles are essential to the creation, maintenance, and improvement of security.

ENHANCING THE SECURITY OF WIRED VERSUS WIRELESS LAN INFRASTRUCTURES

Wired networks offer a form of security that wireless networks lack. That security is the direct physical connection to a wired or cabled network. A hacker requires physical access to your facility or building. Usually, access control of the building can sufficiently prevent most external parties from accessing your private LAN.

Realize, however, that if you allow remote connection via telephone modem, high-speed broadband, or even basic Internet service, then you lose the advantage of a physical access limitation. Once remote access is allowed, the benefits of physical isolation disappear.

The same is true if you allow wireless connections into the network. Wireless networking grants valid and unknown users the ability to interact with the network. This eliminates the need to be physically present in the building to connect to the network. Many organizations are allowing—even encouraging—employees, contractors, and visitors to Bring Your Own Device (BYOD). This is cost effective and often improves efficiency, but it places the end users' devices at risk of control. With the right type of antenna, an attacker could be more than a mile away from your office building and still be able to affect your wireless network.

■ NOTE

An attack known as “Van Eck phreaking” allows an attacker to eavesdrop on electronic devices from a distance. This technique is neither perfect nor simple to perform, but it has been demonstrated on LCD and CRT monitors, as well as on keyboard cables. With minor shielding, you can eliminate most of the risk from such an attack.

To regain some of the security offered by physical isolation, try to incorporate physical isolation into your network design. Isolate all remote access and wireless access points from the main wired network. You can achieve this by using separate subnets and by **filtering** communications using firewalls. Although design does not offer the same level of security as physical isolation, the arrangement provides a significant improvement over having no control over remote or wireless connections.

■ NOTE

You will learn about the Transmission Control Protocol/Internet Protocol (TCP/IP) and the layers within which protocols work later in this chapter.

All remote connections should go through a rigorous gauntlet of verification before you grant access to the internal LAN. Think of a castle design in the Middle Ages that used multiple layers of defense:

a moat, a drawbridge, thick riveted walls, an inner battlement, and finally, a strong keep or inner fortress. Your multilayer defensive design should include multifactor authentication and communication encryption, such as a VPN. Additional checks can include verification of operating system and patch level, confirmation of physical or logical location origin (such as caller ID, MAC address, or IP address), limitations to time of day, and limitations on protocols above the Transport Layer. Any intruder would need to circumvent layer after layer, making intrusion more and more difficult. This is often called a layered defense or **defense in depth**.

INTERNAL AND EXTERNAL NETWORK ISSUES

When deploying a new network or modifying an existing infrastructure, carefully evaluate effects on security and security's impact on the infrastructure. When you or your team overlook or sidestep standard security practices, you can actually decrease security or break the network. Business interruptions can result not only in lost profits, but also in lost opportunities and even in lost jobs. If the compromise is serious enough, the business might not recover.

The threats facing business are numerous and constantly changing. Often, these issues arise daily. Malicious code, information leakage, zero-day exploits, unauthorized software, unethical employees, and complex network infrastructures are just a few of the concerns that every organization and network manager must face.

Malicious code can make its way into a computer through any communication channel. This includes file transfers, email, portable device syncing, removable media, and websites. Precautions against malicious code include network traffic filtering with a firewall, anti-malware scanning, and user behavior modification.

Information leakage often stems from malicious employees who purposefully (maliciously) or accidentally release internal documentation to the public. It can also result from accidents, or can occur when a storage device is lost, recycled, donated, stolen, or improperly thrown away. It also occurs when users accidentally publish documents to peer-to-peer (P2P) file-sharing services or

websites. Precautions against information leakage include conducting thorough background checks on employees, using the principle of least privilege to provide access only to what is essential for completing one's job responsibilities, detailed auditing and monitoring of all user activity, classifying all information and controlling communication pathways, using more stringent controls on use of portable devices, and practicing **zeroization**—purging a storage device to be discarded by filling its space with zeros.

Zero-day exploits are new and previously unknown attacks for which no current specific defenses exist. Think of these as surprise attacks on your networks. *Zero-day* refers to the newness of an exploit, which may be known in the hacker community for some time, but about which vendors and security professionals are just learning.

The zero-day label comes from the idea that work to develop a patch begins the moment a vendor learns of a problem. The moment of discovery of the new exploit is called “day zero.” No specific defenses against zero-day attacks exist, but general security management, use of intrusion detection and intrusion prevention, along with detailed logging and monitoring can assist in quickly discovering and preventing new attacks. Once you know of an attack or exploit, you can begin to take steps to contain damage or minimize the extent of the compromise.

Unauthorized software is any piece of code or program that a user chooses to run on a client system that was not approved by or provided by the organization. Not all unauthorized software is directly problematic, but you should prevent its use for many reasons. Such software could cost the organization money, time, and effort.

Installation or use of the software could be a license violation. The software could include hidden malicious components, known or unknown to the user, which could compromise the security of the network. Steps you can take to prevent the use of unauthorized software should include limiting installation privileges of normal users and using **whitelists** to block the execution of any program not on the approved list.

Users are the final link in security and are often the weakest link in network security. If a user chooses to violate a security policy and release information to the public or execute malicious code, the results could devastate the organization or land the perpetrator in court. Unethical employees purposefully violate the stated rules and goals of the organization. These employees often believe the rules are not important, do not really apply to them, or are not really enforced. Most believe they will never get caught. When users violate the mission and goals of the organization, consequences could be catastrophic.

Methods of preventing unethical employees from doing damage include better background screening, detailed auditing and monitoring of all user activity, and regular management oversight and job performance reviews. Background screening should be more detailed for those employees in positions of trust or who have access to confidential or proprietary information of the organization. When you discover a problematic person, you might be able to grant the employee a second chance after retraining. In many cases, however, the safest choice for the organization is to terminate employment.

Complex network infrastructures often lead to complex vulnerabilities. The larger a network becomes, the more servers,

clients, network devices, and segments it includes. The sheer number of moving parts almost guarantees that something is bound to be misconfigured, improperly installed, lack current firmware or patches, have or induce a bottleneck, or be used incorrectly. Any of these conditions could result in a vulnerability that internal or external attackers can exploit. The larger and more complex a network, the more thoroughly the security team needs to watch over the infrastructure and investigate every symptom, trend, or alert that exceeds the organization's threshold for acceptable use. Preventing complexity from becoming a liability involves detailed planning, careful implementation, regular security management, and constant review of the effectiveness of the infrastructure.

Studies have shown that most threats come from internal sources, but too many organizations focus on external sources and discount the internal threats. A better stance is to count all threats—regardless of the source—as worthy of investigation. Once a potential threat is understood, its risk, potential loss, and likelihood can be better understood and evaluated.

One of the most obvious external threats is the Internet. The Internet is a global network linking people and resources with high-speed, real-time communications. Unfortunately, this wonderful infrastructure is the tool of choice for hackers and malware distributors. Once your organization installs an Internet connection, the world is at your door—or rather, at the fingertips of every employee. That world includes both potential customers and potential attackers.

Without a global communication infrastructure, hackers had to be physically present—wired in—in or near your building to launch

attacks. With the advent of the Internet and wireless technology, any hacker anywhere can initiate attempts to breach your network security. Often, attacks are not targeted at your particular organization, but rather at a specific operating system, communication protocol, or software vulnerability. It is helpful, therefore, to think of the Internet as a threat. It is not one that you should lightly dismiss, nor should you discard the Internet as a powerful tool. The benefits of Internet access are well worth the effort and expense you should expend to defend against its negative features.

Some of the best defenses against Internet threats include a well-researched and written security policy; thoroughly trained personnel; use of firewalls to filter traffic; intrusion, detection, and prevention systems; use of encrypted communications (such as VPNs); and thorough auditing and monitoring of all user and node activity.

Again, perfect security solutions do not exist. Some form of attack, compromise, or exploit can get past any single defense. The point of network security is to interweave and interconnect multiple security components to construct a multifaceted scheme of protection (that is, defense in depth).

The goal is to balance the strengths and weaknesses of multiple security components. The ultimate functions of network security are to lock things down in the best way possible, then monitor for all attempts to violate the established defense. The perfect lock does not exist, so improve on the best locks available with auditing and monitoring. Once you know this, your next step is to understand the common network security components and the best uses of each.

COMMON NETWORK SECURITY COMPONENTS USED TO MITIGATE THREATS

When considering the deployment of a network or the modification of an existing network, evaluate each network component to determine its security strengths and weaknesses. This section discusses many common network security components.

Hosts and Nodes

A **node** is any device on the network. This includes client computers, servers, switches, routers, firewalls, and anything with a network interface that has a MAC address. A **Media Access Control (MAC) address** is the 48-bit physical **hardware address** of a network interface card (NIC) assigned by the manufacturer. A node is a component that can be communicated with, rather than only through or across. For example, network cables and patch panels are not nodes, but a printer is a node.

A **host** is a form of node that has a logical address assigned to it, usually an Internet Protocol (IP) address. This address typically implies that the node operates at or above the Network Layer. The Network Layer includes clients, servers, firewalls, proxies, and even routers. The Network Layer excludes switches, bridges, and other physical devices, such as repeaters and hubs. In most cases, a host either shares or accesses resources and services from other hosts.

In terms of network security, node and host security does vary. Nodes and hosts can both be harmed by physical attacks and denial of service (DoS) attacks. However, a host can also be harmed by malicious code or authentication attacks, and might even be remotely controlled by hackers. Node protection is mostly physical access control, along with basic network filtering against flooding.

Host security can be much more involved, because the host itself should be hardened and you will need to perform general network security. **Hardening** is the process of securing or locking down a host against threats and attacks. This can include removing unnecessary software, installing updates, and imposing secure configuration settings.

Firewalls

Not all traffic on your network is from an authorized source, so you shouldn't allow it to enter or leave the network. Not all traffic is for an authorized purpose, so you should block it from reaching its destination. Not all traffic is within the boundaries of normal or acceptable network activity, so you should drop it before it causes compromises.

All of these protections are the job of a firewall. As its name implies, a **firewall** is a tool designed to stop damage, just as the firewall in an engine compartment protects the passengers in a vehicle from harm in an accident. A firewall is either a hardware device or a software product you deploy to enforce the access control policy on network communications. In other words, a firewall filters network traffic for harmful exploits, incursions, data, messages, or other events.

Firewalls are often positioned on the edge of a network or subnet. Firewalls protect networks against numerous threats from the Internet. Firewalls also protect the Internet from rogue users or applications on private networks. Firewalls protect the throughput or bandwidth of a private network so resources are available to authorized users. Without firewalls, most of your network's capabilities would be consumed by worthless or malicious traffic from the Internet. Think of how a dam on a river works; without the dam, the river is prone to flooding and overflow. The dam prevents the flooding and damage.

Without firewalls, the security and stability of a network would depend mostly on the security of the nodes and hosts within the network. Based on the sordid security history of most host operating

systems, having no firewall would not be a secure solution. Hardened hosts and nodes are important for network security, but these should not be the only components of reliable network security. Firewalls can be installed on client and server computers. These host software firewalls protect a single host from threats from the Internet, as well as threats from the network and other internal network components.

In any case, a firewall is usually configured to control traffic based on a *deny-by-default/allow-by-exception* stance. This means that nothing passes the firewall just because it exists on the network (or is attempting to reach the network). Instead, all traffic that reaches the firewall must meet a set of requirements to continue on its path.

As a network administrator or IT security officer, you choose which traffic is allowed to pass through the firewalls and which traffic is blocked. Additionally, you can also determine whether the filtering takes place on inbound traffic (known as **ingress filtering**), outbound traffic (called **egress filtering**), or both. Filtering inbound traffic is a normal function, as many threats exist outside of the network. Don't forget that sensitive data and company secrets could be sent from a source inside the network to the outside world, which is why outbound traffic filtering is also necessary.

Firewalls are an essential component of both host and network security.

Virtual Private Networks

A **virtual private network (VPN)** is a mechanism to establish a remote access connection across an intermediary network, often over the Internet. VPNs allow for cheap, long-distance connections when established over the Internet, as both endpoints only need a local Internet link. The Internet itself serves as a “free” long-distance carrier.

A VPN uses **tunneling** or **encapsulation** protocols. Tunneling protocols encase the original network protocol so that it can traverse the intermediary network. In many cases, the tunneling protocol employs **encryption** so that the original data traverses the intermediary network securely.

You can use VPNs for remote access, remote control, or highly secured communications within an untrusted network.

Proxy Servers

A **proxy server** is a variation of a firewall. A proxy server filters traffic, but it also acts upon that traffic in a few specific ways. First, a proxy server acts as a middleman between the internal client and the external server. Or, in the case of reverse proxy, this relationship can be inverted with an internal server and an external host. Second, a proxy server will hide the identity of the original requester from the server through a process known as **network address translation (NAT)**. Essentially, a proxy server stands as an intermediary in front of another server.

Another feature of a proxy server can be content filtering. This type of filtering focuses either on the address of the server (typically by domain name or IP address) or on keywords appearing in the transmitted content. You can employ this form of filtering to block employee access to Internet resources that are not relevant or beneficial to business tasks or that might have a direct impact on the business network. This could include malicious code, hacker tools, excessive bandwidth consumption (such as video streaming or P2P file exchange), and inappropriate or offensive websites or material.

Proxy servers can also provide caching services. **Caching** is a data storage mechanism that keeps a local copy of content that is fairly static in nature and that numerous internal clients have a pattern of requesting. Front pages of popular websites, like those of search engines, are commonly cached by a proxy server. When a user requests a page from the Internet that is stored in cache, the proxy server provides that page to the user from cache rather than pulling it again from the Internet site. This provides the user with faster performance and reduces the load on the Internet link.

This type of caching often results in all users experiencing faster Internet performance, even when the pages served to them are pulled from the Internet. Obviously, such caching must be tuned to prevent stale pages in the cache. Tuning is setting the time-out value on cached pages so each expires at a reasonable rate. All expired cached pages are replaced by fresh content from the original source server on the Internet.

Network Address Translation

Network address translation (NAT) converts internal addresses into external public addresses, and vice versa. The network performs this conversion on packets entering or leaving the network to mask and modify the internal client's configuration. The primary purpose of NAT is to prevent internal IP and network configuration details from being discovered by external entities, such as hackers.

FIGURE 1-3 shows an example of a pathway from an internal client to an external server across a proxy or firewall using NAT. In this example, the external server is a web server operating on the default HTTP port 80. The web server is using IP address 208.40.235.38. The requesting internal client is using IP address 192.168.12.153. The client randomly selects a source port between 1,024 and 65,535 (such as 13,571). With these source details, the client generates the initial request packet. This is Step 1. This packet is sent over the network toward the external server, where it encounters the NAT service.

Example Pathway From an Internal Client to an External Server




FIGURE 1-3 An example of how NAT functions.

STEP 1 Initial request:

S: 192.168.12.153: 13571

D: 208.40.235.38: 80

STEP 2 NAT creates entry in translation mapping table:

Internal:

External:

192.168.12.153: 13571 < - - >

72.254.149.76: 27409

The NAT service creates an entry in its translation or mapping table for the request. The table contains the source IP address and port, as well as the translated IP address and port. The NAT server uses its own public IP address as the translated IP address and selects a random currently unused port for the new source port for the packet. This is Step 2.

STEP 3 Translated request:

S: 72.254.149.76: 27409

D: 208.40.235.38: 80

Step 3 is the construction of the new packet with the translated source information, which is then transmitted over the Internet to the external server. The web server receives a request with a source that appears to be the proxy/firewall system.

STEP 4 Response from external server:

S: 208.40.235.38: 80

D: 72.254.149.76: 27409

Step 4 is the response generated and transmitted by the web server.

STEP 5 Response sent to client by NAT:

S: 208.40.235.38: 80

D: 192.168.12.153: 13571

The proxy/firewall receives the response to the web server's request. The NAT service uses the translation table to return the original client's information into the packet header as the destination address. This is Step 5. This process takes place at wire speed. It is usually transparent

to the client and server involved in the communication, because they are unaware that the translation took place.

Another purpose or use of NAT is to reduce the need for leasing a significant number of public IP addresses from an Internet service provider (ISP). Without NAT, you would need a unique **public IP address** for each individual system that would ever connect to the Internet. With NAT, you can lease a smaller set of public IP addresses to serve a larger number of internal users. In many small office/home office (SOHO) networks, all internal network devices can be translated to a single public address.

This consolidation is possible for two reasons. First, the shortage of IPv4 addresses required some networks to consolidate addresses. Second, NAT reserves a specific public address for use by a single internal client only if **Static NAT** is utilized. Instead, NAT randomly assigns an available public address to each subsequent internal client request. Once the client's communication session is over, the public address returns to the pool of available addresses for future communications.

Additionally, NAT usually does not employ IP-to-IP address translation, instead using the more granular option known as **port address translation (PAT)**. With PAT, both the port and the IP address of the client are converted into a random external port and public IP address. This allows for multiple simultaneous communications to take place over a single IP address. This process, in turn, allows you to support a greater number of communications from a single client or multiple clients on an even smaller number of leased public IP addresses.

NAT enables the use of **private IP addresses** while still supporting Internet communications. These addresses are known as RFC 1918 addresses, from the technical specification detailing the protocol. The ranges are as follows:

- **Class A**—10.0.0.0–10.255.255.255/8 (1 Class A network)
- **Class B**—172.16.0.0–172.31.255.255/12 (16 Class B networks)
- **Class C**—192.168.0.0–192.168.255.255/16 (256 Class C networks)

RFC 1918 addresses are for use only in private networks. They are, by definition, unrouteable address spaces. Internet routers drop any packet using one of these addresses. Without NAT, a network using these private IP addresses would be unable to communicate with the Internet. By using RFC 1918 addresses, networks create another barrier against Internet-based attacks and do not need to pay for leasing of internal addresses.

As mentioned, NAT is one of the technologies that has allowed the extended use of IPv4, even after the depletion of available public IP addresses. Now that the transition to IPv6 is underway, NAT is serving a new purpose in translating between addressing schemes, or IPv4-to-IPv6 translation. Many firewall and proxy devices may include IPv6 translation services. This feature is a worthwhile option that you might want to include on a feature list when researching a firewall purchase or deployment.

TECHNICAL TIP

An interesting wrinkle on NAT and PAT is that most devices and software products that support PAT actually

state in the documentation and in the configuration interfaces that NAT is also supported. It seems that NAT has been redefined to include the functionality of port address translation, even though an official acronym for that specific additional technology already exists. Most certification exams combine the two as well. Questions might use the NAT acronym but actually refer to port address translation (PAT).

The Domain Name System

The **Domain Name System (DNS)** is an essential element of both Internet and private network resource access. Users do not keep track of the IP address of servers. So, to access a file server on, say, a social networking site, users rely upon DNS to resolve the **fully qualified domain names (FQDNs)** into the associated IP address. DNS is similar to the address book (contacts) in your smartphone. Instead of remembering people's phone numbers, you just look up their names; the address book (contacts) associates each person's name with the number to dial.

Most users don't even realize that networks rely upon IP addresses to direct traffic toward a destination rather than the domain name that they typed into the address field of client software. However, without this essential but often transparent service, the Internet would function very differently. In such a case, users would need to maintain a list of IP addresses of sites they wanted to visit or always go to a search engine for a site's IP address.

DNS is the foundation of most **directory services** in use today, such as Active Directory and LDAP (see the following section). Active Directory is found in Windows operating systems. Thus, DNS is essential for internal networks, as much as it is for the external Internet.

DNS is vulnerable in several ways. First, DNS is a nonauthentication-query-based system. This allows a false or "spoofed" response sent as a DNS query to appear valid. Second, anyone can request transfers of the DNS mapping data (called the zone file), including external entities if TCP port 53 allows inbound

access. Third, DNS uses a plaintext communication allowing for eavesdropping, interception, and modification.

You can address some of the faults of DNS with local static DNS mapping in the **HOSTS file**, filtering DNS on network boundaries, and using **Internet Protocol Security (IPSec)** for all communications between all hosts. (A HOSTS file is a static file on every IP-enabled host where FQDN-to-IP address resolutions can be hard-coded.) However, these are defenses for DNS, as DNS does not protect the network; rather, it is an essential service that needs protection.

■ NOTE

A few Internet index sites still exist, but they are not as exhaustive (or current) as most search engines. Think of the difference between the Yellow Pages dropped on your front porch each year versus dialing 411: One is more current and relevant than the other.

Directory Services

A directory service is a network index. It helps users locate resources within a private network. A directory service is responsible for keeping track of which servers (sometimes acting as clients) are online and the resources these hosts share with the network. A directory service operates much like a telephone book does for phone numbers and addresses.

Prior to directory services, less efficient methods of tracking or locating available resources were in service. These included local static or dynamic lists, which networks maintained using broadcast announcements. Only workgroup networks still use these outdated methods.

A directory service is another essential piece of modern networks that does not provide obvious security services. This means directory services need the protection provided by other security services and devices. You should limit access to directory services to authorized and authenticated clients and users of the local network. Your network should ignore all external requests for information, with the exception of valid (organization identified) remote access or VPN links. If possible, you should deploy IPSec so that you protect every internal network communication.

Intrusion Detection Systems and Intrusion Prevention Systems

Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) have become popular topics of discussion in the IT and security fields. The concept of an IDS is a system that watches internal hosts or networks for symptoms of compromise or intrusion. Effectively, an IDS is a form of burglar alarm that detects when an attack is occurring within the network.

An IDS often serves as a complementary mechanism to firewalls. Once an IDS detects an intruder, it can send commands or requests to the firewall to break a connection, block an IP address, or block a port or protocol. You must configure the firewall to receive these commands and authorize the IDS to send them. Not all IDSs and firewalls are compatible in this manner.

Eventually, the IPS concept appeared, which strives to detect the attempt to attack or intrude to prevent the attack. Once an attempt is detected, the IPS can respond to prevent the success of the attempt rather than waiting until after a successful breach to respond. IPSs do not replace IDSs. Instead, these are often used as a first or initial layer of proactive defense, relegating the IDS to a detection posture against those events that the IPS misses or that internal personnel perform.

IDSs and IPSs are important components of a complete network security solution. However, these are not without fault. IDSs and IPSs can create a false sense of security under certain conditions. Two commonly discussed conditions include unknown zero-day attacks and false positives.

When unknown zero-day attacks threaten the network, an IDS or IPS might not have the mechanism to detect it. Thus, the lack of an

alarm could cause administrators to assume that no attacks are occurring. In most cases, the lack of an alarm does mean that nothing malicious is happening. So, this assumption is not completely unwarranted. But if an IDS never triggers an alarm, then you should suspect a poor detection system rather than the complete lack of compromise attempts or attacks. This is a false negative, as no threat was detected even though a threat was present.

False positives present the same issue for the opposite reason, namely too many alarms from benign (or valid) occurrences. After the first initial alarms turn out to be triggered by benign activity, the urgency of responding to alarms diminishes. After additional false positives, an administrator might put off investigating alarms. Eventually, you ignore the alarms altogether. Once this situation arises, you may even treat alarms for malicious events as false positives, once again establishing a false sense of security. But the administrator should instead “tune” the device to reduce the number of false positives on the network. A situation where false positives exist is worse than a false negative because it is easy to be lulled into compliance, believing there is no issue on the network because there is no warning present.

Network Access Control

Network access control (NAC) limits access or admission to the network based on the security compliance of a host. NAC is essentially an enforcement tool to make sure that all hosts connecting to the network have current and compliant security components. With NAC, you can block or restrict access if a computer does not have the latest antivirus update, a certain security patch, or a host firewall, or if the device is unauthorized.

NAC usually operates by placing an agent on each authorized host. When the host connects to the network, the agent contacts the master control program to find current host requirements. If the host fails to meet those requirements, then the NAC prevents the host from accessing the network. A noncompliant host may be granted access only to remediation servers. These remediation servers can provide the patches and updates needed to bring the host into full compliance.

Effectively, NAC is a method to enforce host-hardening rules through an automated systemwide mechanism. With NAC, systems must be in compliance, or attempts to access general network resources are blocked. Access is allowed to the network only after noncompliant systems are updated.

TCP/IP BASICS

To fully understand the mechanisms of network security, and especially filtering employed by firewalls, you need a solid understanding of the **Transmission Control Protocol/Internet Protocol (TCP/IP)** protocol suite. Thorough knowledge of TCP/IP benefits a security administrator not just in the area of firewall management, but also in routing, switching, maintaining availability, improving network performance, managing network traffic, analyzing protocol, understanding vulnerabilities and exploits, and even performing penetration testing or ethical hacking.

Most networks, including the Internet, use the TCP/IP protocol. The most prevalent version in use is **Internet Protocol version 4 (IPv4)**. However, **Internet Protocol version 6 (IPv6)** is gaining wider use across the globe. During this transitional period, you should learn about both versions of IP.

OSI Reference Model

Most protocol discussions begin with the **Open Systems Interconnection (OSI) Reference Model**. The OSI model is a standard conceptual tool used to discuss protocols and their functions. The OSI model has seven layers (**FIGURE 1-4**). Each layer communicates with its peer layer on the other end of a communication session. Although the OSI model is helpful in understanding protocols, most protocols are not in full compliance with it.




FIGURE 1-4 The process of encapsulation and the names of header and payload sets at each OSI layer.

Each layer of the OSI model has unique responsibilities, functions, and features. The OSI model defines what needs to take place at each layer and leaves the actual process of accomplishing

those tasks to the protocols and, ultimately, to the protocol programmers.

■ NOTE

The OSI model is the documented standard for discussing and describing network protocols. However, TCP/IP is the de facto or practical standard, as it was already in use before the OSI model was developed into actual protocols. Few products can directly support the OSI model or its derived protocols. Instead, most products support TCP/IP, although it is not the official, documented standard from the International Standards Organization.

- **Application Layer (Layer 7)**—The **Application Layer (Layer 7)** enables communications with the host software, including the operating system. The Application Layer is the interface between the host software and the network protocol stack. The sub-protocols of this layer support specific applications or types of data.
- **Presentation Layer (Layer 6)**—The **Presentation Layer (Layer 6)** translates the data received from the host software into a format acceptable to the network. This layer also performs this task in reverse for data going from the network to the host software.
- **Session Layer (Layer 5)**—The **Session Layer (Layer 5)** manages the communication channel, known as a session, between the endpoints of the network communication. A single Transport Layer connection between two systems can support multiple, simultaneous sessions.

- **Transport Layer (Layer 4)**—The **Transport Layer (Layer 4)** formats and handles data transportation. This transportation is independent of and transparent to the application.
- **Network Layer (Layer 3)**—The **Network layer (Layer 3)** handles logical addressing (IP addresses) and routing traffic.
- **Data Link Layer (Layer 2)**—The **Data Link Layer (Layer 2)** manages physical addressing (MAC addresses) and supports the network topology, such as Ethernet.
- **Physical Layer (Layer 1)**—The **Physical Layer (Layer 1)** converts data into transmitted bits over the physical network medium (network wiring or the air in the case of wireless communication).

WHAT YOU SHOULD KNOW ABOUT ADDRESSES

Logical addresses, such as IP addresses, create global distinction. At least for public IP addresses, all addresses used on the Internet are unique. Assignment of a logical address is independent of physical location. Logical addresses enable communications between two hosts, regardless of the physical proximity of the devices.

Physical addresses, such as MAC addresses, create local distinction. MAC addresses must be unique only within the local subnet. This enables a distinction between physically proximate systems that are able to receive the same electronic signal. The manufacturer usually performs physical address assignment of the

NIC and is thus dependent on the physical hardware component itself.

MAC addresses are dependent on the physical NIC; the manufacturer assigns these as a “permanent” hardware address. However, modifying or spoofing the effective MAC address is usually possible on most systems. MAC changes are possible using native commands, as with Linux, Unix, and Mac OS, or with third-party utilities, such as required by Windows. Tools for this include Change MAC Address, SMAC MAC Address Changer, and MAC Spoof.

As data moves from a software application for transmission over the network, it traverses the layers of the protocol stack from top to bottom. As each layer receives data from the layer above it, that data becomes the payload with a layer-specific header. At the Data Link Layer (Layer 2), where Ethernet resides, the data receives a footer as well. This process is known as encapsulation. The inverse, known as **de-encapsulation**, occurs when a network communication is received. As this process takes place, the data set being manipulated receives unique names, depending on the layer it traverses. The steps in the process are:

- **Application Layer (Layer 7)**—Payload data unit (PDU), that is, data obtained from the host software application.
- **Presentation Layer (Layer 6)**—PDU
- **Session Layer (Layer 5)**—PDU
- **Transport Layer (Layer 4)**—Segment

- **Network Layer (Layer 3)**—Packet
- **Data Link Layer (Layer 2)**—Frame
- **Physical Layer (Layer 1)**—Bits of data

The encapsulation process of adding headers (and a footer at the Data Link Layer) enables data exchange between layers on different systems. This is known as peer-to-peer communications. The content of a header includes information to be processed by the corresponding layer on the receiving end of a network link.

The content of the headers, mainly from Layers 2–4, are the greatest concern and focus of a firewall. Application proxy firewalls and stateful inspection firewalls can also examine the headers and the payload content of Layers 5–7.

Sub-Protocols

TCP/IP is not a single protocol, but a collection of protocols. Often referred to as the TCP/IP suite, this collection includes several core protocols, such as **Internet Protocol (IP)**, **Transmission Control Protocol (TCP)**, and **User Datagram Protocol (UDP)**, as well as several commonly used protocols, such as Address Resolution Protocol (ARP), Internet Control Messaging Protocol (ICMP), Hypertext Transfer Protocol (HTTP), and Transport Layer Security (TLS). In addition to these, there are tens of thousands of proposed Request for Comments (RFCs) introducing potential protocols, but there are only a handful of protocols that typically operate within the network infrastructure created by TCP/IP.

A firm knowledge of the sub-protocols of TCP/IP is assumed but not essential to grasping the scope of this text.

Headers and Payloads

Firewalls, specifically packet-filtering firewalls, inspect the contents of headers to allow or deny **frames**, packets, or segments. Depending on the type of filtering and the layer or protocol focus of the filtering, the item examined and the header it comes from can vary.

TECHNICAL TIP

For a more complete list of the additional sub-protocols, along with the corresponding port assignments, view the document hosted by Internet Assigned Numbers Authority (IANA) at <https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>. Nearly every one of the 65,536 ports has one or more protocols and applications associated with it.

For a list of ports used by malicious software, visit <https://attack.mitre.org/techniques/T1065/>.

Packet filtering often focuses on four main headers:

- The Ethernet header of the frame from the Data Link Layer
- The IP header of the packet from the Network Layer
- The TCP header of the segment from the Transport Layer
- The UDP header of the segment from the Transport Layer

Each of these four headers has numerous details that affect filtering action. These include MAC addresses, IP addresses, TCP header flags, port numbers, and more.

Filtering on Addresses

With regard to the basic differences in IPv4 and IPv6 addressing, the issues to focus on are not subnet masks or the length of an address—both of which are important—but rather address length, available address capacity to meet the need for unique addresses, and subnet masking.

Many forms of filtering focus on the IP address and/or port number to make an allow or deny decision. Such decisions or rules can focus on either the source address or the destination address found in the header of a segment, packet, or frame. In other words, port number, IP address, or MAC address can support filtering as a source and/or destination concern.

Firewall filtering can focus on whether an address is public or private in both the source or destination position in the IP packet header. Generally, private addresses function on the private network and do not reach the outside. NAT will translate source-private addresses into a public address if the packet is heading toward an external destination. On the Internet, both routers and firewalls drop any IP packet that has an RFC 1918 address in the header. Firewalls can filter on whether an address is known or unknown. This function filters against addresses used for the source and/or the destination. Known addresses are usually trusted addresses. Only packets with trusted and known addresses are allowed to reach their destination. Unknown addresses are not likely to be trusted. Packets with unknown addresses can be stopped in every case or further inspected before an allow or deny decision is made.

Firewalls can filter on whether an address is known to be benign or commonly sends malicious content, a variation of known and

unknown filtering. Instead of known addresses being trusted and unknowns not being trusted, this method twists the idea. A benign address is a known, trusted address and a malicious address is a known, not trusted address. In both cases, the addresses are known. Any unknown address encountered using this form of filtering will require additional subsequent filtering that looks at other aspects of the traffic.

Firewalls can also filter on whether an address is real or spoofed. This form of filtering is not as clear and distinct as the previous types. Whether a segment, packet, or frame's address is the real correct address or a spoofed falsified address can be difficult to determine in every situation. When an address is a real address, in the sense that it's within the subnet ranges of a network, the methods for determining whether it's a spoofed address are few.

One method is to compare the address against a use table, such as that maintained by the Dynamic Host Configuration Protocol (DHCP) server, to see if the address is not currently assigned to any authorized system. Another method is to check the route, communication path, receipt vector, or receiving interface of the source address against what that aspect should be for a given address. For example, if a source address typically arrives on port 4 but it shows up on port 7, it is likely to be a spoofed source address.

Spoofed addresses can also be detected at border sentry points. When a source address comes from the opposite side of the firewall, then it is obviously a spoofed address. One clear example is when an internal LAN address appears as a source address in a packet on its way into a network from outside. This form of spoof filtering can be part of ingress filtering.

Likewise, the same process can be used for packets leaving a network. If a packet's source address is from the outside, such as an Internet address, but the packet is received by a firewall from an interface in the private LAN, it must be from a spoofed address. This form of spoof filtering can be part of egress filtering.

CHAPTER SUMMARY

Network security is complex, which is why we began with the fundamentals of network security. Once you have a firm grasp of these basic issues, you will be able to put security into practice on your own network. Network infrastructure design can have a significant impact on security.

A written security policy is the foundation of a successful security endeavor. Without a written policy, security will be chaotic and uncontrolled. A security policy defines and assigns roles and responsibilities to personnel within the organization. The seven domains of a typical IT infrastructure aid in identifying how users interact with each other and the resources on a network. These domains include the User Domain, Workstation Domain, LAN Domain, LAN-to-WAN Domain, Remote Access Domain, WAN Domain, and System/Application Domain. Administrators can employ numerous network components and devices to support a network security policy. These include firewalls, VPNs, and IDSs/IPSSs.

KEY CONCEPTS AND TERMS

Access control

Application Layer (Layer 7)

Asset

Auditors

Authentication

Authorization

Availability

Bottlenecks

Business tasks

Caching

Confidentiality

Countermeasure

Data Link Layer (Layer 2)

De-encapsulation

Defense in depth

Directory services

Domains

Domain Name System (DNS)

Downtime

Egress filtering

Encapsulation

Encryption

Exploits

Filtering

Firewall

Frames

Fully qualified domain names (FQDNs)

Hardening

Hardware address

Host

HOSTS file

Ingress filtering

Integrity

Internet Protocol (IP)

Internet Protocol Security (IPSec)

Internet Protocol version 4 (IPv4)

Internet Protocol version 6 (IPv6)

Intrusion detection systems (IDSs)

Intrusion prevention systems (IPSs)

Local area network (LAN)

Logical addresses

Malicious code

Media Access Control (MAC) address

Network access control (NAC)

Network address translation (NAT)

Network Layer (Layer 3)

Network security

Node

Nonrepudiation

Open Systems Interconnection (OSI) Reference Model

Permission

Physical addresses

Physical Layer (Layer 1)

Port address translation (PAT)

Presentation Layer (Layer 6)

Privacy

Private IP addresses

Privilege

Proxy server

Public IP address

Resources

RFC 1918 addresses

Risk

Roles

Security goals

Security objectives

Security policy

Senior management

Server

Session Layer (Layer 5)

Static NAT

Telco

Threats

Transmission Control Protocol (TCP)

Transmission Control Protocol/Internet Protocol (TCP/IP)

Transport Layer (Layer 4)

Trust

Tunneling

User Datagram Protocol (UDP)

Vulnerability

Virtual private network (VPN)

Whitelists

Wide area network (WAN)

Zero-day exploits


Zeroization

CHAPTER 1 ASSESSMENT

1. Which of the following best describes the principle of least privilege?
 - A. Allow the user to have access to only the items on their payment plan
 - B. Allow the user to access items only during their normal scheduled work hours
 - C. Develop a patch only after the exploit has been discovered and publicized
 - D. Allow the user access to only what is essential for their job responsibilities
2. Which of the following best describes the concept of risk?
 - A. The unsecured content on a network that may be used to gain access to the network
 - B. The surface that is exposed when a network is live that hackers can use for access
 - C. The solution that incorporates the best solution for a particular network
 - D. The likelihood that a threat will take advantage of vulnerability on the network
3. What is the most important characteristic of an effective security goal?
 - A. It is inexpensive.
 - B. It is possible with currently deployed technologies.
 - C. It is written down.
 - D. It is approved by all personnel.
4. Which of the following is true of network security?
 - A. Use of a framework ensures success in securing a network.

- B. Network security includes elements preventing unwanted access and action.
 - C. It is easy to apply the latest fixes to any network setup.
 - D. Network security can often be ensured with no monetary outlay.
5. What is the primary security concern with wireless connections?
- A. Redundancy
 - B. Signal range
 - C. Support for IPv6
 - D. Connection speed
6. What does Van Eck phreaking allow?
- A. Attackers to encrypt the hard drive of a device until a ransom is paid
 - B. Submission of enough packets that a denial of service attack persists
 - C. Attackers to eavesdrop on electronic devices from a device
 - D. Limitation of access to any physical device on the network
7. Which of the following best describes the concept of hardening?
- A. The removal of items from the cloud to physical storage locations
 - B. The allowance of users to access networks through their own devices
 - C. The process of securing or locking down a host against threats and attacks
 - D. The removal of unnecessary software from workgroup computers and devices
8. Which of the following best describes network availability?
- A. Protection against unauthorized changes while allowing authorized changes
 - B. Protection against unnecessary costs while allowing reasonable expenditures

- C. Protection against downtime while supporting authorized access to resources
 - D. Protection against unauthorized access while providing authorized access
9. For what type of threat are there no current defenses?
- A. Zero-day
 - B. Malicious code
 - C. Unauthorized software
 - D. Hardware failure
10. Which of the following best describes nonrepudiation?
- A. It prevents a user from being able to deny having performed an action.
 - B. It controls what users are allowed to do and not allowed to do.
 - C. It is proof of a user's identity prior to granting access to a secured area.
 - D. It ensures that malicious code is not rampant on a corporate network.



© Digital_Art/Shutterstock

Network Security Threats

NETWORK SECURITY IS UNDER CONSTANT ATTACK by threats both internal and external, ranging from disgruntled employees to worldwide hackers. There is no perfect defense, because hackers are able to bypass, compromise, or evade almost every safeguard, countermeasure, and security control. Hackers are constantly developing new techniques of attack, writing new exploits, and discovering new vulnerabilities. Network security is a task of constant vigilance, not a project to complete. It is a job that is never done.

Why is it critically important to understand hacking, exploitation, vulnerabilities, and attacks? As the sixth century BC Chinese military strategist and philosopher Sun Tzu stated in his famous military text *The Art of War*: “If you know the enemy and know yourself, you need not fear the results of a hundred battles.” Once you understand how hackers think, the tools they use, their exploits, and their attack techniques, you can then create effective defenses to protect against your network. Understanding hacking not only improves network security; it also maintains security at a high level of readiness.

CHAPTER 2 TOPICS

This chapter covers the following topics and concepts:

- Types of hackers and their motivations
- Frequently targeted network assets
- Internal and external network threats
- Hacking process
- Common IT infrastructure threats
- Various types of malicious code (malware) and the corresponding security concerns
- Dangers of fast growth and overuse
- Dangers posed by wireless and wired connections
- Risk of eavesdropping
- Different attack types, including hijack and replay attacks, insertion attacks, fragmentation attacks, and buffer overflows
- Preventing session hijacking, spoofing attacks, and man-in-the-middle attacks
- Covert channels and associated attacks
- Threats to network and resource availability, including denial of service (DoS) and distributed denial of service (DDoS) attacks
- Risks of social engineering

CHAPTER 2 GOALS

When you complete this chapter, you will be able to:

- Describe the motivations of hackers and other malicious computer network intruders
- Compare and contrast threats from internal and external sources
- Describe how accidents, natural disasters, and ignorance affect network security
- Explain the risk posed by malicious code
- Identify the effects of wired and wireless connectivity on network security
- Describe common network security exploits and attacks, including hijack attacks, replay attacks, insertion attacks, fragmentation attacks, buffer overflows, XSS attacks, man-in-the-middle attacks, spoofing attacks, covert channels, DoS, DDoS, botnet attacks, and social engineering attacks
- Describe how hacker tools exploit vulnerable targets

HACKERS AND THEIR MOTIVATION

Hacking originally meant tinkering with or modifying systems to learn and explore. However, the term now refers to malicious—and possibly criminal—intrusion into and manipulation of computers. In either case, a malicious or criminal hacking is a serious threat. Every network administrator should be concerned about hacking.

Hackers are people—often misguided, unethical, ingenious, and criminal—but still people. When you design security, keep in mind that the threat is ultimately and mostly human. This awareness should lead to better selections of deterrent, detection, and response tools and techniques. A **deterrent** is any tool or technique that makes hacking your network less attractive than hacking another network. This can be done by enhancing security, which may increase the time, effort, and investment required of the hacker to gain access. Locating the network in an area with strict laws and punishments for illegal or unauthorized access is another example of a deterrent.

Some important aspects of security stem from understanding the techniques, methods, and motivations of hackers. Once you learn to think like a hacker, you may be able to anticipate future attacks. This enables you to devise new defenses before a hacker can successfully breach your organization's network.

So, how do hackers think? Hackers think along the lines of manipulation or change. They look into the rules to create new ways of bending, breaking, or changing them. Many successful security

breaches have been little more than slight variations or violations of network communication rules.

Hackers look for easy targets or overlooked vulnerabilities. Hackers turn things over, inside out, and in the wrong direction. Hackers attempt to perform tasks in different orders, with incorrect values, outside the expected boundaries, and with a purpose to cause a reaction. Hackers learn from and exploit mistakes, especially mistakes of the network security professionals who fail to properly protect an organization's assets.

What motivates hackers to attack computer networks? Why does anyone get involved in illicit activity? The motives are as numerous as the number of ways to conduct the computer attacks. Some do it for the sheer thrill of hacking—the sport of it. Some hackers consider hacking their hobby. Some love a challenge. Some are victims of peer pressure or are seeking social validation. For others, hacking is a way to earn “street cred” with peers or attack a perceived social injustice. Finally, many hackers pursue their exploits for power and financial gain. As it can be difficult to understand the motivation of a hacker, this section offers more detail to give insight into the hacker mentality.

Many criminal hackers are in it simply for the money. Hackers have many ways to gain financially from attacking computer networks, ranging from theft of credit cards and financial statements to blackmail, black market sales, and trafficking. Some monetary gain is immediate; for example, hackers might be able to transfer funds out of a target account. Other methods are more involved, such as stealing corporate documents and selling them to competitors, or hijacking data using encryption and holding it for ransom. Hackers

can become involved in selling their services, including distributing spam, eavesdropping, cracking passwords, and generating DoS events.

FYI

Spam is unwanted and unrequested email. It is often sent in large quantities to an email address in the hopes of causing system delays or denial of service (DoS) events. Spam is not technically malicious software, but spam can have a serious negative effect on IT infrastructures. Experts estimate that 80 to 95 percent of email Internet traffic is spam and other forms of malicious messages. Hackers can easily use DoS attacks through **flooding** a network with traffic or ping requests. As a result, the network servers cannot keep up with the immense amount of traffic; it often renders a website or network access inaccessible.

A more recent type of attack for financial gain is **ransomware**. Historically, a hacker gains access to a system, identifies assets, removes valuable information, and possibly remains resident on the network or creates a backdoor to allow reentry into the network. In a ransomware attack, however, the hacker typically tricks users into clicking a malicious link or email attachment, which encrypts files, folders, or entire drives, preventing users from viewing, copying, or accessing them. The hacker informs the victim how access can be returned, such as with a cryptocurrency payment like Bitcoin.

Payment of the ransom often results in return of access for the victim, but it is not a guarantee, nor is it protection against reencryption.

FYI

Cryptocurrency is electronic currency where the existence of the currency is a mathematical formula stored on the systems of the participants with a rate that fluctuates. The hacker lists the demand as “4 Bitcoin” or some other amount and the victim purchases the cryptocurrency at the going rate. Some ransoms have been paid through donations to charitable organizations if the hacker wanted to punish the victim or right a perceived wrong.

For many hackers, it is an exciting challenge to find new vulnerabilities, develop new attack code, or discover new security breaches. A **breach** is any successful attempt to get past a network’s defenses, such as “breaching” the perimeter. Some hackers attack targets for the same reasons some people climb mountains—because it’s there, because they can, and because they want to beat it. Some hackers continually seek out new and more difficult targets to improve their skills and to increase the level of the challenge.

Demonstrating the ability to compromise a target, especially if the target is reasonably secure, is a way for hackers to prove they are more powerful than the defenders. If you can control something, you have power over it. If a hacker can control your computer remotely, then the hacker has power over your computer; in some ways, the

hacker has power over you, as well. Hackers sometimes hack to boost their image and feed their ego. If they can wage a successful attack, then they are showing dominance over their target.

Hacking can be thrilling. Think of the way a treasure hunter feels on someone's private property. The combination of power, challenge, risk, and potential pursuit, along with discovering potentially valuable assets, is a thrill for some hackers. The thrill of getting away with it motivates many hackers to continue their attacks once they figure out how to successfully hack into a network.

Some hackers enjoy the sheer risk of attacking a network. As with any crime, the risk of being caught is always a possibility. A target network might have **honeypots**, intrusion detection systems (IDSs), intrusion prevention systems (IPSs), firewalls, and other technical defenses that the hacker will need to detect and evade. Security professionals are always looking to discover new attack techniques and to learn the identity of hackers. With every attack, therefore, the hacker is at risk of getting caught and prosecuted.

Some hackers resort to hacking as an attempt to be entertained or distracted from their boredom. A successful system breach will usually initiate a response, often a cat-and-mouse game that is entertaining for hackers. With a ransomware attack there is also a power position: the victim must comply with the hacker's demand, or access is not returned.

Hackers have peers and social groups just like everyone else. Peer pressure can be a powerful motivator to fit in, show off, or demonstrate loyalty to a group, especially for those in the lower levels or rankings of hacker communities. New or inexperienced hackers feel encouraged to perform attacks to maintain their

membership in their peer group. Some attacks are performed as an initiation or rite of passage from a potential hacking group member to an actual full member. Hacking can help the hacker achieve or maintain status. Some hackers try to cause a media story and to get their name in the news. Other hacks are performed to instill fear and compliance in current and future targets.

Hackers who believe that a company is not a good local or global citizen may undertake **hacktivism**, which is hacking as a form of righting a social injustice or to exact revenge for perceived wrongs committed by the company or specific individuals. This type of hacker may steal but often donates the financial gain to a cause deemed worthy or one that works to correct the wrong that initiated the hack. This hacker is more difficult to catch, as the funds are usually directed immediately to a third-party nonprofit, which does not employ or sponsor the hacker. There is no “follow the money” opportunity to capture the individual, as there is when the hacker keeps the proceeds of their hacking.

Three primary types of hackers are recreational, opportunistic, and professional:

- **Recreational hackers** are those who enjoy learning and exploring, especially with computing technology. However, they might make poor choices as to when to use their newfound skills. Bringing in unapproved software from home, experimenting on the company network, or just trying out an exploit to see if it works are all potential actions of recreational hackers. Some hackers might not fully consider that their hobby can be dangerous and their actions may be in violation of the company security policy. For example, a **script kiddie** is a hacker who relies on preexisting programs or code written by others. This type of hacker runs scripts by simply pointing and clicking a tool and may not fully understand the ramifications of these actions.

- **Opportunistic hackers** are timid and not likely to initiate an attack. For whatever reason, they are unwilling to purposefully plan out and wage intrusions. However, if the circumstance presents itself for an attack that can be easily performed with little potential for discovery or consequence, the opportunistic hacker may take advantage of the fleeting opportunity. That moment could arrive when someone happens to work late and is the last employee left in the building. The opportunity could be when a fire drill drives everyone else out of the facility, or when a random power outage occurs and half the workforce leaves for home. It could even happen when they notice a certain office door is left open and no one else is watching.
- **Professional hackers** are criminals whose sole career objective is to compromise IT infrastructures. Whether operating as individuals, offering mercenary hacking services, functioning as a member of a criminal enterprise, or operating in a state-sponsored attack, professional hackers focus their time and energy on conducting the best security assault possible. When people spend years learning and practicing in one primary area of interest, they can develop expertise and skills to rival all defenses. The perfect unbreachable security solution does not exist. Professional hackers have the time, stamina, skill, patience, and backing to keep up an assault against a target until they succeed. They are to a network what termites are to a wooden building. You can deter them; you can keep them out most of the time. But they will always be nearby and eager to gnaw into the foundations if you drop your guard for a moment.

Hackers perform hacks for many reasons, and this list is likely incomplete. Whatever their motivations, they are driven to do unethical activities. You might not completely understand the hacker mindset, but it is important for you to try so you can better anticipate how hackers could attack your network.

FAVORITE TARGETS OF HACKERS

In terms of security, the things you want to protect are known as assets. An asset is anything used to conduct business over a computer network. Any object, computer, program, piece of data, or other logical or physical component employees need to complete a task is an asset.

What valuable assets do attackers most frequently target? Every hacker is different, just as with any individual member of any subculture. While many agree on what is valuable, such as money and access, others might find corporate finances, secret formulas, medical history, credit reports, court records, accounting logs, and so on preferable as their target of choice. For example, why do some people collect toys while others collect cars, soda bottles, or pictures of dogs playing poker?

Among a hacker's favorite targets are easy assets—those that pay off quickly. Easy targets are IT infrastructures and elements not properly secured. Systems with well-known, exploitable holes, known as vulnerabilities, exposed to the world are “easy pickin's” for hackers. An item's cost does not classify it as an asset; instead, it is the usefulness and value to the company network. Some of the most recent ransomware attacks have been against healthcare systems and providers, including hospitals, doctor's offices, and insurance companies. The medical providers have an urgent need to regain

access, and the ransom is often low enough that it is an inconvenience rather than a hardship to pay.

Targets that pay off quickly are those that earn the hacker some form of monetary or barter gain, such as funds or cryptocurrency. Credit cards and bank accounts can be quickly turned into the hacker's desired form of currency. In other cases, gaining access and control of networks, especially those with high-speed Internet links, is valuable. Such a network could be loaded with malicious code that transmits spam, and it can be conducting eavesdropping. Using the hacked network, a hacker can use the servers and processing power to attempt to break encryption. The hacker can then trade or sell these services to others.

Some hackers, however, do not seek the easy targets; instead, they look for unique targets, new challenges, and complex infrastructures to test and improve their skills. Expert and highly experienced hackers often want to continue to improve their abilities rather than waste their time on targets that any amateur hacker, or script kiddie, could compromise.

Assets do not have to be expensive, complicated, or large. In fact, many assets are relatively inexpensive, commonplace, and variable in size. For most organizations, including SOHO (small office/home office) environments, the assets of most concern include business and personal data. If this information is lost, damaged, or stolen, serious complications can result. Businesses can fail, individuals can lose money, identities can be stolen, government regulators can apply penalties or criminal charges, and lives can be ruined.

Valuable resources abound on individual computers, as well as on IT infrastructures comprised of interconnected LANs. Hackers seek

out targets based on a variety of goals and motivations, as well as perceived value of a resource. But who are the hackers?

THREATS FROM INTERNAL PERSONNEL AND EXTERNAL ENTITIES

While historically two-thirds or more of security violations have been the direct result of outsiders, this is changing dramatically. As a result of more automated breaches, more hacking, and increased criminal and terrorist activity, the percentage of “inside” computer crime has been reduced as low as single digits, which still translates to small and medium-sized businesses and small government agencies experiencing thousands or more attacks daily. This has markedly increased the need for trained computer security professionals in organizations of all sizes. According to the Verizon *2019 Data Breach Investigations Report*

(<https://enterprise.verizon.com/resources/executivebriefs/2019-dbir-executive-brief.pdf>), 69 percent of data breaches stemmed from external agents. Internal breaches were responsible in 34 percent of the events categorized.

An on-site employee has physical access to the facility and logical access to the network. An employee may have only a standard or normal user account, but some level of logical access can be parlayed into greater levels of access (known as privilege escalation). Threats from insiders, whether physical or logical, are serious. You must address such threats in your security policy, network design, infrastructure deployment, and ongoing system and security management. The people who represent the most common threats to

an organization's network security include disgruntled employees and contract workers (insider threats), as well as a host of external hackers. **FIGURE 2-1** depicts internal and external hackers.




FIGURE 2-1 Internal and external hackers.

Disgruntled employees believe that they have been wronged somehow by the organization. Whether this “wrong” is real or perceived, their actions can cause severe disruption of mission-critical operations. Disgruntled employees may attempt to embezzle, steal supplies, waste time, infect a system or network with malicious code, leak confidential data, interrupt other workers, or derail projects.

Contract workers are outsiders brought in to an organization to perform work on a temporary basis. Contract workers can be consultants, temporary workers, seasonal workers, contractors, even day laborers. Usually, outsiders do not share the same loyalty to the organization that most full-time employees exhibit. Thus, if the opportunity affords itself to compromise the organization for personal gain, contract workers are more likely to act unethically. Such criminal outsiders don’t worry about long-term stability or viability, or the fate of regular employees. Instead, they take advantage of these ways to personally profit at the expense of others.

Employees and external entities represent differing levels of risk that should be addressed in planning network security. Initially, the risk from an external attacker is low. Outside attackers must seek out and discover methods of gaining logical or physical access to the network or the facility. Without access, hackers must resort to attacking external interfaces such as web servers, VPN devices, firewalls, and remote access servers (RASs). If these are secure, then the default or fallback position of a hacker is to wage a DoS attack, attempt a burglary, or perform social engineering against insiders.

Another threat to network security is intentional damage or sabotage. Disgruntled employees, dismissed contract workers,

opportunistic janitorial staff, unhappy managers, and even careless visitors can wreak havoc in moments if they have access to sensitive equipment or information. Proper pre-employment applicant screening, perceptive supervisory oversight, escorts for nonemployees, background checks, security cameras, training, paying attention to the corporate culture, and providing competitive pay scales or other incentives and courtesies to employees can go a long way toward preventing intentional disruption or destruction of IT equipment and resources.

THE HACKING PROCESS

As much as cybersecurity professionals dislike the practice, hacking can be a fascinating process. Hackers' activities often appear chaotic and random, at least when observed from the mainstream IT industry. Hackers do not have to follow any fixed procedures or recognize any established boundaries. Instead, they seek vulnerabilities on a selected target, using any and all means at their disposal. For them, chaos is both a methodology and a defense mechanism.

Generally, hacking falls into five main subgroups of events or activities referred to as phases. This categorization can represent hacking, but it does not actually control or prevent it. The five categories are reconnaissance, scanning, enumeration, attacking, and post-attack activities. (See **FIGURE 2-2.**) This order of phases occurs if an attack is successful. If an attack is not successful, the hacker can instead attempt a fallback position.




FIGURE 2-2 Five phases of hacking.

FALLBACK ATTACKS

Fallback attacks are the other options for mayhem a hacker can deploy after unsuccessful breach attempts against a target. Common alternatives to intrusion include DoS, eavesdropping, breaking and entering, social engineering, malicious code, session hijacking, man-in-the-middle attacks, wireless hacking, SQL injection, website attacks, and more.

Reconnaissance

Reconnaissance is the initiation of the process of hacking.

Reconnaissance means the act of inspecting or exploring and can also be called **footprinting**, discovery, research, and information gathering. This phase is the first of three pre-attack phases in which hackers learn as much as possible about a target before attempting the first actual attacks. Reconnaissance consists of collecting data about the target from all possible sources online and offline. The hacker is careful to avoid alerting the target that it has probed the network for information.

Reconnaissance can include:

- Researching old versions of a target organization's website at *archive.org*
- Examining search engine contents
- Reviewing the organization's website
- Investigating the background of personnel
- Performing location mapping
- Reading job postings
- Checking insider information leak sites
- Looking at newspaper and magazine articles for mentions of the organization
- Perusing press releases
- Searching newsgroups, chat archives, blogs, and forums
- Auditing financial records or reviewing public filings
- Reviewing court cases and other public records

- Querying **whois, domain registrations**, and public IP assignments
- Eavesdropping on email and other conversations
- Visiting a physical location for the organization

These are just some of the possible reconnaissance activities a hacker can perform. Information gathering is limited only by the time, resources, and imagination of the hacker. Once the hacker has built and organized a reasonable portfolio of information about the target, the next step is scanning.

Scanning

Scanning is the activity of using various tools to confirm information learned during reconnaissance and to discover new details. Scanning is aimed at discovering live and active systems. Scanning can include **wardialing**, **wardriving**, **ping sweeps**, and **port scanning**.

Wardialing is an older tactic that uses the telephone system to locate any active and answering modems. Using a modem and often a dialing program, a hacker's computer automatically dials target phone numbers. The hacker obtains phone numbers during reconnaissance or by dialing all the numbers in an area code or prefix.

Wardriving, a term derived from *wardialing*, is the technique of using a wireless detector (sniffer) to locate wireless networks. Hackers used to drive through cities or neighborhoods with laptop computers to discover wireless networks. Today, driving is not necessary and smartphones detect wireless networks automatically.

Ping sweeps are used to discover systems over network connections that will respond to Internet Control Messaging Protocol (ICMP) echo requests. Hackers commonly use ICMP for network health and testing. The ping command typed from the command line of a computer or network mapping utility sends ICMP echo requests to all possible recipients within an Internet Protocol (IP) address range or subnet. ICMP echo responses from systems indicate their IP addresses and that they are up and running.

Hackers perform port scanning by sending various constructions of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets to ports (**FIGURE 2-3**). If the hacker does not know the computer system exists, then a port scan can determine the

existence of a system at a specific IP address, as well as whether a port is open, closed, or filtered. A TCP port is known to be open if a full TCP three-way handshake can establish a virtual circuit. A UDP port cannot be confirmed as open because the default response from an open port is always silence.

TCP Scan



SYN Request Packet




SYN/ACK Response Packet




Open Port

ACK Response Packet



SYN Request Packet



RST Response Packet



Closed Port

Datagram



Open Port

No Response


UDP Scan



Datagram



ICMP Type 3 Response



Closed Port

FIGURE 2-3 Basic TCP and UDP port scanning.

TECHNICAL TIP

Ports exist at the Transport Layer (Layer 4) of the Open Systems Interconnection (OSI) Reference Model. TCP and UDP use ports to support multiple simultaneous communications, connections, or sessions over a single Layer 3 (Network Layer) IP address. There are 65,535 ports, but most systems can only support a few hundred concurrent transactions.

A port is open if an active service is ready to process data through the specific port. A port is closed if no service is associated with that specific port.

When communications elicit error messages or abnormal responses from a port, a firewall can filter out these responses. This will result in a port being visible for valid and normal communications. A firewall can block any attempt to elicit errors or abnormal responses.

During scanning, often referred to as probing, a hacker sends packets to elicit responses. When a hacker performs scanning, it is detectable. Reconnaissance is generally silent, secretive, and unobtrusive. You are also unlikely to detect scanning to verify individual data items, such as a single open or closed port. But when hackers scan to discover all possible IP addresses in use and all possible open and closed ports, it is noticeable to network administrators and usually captured in system logs. Reconnaissance

activity might or might not draw attention, but scanning—if not conducted carefully and on single ports, as opposed to all ports at once in a port sweep or blast—will be hard to ignore.

Hackers perform scanning until they discover one or more targets. Because scanning uncovers only a system and the open ports, hackers learn very little about the targets. Hackers stop scanning and move to the enumeration phase, the next step toward attacking, whenever they have enough information or their scans fail to reveal new information, based on the purpose of their attacks. Keep in mind that an attacker needs only a single vulnerability to gain access to a system or network. Once hackers can access a machine, moving on toward attacking it quickly is the most common escalation.

Enumeration

Enumeration is the discovery and listing of potential attack targets. Enumeration is the hackers' process of discovering sufficient details about a potential target to learn whether a vulnerability exists that they can successfully attack. Enumeration often starts with operating system (OS) identification, followed by application identification, then extraction of information from discovered services.

Hackers perform OS identification by probing an open and closed port of a target. The responses from these ports identify the OS. This identification is possible because of the idiosyncrasies of different programmers writing interoperable code. Each OS uses a different group of programmers to write the network protocol stack. Even though the resultant protocol stack may be in compliance with IEEE standards, the defaults and reactions of the stack often differ from one OS to the next. These differences are known and maintained in a small database, which is coded into most network scanning and probing tools, such as **nmap**.

Each open port has a service running behind it. **Banner grabbing** is the activity of probing those services to obtain information (**FIGURE 2-4**). Once a connection has begun, a service may send an announcement of connection or communication confirmation. This announcement is called the banner. The **banner** may contain additional details, such as the product name and version number of the service.




FIGURE 2-4 A banner grabbed from a web server.
Courtesy of Zenmap.

The information in the banner can be used by a hacker for reconnaissance purposes—to gain the intelligence needed to form the basis of an attack. One low-cost method of protection is to change the banner to mislead hackers into using the wrong exploits against the server, but still allow the service on the given port to perform its duties. In some cases, banners are not actually used, so these may be greatly shortened; in other cases, advanced features that are not used may be abandoned in a trade-off that results in better security.

Once hackers have identified the service, they can request additional information. The information request may be ignored and is secure, as no response is sent. But if insecure, the service may return volumes of data to the hacker. Depending upon the service and the queries the hacker performs, the extracted information could include system name, network name, user names, group names, share names, security settings, resources available, access control settings, and more.

Enumeration provides the hacker with identified, potential attack points. After reviewing vulnerability databases, such as **MITRE** (<http://cve.mitre.org>) or **National Institute of Standards and Technology (NIST)** (<https://nvd.nist.gov>), hackers evaluate the potential vulnerabilities. Once the hacker selects an attack target, they collect exploit tools and wage the attack.

One aspect that characterizes modern attacks is the amount of time that passes between the reconnaissance phase and the attack phase. In the past, a matter of minutes may have separated the two, but the more modern attack might take weeks, months, or longer to gather information before the attack is unleashed. Such an attack, the

highly targeted advanced persistent threat (APT), will be discussed in more depth later in this chapter.

Attacking

Attacking is the fourth phase of hacking. Although this seems to be the phase that attracts most of the hype about hackers, in fact, it is the briefest phase of the overall hacking process. A successful attack, based on solid research and preparation, can take just seconds.

If an initial attack fails, hackers can modify their exploits, tune their payloads, adjust their **shell code**, reset their vectors, and relaunch the attack. Once hackers figure out that an assumed vulnerability does not exist or has been secured, they return to their enumeration results to select a new point of assault. Again, think of termites: If they can't get into the house structure through a doorjamb, they will just as eagerly try to enter through a windowsill. Both pests are relentless.

Repeated attacks will lead either to an eventual successful breach or to the frustration of a hacker due to successful defense by the target. If successful, the attacker moves on to post-attack activities. If unsuccessful, the attacker can elect to move to alternative fallback attacks.

Post-Attack Activities

In a successful attack, the hacker usually has breached the target's security to gain some level of logical access. This could be the credentials of a standard user account or access to a **command shell** accessed through a buffer overflow exploit. In any case, some common post-attack activities usually take place. These include **privilege escalation**, depositing additional hacker tools, pilfering data, creating a re-entry point (backdoor) and removing evidence of the hack.

Privilege escalation is attempting to gain higher levels of access or privilege from the target. This can occur using a keystroke logger, known OS exploits to steal administrator or system access, manipulation of scheduled tasks, social engineering, Trojan horses, remote control programs, and other mechanisms. The result is that the hacker gains access to a user account or a command shell that operates as an administrator, root, or the system itself. With privileged access, the remaining post-attack activities are much easier.

Depositing additional hacker tools gives the hacker more power over the compromised system. Tools may enable additional abilities unavailable through the current connection method. Tools may assist in pilfering data, removing evidence of the hack, or maintaining or regaining access in the future. To regain access, hackers often create new accounts on the system. Scanning for new accounts is an easy way to identify this technique.

Pilfering data is just that: scouring storage devices looking for files of interest. Hackers look for anything they can turn into cash quickly. They also look for things that would be fun, interesting, or

damaging to disclose to the public. They look for items they can use to blackmail or coerce users. And, of course, they are on the lookout for potentially valuable information for bartering or trading with other hackers or criminals. Dumping the user account database and password hashes is often a priority as well. Cracking the passwords of other users will help in accessing the system in the future.

Removing evidence of the compromise and subsequent activities is an important step for the hackers. Failing to cover their tracks could lead to apprehension and prosecution. Allowing the IT and security staff to discover the intrusion will lead only to heightened levels of security. Discovery makes future returns more difficult, if not impossible, depending on the hacker's skill set.

Once the hacker performs evidence cleanup, the attacker can claim to have owned (or **pwned**, in **leetspeak**, hackers' code language) a system. The hacker has demonstrated skills through the discovery, tracking, and penetration of a target. This is often the goal of a hacker: to successfully penetrate a target. However, such successes are not always easy, or common. If the attacks fail, the hacker may have fallback attack options.

COMMON IT INFRASTRUCTURE THREATS

The more you understand about the various threats and risks to network security, the more defenses you can mount against these attacks. Threats to network security include hacker exploits, as well as Mother Nature, device failures, and users performing normal business activities.

Hardware Failures and Other Physical Threats

Network security is not just about protection against hacking. Computer systems and networks face many other threats on an ongoing basis. Computer equipment is complex and sometimes fragile. Hardware failures are the most common cause of unexpected downtime. Most equipment is expected to operate well beyond its expected lifetime in normal environments. Some forms of technology, however, are more prone to failure than others.

One of the most commonly discussed causes of unexpected downtime is hard drive failure. A hard drive is one of the few common computer components that has moving parts. While optical drives, tape drives, mice, and keyboards have moving parts, these devices seem to outlast hard drives by a significant margin. Hard drive failure can occur unexpectedly or with reasonable warning. The warning is usually a grinding, whining, or clicking noise coming from the drive as it begins to fail. These noises are clear signs that the end of the device's usable life is near and it should be replaced immediately.

The best defense against hard drive failure, as well as hardware failure in general, is to be prepared. Being prepared includes consistent periodic backups, preferably using a redundant array of independent disks (RAID); performing general cleaning and maintenance; and having spare parts on hand for the inevitable. Another method to avoid downtime and a loss of availability is to replace equipment before it fails. Most devices have a mean time to failure (MTTF) or mean time between failures (MTBF) that can determine the statistical likelihood of a failure. It is a good practice to

replace the device before that period expires. While you will be replacing some devices long before an actual failure, this technique keeps the statistics on your side. Though planned downtime is costly, it is less costly than unplanned downtime. **Solid-state drives (SDDs)** eliminate many of the problems associated with hard disk drives but at a greater initial cost for the hardware.

Another physical threat is heat. Too much heat damages computer equipment. Systems that experience severe temperature cycles, such as very hot to very cold, can have incidents of chip creep or warping and cracking of materials. **Chip creep** is caused by the expansion and contraction of metal because of temperature changes. Severe temperature cycling can even break soldered connections. The traditional debate about turning computers off at night or over weekends or leaving the equipment in a running state speaks to the effects of these temperature changes. A running computer is at its optimum functioning temperature.

Static electricity discharge (SED) from dry conditions can destroy most circuits. Frayed wires caused by rubbing against sharp metal edges or rough surfaces, such as a concrete block wall, can cause a short circuit. Moisture, due to high humidity in the air or liquid spills, is always bad for electrical devices. Excessive vibration can be damaging to computer equipment. Vibrations—caused by nearby heavy construction or regular passing of trains, subways, and airplanes, or even forklifts and large equipment on a shop or manufacturing floor—can cause damage over time. Any obvious physical damage as a result of dropping the device, knocking it over, dropping heavy objects on the device, and so on can result in broken equipment. You can eliminate or significantly reduce most physical

risks and threats with reasonable precautions and common sense in proper handling, care, and storage of electronic equipment.

One such risk is **intentional electromagnetic interference (IEMI)**, in which an intentional discharge is made that damages or destroys electronic equipment from cell phones and video surveillance to computers and servers. IEMI discharges have been recorded up to two miles away and pose no risk of damage to living creatures. Although IEMI and related technologies have been used by the military for years, IEMI is just now becoming a threat to computer security.

One final physical threat is theft. Physical facility protections should ensure that an IT infrastructure is not threatened by unauthorized outsiders (or insiders) walking away with storage devices or other critical components. A hacker who has stolen a device can spend as long as it takes to attempt to gain access to it.

Natural Disasters

Mother Nature is powerful and unpredictable. All sorts of serious weather events can damage or destroy IT infrastructures. Knowing the types of severe weather common in your area will suggest the correct precautions, such as special insurance, structural reinforcements, lightning protection, surge protectors, bilge pumps, and so on. No matter what the potential disaster, the best protection for data is a reliable regular backup stored in a secured, off-site facility.

Accidents

Accidents happen. Whenever humans are involved, things will go wrong. Murphy's Law states, "Anything that can go wrong will." Preparing for accidents can significantly lessen the impact if one occurs or aid in avoidance.

An IT infrastructure is a large, complex, but fragile thing and it is completely at the mercy of human beings. Accidental damage in the wrong location or at the wrong time can have devastating results. Accidents include spilling liquids on equipment, tripping over cables, pulling out the wrong power cord or cable, tripping the building's circuit breaker, setting off the water sprinklers, knocking over a computer, turning off a system prematurely, installing the wrong driver, and so on.

The best precautions and protections against accidents are backups, configuration documentation, and training. With some commonsense adjustments to worker activity, paying closer attention to activities they perform, and watching out for precarious circumstances, you can avoid many "common" accidents. Training employees what to do in the event of an accident or emergency significantly increases the chances of successful recovery.

MALICIOUS CODE (MALWARE)

Malware is the shortened term for malicious software. Malware is unethical code hackers write to cause harm and destruction. Malware gains access to a system in myriad ways, usually without the consent or knowledge of the user. The most common vectors of this computer contaminant are portable storage devices and Internet communications, like email or social media. A wide range of malware exists, including viruses, worms, Trojan horses, keystroke loggers, spyware, adware, rootkits, logic bombs, trapdoors, backdoors, dialers, URL injectors, and other exploits. The number of unique malicious code examples is astounding. Some recent research of malware attacks has estimated the potential number of unique **signatures** (strings of code used to detect and identify specific malware, much like a fingerprint) as into the billions. This is a huge increase since Symantec noted more than 17 million unique signatures less than 10 years ago.

Just like a biological virus, a computer **virus** needs a host object to infect. Most viruses infect files, such as executables; device drivers; dynamic link libraries (DLLs); system files; and sometimes even document, audio, video, and image files. Some viruses infect the boot sector of a storage device, including hard drives, and USB drives. Viruses spread through the actions of users. As users open infected files, the virus spreads to other files. As users send infected files to other systems, the virus spreads there, as well.

Unlike viruses, which spread from file to file, worms spread from system to system. Because human interaction isn't necessary for propagation, these files can (and do) spread much more quickly than viruses. Today, nearly every threat described as a virus is really a worm. Hackers design worms around specific system flaws. The worm scans other systems for this flaw, and then exploits the flaw to gain access to another victim system. Once hosted on another system, the worm spreads itself by repeating the process. Worms can be carriers to deposit other forms of malicious code as they multiply and spread across networked hosts.

A Trojan horse is actually a mechanism of distribution or delivery more than a specific type of malware. During the Trojan War, the Greeks built a huge, hollow wooden horse; hid warriors inside; delivered the horse statue to the Trojans; and seemingly departed the area. The Trojans took the horse into their citadel and were massacred overnight when the Greek warriors emerged from hiding. The concept refers to an embedded malicious payload within a seemingly benign carrier or host program, a program that the user wanted—such as game or video clip—and intentionally downloaded or copied. Then, when the host program runs the game or video clip, the malware is delivered.

The gimmick of a Trojan is fooling someone (a type of social engineering attack) into accepting the Trojan program as safe. Any program can be converted into a Trojan by embedding malware inside it, in the same way that any food can be poisoned by adding a toxic substance to it. In fact, hackers have specialized tools designed for the express purpose of building Trojan horses called wrappers or Trojan construction kits.

Keystroke loggers record the keyboard activity of a user.

Hackers can deposit software keystroke loggers onto a victim's system through a variety of techniques, including a worm or a Trojan. Once a system is infected, the keystroke logger periodically transmits key logs to the originating hacker through email, File Transfer Protocol (FTP), or **instant messaging (IM)**. Hardware keystroke logger attacks can come through the keyboard cable. These are hard to detect because the device is very small; while often linked to the keyboard cable, it is usually below the desk or behind the computer where users typically would not look.

Spam email is often used to perpetrate **phishing** attacks where an email appears legitimate or enticing, even if it was never requested, causing the receiver to click on it, thus launching malware into their computer system.

Spyware is an advancement of keystroke logging to monitor and record many other user activities. Spyware varies greatly, but can collect a list of applications launched, URLs visited, email sent and received, chats sent and received, names of all files opened, recording of network activity, periodic screen captures, and even recordings from a microphone or images from a webcam.

Adware infiltrates advertisements. Spyware and adware are often linked together in a symbiosis, because the information learned about a target from spyware helps in selecting materials the adware will push through to the user. Adware can push advertisements as pop-ups, as email messages, or by replacing existing legitimate ads on websites as each is displayed in the browser.

Rootkits are malicious camouflage that function as invisibility shields for anything a hacker wants to hide on a computer. A rootkit

acts like a device driver and positions itself between the kernel (the core program of an operating system) and the hardware. From there, the rootkit can selectively hide files on storage devices and keep active processes in memory from being viewable, accessible, or detectible by the OS. Rootkits hide other forms of malware or hacker tools and can include other malware functions in addition to their stealth abilities.

A **logic bomb** is an electronic land mine. Once a hacker embeds a logic bomb in a system, it remains dormant until a triggering event takes place. The trigger could be a specific time and date, the launching of a program, the typing of a specific keyword, or accessing a specific URL. Once the trigger occurs, the logic bomb springs its malicious event on the unsuspecting user.

Trapdoor and backdoor malware are two terms for the same type of malware. A **backdoor** or **trapdoor** program opens an access pathway for a hacker to gain easy access into a compromised system. Backdoor malware is commonly delivered by the Trojan horse method. Unlike malware such as ransomware or a worm, it is most important to the malicious portion of the Trojan to remain hidden. The access could be the creation of a new user account with credentials the hacker has defined; a rogue web, Telnet, or Secure Shell (SSH) server that gives the hacker remote command prompt access; or a source that enables full remote control over the victim's machine (sometimes just by turning on Remote Desktop on a Windows host). Many other possible trapdoor or backdoor manipulations can grant access to external hackers.

A **dialer** is a rogue program that automatically dials a modem at a predefined number. Sometimes this process auto-downloads

additional malware to the victim system or uploads stolen data from the victim. In other cases, the dialer calls premium rate telephone numbers from the victim system to rack up massive long-distance charges. If the user normally connects to the Internet over a dial-up link, the dialer could dial a rogue proxy site instead of the Internet service provider (ISP). This site would act as a man-in-the-middle and then eavesdrop on all communications.

URL injectors replace URLs in HTTP GET requests for alternative addresses. These injected URLs cause a different webpage to appear in the browser than the one requested by the user's click. These replaced webpages could present advertisement sites, generate traffic to falsify search engine optimization (SEO), or lead to spoofed sites.

Exploits are any form of malware designed to take advantage of a flaw in programming, timing, communication, or storage. Hackers often embed exploits into other forms of malware to assist in infection and distribution. Exploits also exist independently, usually as tools employed by hackers to wage attacks, cause damage, and perform intrusions.

Malware is spread through the same communication channels as legitimate, benign data. The difference is that hackers design malware to cause distress and destruction. A growing area of risk for the spread of malware is mobile code. **Mobile code** is software that hackers write for easy distribution over communications networks, such as the Internet and mobile phone networks. Hackers design mobile code to download to a host, and then execute on the host. Malware, under the guise of mobile code, is spreading more rapidly than ever. The general population—and even IT professionals—need

to be more aware, use proper precautions, and use anti-malware applications.

Advanced Persistent Threat

For many years, it was common for the general public, and even many knowledgeable security professionals, to call every type of malware a virus. It has become just as common for every type of malware to now be called an **advanced persistent threat (APT)**. Although the occurrence of APTs has increased dramatically, they still account for a small percentage of attacks. It is true that APTs represent the next generation of malware in that the threat quietly resides on a target machine until activated (hence the term “persistent”). However, APTs are highly targeted, with the targeting intelligence often gleaned from other types of attacks, from phishing to social engineering. Historically, most attacks have been opportunistic attacks seeking the weakest systems to break into: the lowest-hanging fruit. As we see a shift from financially motivated attacks to state-sponsored espionage and hacktivism, or politically motivated hacking, it follows that we are likely to see the continued growth of targeted attacks such as APTs.

Even with all the variations of malware that currently exist and those that will exist in the future, you can choose from only a few common defenses: antivirus software, anti-malware scanners, integrity checking scanners, and user awareness. Antivirus software actively searches for viruses, worms, Trojans, and other similarly destructive forms of malware in memory and on storage devices. Anti-malware scanners look for spyware, adware, dialers, and so forth that an antivirus software might not address. An integrity checker keeps a database of hash values for all system and application files and reports when unauthorized changes occur to those files. You can improve user awareness by offering training that

encourages responsible action with regard to security. Training will also encourage users take reasonable precautions against infection and attack both at work and at home. Training should include awareness of what to do and whom to contact, such as the IT team or the helpdesk, to report incidents within an organization.

FAST GROWTH AND OVERUSE

Network security is not always an organizational priority. Some organizations are more concerned with profits and rapid growth than with spending time on network security. Security is sometimes viewed as an annoying overhead expense that consumes resources without providing any return to compensate for the outlay. Although this mindset is common among senior management, it is a poor and incorrect understanding of the return on investment (ROI) of the crucial investment in network security.

Network security—in fact, all forms of security—protect the organization so that its profit centers can function without interruption or interference. Without network security, the capability and availability of the IT infrastructure would be unstable, especially during periods of accident, infection, attack, or hardware failure. Network security reduces the occurrences of downtime and damaged or lost resources. What could be more important to an organization's bottom line?

Organizations that fail to address security issues in periods of explosive growth are more likely to experience catastrophic failure. By failing to protect assets (communications, data stores, intellectual property, customer data, financial records, and private personnel data), any level of hacker breach could result in organizational implosion.

Racing to get ahead, without proper planning and preparation, usually ends in failure. When constructing a skyscraper, the top of the

structure is not built until each floor below it is properly erected. Growing a company too fast without adequate network security protections is like attempting to build the penthouse before the lower floors are completed. The rise to the top floor may be exhilarating, but the subsequent crash will be unavoidable. What goes up too fast will inevitably come down faster.

A slightly slower growth rate to build network security concurrently with the expansion of the organization is a much smarter plan. Such a deliberate approach is more likely to provide sustained growth and longevity than one based on an unbridled push for forward momentum without considering the risks.

Another potential oversight is when an organization pushes equipment, software, and connectivity beyond a reasonable load level. Trying to pull a yacht with a sports car comes to mind. It looks sexy, but what's at risk? Modern IT equipment is able to perform at astounding levels. But even the best equipment can do only so much before it exceeds its peak operational limitations—and starts trending toward failure.

The growth of most organizations is predictable. Predictable growth can help plan for expansion of infrastructure before the infrastructure becomes a bottleneck. As growth passes 60 percent capacity of the current infrastructure, you should already be planning for expansion. As growth passes 80 percent capacity, take steps to implement expansion. As growth passes 90 percent, accelerate your efforts to complete the expansion.

If the company reaches 100 percent capacity before it completes expansion, a bottleneck inhibiting growth will result. This obstacle can create a bounce-back effect, in which interrupted growth could

shrink the organization, making the return to growth more difficult. Equipment, storage space, memory capacity, backup capabilities, communication bandwidth, and processing capabilities should never reach maximum use or consumption. You should reserve sufficient overhead for the occasional spike above normal maximum activity.

Managing organizational growth has many benefits, including measured, planned upgrades and improvements versus equipment failures and outages. If growth occurs too quickly, how will you know if the lack of bandwidth is due to overuse and overgrowth of the legitimate network processes or a malware infestation? These two situations have similar effects on the behavior of a network and user satisfaction.

WIRELESS VERSUS WIRED

The security implications of a wireless network compared to a wired network are often exaggerated. The biggest difference is the mechanism and proximity of the attack. With wired networks, a hacker must gain physical proximity to a target to make direct contact with it. Once connected to the wired network, the hacker can attempt various attack and exploits.

With wireless networks, the hacker does not have to be physically close. Hackers can attempt network breaches from a mile or more away from the access point (**FIGURE 2-5**). In most real-world situations, however, the range is often under 1,000 feet with a small, but powerful, directional antenna.




FIGURE 2-5 Wired networks require local attacks; wireless networks allow for remote attacks.

In either case, wired or wireless, the hacker must first obtain a network connection with the target network to attack if the hacker's goal is to gain access to user accounts or data stored on the network. If the hacker is mainly interested in destruction and DoS, then logical network access is not necessary.

EAVESDROPPING

Eavesdropping is listening in on communications. Eavesdropping can be the recording of network traffic using a packet-capturing tool, generically known as a sniffer (**FIGURE 2-6**). Hackers can eavesdrop against data packets or against voice traffic. Eavesdropping can occur over wired or wireless connections.




FIGURE 2-6 Eavesdropping on an existing session between client and server.

Any communication performed in plain and directly usable data forms is subject to interception and recording. You can prevent eavesdropping by using encrypted protocols. Only strongly cryptographically encoded messages are safe from outsiders learning the content of the conversation. But as super computers gain speed and efficiency, the ability to crack encryption is only a matter of time once the packets are acquired.

HIJACK AND REPLAY ATTACKS

A **hijack attack** occurs when a hacker uses a network sniffer to watch a communications session to learn its parameters. The hacker then disconnects one of the session's hosts, impersonates the offline system, and begins to inject crafted packets into the communication stream. If successful, the hacker takes over the session of the offline host, while the other host is unaware of the switch.

Replay attacks are also known as **playback attacks**. A **replay attack** occurs when a hacker uses a network sniffer to capture network traffic, and then retransmits that traffic back on to the network at a later time. The goal of a replay attack is to gain interactive or session access to a system.

The traffic captured and retransmitted for a replay attack is authentication packets (**FIGURE 2-7**). In this type of attack, the hacker captures traffic between a client and server, and then later retransmits it against the same server as the original communication. Replay attacks often focus on authentication traffic in the hope that retransmitting the same packets that allowed the real user to log onto a system will grant the hacker the same access.




FIGURE 2-7 Replay attacks collect authentication packets, and then retransmit the packets at a later time.

Fortunately, you can thwart most replay attacks by using one of several common communication improvements. Many authentication transactions include a random challenge-response dialogue that cannot be replayed. This dialogue consists of one endpoint generating a random seed value sent to the other endpoint. The second endpoint uses a mutual secret known by both endpoints to compute a response using a one-way computation. The response returns to the original

endpoint, where the response was predicted. If the received and predicted responses match, the user is authenticated.

Time stamps are another defense against replay attacks. Some authentication exchanges have encoded time details that are difficult to reproduce or modify without detection. Additionally, the use of one-time pad or session-based encryption can make replay attacks impossible.

INSERTION ATTACKS

Insertion attacks come in many forms, but all involve the introduction of unauthorized content or devices to an otherwise secured infrastructure. Three common insertion-based attacks include SQL injection, IDS insertion, and rogue device insertion. Knowing what should be on the network and dropping unexpected or improperly formatted data packets greatly diminishes these types of attacks.

SQL injection is an attack that inserts a hacker's code into a script hosted on a website or through the querying of a database. SQL injection attacks can give the hacker access to the back-end database of a web application. The technique exploits a weakness in common web communications that treats certain characters differently because they are assigned a special meaning or purpose rather than just treated as text. These are called **metacharacters** and act as programming markup. If you don't write a script defensively to block out or ignore metacharacters, then injection attacks can effectively rewrite the script based on content a hacker submits. The injected code can perform just about any possible command line task imaginable.

Cross-site scripting (XSS) is similar to SQL injection, but the results attack future visitors to a webpage rather than grant the hacker access to the back-end database. An XSS attack submits script code to a website. XSS can result in persistent malicious modification of web source files. This causes all future visitors to the site to receive compromised content.

XSS attacks can include emails to victims with falsified hyperlinks that point the script injection to a target site when the victim clicks on the email's embedded links. Such an attack can grant the hacker access to the seemingly secured web transaction of the victim. This form of attack is non-persistent because it affects only those who click on the links in the malicious email.

IDS insertion is a form of attack that exploits the nature of a network-focused IDS; it collects and analyzes every packet to trick the IDS into thinking an attack took place when it really has not. The common purpose of IDS insertion attacks is to trick signature- or pattern-matching detection of malicious network events. By interspersing attack traffic with packets that the target host will reject but the IDS will view, the IDS fails to see the attack pattern, but the attack still takes place. For example, suppose an attack is composed of four packets—A, B, C, and D—and the IDS signature is a packet stream of ABCD. If the hacker transmits the attack as AXBCYD, where X and Y are invalid packets rejected by the target, then the IDS doesn't recognize the pattern. After X and Y are discarded, the ABCD attack occurs against the target.

Rogue device insertion is a physical form of insertion attack where a hacker inserts an imposter device into an infrastructure. The most common example of this is the insertion of a rogue wireless access point (WAP) configured similarly to the real, authorized access point. Some users might be fooled into connecting to the rogue access point. This would constitute a man-in-the-middle attack where the hacker would intercept all transactions from the compromised system. A smartphone can mimic a WAP. Another type of insertion attack is the physical key logger discussed earlier.

Each insertion attack method requires that you create a unique defense. You can prevent SQL injection attacks by defensive programming and filtering input. XSS attacks are generally preventable if you use defensive coding techniques, metacharacter filtering, and input validation. For end users, defenses include cookie management and disabling scripting support in browsers and email clients. Squelch IDS insertion attacks by using modern IDS techniques such as anomaly, behavioral, and heuristic detection. You can derail a rogue device insertion through encrypted communications, preconfigured network access, prohibited wireless networking, user training, and regular site surveys.

FRAGMENTATION ATTACKS

Fragmentation attacks are an abuse of the fragmentation offset feature of IP packets. **Fragmentation** occurs when there are many different network links connected to construct a global infrastructure. Some network segments support smaller datagrams (another term for packet or frame) than others, so larger datagrams are fragmented into the smaller, more compatible size. When the fragmented elements of the original datagram reassemble, manipulations of fragmentation can cause several potentially malicious reconstructions, such as overlapping and overrun. Think of the transporter on *Star Trek*: If anything gets in the way of the reassembly of the person being transported, you might end up with an evil Mr. Spock with a goatee.

Overlapping can cause full or partial overwriting of datagram components, creating new datagrams out of parts of previous datagrams. Overrun can result in excessively large datagrams. Other fragmentation attacks cause DoS or confuse IDS detection and firewall filtering.

Protections against fragmentation attacks include modern IDS detection and firewall- filtering features, as well as performing sender fragmentation. Sender fragmentation queries the network route to determine the smallest **maximum transmission unit (MTU)** or datagram size. The sender then pre-fragments the data to ensure that no fragmentation needs to occur en route. “Beam me up, Mr. Scott—and make sure I get back all in one piece.”

BUFFER OVERFLOWS

A buffer is an area of memory designated to receive input. Buffers are of a specifically determined size set by the programmer, as only a finite amount of memory resides on a host. A buffer overflow is an attack against poor programming techniques and a lack of quality control. Hackers can inject more data into a buffer than it can hold, which may result in the additional data into the next area of memory.

This overflow could be totally ignored, could trigger an overflowing crash or freeze, or could result in arbitrary code execution. In the last case, the hacker crafts the input stream so the overflowed data is a command-line code statement executed with system-level privileges.

Programmers can prevent buffer overflows. Using defensive programming techniques, such as input limit checks and avoiding programming language functions that do not check boundary limitations, buffer overflows become a useless form of attack.

SESSION HIJACKING, SPOOFING, AND MAN-IN-THE- MIDDLE ATTACKS

Attacks on systems and networks can involve falsification of credentials or misrepresentation. This collection of attacks involves a hacker posing as another entity or sending messages that their system is actually a different machine. These attacks include session hijacking, spoofing, and man-in-the middle attacks.

Session Hijacking

Session hijacking occurs when a hacker is able to take over a connection after a client has authenticated with a server (**FIGURE 2-8**). To perform this attack, a hacker must eavesdrop on the session to learn details, such as the addresses of the session endpoints and the sequencing numbers. With this information, the hacker can desynchronize the client, take on the client's addresses, and then inject crafted packets into the data stream. If the server accepts the initial false packets as valid, then the session has been hijacked.





FIGURE 2-8 Session hijacking steals a connection from a client.

In a session hijack, the attacker does not directly learn the credentials of the client. If the hacker loses the connection, he will have to look for another session to hijack. The client who lost the session will be aware that the connection was lost, but will not necessarily be aware that the disconnect was a hijack attack.

Session hijacking sometimes employs DNS spoofing, poisoning, ARP spoofing, ICMP redirects, and rogue DHCP to alter the route or pathway of a session. The hacker uses this pathway alteration to make the session hijacking attack easier by forcing the target session to travel over a more accessible network segment.

Any host that uses TCP/IP without encryption is vulnerable to session hijacking. Even with complex or pseudo-randomized packet sequence numbering, a little eavesdropping is all that is necessary for hackers (or the hackers' tools) to predict future sequence values. The only true protection against session hijacking is encryption, such as a VPN.

Spoofing Attacks

Spoofing is falsification of information. Most spoofing is a falsification of the identity of a source. Email addresses, MAC addresses, and IP addresses are all easily spoofed. Spoofing tricks a user or a host into believing a communication originated from somewhere other than its real source. This is a common tactic in the transmission of spam. Spoofing impersonates an authorized entity, such as MAC spoofing to bypass wireless access-point MAC filtering.

Spoofing is difficult to prevent and somewhat hard to detect. Most spoofing detection occurs when you watch normal traffic and look for addressing anomalies. For example, if a switch sees that a specific MAC address is the source address for frames received on switch port 6, and that MAC address also appears as the source address for frames received on switch port 9, that is a symptom of MAC spoofing (**FIGURE 2-9**). In another example, if a firewall receives a packet on its external interface and the source IP address is an internal LAN address, spoofing could be going on.




FIGURE 2-9 Spoofing of a client's MAC address by a hacker's computer.

TECHNICAL TIP

The **alternate data streams (ADS)** of NTFS are a feature added to this file system to support files from POSIX, OS/2, and Macintosh. This feature was added to **New Technology File System (NTFS)** in the mid-1990s to drive government purchase of Windows NT. However, even with POSIX and OS/2 support now dropped from Windows and Macintosh **hierarchical file system (HFS)**, NTFS has retained this feature. ADS is the ability of a file to contain multiple resource forks. The result of NTFS support for ADS is that complete additional files—not only additional resources—can hide below any normal file object.

A normal file object, including directories, can contain numerous additional files underneath itself. The number of additional files is limited only by the total amount of free space on the drive and the size of the hidden files. Once a file stores as an ADS, it is no longer visible or easily accessible by the OS itself. Several hacking tools can create and manipulate ADS. Only a few scanning tools, such as Streams from **sysinternals.com**, and only a handful of malware scanners can specifically explore a drive for ADS hidden code.

Spoofing is something to watch and filter for, but no real or direct prevention of spoofing exists. Additionally, hackers can intercept and modify data already in transit from a real source if it is not encrypted. Thus, spoofed data does not always originate as falsified communications.

Man-in-the-Middle Attacks

Man-in-the-middle (MitM) attacks occur when a hacker intervenes in a communication session between a client and a server. The attack usually involves fooling or tricking the client into initiating the session with the hacker's computer instead of with the intended server (**FIGURE 2-10**). The hacker performs a spoofing attack in order to trick the client. The result is that the connection between the client and server is proxied by the hacker. This allows the hacker to eavesdrop and manipulate the communications. An MitM attack is also called an **interception attack**, a **proxy attack**, or a **monkey-in-the-middle attack**.




FIGURE 2-10 Man-in-the-middle attacks fool clients into initiating sessions with the hacker instead of the target server.

MitM attacks involve a pre-attack element, in which the client is given false information that leads the client to request a session with the hacker's computer rather than with the real server. The hacker can accomplish this using one of several methods:

- **ARP spoofing**—Address Resolution Protocol (ARP) is a nonauthenticating broadcast query service that requests the Media Access Control (MAC) address from a system using a specific IP. If a hacker running an **ARP spoofing** tool sends a false response to the requester before the real response returns, then the sender will use the false MAC address. Subsequent frames go to the rogue MAC address, which the hacker's computer uses. Because the ARP protocol is Layer 2 and nonroutable, ARP spoofing must occur within a subnet.

- **MAC spoofing**—The hacker's computer uses a server's MAC address; while the server is flooded, the hacker's system receives traffic instead of the intended server. Similar to ARP above, **MAC spoofing** must occur within a subnet.
- **DNS poisoning**—To perform **DNS poisoning**, a hacker compromises a Domain Name System (DNS) server and plants false FQDN-to-IP mapping records. The DNS source will feed subsequent user queries false data.
- **DNS spoofing**—DNS is a nonauthenticating query service that requests the resolution of a FQDN into its related IP address. A hacker hosting a rogue **DNS spoofing** tool can send back false DNS responses.
- **ICMP redirect**—On subnets with multiple routers, **ICMP redirects** can cause a host to alter its routing table. This attack could redirect traffic along a different route than the default, expected, or optimal one.
- **Proxy manipulation**—To perform **proxy manipulation**, a hacker reconfigures a client's proxy configuration. Requests for services go to the hacker's system that acts as a MitM proxy.
- **Rogue DHCP**—A **rogue DHCP** is a false Dynamic Host Configuration Protocol (DHCP) server that can provide IP address configuration leases for a unique subnet and define the default gateway, because the hacker's computer acts as a MitM router/proxy.
- **Rogue access point**—To create a **rogue access point**, a hacker configures a rogue wireless access point similar to the real authorized access point that can fool users into connecting, which then serves as a MitM proxy.

TECHNICAL TIP

When a **nonauthenticating query service** is in use, it does not confirm the source or the validity of any

response received. Thus, if you receive a fake, spoofed, or rogue response, the system accepts it as genuine and the query session ends. If the real response arrives, the system rejects it as an invalid or stray packet because it will no longer correspond to any open query session.

Defenses against MitM attacks include IDS and IPS solutions that monitor for common network abuses or abnormal network activity. Additionally, strong multifactor authentication and mutual authentication can reduce the success of MitM attacks.

COVERT CHANNELS

Covert channels are hidden, unknown, unique, atypical pathways of information transfer. The channel is covert because it is unknown and unseen. Hackers use covert channels for secretive communications, often to leak data (transport it) out of a secured environment, such as an air-gapped system. Covert channels are insecure pathways of transmission. If the pathway were known, it would be an overt channel and likely blocked, filtered, or otherwise secured.

Two main forms of covert channels exist: timing and storage. A timing channel conveys information through timed and synchronized activities. A few potential examples of timing covert channels include:

- Blinking lights to distribute information in Morse code
- Manipulating a fan's speed so the higher and lower pitched noise creates binary transmission
- Throttling the bandwidth consumption on an Internet link so that at a specific interval a utilization measurement reads a value below 60 percent as a zero and a value above 60 percent as a one for binary communications. This does results in a lower **baud rate** (the rate data is transmitted), which will still work for small communication streams such as user credentials, but it will not work for large amounts of packets, like a database.

A storage covert channel conveys information through unseen or undiscovered storage locations. A few potential examples of storage covert channels include:

- Using the unpartitioned space of a hard drive to store data written via a hex editor
- Using a firmware flash memory onboard chip to store data
- Using the alternate data streams of NTFS to hide files
- Using the slack space of a hard drive to store data

The best defenses against covert channels include IDS and IPS, as well as thorough monitoring of all aspects of an IT infrastructure for aberrant or abnormal events of any type. Predicting covert channels is difficult because, by their very nature, they are unknown and unseen.

HOW HACKERS HIJACK SLACK SPACE

A hard drive contains segments known as sectors. A sector is the smallest fixed-size block of storage space of a drive and is 512 bytes. When a file system is applied to a partition, clusters are created out of one or more sectors. Slack space is the unused portion of the last cluster only partially consumed by a stored file. The cluster-to-sector ratio typically ranges from 1:1 to 1:128 for clusters of 512 bytes to 64 KB.

A file system has a fixed maximum number of addresses assigned to clusters. Larger drives have the same number of clusters as smaller drives, but the clusters are larger. It's a little like shoes: kids' shoes are one pair to a box, just like adults' shoes, but the box containing adults' shoes is bigger. A cluster is the smallest consumable element of storage space and can

contain data from only a single file. No native mechanism addresses subcluster divisions in standard file system formats.

When a file writes to a drive, it consumes as many clusters as necessary to contain all of the data of the file. All clusters but the final or last cluster containing the file are filled. The last cluster is fully consumed only if the file happens to be an exact multiple of the cluster size (which is not very common). The unused portion of the last cluster is known as slack space (**FIGURE 2-11**). Slack space is effectively unusable, wasted storage space.




FIGURE 2-11 Slack space is the unused portion of the last cluster only partially consumed by a stored file.

Hackers have developed special file manipulation tools that can locate and hijack the slack space and use it to create hidden volumes on a hard drive. These volumes are nearly impossible to detect because they are not contained, referenced, or addressed by the file system of the storage device. Instead, the slack space drive exists only because special software can address subcluster storage locations. The slack space drive software operates independently of the OS and the file system.

TECHNICAL TIP

Web servers, or at least websites, appear on the Internet in a least four different architectural deployment options: reverse proxy, demilitarized zone (DMZ), co-location, and hosting. Reverse proxy uses a static network address translation (NAT) mapping or port forwarding to allow outside visitors to initiate communications with an internal server. This is the poorest security choice because it could grant hackers access to the entire intranet if the web server is compromised.

Hosting a web server in a DMZ is a more secure solution. However, if compromised, a DMZ web server gives the hacker a weapons platform just outside of the private network's front door. Co-location is placing a web server host directly on an ISP network within a facility. Hosting is

leasing access to space on an existing web server owned and managed by the hosting entity. These last two options are the most secure; if the web server or site is compromised, the hacker gains no location benefit to breach the private network.

NETWORK AND RESOURCE AVAILABILITY THREATS

To be successful, many exploits and attacks require special access on a private network. Some exploits will function against an Internet-facing web server, but such a server might not directly connect to a private network. If a hacker is unable to find an exploitable vulnerability that gains access or control over the targeted systems, a fallback or final resort option is to launch availability attacks.

An availability attack aims at preventing legitimate access or use of resources to delay or interrupt business. Generally, this is known as a denial of service attack.

Denial of Service (DoS)

A **denial of service (DoS) attack** interrupts the normal patterns of traffic, communication, and response. A DoS attack interferes with timely processing and reply to legitimate requests for resources. A DoS attack can be of two primary forms: flaw exploitation or traffic generation.




FIGURE 2-12 Denial of service flooding attack against a client.

Flaw exploitation DoS attacks take advantage of a programming bug, flaw, or convention. The DoS exploit results in the system freezing, crashing, rebooting, or failing to respond to external

communications. You can mitigate flaw exploitation DoS attacks through the application of a patch and the use of an IDS or IPS system. Once you apply a patch, the DoS will no longer be effective.

Flaw exploitation attacks are usually specific to a software version.

Traffic generation DoS attacks flood a target with traffic (**FIGURE 2-12**). The traffic consumes available bandwidth and processing, preventing legitimate communications. No patches exist to mitigate traffic generation DoS attacks. Instead, traffic filtering is the only effective response. **Upstream filtering**, however, is more effective than edge device filtering.

Upstream filtering occurs when a parent network, usually the ISP, provides filtering for traffic before it enters the child network to which individual and business customers connect. Edge device filtering will prevent malicious traffic from entering the private network, but it will not prevent a successful DoS. Only upstream filtering will reduce or eliminate the DoS traffic and allow legitimate communications to continue.

Distributed Denial of Service (DDoS)

Distributed denial of service (DDoS) attacks advance DoS attacks through massive distributed processing and sourcing. The foundations of a DDoS attack are **agents**, **bots**, or **zombies**, which are malicious code implanted on victim systems across the Internet. These mobile agents may create their own peer-network interaction or connect into a public communication medium, such as an **Internet relay chat (IRC)** channel. The resultant network is known as a **botnet**, **zombie army**, or **zombie array**.

The hacker remotely controls the botnet and directs it to perform various malicious activities—including flooding attacks—against selected targets. Generally, the main targets of the botnet are known as primary victims, while the compromised systems hosting the botnet's agents are known as secondary victims. This form of DoS is distributed because the bots are disseminated across numerous secondary victims, and the resulting attacks originate from a plethora of source vectors.

A hacker distributes the bots, agents, or zombies to many secondary victims located throughout the Internet. The bots then connect to some form of communication server, commonly a chat service like IRC, where each can receive instructions from the hacker. Once the hacker sends attack instructions, the bots launch attacks against the primary target (**FIGURE 2-13**).




FIGURE 2-13 Distributed denial of service flooding attack against a primary target.

A botnet can perform a wide range of malicious actions including flooding, spamming, eavesdropping, intercepting, MitM, session hijacking, spoofing, packet manipulating, malware distributing, phishing site hosting, password stealing, encryption cracking, and more.

Several botnets, such as Storm and Conficker, have appeared in the last decade; these two worms had an estimated secondary victim base of 75 to 100 million systems.

Defenses against DDoS focus on either avoiding becoming a secondary victim or protecting against primary victim onslaughts. To avoid becoming a secondary victim, measures include current antivirus and anti-malware scanning, user behavior modification, firewall filtering, and IDS/IPS solutions. Protection against primary victim onslaughts includes firewall filtering, honeypots, IDS/IPS solutions, and DDoS mitigators. A **DDoS mitigator** is a purpose-built appliance that inspects traffic entering a network; it blocks attacks while allowing legitimate traffic to flow. Hybrid solutions are available that combine on-premises appliances with a cloud-scrubbing service to filter traffic.

HACKER TOOLS

Hacker tools include a wide variety of software: from mundane native OS utilities, to commercial applications, to custom-coded exploits. Generally, any software put to a malicious or unauthorized (according to company security policy) use is a hacking tool. Hacking tools perform hacking activities. All of the exploit concepts mentioned in this chapter are possible through a wide variety of hacking tools and utilities.

No master list of hacker tools exists to search for or block access to protect IT systems. Just about every legitimate program can be put to some illicit task. To defend generally against hacker tools, consider using a whitelist restriction system.

A whitelist restriction system incorporates a list of software executables authorized for use. A user can launch any application on the list, but all executables not on the list can be blocked from running. A whitelist cannot focus only on authorized filenames; this process also uses a hash value to prevent a hacker from easily bypassing the limitation through simple file renaming. In addition to whitelisting, you can reduce the threat of hacking tools by limiting Internet downloads and file exchanges, controlling use of portable storage devices (especially those used on external systems), filtering email attachments, installing IDS/IPS solutions, and providing user education.

SOCIAL ENGINEERING

Social engineering is the art of manipulating and exploiting human nature by getting people to perform tasks or release information that violates security. Social engineering is an exploit that can almost always be performed against a target organization due to the presence of humans. Humans are the primary targets of social engineering.

Humans are the weakest link in most security solutions because humans are the only element in an organization with free will. Every other element can perform only within its programming and design. In addition, humans can be tricked or fooled, while hardware and software can perform only in accordance with their design and programming. Social engineering can take place over any communication method, including face-to-face, telephone, email, IM, and websites. Social engineering may focus on extracting information from a target or convincing the target to take action that alters the security status of a host or network.

Many social engineering attacks stem from some form of relationship, from initial and casual to business professional to long-term and highly developed. The more in-depth and long-term the relationship, the more leverage the hacker can use to turn, trick, or abuse the target. Social engineering can involve a wide range of techniques, including impersonating a position of authority, reciprocating favors, using social validation, and creating urgency through scarcity. Often these attacks become more successful if the hacker can impersonate an insider.

Gaining access to inside information is often the first element of a social engineering attack. **Dumpster diving**, using reconnaissance, and **cold calling** are techniques to learn about the internal culture of the target. As a hacker learns more and more terminology, processes, organizational hierarchy, policies, events, gossip, social occurrences, calendars, project scheduling, and so on, the more the hacker is able to simulate being an insider. Once a target is fooled into believing that the hacker is just another employee, the initial attack is successful. With the standing gained, the hacker manipulates the target into revealing more internal information, reconfiguring systems, or downloading tools from questionable Internet locations.

Social engineering may be the first wave of hacker attacks, or it could be the last-resort fallback plan if attempts to perform logical intrusion or physical burglary fail. Some hackers are naturally gifted at social engineering, while others must practice to obtain workable competency at the craft. These skills of social engineering are not unique to this unethical activity; instead, they are the same skills most people use in normal social situations when trying to get their way, convince someone to go out on a date, ask for help, improve social status, get out of trouble, lead a group, sell and market, create advertising, and so on. The difference is that hackers have an unethical goal in their use of these skills.

Social engineering, primarily attacks against people, is invulnerable to typical IT countermeasures. Instead, the best defense against social engineering is thorough user training and awareness. Once personnel are aware that they are, have been, and will be targets of attack, they can adopt a slightly suspicious and cautious outlook.

Employees should skeptically evaluate any activity, question, interaction, or relationship that seems odd or out of place.

You can help reduce the threat of social engineering by using security policies that employ information classification with related restrictions on communication methods. If you limit the communication channels that specific classes of information traverse, you will reduce information leakage caused by social engineering. For example, if you restrict the use of passwords over the telephone or by email, then anyone who requests a password will be obviously attempting to violate security. Employees should be trained to report all such requests to the network security staff.

CHAPTER SUMMARY

Hackers are consistently seeking to take advantage of anyone or any system not prepared or properly secured. Understanding the various means of attacks hackers commonly employ directly improves awareness and overall network security.

Hackers often seek monetary gain through attacks against individuals and organizations. Hackers can be employees or outsiders. Compromising situations are not limited to hacker attacks, but they can also include accidents, oversights, hardware failure, rapid growth, and severe weather. Hacker tools and techniques include malicious software, exploiting wireless connections, eavesdropping, replay, insertion, fragmentation, buffer overflow, XSS, man-in-the-middle, session hijacking, spoofing, covert channels, and the availability attacks of DoS and DDoS. You should take action to restrict or limit hacker tools and use caution and training to avoid social engineering.

Hacking attacks and techniques other than those listed here exist. This chapter offers a generic description of the hacking process, not a definitive or

exhaustive examination. However, from this foundation, you can develop a greater understanding of hacking and the threats posed by hackers (as well as other sources of threats and risk), leading to improved security design, policy, and implementation.

KEY CONCEPTS AND TERMS

Advanced persistent threat (APT)

Adware

Agents

Alternate data streams (ADS)

Arbitrary code execution

ARP spoofing

Attacking

Backdoor

Banner

Banner grabbing

Baud rate

Bots

Botnet

Breach

Buffer

Buffer overflow

Chip creep

Clusters

Cold calling

Command shell

Contract workers

Covert channels

Cross-site scripting (XSS)

Cryptocurrency

DDoS mitigator

Denial of service (DoS) attack

Deterrent

Dialer

Disgruntled employees

Distributed denial of service (DDoS) attacks

DNS poisoning

DNS spoofing

Domain registrations

Dumpster diving

Eavesdropping

Enumeration

Fallback attacks

Flaw exploitation attacks

Flooding

Footprinting

Fragmentation

Fragmentation attacks

Hackers

Hacking

Hacktivism

Hierarchical file system (HFS)

Hijack attack

Honeypots

ICMP redirects

IDS insertion

Insertion attacks

Instant messaging (IM)

Intentional electromagnetic interference (IEMI)

Interception attack

Internet relay chat (IRC)

Keystroke loggers

Leetspeak

Logic bomb

MAC spoofing

Man-in-the-middle (MitM) attacks

Maximum transmission unit (MTU)

Mean time between failures (MTBF)

Mean time to failure (MTTF)

Metacharacters

MITRE

Mobile code

Monkey-in-the-middle attack

**National Institute of Standards and Technology
(NIST)**

New Technology File System (NTFS)

Nmap

Nonauthenticating query service

Opportunistic hackers

Partition

Phishing

Ping sweeps

Playback attacks

Port scanning

Privilege escalation

Professional hackers

Proxy attack

Proxy manipulation

Pwned

Ransomware

Reconnaissance

Recreational hackers

Replay attack

Return on investment (ROI)

Rogue access point

Rogue device insertion

Rogue DHCP

Rootkits

Scanning

Script kiddie

Sector

Session hijacking

Shell code

Signatures

Slack space

Social engineering

Solid-state drives (SDDs)

Spam

Spoofing

Spyware

SQL injection

Static electricity discharge (SED)

Trapdoor

Trojan horse

Unpartitioned space

Upstream filtering

URL injectors

Virus

Wardialing

Wardriving

Whois

Worms

Wrappers

Zombies


Zombie army

Zombie array

CHAPTER 2 ASSESSMENT

1. Which type of hacker represents the greatest threat because they likely already have physical access to a target?
 - A. Consultant
 - B. Competitor
 - C. Cleaner
 - D. Customer
2. Most exploits are based on the existence of which of the following?
 - A. Bandwidth speed
 - B. Human beings
 - C. Filtering protocols
 - D. System anomalies
3. What is the first stage or step in the hacking process?
 - A. Enumeration
 - B. Penetration
 - C. Reconnaissance
 - D. Scanning
4. Which form of attack captures authentication packets to retransmit them later?
 - A. Hijacking
 - B. Insertion
 - C. Interruption
 - D. Replay
5. Which form of attack submits excessive amounts of data to a target to cause arbitrary code execution?
 - A. Buffer overflow
 - B. Fragmentation

- C. Insertion
 - D. Interruption
6. Which attack is based on the impersonation of a legitimate host?
- A. DoS
 - B. Fragmentation
 - C. Hijacking
 - D. Spoofing
7. Which method of communication is unseen, unfiltered, and based on timed manipulations?
- A. Buffer overflow
 - B. Covert channel
 - C. IDS insertion
 - D. Man-in-the-middle
8. Which attack uses nontechnical means to achieve results?
- A. Buffer overflow
 - B. Covert channel
 - C. Social engineering
 - D. SQL injection
9. What does a hacker exploit in a target system?
- A. A botnet
 - B. A vulnerability
 - C. Multifactor authentication
 - D. Traffic filtering
10. Which of the following might a hacker launch if the other attempts are not successful?
- A. Buffer overflow
 - B. Fallback attack
 - C. Covert channels
 - D. Zombie attack



© Digital_Art/Shutterstock

Common Network Topologies and Infrastructures

COMPUTER NETWORKS allow organizations to share information and resources in order to allow for and enhance employee productivity. Instead of physically going to the location of a resource, if the desired item (such as a file, folder, software program, image, or video) is loaded on the network, and the user requesting the item has authority and is granted access to that item by the administrator, then the item can be accessed more quickly and efficiently from the user's desktop. This streamlines inventory management of data files and resources, and it can enhance communication and collaboration. Networks also provide centralized

storage for important information and enables authorized users to access that information quickly and efficiently.

Networks can range in size from a few devices to thousands of interconnected computers, devices, and peripherals. Some networks are located in the same room or the same building, while other networks span the globe. Nontraditional devices, including gaming systems, smartphones, and Internet of Things (IoT) devices, are connectable and may constitute a network. How those networks are constructed, and the physical and logical ways in which each is laid out and communicates, is the network's topology.

Networks are composed of different devices connected together using a common addressing scheme; in this way, each device can be easily accessed, and messages to it can be quickly delivered. The primary methods of addressing are the use of Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) addressing. Addressing is how communication packets on a network can be sent and received by devices on the network.

CHAPTER 3 TOPICS

This chapter covers the following topics and concepts:

- Common network topologies, including bus, ring, star, mesh, and wireless
- Various network devices that are part of a network topology
- Relationship between physical and logical topologies

- Differences between IPv4 and IPv6
- Examples of network infrastructures

CHAPTER 3 GOALS

When you complete this chapter, you will be able to:

- Describe common network topologies—bus, ring, star, mesh, and wireless—and the pros and cons of each
- Describe differences between wired and wireless configurations
- Differentiate between physical and logical topologies
- Identify scaling concerns for each major topology
- Explain the significance of IPv6 to modern networks
- Describe common network infrastructures and related security concerns

WHAT IS A NETWORK TOPOLOGY?

Computer networks are an essential part of most businesses today. Whether that business is a small home-based office or a multinational organization, computer networks often reduce the cost and time of sharing resources and increase business efficiency. Understanding how a computer network functions is essential to understanding how to secure such a network environment.

In the context of networks, the term **topology** references the order and arrangement of the elements of a communications network. Topology has often been described as what the computer network would look like if the roof were open and you could see all the computers and peripherals clearly connected together below. The way in which computers and network peripherals are connected determines which communication methods may be used on the network, what hardware is required, and whether redundancy exists.

Obviously, it is not easy to look down into a building from the roofline, but you can create network diagrams and network maps that aid in understanding the connections and communication pathways that exist. These diagrams are an essential tool in troubleshooting network problems and issues, because the difference between what you expect to find and the way the network is configured is often the issue to be addressed or corrected.

Before continuing into the various topology options available for configuring a network, it is logical to define the different types of

devices and appliances that are commonly found in networks, which helps make a network topology more understandable.

TYPES OF NETWORK DEVICES

Computers and servers, including blade servers and special-purpose computers, are typically found on networks, but there are many other devices used for efficiency, segregation, or security.

A **hub** is a device that consolidates or collects the wires and connections of other computers and devices. Often, a server supports more than a few computers, which would require each computer to connect to the server individually. Most servers do not come equipped with enough individual ports to accommodate this type of connection. A hub simply collects the individual physical wires and provides a single connection to the server (**FIGURE 3-1**), limiting the physical wiring and streamlining the connections. A simple collection hub is commonly considered a **dumb hub** or **passive hub** because there is no logic or power within the box—it is only a collection point for wires.




FIGURE 3-1 A hub on a simple local area network.

A **punchdown block**, also known as a **punch panel**, is an example of a hub, although it is not shaped like a traditional passive hub. A hub is an Open Systems Interconnection (OSI) Layer 1 (Physical Layer) device because it physically connects the wires together into a single appliance. A hub is nondiscriminating, and many data collisions occur due to the pass-through behavior of the device. Traffic coming in on an inbound port is sent to all other ports connected to the device, whether or not the data is intended for that port and the connected devices, which makes for additional traffic generation and bandwidth usage.

Another type of hub is the **smart hub** or **active hub**. This type of hub has some intelligence built into it. These are the most common hubs on the market and often do not include “smart” or “active” in

the naming convention of the product. Other common terms for these devices include **concentrators** and repeaters. These terms are applied because this type of hub often concentrates the signal as it enters the device, and it can strengthen the signal through the repeating process as the signal exits the device.

Historically, smart hubs included some simple management software that enabled network administrators to perform basic commands and elicit statistics about the network's performance, such as throughput measures, uptime percentages, routing table entries, and other valuable information for troubleshooting and baselining. Many of these same measures can be performed from the various servers and devices on the network. However, because a hub is a consolidation point, information consolidation typically performed by the user is done by the smart hub instead.

■ NOTE

Internet of Things (IoT) home automation hubs are more device-specific and are not the same type of hub used to control a computer network. However, an IoT hub may be implemented as a device on a home network.

A **repeater** is a type of hub that rebuilds the signal received through a bit-by-bit process. It does not amplify the data; it rebuilds it to pass it on. This function is necessary because a natural consequence of long segments of network cable, or distance between wireless devices, is that the signal can become weak or degraded. A repeater prevents data loss by offering the newly reconstructed packet

at a stronger signal strength so it can continue through the network to its destination. Active hubs extend the distance available between devices or network segments.

A **switch** is a Layer 2 (Data Link Layer) device that uses Media Access Control (MAC) addresses to differentiate traffic. A typical switch is shown in **FIGURE 3-2**. The MAC address is a unique hardware identification number assigned by the manufacturer under guidance of the Federal Communications Commission (FCC). Each node on the network has a network interface controller (NIC), which has a unique MAC address. Because the switch reviews the MAC address in determining where each data packet is meant to go, there are no collisions; the data packet is sent only to the intended destination (if it exists). If the destination address does not exist, the packet is dropped. Switches are typically used in large networks. They become a point of connection within a network because they switch to the proper port when sending packets to the intended recipient device. A switch can perform some error checking and discard any packets with detected errors.



FIGURE 3-2 A switch.
© VtIs/Shutterstock

Switches provide network segmentation through hardware. Across a switch, temporary dedicated electronic communication pathways connect the endpoints of a session (such as a client and server). This switched pathway prevents collisions. Additionally, switches allow you to use the full potential throughput capacity of the network connection by the communication, instead of wasting 40 percent or more by collisions (as occurs with hubs). See **FIGURE 3-3**.




FIGURE 3-3 An example of switch deployment.

You can see this basic function of a switch as a security benefit once you examine how this process takes place. A switch operates at Layer 2, the Data Link Layer, of the OSI model, where the MAC address is defined and used. Switches manage traffic through the use

of the source and destination MAC address in an Ethernet frame. The Ethernet frame is a logical data set construction at the Data Link Layer (Layer 2), consisting of the payload from the Network Layer (Layer 3) with the addition of an Ethernet header and footer **(FIGURE 3-4)**.

Example of Switch Deployment




FIGURE 3-4 An IEEE 802.3 Ethernet Type II frame.

Switches employ four main procedures, labeled as “learn,” “forward,” “drop,” and “flood.” The learn procedure is the collection of MAC addresses from the source location in a frame header. The source MAC address goes into a mapping table, along with the number of the port that received the frame. Forwarding occurs once a frame’s destination MAC address appears in the mapping table. This table then guides the switch to transmit the frame out the port that originally discovered the MAC address in question.

If the frame goes on the same port that the mapping table indicates is the destination port, then the frame is dropped. The switch does not need to transmit the frame back onto the network segment from which it originated. Finally, if the destination MAC is not in the mapping table, the switch reverts to flooding—that is, a transmission of the frame out every port—to ensure that the frame has the best chance of reaching the destination. By monitoring the activity of the switch—specifically, watching the construction and modification of the mapping table and the variation of MAC addresses seen in frame headers—you can detect errors and malicious traffic. Intelligent or multilayer switches or external intrusion detection system (IDS)/intrusion prevention system (IPS) services can perform this security monitoring function.

When the switch procedures fail, the network suffers. This could mean a reduction in throughput, a blockage causing a denial of service (DoS), or a redirection of traffic allowing a hacker to attempt to modify or eavesdrop. Fortunately, such attacks are “noisy” in terms of generating significant abnormal network traffic, making it easier to

detect with basic security sensors. This benefit alone is a significant reason to use switches rather than hubs.

A **router** is a wired or wireless device that works with Layer 3 (Network Layer) Internet Protocol (IP) addresses (see **FIGURE 3-5**). These are the addresses used to identify and move packets destined for and received from the Internet. A router is usually positioned between different segments of a network, allowing for a connection between the segments and routing traffic from one segment to the destination segment. A router has several ports that are consolidated, along with a port that connects to the Internet or to the network segment that eventually connects to the Internet. Consolidated router ports take in many streams of data and merge them to send to various parts of the network. A router is most often used to connect networks or network segments to ease the process of packet transmission.



FIGURE 3-5 A wireless router.
© Natalia Siverina/Shutterstock, Inc.

■ NOTE

A **sniffer**, also known as packet analyzer, network analyzer, and protocol analyzer, is a software utility or hardware device that captures network communications for investigation and analysis. To help prevent infrastructure level attacks, ensure that routers are using authentication to exchange routing data and that they are protected against unauthorized physical access.

A router's primary purpose is to direct traffic toward its stated destination along the best-known currently available path (see **FIGURE 3-6**). Routing protocols such as RIP, OSPF, IGRP, EIGRP, BGP, and others dynamically manage route selection based on a variety of metrics. A router supports security by guiding traffic down preferred routes, rather than routes that might not be as logically or physically secure. If a hacker can trick routers into altering the pathway of transmission, network traffic could traverse a segment where a hacker has positioned a sniffer.




FIGURE 3-6 An example of a routed network deployment.

Although in some instances switches and routers may be used interchangeably, the most common use of each is to place routers prior to the Internet connection and to use switches at the connection between different groupings of computers. Routers used to be slower than switches, but that is no longer the case. Recent additions to the market offer the same speed and performance as switches.

A **bridge** is a Layer 2 (Data Link Layer) device. This device acts like a smart hub or repeater but includes the ability to filter the packets based on MAC address, allowing connection of two local area networks (LANs) if those LANs are using the same protocol for communication. Unlike hubs, routers, and switches, bridges have only two ports because the device connects only two LANs.

Bridges link between networks by creating a path or route between networks, used only when the destination of a communication actually resides across the bridge on the opposing network. An analogy of this would be a city split in two by a river. When people in that city want to visit a location on their side of the river, they have no need to cross the bridge to get there.

The same is true of a network bridge. Bridges work in a similar manner to switches in that they use several basic procedures to manage traffic. Bridges are used to connect networks, network segments, or subnets whenever you desire a simple use-only-if-needed link and the complexity of a router, switch, or firewall is neither specifically required nor desired.

Bridges use the process of learn, forward, and drop. With these processes, only traffic intended for a destination on the opposing network will travel across the bridge to that network. All traffic intended for a destination on the same side of the bridge that received it will be dropped, because the traffic is already on the correct network. See **FIGURE 3-7**. Bridges can also link networks with variations in most aspects of the infrastructure, barring protocol. This includes different topologies, cable types, transmission speeds, and even wired versus wireless.

An Example of Switch Deployment




FIGURE 3-7 An example of a MAC-layer bridge deployment.

You can use bridges to detect the same types of traffic abuses that switches do. In addition, both bridges and switches can impose filtering on MAC addresses. Here you can use a **blacklist** or whitelist concept, where you allow all traffic except for that on the black/block list, or you block all traffic except for that on the white/allow list.

You can arrange to **log** activities across a router, switch, and bridge. You should review traffic logs regularly, typically on the same schedule as firewall and IDS/IPS logs. You will note the signs of abuse, intrusion, denial of service, network consumption, unauthorized traffic patterns, and more in these common network devices not typically considered to be security devices. In large networks, it is preferable to have an automated monitoring solution, as logs are generated in real time based on the actions occurring on a network. Keeping up with new entries and anomalies can be very difficult.

Two types of bridges exist: transparent bridges and source routing bridges. In a **transparent bridge**, the LANs that are connected and the individual devices on each LAN are unaware that there is a bridge on the network. It usually works seamlessly and is therefore transparent to the actual user or device. This is a great benefit to network administration, as the devices on each LAN do not require configuration to recognize the bridge.

In a **source routing bridge**, routing occurs based on the use of a discovery frame that travels the network to determine all possible paths to the intended destination, and then it selects the best one to use. This causes additional network traffic but often results in a faster delivery of the packets. Over the course of a day, the traffic on a

network can change drastically based on resources requested and the users accessing it. As the discovery frame is sent out right before the actual packets, the most optimal route is always selected.

A **brouter** is a combination of a *bridge* and a *router* and can be referred to as a bridging router. This device performs the functions of both a bridge and a router in that it can effectively route packets to the different segments of a network and work as a bridge to send those packets to a different network. Depending on the specific functions performed on the network, the device can work at either Layer 2 (Data Link Layer) or Layer 3 (Network Layer).

A **gateway** is a device that connects two networks that use dissimilar protocols for communication (**FIGURE 3-8**). A gateway is similar to a bridge and can function as a bridge; however, a bridge cannot function as a gateway due to dissimilar protocols in use. This device essentially takes data packets from one network in a specific protocol and converts those packets to the protocol of the receiving network. Because of this action, the device is sometimes referred to as a **protocol converter**.




FIGURE 3-8 A gateway connecting two networks.

FYI

Modem is the name originally given to the device that translated digital signals from the computer into analog signals on traditional telephone lines so the information could be transmitted. The term is derived from the words ***mod***ulate and ***dem***odulate to relay the transformation the data undergoes in the process. Modems are used to connect to **Internet service providers (ISPs)** using cable wire, fiber optic wire, or telephone wire, as was traditionally used.

Historically, the router and modem were different devices that performed different tasks for the network environment. Today it is possible to purchase devices that offer the functionality of both. Although the telephone connection is rarely used in most settings, the term *modem* is still used to refer to the device offered by the cable company to bring television, computer, and video access to the home.

WHAT DIFFERENTIATES LOGICAL AND PHYSICAL TOPOLOGIES?

Computer networks are arranged in a manner that includes the way each is physically connected, which is the **physical topology** (wired or via the air if wireless), and the logical manner, which pertains to how resources are segregated and allocated. All computer networks have physical and logical topologies, and it is possible for the topologies to be similar. From the user perspective, neither topology matters unless the user is unable to access the resources needed to complete tasks.

Because it is easiest to consider a topology as the way computers and devices are connected, let's cover wired topologies first. **Wired topologies** have a physical wire between devices, allowing for communication among those devices. A **route** is the pathway that communication packets travel within the wired connections of a network. The route for network communication is not always easily identified or managed if there is not appropriate documentation to aid the network administrator or engineer. Network equipment, including routers, switches, and gateways, may alter what appears to be a straightforward path of communication to better balance the communication load or to restrict access to certain devices or parts of the network.

The **logical topology** is how the network appears from any device or user and is governed more by policy and access than by physical

connectivity. Instead of how the computers and components are physically arranged, think of the logical topology as how those computers access and share data and resources. It may have little resemblance to the physical connectivity of the devices.

Whether identifying the physical or logical topology, some elements should be kept in mind. Due care should be exercised when planning a network installation or upgrade, or **scaling** (whether making it bigger or smaller) to ensure that the appropriate physical and logical topologies are selected to support the needs of the business. Both topologies should be documented clearly in both the security policy and in a standalone network document, so it is accessible for troubleshooting networking problems by the network administrators and engineers.

Types of Physical Topologies

A **bus topology** is considered a linear network because it has computers in a line (although not necessarily a straight line). See **FIGURE 3-9**. Each computer or device in the topology is connected to the computer or device that is next in line. The actual configuration could be wavy (weaving up and down rows of computers in an office space) or appear circular in formation, although it is not a ring or circle, as the ends never meet. A bus topology requires equipment called **terminators** to be placed at each end of the physical cable. The terminators stop the signal on the line from bouncing back and forth down the line.




FIGURE 3-9 A bus topology.

In this topology, a single cable is used as the backbone of the network. The computers and devices that are part of that network are connected to that cable. This network topology is easier to implement and troubleshoot in home office and small business settings, but it is often slower, especially when there are more than 10 devices on that network. When a system fails in a bus topology, the entire network goes down.

Connectivity issues in the network affect all other computers and devices connected to the network in a bus configuration. A network with a bus topology is often more susceptible to performance issues as the number of computers and devices on the network increases beyond 10. There are also length limitations on the cabling—the signal degrades the further from the origination point the packet gets. Repeaters often solve some of the problems with long distances, but if the desired quality of throughput is not achieved, wireless may be added to reach additional distances. However, wiring and installation is easier and usually less expensive than other configurations, which is why the bus topology is the most common network design in small office/home office installations. Another advantage of the bus network is the relative ease of scaling the network up to the maximum number of devices or removing devices to scale down.

The **ring topology** uses a cabling scheme that is closed; that is, it connects one end of the cable/network to the other end to form a ring (see **FIGURE 3-10**). Each computer or device is connected to two devices, one on each side. This topology requires a special packet called a **token**, similar to a “speaking baton” in childhood games where only the person in possession of the baton may communicate. The use of a token is the reason this topology is often referred to as “token ring.” Because this network scheme uses a token to control access to the network cable for communication and transmission, it has fewer problems with data collisions than other physical topologies; however, it will still experience slower speeds when multiple devices on the network attempt to communicate at the same time. Ring networks often use repeaters when there are a large number of nodes on the network.




FIGURE 3-10 A ring topology.

If the ring network is **unidirectional**, which means the token moves in only one direction around the ring, there could be delays.

Every computer and device on the network is polled (asked or queried) if it has packets to send before the token moves to the next device, in order, around the ring.

A **bidirectional** ring requires two connections to each device and provides redundancy not available in a unidirectional token ring network. This type of ring has an entire cabling setup for a token to go in either direction. That means it is not actually a bidirectional ring; it is two separate rings, each going in a separate direction.

Essentially, the ring network is implemented twice with flow in opposite directions. The wait time is not necessarily cut in half due to the addition of the second network, but it is a faster communication system than a single ring. Cost is a deterrent for many types of ring networks, and troubleshooting is more complicated because the administrator must first determine which network is experiencing the problem.

Benefits of a bidirectional ring include redundancy because of the duplication of the entire network. However, if a computer goes down, it is possible the entire network could go down in both directions. Speed is potentially improved, due to load balancing across the two ring networks, but if the network is at capacity, it would still seem slow to users. The token is also not directly opposite on each network. Tokens work through a polling process. Each device is polled to identify if there are communication packets to send across the network. If not, the token moves to the next device in the ring. If there are packets to send, the device sends those packets, and the token moves to the next device in order. Depending on how much data is sent, the system could run slower than anticipated.

The **star topology**, shown in **FIGURE 3-11**, provides a layer of performance protection that is not present in the bus or ring topologies. The network design has an independent cable to each device. This configuration often uses much more cabling than a bus or ring topology, but it benefits from exclusive use of each cabling segment. If one segment is disrupted, it usually does not bring down the entire network. Because the cable to each device is isolated from other devices, it offers redundancy from element breakage. But it does not offer full redundancy to all devices, as the affected device will not have connectivity to the network until the issue is resolved.




FIGURE 3-11 A star topology.

A common visual for this design is a circle with a line emanating from the center to devices on all sides. This makes the image resemble the spokes in a wheel or outlined pieces from a pie. In application, the “spokes” do not need to extend to all sides. It is common to have the shape appear more like a squid where the server and hub are at the top and tentacle-like cables drop down to each

device. Either way, the inclusion of a central “hub” and segregated cables to each device identifies the topology, even if the physical design bears little resemblance to an asterisk. The star topology is more fault-tolerant than the bus topology, as the independence of each wired line allows for one or more devices to lose connectivity; the remaining devices maintain connectivity to the rest of the network.

Mesh topology is the most redundant of the wired topologies, as the design provides isolation for devices and redundancy in the event that a segment of the network no longer functions to allow communication through. Mesh topologies can be enacted in two forms: full and partial connections. In a full connection mesh network (**FIGURE 3-12**), all devices are connected to all other devices. This essentially means that before every device on the network is a hub to facilitate and streamline the connections.




FIGURE 3-12 A full mesh topology.

A more common and less cost-intensive method is the partial connection mesh (**FIGURE 3-13**), where all high-demand, resource-intensive devices have redundant connections, but not every device to every other device. In this type of application of the mesh topology, there is redundancy for all servers and the main systems in each area. For example, in a school that offers online courses, the main servers would have redundant lines between each, and the online learning management system (LMS) that hosts the online class environment

would have cables to several other servers and devices. If one path to the Internet was inaccessible, there would be others, so the classroom environment would still be available to students. This is common in mission-critical areas of a business and in the military, as communication is essential to survival.




FIGURE 3-13 A partial mesh topology compared to a full mesh topology.

Wireless topology is a bit of a misnomer. In a wireless installation there is very little wire, as the devices connect using the air as the medium of exchange for communication packets. Some wiring may exist in areas that require additional redundancy, so that wired and wireless provide that alternative in the event that one method fails to function properly. There is at least one wireless router in a wireless network topology. This is the connectivity point that allows for interconnection and access to the Internet. See **FIGURE 3-14**.




FIGURE 3-14 A wireless network topology.

Logical Topology

All of networking is the communications between devices, but logical networking is more about connections and permissions than arrangement. In a typical office environment, workers may be able to see each other as they work because the desks in the work area are in close proximity. The computers, however, may not be able to detect the existence of the other computers and devices. Implementation of a logical network is important for separation of duties and access. Usually, the employee responsible for marketing an organization's products or services is not the same person who completes the organization's payroll or pays the bills. Those job functions are usually completed by two or more people who would not typically access the same programs and tools. An example of a logical topology is shown in **FIGURE 3-15**.




FIGURE 3-15 A logical topology.

For those tools that the entire company needs to access, such as email, all employees should be granted access. For the other resources, like payroll applications or inventory management, only those employees with job responsibilities requiring access to those programs are granted access to those systems. If there are several employees in each department or division, access could be granted through the use of a template or group.

Logical networks limit access to data and resources by allowing only those individuals and devices that require such access, if the

network is set up correctly. This segmentation protects employees from accidentally accessing, altering, or deleting information. Segmentation aids in legal compliance efforts by supporting least privilege and separation of duties. When access to network resources is denied to those persons and systems without a legitimate need, it can also aid in slowing or stopping the spread of malware, as an infected system or account is unable to spread the malware past its privilege position.

Logical networks also aid in hiding assets, as not all devices can see or access every asset on a network, only those to which permission is granted. This is helpful when there are several servers performing the same function, such as load balancing a web server. Limiting access also potentially limits network traffic through certain pathways and communication channels, or it may redirect traffic to a more optimal connection stream. This alteration can aid in network performance—or the perception of network performance—if network traffic and communication appear to be faster because of limited access to specific segments.

Despite the many benefits to using logical network topologies, there are some drawbacks to consider prior to implementation. The logical network can complicate network scaling, making it bigger or smaller with the addition or removal of devices, as placement of each device does not imply the permissions granted. Visually identifying devices in close proximity does not mean those devices have access to or recognize the existence of each other. It can complicate troubleshooting problems because network engineers cannot rely on appearance of the network; they must understand the logical pathways and accesses in order to determine where an issue exists.

The need to grant permissions and deny access is a complicated process that requires a thorough understanding of how the network should function. Use of a template or group aids in this effort, but customized access must often be granted on an individual account or system basis. Failure to properly document or implement access could negate any benefit derived from the initial segmentation. There can also be a false sense of security when utilizing logical network strategies to restrict access and hide assets; one may believe that use alone is sufficient protection, instead of considering the need for updating all elements of a network security plan, such as patching, reviewing accounts and access regularly, and removing any unnecessary accommodations.

Creating Logical Topologies

There are two main ways to create logical networks: subnetting and VLANing. **Virtual local area networks (VLANs)** are created at OSI Layer 2 (Data Link Layer), and subnetting occurs at OSI Layer 3 (Network Layer).

A VLAN is an example of a logical network. A VLAN can be created from devices from different LANs that function as if they are on the same network (**FIGURE 3-16**). In many applications, the creation of the logical network is the same regarding OSI layers, but it is possible to construct the network independent of Layers 2 and 3 matching. Traffic configured to a subnet from a physical network does create a logical network that uses the same switch, which can slow traffic. If a VLAN is used for each logical network, the load on the switch is relieved because traffic does not cross the switch from one VLAN to another.




FIGURE 3-16 VLANs are logical partitions on a network that define groupings of hosts based on logical associations rather than physical connections to a switch.

The process of logically segmenting a network into smaller groups to manage access and resources is also called subnetting. **Subnetting** is the process where a collection of IP addresses is broken into smaller subsets to better identify systems on the network and to distribute resources. Segregating collections of computers and devices into departments or areas is a common use of subnetting.

That way, all IP addresses in a particular range are assigned to a group with common objectives, location, or needs.

With IPv4, subnetting is necessary in very large networks because there are not enough available IPv4 addresses to assign one to each device. By subnetting, the pool of available addresses is expanded to allow more devices to have concurrent access to the network. This works by applying a subnet mask, which applies to a specific class of addressing. The address class differentiates how the address can be used. Use includes public, private, experimental, and multicast, which is when a single host sends messages to multiple network devices. Globally, IP addresses are managed by the Internet Assigned Numbers Authority (IANA).

IPv6 does not require the use of subnetting to expand the available pool of IP addresses for use in an organization, as there are many more addresses available under that scheme. Although not imposed by the need to make additional addresses available, logical network segregation is still important to allow for restriction of access, minimization of malware access, legal compliance, and other business reasons. IPv4 and IPv6 are covered in detail next.

DIFFERENCES BETWEEN INTERNET PROTOCOL VERSION 4 (IPV4) AND INTERNET PROTOCOL VERSION 6 (IPV6)

IPv4 has been in use as the predominant protocol on the Internet for nearly three decades. It was originally defined in 1981 in a technical specification known as RFC 791. IPv4 is a connectionless protocol that operates at the Network Layer (Layer 3) of the OSI model. IPv4 is the foundation of the TCP/IP protocol suite as it exists today.

IPv4 was designed with several assumptions in mind, many of which have been proved inaccurate, grossly overestimated, or simply nonapplicable. While IPv4 has been the predominant protocol on the Internet, a replacement is long overdue. Some of the key issues of concern are exhausted address space of only 32 bits, subnetting complexity, and a lack of integrated security. Some of these issues have been minimized with the advent of network address translation (NAT), classless inter-domain routing (CIDR), and Internet Protocol Security (IPSec). But in spite of these advancements, IPv4 is being replaced with IPv6.

IPv6 was defined in 1998 in RFC 2460. The new version was designed specifically as the successor to IPv4, mainly due to its dwindling availability, which was foreseen in the late 1990s. IPv6 uses a 128-bit address, which is significantly larger than IPv4's 32-bit

addressing. **FIGURE 3-17** compares an IPv4 address to an IPv6 address.




FIGURE 3-17 Comparing a typical IPv4 address to an IPv6 address.

■ NOTE

If an IPv6 address has one or more consecutive four-digit sections of all zeros, the sections of zeros can be dropped and replaced by a double colon. For example: 2001:0f58:0000:0000:0000:0000:1986:62af can be shortened to 2001:0f58::1986:62af. However, if there are

two sections of zero sets, only a single section can be replaced by double colons.

TECHNICAL TIP

The assignment of both IPv4 and IPv6 addresses is the domain of suborganizations of a global organization called the Internet Assigned Numbers Authority (IANA). To find out more about IANA, go to www.iana.org. In the United States, the suborganization is the American Registry for Internet Numbers (ARIN); for more information, visit www.arin.net.

Additionally, changes to subnetting, address assignment, and packet headers, and simpler routing processing, make IPv6 much preferred over its predecessor. Some benefits of IPv6 include:

- **Increased address space**—IPv6 supports 340 undecillion (that's 1 followed by 36 zeros) IP addresses for network devices.
- **More efficient routing**—The routing functionality of IPv6 has been enhanced.
- **Reduced management requirement**—A more robust protocol, IPv6 requires less management. In contrast, IPv4 required significant management due to requirements added after the introduction of the protocol.
- **Better quality of service (QoS) support for all types of applications**—Support for QoS is built into IPv6, whereas it was an add-on in IPv4.
- **Security**—IPv6 includes a native information security framework (IPSec) that provides for both data and control packets.

- **Plug-and-play configuration with or without DHCP**—IPv6 permits hosts to automatically configure during connection to an IPv6 network by querying the local routers with a multicast message. If the routers are configured correctly, a response with the appropriate configuration information is sent. If this mechanism is not appropriate for a specific environment or application, support is available for an IPv6 version of DHCP. It also supports static configurations.

Three mechanisms currently exist for the transition from IPv4 to IPv6. It is important to understand that when the IPv6 standard appeared, the expectation was that the two protocols would need to coexist in the network for 20 to 30 years, allowing for a gradual transition period. The migration strategies are:

- **Dual-stack**—With this transition solution, both IPv4 and IPv6 protocol stacks coexist in the same terminal or network equipment. This allows the network to communicate using both protocols, but it adds significant overhead to the network infrastructure.
- **Tunneling**—This solution allows two IPv6 hosts to create a tunnel for traffic between two IPv6 hosts through an IPv4 network, or vice versa, as IPv4 is phased out. This could add significant configuration overhead.
- **Translation**—This solution enables an IPv4 host to talk to an IPv6 host. This solution will require additional development, as currently IPv6 does not include this capability.

IPSec and IPv6

IPSec is a mandatory component for IPv6, and it is used to natively protect IPv6 data as it is sent over the network. The components of IPSec in IPv6 are not dramatically different from IPSec in IPv4, which industry has been using since the 1990s. In IPv6, IPSec uses the Authentication Header (AH) and the Encapsulating Security Payload (ESP) extension header.

The IPv6 IPSec is a set of Internet standards that uses cryptographic security services to provide the following:

- **Confidentiality**—IPSec traffic is encrypted and cannot be deciphered without the appropriate encryption key. This should be easier to use in conjunction with IPv6 because it is a native component of the protocol.
- **Data origin authentication**—IPv6 IPSec uses a cryptographic checksum that incorporates a shared encryption key so that the receiver can verify that it was actually sent by the apparent sender. This prevents spoofing of transactions.
- **Data integrity**—The cryptographic checksum can be used by the receiver to verify that the packet was not modified in transit.

You can find more information about IPv6 in RFC 2460 at <http://tools.ietf.org/html/rfc2460>.

THE TCP/IP SUITE

The TCP/IP suite is composed of a large number of protocols, defined in a series of RFCs. Each of the protocols in the TCP/IP suite provides a different function and together provide the functionality known as TCP/IP.

The TCP/IP suite got its name from the two main protocols in the suite: TCP (Transmission Control Protocol) and IP (Internet Protocol). TCP is responsible for providing reliable transmissions from one system to another, and IP is responsible for addressing and route selection. IP defines how computers communicate over a network.

From Harwood, Mike. *Comp TIA Network+ N10-004 Exam Prep*. Pearson Education.

EXAMPLES OF NETWORK INFRASTRUCTURES AND RELATED SECURITY CONCERNS

As you design a network, you need to evaluate every aspect in light of its security consequences. With limited budgets, personnel, and time, you must also minimize risk and maximize protection. Consider how each of the following network security aspects affects security for large corporations, small companies, and even home-based businesses.

Workgroups

A **workgroup** is a form of networking in which each computer is a peer or equal. Peers are equal in how much power or controlling authority any one system has over the other members of the same workgroup. All workgroup members are able to manage their own local resources and assets, but not those of any other workgroup member.

Workgroups are an excellent network design for very small environments, such as home family networks or very small companies. In most cases, a workgroup comprises fewer than 10 computers and rarely contains more than 20 computers. No single rule dictates the size of a workgroup. Instead, the administrative overhead of larger workgroups encourages network managers to move to a client/server configuration.

Workgroups do not have a central authority that controls or restricts network activity or resource access. Instead, each individual workgroup member makes the rules and restrictions over resources and assets. The security defined for one member does not apply to or affect any other computer in the workgroup.

FIGURE 3-18 shows a typical workgroup configuration. In this example, a switch interconnects the four desktop workgroup members, as well as an Internet connection device and a wireless access point (WAP). Additional clients can connect wirelessly via the access point or wired via a cable connecting to the switch.

Seven Domains of a Typical IT Infrastructure




FIGURE 3-18 An example of a typical workgroup.

Due to system-by-system–based security, a worker or a workgroup member needs to have a user account defined for each of the other workgroup members to access resources on those systems. Each of these accounts is technically a unique user account, even if it is created by using the same characters for the username and password.

This results in either several unique user accounts with different names and different passwords or several unique user accounts with the same name and same password. In either case, security is poor. In the former case, the user must remember several sets of credentials. This often results in the user writing down the credentials. In the latter case, an intruder need compromise only one set of credentials.

This lack of a central authentication authority is both a strength and weakness of workgroups. This characteristic is a strength in that each user of each computer can make choices about sharing resources with others. However, this is also a weakness because of the inconsistent levels of access.

Workgroups are easy to create. Often, the default network configuration of operating systems is to be a member of a workgroup. A new workgroup is created by defining a unique name on a computer. Once one computer names the workgroup, it now exists. Other computers become members of the new workgroup just by using the same workgroup name. Because workgroups lack a central authority, anyone can join or leave a workgroup at any time. This includes unauthorized systems owned by rogue employees or external parties.

Most workgroups use only basic resource-share protections, fail to use encrypted protocols, and are lax on monitoring for intrusions or intrusion attempts. Although it is possible to impose some security on workgroups, usually each workgroup member is configured individually. Fortunately, because workgroups are small, this does not represent a significant amount of effort.

SOHO Networks

SOHO (small office/home office) is a popular term that describes smaller networks commonly found in small businesses, often deployed in someone's home, garage, portable building, or leased office space. A SOHO environment can be a workgroup or a client/server network. Usually a SOHO network implies purposeful design with business and security in mind.

SOHO networks generally are more secure than a typical workgroup, usually because a manager or owner enforces network security. Security settings defined for each workgroup member are more likely to be consistent when the workgroup has a security administrator. Additionally, SOHO networks are more likely to employ security tools, such as antivirus software, firewalls, and auditing.

Client/Server Networks

A client/server network is a form of network where you designate some computers as servers and others as clients. Servers host resources shared with the network. Clients access resources and perform tasks. Users work from a client computer to interact with resources hosted by servers. In a client/server network, access is managed centrally from the servers. Thus, consistent security is easily imposed across all network members.

FIGURE 3-19 shows a possible basic layout of a client/server network. In this example, three servers host the resources, such as printers, Internet connectivity, and file storage shared with the network. Both wired and wireless clients are possible. Switches interconnect all nodes. Client/server networks are more likely to use hardware or appliance firewalls.

Possible Basic Layout of Client/ Server Network




FIGURE 3-19 An example of a typical client/server network.

Client/server networks also employ **single sign-on (SSO)**. SSO allows for a single but stronger set of credentials per user. With SSO, each user must perform authentication to gain access to the client and the network. Once the user has logged on, access control manages resource use. In other words, client/server authentication with SSO is often more complex than workgroup authentication—but it's more secure. Users need to log on only once, not every time they contact a resource host server.

Because of their complexity, client/server networks are invariably more secure than SOHO and workgroup networks. But complexity alone is not security. Instead, because they are more complex, client/server networks require more thorough design and planning. Security is an important aspect of infrastructure planning and thus becomes integrated into the network's design.

Client/server networks are not necessarily secure, because you can deploy a client/server network without any thought toward security. But most organizations understand that if they overlook network security, they are ensuring their ultimate technological downfall. Security is rarely excluded from the deployment process. And some networks are by nature more secure than others.

LAN Versus WAN

LAN stands for local area network. A LAN is a network within a limited geographic area. This means that a LAN network is located in a single physical location rather than spread across multiple locations. Some LANs are quite large, while others are very small. A more distinguishing characteristic of a LAN is that all of the segments or links of a LAN are owned and controlled by one organization. A LAN does not contain or use any leased or externally owned connections.

WAN stands for wide area network. A WAN is a network not limited by any geographic boundaries. This means that a WAN network can span a few city blocks, reach across the globe, and even extend into outer space. A distinguishing characteristic of a WAN is that it often uses leased or external connections and links. Most organizations rely on telecommunication service providers (often referred to as telcos) for WAN circuits and links to physical buildings and facilities, including the last-mile connection to the physical demarcation point. Both LAN and WAN networks can be secure or insecure. With a LAN, the owner of the network has the sole responsibility of ensuring that security is enforced. With a WAN, the leasing entity must select a telco that has a secure WAN infrastructure and incorporate service level agreements (SLAs) that define the level of service and performance that is to be provided on a monthly basis for the customer. In most cases, WAN data is secure only if the data sent across leased lines is encrypted before transmission. This service is the responsibility of the data owner, not the telecommunications service provider, unless this option is offered as a value-added service.

Thin Clients and Terminal Services

Thin client computing, also known as **terminal services**, is an old computing idea that has made a comeback in the modern era. In the early days of computers, the main computing core, commonly called a mainframe, was controlled through an interface called a terminal. The terminal was nothing more than a video screen (usually monochrome) and a keyboard. The terminal had no local processing or storage capabilities. All activities took place on the mainframe, and the results appeared on the screen of the terminal.

With the advent of personal computers (PCs), a computer at a worker's desk offered local processing and storage capabilities. These PCs became the clients of client/server computers. Modern networking environments can offer a wide range of options for end users. Fully capable PCs used as workstations or client systems are the most common. PCs can run thin-client software, which emulates the terminal system of the past. That means they perform all tasks on the server or mainframe system and use the PC only as a display screen with a keyboard and mouse. Even modern thin-client terminals can connect into a server or mainframe without using a full PC.

Remote Control, Remote Access, and VPN

Remote control is the ability to use a local computer system to remotely take over control of another computer over a network connection. In a way, this is the application of the thin-client concept on a modern fully capable workstation to simulate working against a mainframe or to virtualize your physical presence.

With remote control connection, the local keyboard and mouse control a remote system. This looks and feels as if you are physically present at the keyboard of the remote system, which could be located in another city or even on the other side of the world. Every action you perform locally takes place at that remote computer via the remote control connection. The only limitations are the speed of the intermediary network link and the inability to physically insert or remove media, such as a flash drive, and use peripherals, such as a printer.

You might consider remote control as a form of software-based thin client or terminal client. In fact, many thin client and terminal client products are sold as remote control solutions.

Many modern operating systems include remote control features, such as the Remote Desktop feature found in most versions of Windows. Once enabled, a Remote Desktop connection remotely controls another Windows system from across the network.

Remote access is different from remote control. A remote access link enables access to network resources using a WAN link to connect to the geographically distant network. In effect, remote access creates a local network link for a system not physically local to the network. Over a remote access connection, a client system can technically perform all the same tasks as a locally connected client, with the only

difference being the speed of the connection. Network administrators can impose restrictions on which resources and services a remote access client can use.

Remote access originally took place over dial-up telephone links using modems. Today, remote access encompasses a variety of connection types, including ISDN (Integrated Services Digital Network), DSL, cable modem, satellite, mobile broadband, and more.

In most cases, a remote access connection links from a remote client back to a primary network. A **remote access server (RAS)** accepts the inbound connection from the remote client. Once the connection goes through, the remote client now interacts with the network as if it were locally connected.

Another variant of remote connections is the virtual private network (VPN). A VPN is a form of network connection created over other network connections. In most cases, a VPN link connects a remote system and a LAN, but only after a normal network connection links to an intermediary network.

In **FIGURE 3-20**, a LAN has a normal network connection to the Internet, and a remote client has established a normal network connection to the Internet. These two connections work independently of each other. The LAN's connection is usually a permanent or dedicated connection supporting both inbound and outbound activities with the Internet. The remote client's connection to the Internet can be dedicated or nondedicated. In the latter case, the connection precedes the VPN creation. Once both endpoints of the future VPN link have a connection to the intermediary network (in this example, the Internet), then the VPN exists.

Before VPN Connection Established




FIGURE 3-20 Common resource access connections.

In **FIGURE 3-21**, a new network connection connects the remote client to the LAN across the intermediary network. This new network connection is the VPN.

After VPN Connection Established




FIGURE 3-21 A VPN established between a remote client and a LAN over the Internet.

A VPN is a mechanism to establish a remote access connection across an intermediary network, often the Internet. VPNs allow for cheap long-distance connections over the Internet, as both endpoints only need a local Internet link. The Internet itself serves as a “free” long-distance carrier.

A VPN uses tunneling or encapsulation protocols. Tunneling protocols encase the original network protocol so that it can traverse

the intermediary network. In many cases, the tunneling protocol employs encryption so that the original data securely traverses the intermediary network.

Boundary Networks

A **boundary network** is a subnetwork, or subnet, positioned on the edge of a LAN. A boundary network isolates certain activities, such as programming or research and development, from the internal environment (private LANs), which does not need to be externally accessible. External users can also access resources hosted in a boundary network. Such a subnet is known as a **demilitarized zone (DMZ)**.

A DMZ is a boundary network that hosts resource servers for the public Internet. An **extranet** is a boundary network that hosts resource servers for a limited and controlled group of external users, such as business partners, suppliers, distributors, contractors, and others.

A DMZ subnet is a common network design component used to host public information services, such as websites. A DMZ allows for anyone on the public Internet to access resources. But at the same time, the DMZ provides some filtering of obvious malicious traffic and prevents Internet users from accessing the private LAN.

An extranet subnet provides companies that need to exchange data with external partners a safe place for that activity. An extranet is secured so only the intended external users can access it. Often, accessing an extranet requires VPN connections. An extranet configuration keeps public users from the Internet out and keeps the external partners out of the private LAN. An extranet example you may be familiar with is tracking the progress of a package by accessing a delivery company's website and entering in your package's unique tracking number.

You deploy DMZs and extranet subnets in different ways. Two of the most common designs are a screened subnet using a three-homed firewall (**FIGURE 3-22**) and N-tier deployment (**FIGURE 3-23**). A screened subnet using a three-homed firewall creates a network segment for the private LAN and a network segment for the DMZ or extranet. The three-homed firewall serves as a filtering device, as well as a router, to control traffic entering these two segments.

VPN established between a remote client and a LAN over the Internet




FIGURE 3-22 A DMZ or extranet deployed as a screened subnet with a threehomed firewall.




FIGURE 3-23 A DMZ or extranet deployed in an N-tier configuration.

The N-tier deployment creates a series of subnets separated by firewalls. The DMZ or extranet subnet serves as a buffer network between the Internet and the private LAN. The “N” in N-tier is a variable to be replaced with whatever number of tiers you actually build in your network.

CHAPTER SUMMARY

There are several common network topologies, all with distinct benefits and negative considerations. Knowing the organization's goals, the ability to employ each physical design, and how the network should be configured logically can aid in making the right topology decision. Although some small organizations could be served by the bus topology, it is usually reserved for very small installations. A mesh topology offers the most redundancy, but it is also the most expensive.

Although most topologies can be implemented to follow specific logical topographical frameworks, it can be less secure to implement some logical designs on physical networks that require all devices to know and recognize all other devices. It can also be less secure where flow from one system to another is precluded by physical connectivity to the next device, as in bus and ring implementations. Even if those devices should not communicate with the next device in the network, the device is still known to the computer; therefore, an attempt may be made to access it. This is why security is a necessary consideration prior to constructing a network, and it continues to be an important element as that network is managed and maintained.

This chapter also took a detailed look at IPv4 and IPv6, as well as various network infrastructures.

KEY CONCEPTS AND TERMS

Active hub

Appliance

Bidirectional

Blacklist

Boundary network

Bridge

Router

Bus topology

Clients

Client/server network

Concentrators

Demilitarized zone (DMZ)

Dumb hub

Extranet

Gateway

Hub

Internet service providers (ISPs)

Log

Logical topology

Mesh topology

Modem

Passive hub

Physical topology

Protocol converter

Punch panel

Punchdown block

Remote access


Remote access server (RAS)
Remote control
Repeater
Ring topology
Route
Router
Scaling
Single sign-on (SSO)
Smart hub
Sniffer
SOHO (small office/home office)
Source routing bridge
Star topology
Subnetting
Switch
Terminal services
Terminators
Thin client computing
Token
Topology
Transparent bridge
Unidirectional
Virtual local area networks (VLANs)
Wired topologies
Wireless topology
Workgroup

CHAPTER 3 ASSESSMENT

1. A protocol converter is another name for what device?
 - A. Bridge
 - B. Gateway
 - C. Hub
 - D. Router
2. Which of the following is a downside of using a workgroup for business network activities?
 - A. Workgroups do not have a central authority that controls or restricts network activity.
 - B. Workgroups offer each member the ability to secure their resources from the rest of the devices.
 - C. Workgroups allow the security defined for one member to apply to that member only.
 - D. Workgroups are often more cost effective for small companies and SOHO installations.
3. Which of the following is a type of passive hub?
 - A. Punch panel
 - B. Brouter
 - C. Repeater
 - D. Gateway
4. Which of the following is a type of smart hub?
 - A. Brouter
 - B. Concentrator
 - C. Firewall
 - D. Gateway
5. Which device works at Layer 2 (Data Link Layer) and uses MAC addresses to differentiate traffic?

- A. Router
 - B. Firewall
 - C. Switch
 - D. Punch panel
6. What is the term for the unique address identifying hardware assigned by the manufacturer under the guidance of the FCC?
- A. MAC address
 - B. IPv4 address
 - C. Loopback address
 - D. IPv6 address
7. Which of the following is true about a star topology?
- A. It requires less wiring than a traditional bus network.
 - B. It is more secure than a mesh network.
 - C. It requires less wiring than a wireless network.
 - D. It is more fault-tolerant than a bus network.
8. Which of the following is *not* true of a logical network?
- A. It is possible for a network to have the same logical and physical topologies.
 - B. It is possible for computers physically located side by side to be on different subnets in the logical network.
 - C. It is possible for the physical network to be a star and the logical network to be a ring.
 - D. It is possible for the physical network to employ a gateway.
9. Which of the following networks provides the most redundancy?
- A. Star
 - B. Mesh
 - C. Bus
 - D. Ring
10. Which of the following network topologies requires the use of terminators?

- A. Bus
- B. Star
- C. Ring
- D. Mesh



© Digital_Art/Shutterstock

Network Design Considerations

NETWORK SECURITY IS ESSENTIAL to every organization, from an individual computer at home to a multinational corporation. The foundations of security apply universally, no matter the size, purpose, or function of the computers and network.

Implementing network security requires you to work through the installation, deployment, and configuration of all components, but it must begin with thorough planning and design. That means considering layered defenses, proper use of protocols, communication management, intrusion detection and prevention systems, system hardening, and more. You can significantly improve the security of your organization's computer systems and network based on some common, often simple, principles.

Following the suggestions in this chapter will reduce the risk of system compromise from accident, oversight, Mother Nature, or malicious intent.

CHAPTER 4 TOPICS

This chapter covers the following topics and concepts:

- Network design and defense in depth
- Protocols
- Common types of addressing
- How to control communication pathways
- Intrusion detection systems (IDSs) and intrusion prevention systems (IPSs)
- How to harden systems
- Which method to use for selecting equipment
- Authentication, authorization, and accounting
- Communication encryption
- The best architecture for local hosts only or remote and mobile hosts
- Examples of redundancy
- What constitutes node security
- How to perform a risk assessment
- Network zones of risk

CHAPTER 4 GOALS

When you complete this chapter, you will be able to:

- Describe elements of network security design

- Compare and contrast public and private addressing, as well as static and dynamic addressing
- Understand the difference between an IDS and an IPS
- State the importance of system hardening
- Describe why authentication, authorization, accounting, and encryption are essential for network security
- Identify the security issues of local hosts, as well as remote and mobile hosts
- Define the elements of node security
- Assign values to network risks

NETWORK DESIGN AND DEFENSE IN DEPTH

Every network is different. However, common security principles apply to every network, regardless of its unique elements. One of these common principles is secure network design. Secure network design embeds core protections and improvements into an IT infrastructure before it is implemented. Design comes from planning. Planning comes from sufficient knowledge and understanding of business needs, user behavior, and network vulnerabilities.

Common security goals include confidentiality, integrity, availability, privacy, authentication, authorization, nonrepudiation, and accounting. To accomplish these goals efficiently, informed planning assists you in designing the network before deployment. Understanding common network vulnerabilities and prescribing solutions aids in the design of a secure network.

Designing network security is neither simple, nor is it a short-term task. Thorough network security design must include adequate research, thorough planning, and extensive **modeling** and testing. The process of security design must evaluate a wide range of technologies performing an astounding number of functions.

Ultimately, good network security design produces a blueprint to guide the construction of a securely functioning network infrastructure. The blueprint is the foundation for your organization's security policy. Most network designs have limitations. These limitations include budget, internal politics, regulations, standards,

and industry practices. Network design should focus on providing the best security possible within these prescribed budgetary boundaries.

Achieving Defense in Depth through Layering

An underlying fundamental of network security design is that no security solution is perfect. Any single security protection, countermeasure, and safeguard is insufficient. Hackers will use some method, technique, or exploit to bypass, evade, or render useless a security protection. The potential concerns include placement, programming flaws, default settings, maximum values, processing capabilities, memory capacity, backdoors, malicious code, social engineering, and physical attacks. This list is not exhaustive, but it represents the key issues. In theory, no security solutions are sufficient and complete.

Thus, you need to use multiple security components. This is known as defense in depth or multiple layers of defense (**FIGURE 4-1**). If you follow a defense-in-depth design concept, numerous safeguards will protect each asset. As one defense tool interlocks with another, they overlap and improve the overall security. The strengths and benefits of one countermeasure supplement or compensate for the weaknesses and limitations of another.




FIGURE 4-1 An example of defense in depth around an asset.

Defense in depth leads many security professionals to two additional guidelines. First, avoid single points of failure. Second, divide and conquer. A **single point of failure** is any element, component, or aspect of a system that could lead to failure or compromise of the entire system. Divide and conquer is the process of separating a large project into multiple smaller and more manageable pieces.

Avoiding single points of failure must take place on multiple fronts. A hacker needs only a single flaw or weakness to exploit a target. Efforts should focus on finding and eliminating as many vulnerabilities as possible to remove the single points hackers seek to exploit.

Good design filters every user interaction with an asset multiple times. This filtering should include authentication, authorization, content filtering, and context filtering. Reliance upon a single filter or check system is a form of a single point of failure. Always assume that any one service or function is flawed or will fail.

Effective network design monitors and examines all activities against an asset using multiple techniques. This could include object auditing, server monitoring, client monitoring, network monitoring, and so on. Only using a single monitoring viewpoint could be a single point of failure. Everyone has seen video footage from a single perspective that guides the viewer into seeing or believing one thing, but from another camera angle the truth, the trick, or an alternate explanation becomes clear.

Divide and conquer is not just a tactic for waging actual war; it is also a tactic in the war against network security breaches. By dividing

up a larger project or task into manageable components, you can focus on and care for each component to ensure accuracy and completeness in addressing network security concerns. Attempting to tackle the security of a network as a whole is often a recipe for disaster. Evaluating the big picture is always a good idea, but working exclusively on the whole may lead to overlooking details or missing subtle nuances only perceived upon close detailed inspection.

A layered security approach throughout the IT infrastructure works best: slow, methodical, compartmentalized, and thorough. Properly designed network security should support timely delivery of information and adequate response during transactions. A properly secured network provides reliable and stable communications. Well-designed security adapts to changing conditions and anticipates future growth and expansion.

One method to reduce compromise by hackers is to keep them from finding your network as a target. By staying offline and only using trusted communication pathways, your organization can avoid significant levels of risk. However, this form of technological hide and seek is not perfect, nor does it eliminate all issues. External hackers might not be able to hack a non-Internet-connected intranet from the outside; however, the risk from disgruntled employees and other internal users is still present.

The idea of hiding from danger is commonly known as security through obscurity. It is true that if you are not found, you cannot be attacked; yet, the issue is often a false hope if obscurity—rather than actual countermeasures and safeguards—protects the network. Being obscure or difficult to locate may be a good thing, but it's not a

reliable form of security. Use only direct and real security defenses when you are designing network security.

Not all attacks are targeted attacks. Although many attacks against a system or network are directed at that system by a hacker, other attacks are blanket attacks unleashed on the Internet in the hope of causing loss, harm, and destruction to any system susceptible to that particular attack because of inherent vulnerabilities in the system. Even when an attack is targeted at a specific system or network, there can be collateral damage, depending on how the attack was scripted.

Security is essential to the long-term survival of any modern organization. Without security, logical, physical, and social breaches would render most companies vulnerable to failure. But security cannot work without balance. You can over-secure an infrastructure to the point where security interferes with work tasks. Usability and security must be in balance. Usability will not survive long without security, and too much security can cripple usability. Good network security design balances security and usability.

Planning for Scalability

One goal of most organizations is to expand and grow. They often seek to attract new customers, support more clients, sell more products, offer more services, make more money, and so forth. But growth can be a two-edged sword. While growth can lead to a more reliable and assured future, it can also cause growing pains. For any organization, growing pains occur when the existing infrastructure, facilities, and even personnel are pushed to the limit or beyond to support the additional workload caused by growth.

Growth can be expected, unexpected, gradual, or abrupt. A proper network design process evaluates and predicts potential growth scenarios and plans contingencies for each. One contingency for growth is to build additional capacity into the current infrastructure. If slow growth is expected, then 20 percent additional capacity may be sufficient, while rapid growth may consume 50 percent additional capacity in a short time (**FIGURE 4-2**).




FIGURE 4-2 Rate of growth used to predict needed additional capacity.

Growing too fast is as much of a burden as shrinking. As an organization stretches beyond its capacity to support, sell, create, maintain, respond, and produce, small problems quickly snowball into avalanches. Steady, controlled, limited growth is often a method of ensuring long-term viability and stability. This is true in general business management, and it's true in network security design.

When an organization shrinks, there may be systems, processes, and resources that are no longer used or monitored. If these items are not secured, a hacker can use this as an entry point into the network and potentially remain undetected as they scan the network for items they deem valuable. If an item is connected to the network, it should be secured and monitored. If the item is no longer necessary, it should be disconnected from the network and stored for later use or disposed of properly (such as proper destruction of hard drives in systems or printers that may contain sensitive data).

During the network design phase, consider the scalability of all technologies you select for deployment. Does a component or system have a maximum value or limitation it will quickly reach? Can the component or system expand without compromising efficiency, cost effectiveness, and security? Will the component or system need replacement by a scalable solution once moderate growth occurs? If so, why not use the scalable solution now? Planning for growth will reduce problems associated with outgrowing function and security capacity.

Working with Senior Management

No security endeavor will succeed without the active involvement of senior management. In fact, without its explicit approval and support, any security effort is likely doomed to fail. Senior management has the responsibility to dictate the strategic goals and plans of the organizations and, hence, its IT infrastructure and integrated security. Senior management must approve budgetary funding, encourage compliance, and support security, even when problems occur. Senior management sets the corporate expectations that shape the corporate culture of an organization.

Throughout the entire design and implementation process for network security, senior management monitors and approves progress reports. Senior management steers the organization and its security planning through the changing business environment. But secure network design is not only about following the leader; it is also about integrating every employee into the overall security design process. Security is the responsibility of everyone in the organization, not just managers and executives.

The elements of secure network design touch on every aspect of an IT infrastructure. This includes hosts, nodes, communications, encryption, local and remote systems, redundancy, and more. Often, the process of designing security starts by focusing a central or core element found throughout the infrastructure. Important distributed core components are networking protocols.

PROTOCOLS

A significant portion of network security is about making the right technology choices without falling into easy traps or defaults. One common trap is to continue doing the same thing or using the same product. You need to reevaluate old technologies and existing solutions on a regular basis. Most organizations choose to perform a security design evaluation annually. When performing a security evaluation, rethink every aspect of the infrastructure, including network protocols.

Most networks use Transmission Control Protocol/Internet Protocol (TCP/IP) as the primary network protocol suite. TCP/IP is a suite, as it is two communication protocols working together to provide enhanced speed and consistency. Specifically, most networks still use Internet Protocol version 4 (IPv4), as opposed to Internet Protocol version 6 (IPv6). Using IPv4 is not an open invitation for hackers, but it does have numerous commonly exploited weaknesses and concerns. IPv4 typically defaults to a plaintext form of transmission, while IPv6 can be set by default to encrypt transmissions. IPv4 can be encrypted using IP Security (IPSec) or other virtual private network (VPN) protocols.

Other issues to consider include:

- Is the current protocol easy to compromise?
- Are there numerous exploits for this protocol available for novice hackers?
- Can encryption be applied?

- Is the process of adding encryption complex or costly?
- Will encryption interfere with other technologies, such as IPSec and network address translation (NAT)?
- Is there an alternative or replacement available?
- Is the alternative backward compatible?
- Is the alternative supported by all current hosts and nodes?

This is only a partial list of questions you need to consider when assessing the currently deployed networking protocol for possible replacement. The point is to give serious consideration to this issue on a regular basis. Most of the elements of a network's security are based on the protocol in use. If the protocol has changed or improved, this could cause sweeping changes throughout the production environment—most importantly in the security of that environment.

As a general rule at the time of this writing, if most hosts and software are less than five years old, then upgrading to IPv6 is likely possible with minimal complication. However, if IPv4 with IPSec or other forms of encryption are functioning well within performance and security parameters, there's no strong need to upgrade to IPv6.

When you consider a protocol upgrade, you'll need to thoroughly research and test every aspect of the production and security environment to confirm compatibility with IPv6. Any transition is going to have some hurdles, and certainly switching the main network protocol is a candidate for major hurdles. Perform the rollout of a protocol change in stages, only after **piloting** in a lab, and consider running dual protocols for a transition period.

TCP/IP, or at least IP, is not the only protocol in use on most networks. Every single protocol across every theoretical layer of the OSI (Open Systems Interconnection) model network protocol stack

needs reevaluation on a regular basis. Do you still want to continue using **Simple Mail Transfer Protocol (SMTP)** and **Post Office Protocol (POP)**? What about **File Transfer Protocol (FTP)**, **Network News Transfer Protocol (NNTP)**, and **Telnet**? If a protocol operates in plaintext, consider using a protocol with native encryption, or investigate the possibility of encapsulating it inside an encrypting tunneling protocol, such as IPsec or Secure Sockets Layer/Transport Layer Security (SSL/TLS).

Are AppleTalk, Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), Systems Network Architecture (SNA), NetBios Extended User Interface (NetBEUI), or other protocols still present on the network? Are they still necessary? Have the older systems requiring these protocols been replaced or removed? Can newer secure alternatives replace these older and insecure protocols? Can any system still using a legacy protocol be replaced to gain host security and remove the protocol?

Don't let tradition, personal bias, or sunk cost get in the way of making a smart security design decision. The cost of an item is not a reason to keep using it if that continued use poses more security risks and concerns than benefits. Just because something has always been done a certain way does not mean it should continue to be done that way, especially if the old way was insecure. Personal bias in terms of likes, dislikes, comfort, and familiarity are not good business reasons to avoid moving on to more secure solutions.

Sunk cost is the money or investment already made in the past, as opposed to prospective costs that are investments likely in the future. Often, those managers overly concerned with wasting resources already invested will make poor decisions. Avoiding loss is a good

principle of business, but just because you have already spent some money, time, and effort does not mean you should continue on a set path. Often, throwing good money after bad is compounding the loss due to sunk cost. If the future benefit of an existing system, solution, or product is not assured, then no amount of sunk cost justifies its continued use. In most cases, you should analyze only future costs and benefits in making a decision. Already incurred expense (sunk cost) should not influence future choices.

COMMON TYPES OF ADDRESSING

Addressing is the assignment of a logical numbering system to the hosts on a network for the purposes of efficient traffic routing.

Addressing is more than just a system imposed by a network server; it is often a means to control traffic. Traffic managing through routing and traffic filtering are possible through the use of logical addresses.

The most common protocols in use worldwide are TCP/IP. This network protocol suite dictates the most common addressing scheme. The addressing schemes of IPv4 and IPv6 are quite different. Some common elements, security concerns, and management techniques remain consistent, however.

Internal IP addresses can be public addresses, private addresses, or a mixture of both. A public address is an address issued by the IANA, monitored by RIRs, and leased directly through ISPs. The

Internet Assigned Numbers Authority (IANA)

(<http://www.iana.org/>) is the entity responsible for global coordination of IP addressing, DNS root, and other Internet protocol resources. A **Regional Internet Registry (RIR)** is one of five regional organizations that oversee and monitor the use and assignment of IP addresses (both IPv4 and IPv6). An Internet service provider (ISP) may randomly assign or semipermanently lease an IP address to an individual or organization. Public addresses are those obtained from an ISP.

■ NOTE

In the past, it was possible to purchase or own IP addresses, specifically large groups or an entire class of addresses. However, this practice is mostly no longer possible, not because ownership of IP addresses is prohibited, but because of the lack of available IP addresses to sell. Many of the original Class A and Class B subnet owners still own, control, and use their purchased addresses. Some of these are actually ISPs that now lease out subsets of their owned IP address ranges. Today, most public IP addresses are leased rather than sold or owned.

A public address also implies that it communicates directly with resources on the Internet. The Internet uses only public addresses. Without a public address, it is impossible to communicate to or receive responses from an Internet-hosted resource.

Public addresses are assigned from Class A, B, and C ranges of the IPv4 address spectrum (as Class D and E are reserved for multicasting and experimentation, respectively). Public addresses for IPv6 are most of the 2^{128} addresses, except for the `fc00::/7` address block.

In a technical specification called RFC 4193, IANA set aside the `fc00::/7` address block for use as private addresses for IPv6, similar to RFC 1918 for IPv4. RFC 1918, which provides for private IPv4 addresses, sets aside three class ranges for private use:

- **Class A**—10.0.0.0–10.255.255.255/8 (1 Class A network)

- **Class B**—172.16.0.0–172.31.255.255/12 (16 Class B networks)
- **Class C**—192.168.0.0–192.168.255.255/16 (256 Class C networks)

TECHNICAL TIP

A DHCP reservation is the pre-assignment of a specific IP address to a host by reserving it, using the target host's MAC address. Reservations ensure that the same address is always issued to a specific host. It can also simulate static addressing, but retain centralized control of address assignment.

A private address is used only within a private network. Individuals and organizations without approval or a fee from an outside entity can use private addresses. However, using private addresses requires NAT services to communicate with Internet resources. All Internet routers automatically drop any packet with a private address in its header. Private addresses serve as a basic isolation security measure, as external entities with public addresses cannot directly communicate with internal privately addressed hosts. But a NAT server allows communication with Internet resources.

You should review your organization's choice to use private or public addresses internally. The issue is not only about saving money. Private addresses are free, while public addresses are usually leased. Private addresses require translation, while public addresses do not. Private addresses are natively isolated from the Internet, while public addresses are not. It is even possible to mix private and public addresses on an intranet.

Another addressing concern is whether to employ static or dynamic addressing. Static addressing pre-assigns a specific IP address to each host, while dynamic addressing hands out IP addresses to hosts from a pool. Dynamic addressing does not guarantee that a host will always have the same address assigned to it, unless a reservation is created for the host.

Static addressing typically requires that the IP address be configured on each individual host. This ensures that a host always uses the same IP address. However, if changes to the network configuration or topology occur, manual changes to IP addresses on a host-by-host basis involve a significant amount of additional administrative overhead.

Because of this, most organizations use dynamic assignment, typically using a Dynamic Host Configuration Protocol (DHCP) system. If static addressing is preferred, then DHCP reservations can simulate static addressing while maintaining centralized control. With reservation-based static addressing, changes can be made by editing the reservations on the DHCP server, without needing to manually adjust each host individually.

When addresses are assigned dynamically, it may be possible for a rogue system to come online and receive a valid IP address just by asking. If addresses are assigned statically, then the attacker will need to discover a valid but unused IP address and manually configure his or her system to use it. Similarly, if DHCP reservations are used, the attacker will either statically assign his or her own address manually or spoof a Media Access Control (MAC) address to “borrow” an IP address from another offline system.

IPv6

Approximately 25 percent of global Internet addressing is done with IPv6 addresses, and a small percentage of exploits takes advantage of it. Therefore, it is important for the security professional to stay up-to-date, and a familiarity with IPv6 is important. IPv6 varies from IPv4 in two important ways: The packet headers differ in structure, and an IPv6 address has a different size and construction from that of an IPv4 address. There are also a handful of transition mechanisms that allow IPv4 to ride inside of IPv6 packets and vice versa. The transition mechanisms make the job of going between IPv4 and IPv6 easier, but it can be a nightmare for security and policy enforcement.

An IPv6 address consists of 128 bits, as compared to the 32 bits of the IPv4 address. Addresses are classified into various types for applications in the major addressing and routing methodologies: unicast, multicast, and anycast networking. In each of these, various address formats are recognized by logically dividing the 128 address bits into bit groups and establishing rules for associating the values of these bit groups with special addressing features. There is no Class A or Class B in IPv6, as there is in IPv4.

- A unicast address identifies a single network interface. The Internet Protocol delivers packets sent to a unicast address to that specific interface.
- An anycast address is assigned to a group of interfaces, usually belonging to different nodes. A packet sent to an anycast address is delivered to just one of the member interfaces, typically the nearest host, according to the routing protocol's definition of distance. Anycast addresses cannot be easily identified, they have the same format as unicast addresses, and they differ only by their presence in the network at multiple points. Almost any unicast address can be employed as an anycast address.

- A multicast address is also used by multiple hosts, which acquire the multicast address destination by participating in the multicast distribution protocol among the network routers. A packet that is sent to a multicast address is delivered to all interfaces that have joined the corresponding multicast group.

IPv6 does not implement broadcast addressing. Broadcast's traditional role is subsumed by multicast addressing to the all-nodes link-local multicast group ff02::1. However, the use of the all-nodes group is not recommended, and most IPv6 protocols use a dedicated link-local multicast group to avoid disturbing every interface in the network.

Address management is an important concern of network security. Another concern is the management of communication pathways.

CONTROLLING COMMUNICATION PATHWAYS

Controlling the flow of information is a key element of network security. This involves ensuring that data travels along pathways isolated, secured, and controlled, and not along pathways that are public, insecure, and uncontrolled. Part of communication pathway control is about topology selection, but it is also about router configuration, encrypted protocols, physical access management, and filtering.

Router Configuration

Routers are the primary network devices administrators use to control the pathways that communications traverse. Failing to design router configuration and deployment with security in mind is a serious oversight. Routers make real-time determinations of the best available path to a destination. However, the information available to a router to make those decisions can be true and accurate or incorrect, falsified, and misleading.

Secure network design includes protections for routers, routing protocols, and routing information. Physical isolation of a router is important to ensure that only authorized router administrators can access the device. Failing to protect routers physically means that the logical activities and the resulting routes selected will not be trustworthy.

Routers employ routing protocols to exchange information about routes and connected pathways. This information calculates the best path to guide a packet toward its destination. Depending upon the make and model of router, the routing protocol, and the related configurations, a hacker can spoof or manipulate routing data through false Internet Control Message Protocol (ICMP) type 5 redirect messages. Configure routers to accept routing information only from other known routers through authentication of the source. Consider also encrypting all communications between routers.

The path specified by a router is an opportunity for a hacker to steal or hijack information. The hacker passes bad paths (spoofing), so packets are sent along the compromised route and are easily copied or captured. For this reason, router anti-spoofing techniques should be used. **Router anti-spoofing** techniques including source-

address validation and dropping any packets that have a false source address. One of the easiest ways to prevent spoofed packets is to set up firewall or router rules that examine incoming packets. **Router access control lists** are an effective tool to limit network access and traffic, and to increase network performance. An access list is simply a listing of what may or may not have access to a segment of a network behind the router.

Encrypted Protocols

Encrypted protocols are another important aspect of communication pathway security. Even the best design, proper installation, and reasonable physical isolation are not guarantees that a wired or wireless communication channel will not be the target of an eavesdropping, interception, or man-in-the-middle attack. Assuming that physical access is under your control all the time is naïve. The possibility of an internal malicious entity or of the planting of a socially engineered listening device always exists.

To thwart eavesdropping and related attacks based on eavesdropping, you should encrypt all traffic over a network communication link. This especially applies to any traffic traversing a network segment physically accessible from outside your organization's facilities. But it also applies to physically isolated and internal connections, as well. Compromise of every physical connection is always possible, so the best defense against content eavesdropping is encryption.

Physical Access Management

Physical access management should always be a part of communication pathway security. Even with encrypted protocols, hackers can gain significant information by eavesdropping on a network segment. If you encrypt every single packet (which is often not the case), eavesdropping can still glean a wide variety of information about the protected communications.

Such gleaned information can include a count of the number and size of packets. This can allow for estimating the size of the payload delivered, which in turn can extrapolate the likely type of data, such as email transmission, web surfing, file exchange, or database synchronization.

Eavesdropping can also glean the identity of the endpoints of the secured communication. If the transaction is using transport mode encryption, then the endpoints are the actual sending and receiving hosts. If the transaction is using tunnel mode encryption, then the endpoints are either both VPN gateways or one end is a remote host.

Eavesdropping can also glean the general identity or purpose of each endpoint discovered, based on the timing and volume of traffic sent to and from each discovered endpoint. This can allow an outside user to reliably predict which endpoint is a server, a client, or a VPN gateway.

Generally, data gathering through eavesdropping on communications, whether encrypted or not, is known as traffic and trend analysis. Such analysis can reveal many important details about internal processes and the importance, value, or criticality of systems. Thus, even with encryption, prevention of physical access to communication cables and wireless signals is paramount.

Filtering

Filtering is another important part of communication pathway security. The movement of data among departments, subnets, WAN-connected LANs, and the Internet requires that you monitor and filter communications to prevent violations of disclosure, intrusion, and malicious code infection.

Covert channels are a risk for many organizations in communication pathway security and control. Covert channels are pathways of communication unknown to or uncontrolled by security systems or personnel. Covert channels, whether timing- or storage-based, can leak information out or bring malicious content into the network.

Uncovering covert channels is difficult because the very nature of these channels is to remain unknown and unseen. The best defenses against covert channels include intrusion detection systems and intrusion prevention systems, as well as thoroughly watching all aspects of an IT infrastructure for aberrant or abnormal events of any type. To aid in effective implementation, evaluate the protections you'll need for traffic, both inbound and outbound, while planning and designing communication pathway security. Consideration should also be given to managing internal-only, external-only, or border-crossing communications.

You'll need to examine inbound traffic by asking several key questions: Is the inbound communication a response to a previous request from an internal entity, or is it a communication that originates from an outside source? Responses are often allowed, unless the initial request was for a resource that's off limits. Restrictions along these lines include blocked protocols, IP addresses

or domain names, unauthorized services and applications, or users without sufficient or correct authorization.

If the communication has external origins, rather than being a response, is the communication generally allowed or not? If the communication is for a resource offered to the public, then the communication could be allowed. However, if no public resources exist, the communication is more likely unauthorized.

Is the source address in the communication from a known or unknown location, and if known is it known to be malicious or questionable? If the latter, then blocking the traffic is more likely the proper security stance. Is the traffic obviously spoofed? Does the traffic match any known malicious patterns, have any construction anomalies, or have questionable content? Can it otherwise be classified as abnormal or atypical? In most of these cases, the packet should be dropped rather than allowed to continue on to its claimed destination.

An area of concern for communication pathway security is the difference between traffic management based on whether the traffic is internal-only, external-only, or border-crossing. Generally, internal-only traffic is more trustworthy than any other form of traffic. In most cases, internal traffic originates from a trusted internal host and terminates at a trusted internal host. However, because the possibility always exists of a rogue internal host or a malicious insider, blindly trusting traffic just because it originates internally is a security blunder.

TECHNICAL TIP

With the proliferation of telecommuting and outsourcing, external users are often valid users as well. Realizing this, you must design networks to support inbound connections from valid external users. This can involve traditional remote access connections, but more likely will use VPN links. VPNs provide connections for remote users that function like local connections and encryptions to keep the content of those connections confidential.

A good practice is to treat all traffic with caution. Trust nothing until it is proved to comply with security policy and not to match any known malicious patterns. Monitoring and filtering of internal communications is as important as monitoring external and border crossing communications.

You should subject outbound traffic to the same investigations and analysis as inbound. Outbound traffic may appear legitimate if it was initiated from within the network, but it could be from a hacker who infiltrated the network and is sending out the company's trade secrets or from an employee you do not realize is disgruntled. Your best chance to learn of their intention is to secure the traffic and the logs and review the logs regularly. Does the outbound communication take place over an abnormal protocol or port? Does it attempt to communicate with a blocked or prohibited host or service? Is the traffic spoofed, does it have abnormal time stamps, is it a clone of another packet, or is it part of a flood?

Fortunately, inbound and outbound traffic filtering is the primary function and purpose of firewalls using ingress- and egress-focused filters. Secure communications pathway management often requires

the use of firewalls. Firewalls are an essential element in secure network design.

Border-crossing communications are those that either leave the intranet heading to the Internet or enter the intranet from the Internet. In either case, an increased risk of compromise exists. Inbound communications could be carrying malicious code or an intrusion attempt. Outbound communications could be revealing internal secrets or distributing confidential files. These are extreme examples, but they are not uncommon. Most border-crossing traffic is benign, but because the risk of malicious traffic is greater at border crossings, you need additional filtering, monitoring, and blocking.

Naturally, external-only communications are more likely to be malicious; however, they do not end or originate from an internal source, so there is usually little need for concern. Malicious activity that does not attempt to breach your network borders is not really your problem. However, if external-only communications are defined as those that do not interact with your intranet but that may interact with your demilitarized zone (DMZ) or extranet, then you should be concerned.

With this definition of external-only communications, you must filter, monitor, and block the most obvious malicious packets and events, but still allow any conforming communication request—even if the origin or source is unfamiliar.

Controlling communication pathways is an important part of managing and designing network security. But another important piece of the security puzzle is management of hardened systems between which the secured communications take place. All systems

on the network are assets that contain other assets (such as data, processes, and configurations) that are of value to the business.

TECHNICAL TIP

When considering internal user access to network resources, one should scrutinize all user accounts, because hackers can gain access to and remain resident on networks for years. Just because the activity is internal to the network does not mean it is any more secure or legitimate than activities initiated outside the network. Network security is an essential component of network design and management. One must understand the design and flow—both physical and logical—of a network to recognize if an event is a possible malware attack, or if the network has reached its capacity for data throughput, as both instances can result in similar behavior of the network.

INTRUSION DETECTION SYSTEMS AND INTRUSION PREVENTION SYSTEMS

Two security technologies available to secure networks are intrusion detection systems (IDSs) and intrusion prevention systems (IPSs). An IDS detects unauthorized user activities, attacks, and network compromises (**FIGURE 4-3**). IDSs come in two types: host-based and network-based.




FIGURE 4-3 IDS detects an attack and alerts operators—manual intervention needed.

■ NOTE

IDS/IPS is used largely on Internet connections, because those connections typically present the largest threat to the network. You can also deploy IDS/IPS in strategic locations on the internal network. This is an excellent idea if your internal network has connections to third-party networks such as customers, vendors, or business partners.

An intrusion prevention system (IPS) is very similar to an IDS, except that in addition to detecting and alerting, an IPS can also take action to prevent a breach from occurring (**FIGURE 4-4**).




FIGURE 4-4 IPS detects attack, alerts operators, and then modifies firewall and router configuration to address the attack.

Two common deployment methods are used when placing an IDS/IPS for protecting a network from the Internet. Each has advantages and disadvantages.

An unfiltered IDS/IPS installation examines the raw Internet data stream before it crosses the firewall (**FIGURE 4-5**). This provides the highest amount of visibility to attacks, but it also means that a significantly higher volume of data will be monitored, with a higher possibility of false positives. During periods of high traffic, the IDS/IPS might not be able to process all the packets, and attacks can be missed. This placement often causes a bottleneck within the network, and users may notice the delay and complain about network performance.




FIGURE 4-5 Placement of the intrusion detection system to get unfiltered traffic for analysis.

A screened IDS/IPS solution monitors traffic that gets through the screening firewall (**FIGURE 4-6**). The advantage to this model is it dramatically reduces the amount of traffic monitored, reducing the chances of false positives and lost packets during high traffic volumes. A loss of visibility accompanies this model, as you cannot see attacks blocked by the screening firewall. Often you can learn where attacks are coming from by the packets dropped at the firewall. This placement does lessen the bottleneck possibilities on the network, as a smaller number of packets is assessed.




FIGURE 4-6 An intrusion detection system deployed behind a screening firewall.

HARDENING SYSTEMS

Hardening systems focuses on improving security of hosts and nodes. Hardening is the process of reducing the **attack surface** of a potential target by removing unnecessary components and adding in protections. Although each organization usually creates its own custom and internal hardening processes and procedures, most hardening guidelines have common elements and components.

When considering your network and systems, think of an attack surface as the candy coating on a popular button-shaped chocolate. The colorful outer shell is not supposed to melt easily and therefore protects the heat-sensitive inner contents (the chocolate). If the outer candy coating is misapplied, forms cracks due to pressure or dropping, or is exposed to air and heat through a manufacturing defect, the chocolate will begin to melt and leak out. To protect the inner contents, the outer shell must be free from defects and continuous with no gaps. The same is true for your network. Any weakness in the network exposes the entire network.

Some of the common recommendations to improve security or harden a host include:

- Remove all unnecessary protocols.
- Uninstall all unnecessary applications and services.
- Define a complex password for all accounts; do not leave any account with a default password or a blank password.
- Configure account lockout and define a logon warning banner.

- Install all available final release updates, patches, fixes, service packs, and so on for the operating system and every remaining application and service.
- Update all hardware device firmware or BIOS with the latest final release from the vendor.
- Install the latest final releases of all device drivers.
- Install and update antivirus and anti-malware scanners.
- Configure communication encryption.
- Install and configure a host firewall.
- Use a file system that supports file level permissions and **auditing**, and set those permissions.
- Configure system monitoring and auditing.
- Synchronize the clock of all devices and systems across the network for ease in auditing.
- Run vulnerability assessment tools against the host, such as OpenVAS and Nessus.
- Configure regular backups, and test those backups to ensure that each works and data can be retrieved.
- Impose any organization-specific security limitations, such as blocking USB drives or using whitelist execution management. This is often performed using a security template file.
- In Windows, disable the guest account, and rename the Administrator account. In Unix, establish policies whereby the root account is never used directly, but administrators must use the su command to obtain root access (thus creating a log of their events).

In addition to these hardening suggestions, organizations should add additional steps to their securing procedures based on the purpose of

the system, the criticality of the system, and the risk present in the environment.

Once you've hardened your system, you must maintain it over time. On a regular schedule, at least annually, reexamine every host against your organization's hardening policies to ensure compliance. Remove from service any system out of compliance with hardening policies until you are able to bring it into compliance. Then, investigate the cause of the security noncompliance and take countermeasures to prevent a reoccurrence.

EQUIPMENT SELECTION

Equipment selection is a commonly overlooked aspect of secure network design. The general belief that any hardware capable of performing an IT function is suitable for deployment is, unfortunately, not the case. Regardless of cost, all products may have well-known or not-yet-discovered security flaws.

Arbitrarily or automatically choosing the least expensive or the most expensive, or a certain brand of products, is not a winning security strategy. You should carefully evaluate each piece of computer equipment, from network device to host system, for its native security defenses or lack thereof, regardless of its cost or branding.

As you select, purchase, and deploy equipment, consider the vulnerabilities introduced and any protections or improvements to the infrastructure's security stance. Every piece of equipment will either improve security or reduce it in some way. Some equipment adds new weak points or expands the organization's attack surface, while other equipment will act as a countermeasure, protecting weak points of other components and reducing the attack surface.

Whenever possible, select equipment that provides greater improvement to security rather than acts as a detriment. This seems an obvious guideline, but you can only follow it if you are conscientious about evaluating the security profile of each device. Failing to evaluate the security of a new device properly could mean

that you inadvertently introduce new vulnerabilities into the organization's network and systems.

Some of the security concerns regarding equipment include:

- **Removable case**—The easier it is to remove or open the case of a device, the easier it is to hack into the device, plant a listener, or modify its functions.
- **Portability**—Is the device small enough to fit into a pocket, purse, or backpack, making it easy to steal?
- **Rack mountable**—A rack-mounted device is less likely to be stolen once screwed into a rack case, which can be locked.
- **Removable media**—Equipment with removable media bays (such as tapes and optical discs) or external peripheral ports (such as universal serial bus [USB], firewall, and Ethernet) may be easier to compromise than those with fewer or none of these access points.
- **BIOS/firmware flashing**—The ability to change the embedded software of a device is both an advantage and a disadvantage. Flashing to an updated, more secure version of firmware is positive; however, it could be a problem if a hacker can easily replace firmware with a third-party version or an older version with flaws.
- **Remote connection**—If remote connectivity to a device is possible, then risk increases. Limit, encrypt, and monitor remote connections.
- **Plaintext protocols**—The more a device defaults to or only supports plaintext protocols, the less secure it is. Choose equipment that supports encryption.
- **Easy-access management console port or interface**—If device can be reconfigured through an LCD screen and a few buttons (such as a printer) or if a console or terminal port allows quick access to a configuration or management interface (such as a wireless access point, router, or switch), then use caution when

deploying the device and assess the level of physical access security.

- **Reset button**—A reset button is used to return a device to the default factory settings. Any defined security configuration on the device is lost if someone presses the reset button. When possible, select equipment without a reset button or a button that can be disabled to provide physical security for the device.
- **Location**—Consider the location of a device that requires periodic login access or human interaction, as it may be possible for employees nearby to shoulder surf and learn the necessary credentials to access that device. This is true for employee workstations as well. With smartphones in the workplace, all one needs to do is aim the embedded camera from the phone in the direction of the keyboard someone else is using to see or record the keystrokes for later use.
- **Easy-access power switch**—If the power switch is easily accessed or triggered, casual contact with the device could cause power interruption. Additionally, an easy-access power switch allows a malicious person to power off equipment or trigger a reboot.
- **Heat produced**—The more heat a device produces, the more the heating, ventilation, and cooling (HVAC) system must work to keep the temperature of the room within acceptable boundaries. Excessive heat-producing devices can also increase the risk of fire.
- **Electricity consumption**—Excessive energy use can cause not only increased electric bills, but also increased temperature within the device and on electrical distribution systems. A circuit drawing too much power can cause a breaker overload. A tripped breaker causes downtime and may cause data loss and equipment damage.

Many other aspects of equipment security are important when evaluating the deployment of a new device. Whether the device is a

high-end server, a user's notebook, a smartphone, a network router, or anything else, consider the security of the equipment thoroughly; do not make purchasing recommendations solely based on equipment cost.

Keep in mind that devices that are inexpensive or free up-front may cost considerably more to manage and secure over time than a more expensive device. However, just because something has a high cost does not ensure that it has a low security management requirement. Money alone is rarely the true measure of the security of anything.

AUTHENTICATION, AUTHORIZATION, AND ACCOUNTING

Security ultimately is supported and enforced by **authentication, authorization, and accounting (AAA)**. AAA is part of field of security commonly referred to as **identity and access management (IAM)**. Without all three of these security fundamentals properly implemented, real security cannot exist.

Authentication is the verification or proof of someone or something's identity. The most common form of authentication is the use of a password. Although passwords are the most common, they are also one of the weakest forms of authentication. People typically pick passwords that are easy to guess or that are somehow predictable. They often reuse the same passwords on multiple systems.

Passwords reside in account databases in hashed form, which means the original password cannot be recovered from the hash value. However, by hashing large numbers of potential passwords, password-cracking techniques can potentially match a password hash to the target hash. Password-cracking techniques—including **dictionary password attacks**, brute-force password attacks, and hybrid attacks—can often reveal poorly constructed passwords.

TECHNICAL TIP

Password cracking typically focuses on generating and hashing large numbers of passwords with the goal of matching a stolen or captured password hash. **Dictionary password cracking** uses a list of potential passwords, prepared in advance. Each password from the list is hashed using the same hashing algorithm as the target/stolen password. If a match is found before the list is exhausted, the attack is successful. **Rainbow tables** are precompiled lists of hashed passwords used in attacks. A hacker who has access to the physical system may be able to build a dictionary of potential passwords from the hard drive within the device, because most passwords entered at the keyboard are stored on the hard drive in plaintext. That's why hard disk encryption is an important security measure.

A **brute-force password attack** builds potential passwords out of a selected character set, creating ever longer and longer passwords that use every possible valid combination of characters. The hacker hashes each crafted password and compares it to the target hash until a match is found or the attack is abandoned. Given enough time and the right character set, a brute-force attack will eventually be successful.

A **hybrid attack** uses a dictionary list as seed passwords that are then brute-force-modified. First the hacker makes all possible one-character modifications, then two, and then three. This technique is often more successful, because many people pick an easy-to-remember word

and then make only a few character modifications, such as changing an “a” to an “@” or just adding one or two characters.

The best defense against password-cracking techniques is to select a long password, for example, using at least 15-character passwords on Windows systems. Adding complexity (mixing multiple character types: uppercase, lowercase, numbers, and symbols) and using multiple words or phrases instead of a single base word also improve password strength. No matter how long or strong a password appears, a successful phishing attack, social engineering attempt, or shoulder surfing (watching someone type in their password) can all defeat it.

A dynamic password token is a device with a display screen that shows a seemingly random nonrepeating one-time-use password. The password displayed on the token must be included in the logon process, usually with a separate Type 1 PIN or password. A similar technique is a mobile authenticator app that sends the password to your mobile device. The password is temporary, and if the delay is too long, a new request must be sent for another one-time password as the previous password timed out. Google’s Titan Key, Types 1 and 2, uses proximity to verify authorization to a system without the need to enter a numeric code. These two-part mechanisms are forms of multifactor authentication.

Multifactor authentication is significantly more secure than any single-factor form of authentication. Passwords are but one example of the first type of authentication factor. Three commonly recognized authentication factors are:

- **Type 1**—Something you know
- **Type 2**—Something you have
- **Type 3**—Something you are or do

Something you know can be anything you memorize so that you can type, write, or speak it when asked to authenticate. Passwords are the most common example of a Type 1 authentication factor.

Something you have can be anything you must physically carry with you, such as a device or token. These can include metal keys, smart cards, radio-frequency identification (RFID) chips, ID badges, electronic devices known as dynamic password tokens, or texts sent to your preregistered smartphone.

Something you are or do is commonly known as biometrics. Some part of the human body is used as an element in an authentication process. This can include fingerprints, iris scans, retina scans, facial geometry, palm scans, signature dynamics, keystroke dynamics, and voice-pattern analysis.

A mixture of two or more authentication factors is multifactor authentication. Multifactor authentication is much more secure than single-factor. With single-factor authentication, an attacker needs to have only a single skill or exploit to successfully log on with a compromised user account. With multifactor authentication, an attacker needs to have multiple skills or exploits to successfully log on with a compromised user account.

Strong authentication prevents unauthorized entities from gaining easy access to the internal workings of an organization's infrastructure. Apply strong authentication to the logical environment, as well as to the physical environment.

Authorization, commonly known as access control, defines which actions a user can and cannot perform. Proper granular use of authorization ensures that authorized users perform only authorized activities.

The principle of least privilege is often a good guideline on which are the most appropriate authorization settings to make. This principle states that you should grant users the fewest capabilities, permissions, and privileges possible to complete their assigned work, without unnecessary access and additional capabilities. In other words, you grant users enough power and access to perform their assigned work, but no additional capabilities beyond their job descriptions. If a user has access to a device, the potential exists for intentional or unintentional alteration of the device or its contents.

Accounting is the activity of logging, monitoring, and auditing the environment, focusing both on users as well as system activities, to check for security policy compliance. This process holds users and systems accountable for their actions and activities. Through the use of thorough accounting and detailed auditing, you can detect and respond to any violations, attempts to violate, or trends toward violations.

Remember that security is locking things down and then watching for attempts to breach the lock. Authentication and authorization are forms of locking, and accounting is watching. However, these are not

the only forms of locking and watching on a secure network. Another important area of network security is communication encryption.

COMMUNICATION ENCRYPTION

Communication encryption is the use of encryption protocols to secure the contents of communications. Use encryption any time a transaction occurs with an outside entity or when a communication crosses a segment that is at risk of eavesdropping. You should also use encryption internally whenever the potential for loss or compromise, even by internal personnel, would cause significant harm to the organization.

You can protect transactions by encryption in two main ways. One is to use encapsulating intermediary protocols that provide encryption, such as IPsec and SSL/TLS; another is to encrypt data before it goes to a network protocol.

Protocol encryption ensures that all or most data sent over the network is safe from eavesdropping, modification, and other forms of compromise. A widely used mechanism of protocol encryption is that of a VPN. VPNs can function between individual systems or between entire networks or any other combination of endpoints.

Data encryption, performed either by the client or server software, ensures that even if the protocol encryption fails or is compromised, the data is safe by its own encryption. Software or data encryption is not as interoperable as protocol or VPN encryption, but it can be a viable option when both endpoints of the transaction are using compatible software components.

When in doubt, always encrypt. Although it is not a perfect solution (indeed, no perfect security solution exists), encryption offers significant protections from outside eavesdropping and modification. However, encryption fails if the selected algorithm is

poor, key management is insecure, either endpoint of the communication has been compromised (such as by hacker intrusion or planting of malicious code), or intermediary network nodes that decrypt and reencrypt fail.

Choosing to encrypt by default is an excellent network security rule of thumb. This guideline does not imply that the same encryption protects both internal-only communications and those that cross the network's boundaries. In deciding which encryption to use, it's important to consider where the hosts are located.

HOSTS: LOCAL-ONLY OR REMOTE AND MOBILE

When designing network security, focus on the function—as well as the physical and logical location—of each device. When a host is local and communicates only with other local hosts, then you can use slightly less stringent security on internal communications. However, when a host is remote, mobile, or otherwise outside of the private network, use significant additional precautions.

Once you allow remote access, you lose the benefit of the physical access controls. Once you support remote connectivity, a hacker or intruder need no longer be physically present in a facility to launch an attack. Thus, remote access is a risk, and it lowers the overall security of an environment. To compensate for this security reduction, you should impose more rigid limitations in terms of authentication, authorization, and accounting.

As security is designed and deployed, consider the purpose or function of each device, especially clients and servers, as well as its location in one (or more) of the seven domains of a typical IT infrastructure (**FIGURE 4-7**). Then, consider which other devices will need to communicate with it. The more vulnerable the communications pathways, the more security you will need to impose.

7-Domains of a Typical IT Infrastructure




FIGURE 4-7 The seven domains of a typical IT infrastructure.

If both endpoints are physically located in the same facility but the network pathway linking the devices involves any exterior segments, then the transaction demands greater security than a local-only communication. If the endpoints are geographically distant, then you need to use encrypted communications.

Remote and mobile devices are inherently riskier, as they are exposed to a greater number of potential threats. Inside the office, most threats are known, controlled, and monitored. But mobile and remote devices are potentially exposed to unknown, uncontrolled, and unmonitored situations, data, software, and users. A location outside of the organization's facility also means less strictly controlled physical access to the device.

TECHNICAL TIP

Location-aware anti-theft software periodically collects IP information and any available location data and uploads it to a centralized site. If the owner of a system reports the item lost or stolen, this automatically posted data might help recover the device. However, if the thief reformats the hard drive, then the anti-theft software might be removed.

Changes to security and potential exposure to risks require additional protections on remote and mobile devices. These devices should continue to host antivirus scanners, anti-malware scanners, and firewalls. But they can also benefit from whole hard-drive

encryption, multifactor authentication, and location-aware anti-theft software.

Remote and mobile devices are most likely to be out of compliance with security requirements, such as patch levels or application of security templates. A network access control (NAC) system can isolate and quarantine devices until they have installed all the necessary patches and updates.

Even with good host management—whether local, remote, or mobile—consider redundancy requirements. Is one of something enough or just enough to cause a significant problem when it goes offline?

REDUNDANCY

Security is not complete without adequately addressing preparedness. Redundancy is at the heart of preparedness. Can your organization survive downtime, blackouts, communication loss, server crashes, hard drive failure, floods, building eviction, virus infection, or any other potential threat? In most cases the ability to answer, “Yes, the organization can survive,” depends on the level of preparedness of that organization.

Preparedness is also known as business-continuity planning or disaster-recovery planning. The purpose of such planning is to ensure that a strategy exists to recover from any realistic threat. The difference in the type of planning is based upon the severity and continuity or the threat. To determine what is needed in your organization, assess each serious threat to the stability and function of the organization. Then, develop procedures to respond to the threat. The procedure can focus on prevention or recovery (or a combination of both).

A core element throughout this form of planning is redundancy. **Redundancy** is the act of avoiding single points of failure by building in multiple elements, pathways, or methods of accomplishing each mission-critical task. This is an area that is often ignored in security because it is viewed as an additional cost or expense to complete the exact same work or process.

Redundancy works in many ways. A **redundant array of inexpensive disks (RAID)** is a form of redundancy for hard drives.

RAID protects against drive failure. Hard drives are mechanical devices that are engineered to work for only a set amount of time, expressed as the mean time between failure (MTBF) rating in hours of use. All hard drives will eventually fail. Having a plan and sufficient backups can ease the problems of drive failure during a peak productivity period on the network, or worse, failing when no one recognized it has failed and no replacement is planned.

An uninterruptable power supply (UPS) is a form of redundancy for power. UPS devices provide temporary power in the event of a brownout or blackout. A UPS can trigger a graceful shutdown to prevent loss of data if facility power is not restored promptly. These devices also offer power conditioning if there are issues with building power due to periodic lows and spikes, such as when an air conditioning unit cycles on. This is especially important in cities where there are imposed rolling brownouts or blackouts, or in older buildings, to prevent data loss and allow for maintenance of system integrity.

Redundancy can apply to communication pathways. Good network design requires at least two pathways to every important resource. Redundant links to and from clients may not be a strong need in most cases, but having multiple pathways to reach resources is often essential.

TECHNICAL TIP

Try to avoid network designs that have a single point of failure; it is important to have redundant options to ensure that your mission-critical functions can take place. Keep in mind that bottlenecks are still likely to happen, even with

redundant pathways. Monitor traffic and use levels on every segment across the network to look for trends toward reduced throughput or productivity over time. A bottleneck may at first be a slight hindrance to high performance and productivity, but it can later become a form of a DoS attack.

Redundant communication links for voice and data are also important considerations. Construction workers could accidentally dig in the wrong location and sever the primary wire bundle supporting the organization. Redundant lines to service providers could be the difference between a nuisance and a company-ending event. Do you want to leave the power of terminating your organization in the hands of outsiders?

Redundant firewalls, proxies, routers, switches, servers, and databases begin to make security and financial sense once each becomes an essential or mission-critical element of the IT infrastructure. Do not entrust the viability of an organization to any single component. Always have a secondary option or backup plan.

Failing to plan is planning to fail. Hardware failures, mistakes, accidents, or intentional damage will occur in every system in an organization at some point. No organization is immune. What differentiates successful businesses is whether an organization is prepared for the problem with a ready-to-use response solution, or will be caught off guard, resulting in scrambling to recover.

ENDPOINT SECURITY

Implementing network security is about both the big picture and the granular details. Infrastructure design, topology, and redundancy are all important big-picture items. But do not overlook the details that need attention on a node-by-node basis. Node security focuses on the tasks necessary for each type of networking device to improve its security. Node security or node hardening takes the generic recommendations of system hardening and expands those with additional node/host-specific improvements.

Clients

Clients are the devices directly controlled by people. Thus, clients must protect the network from the user and vice versa. Thinking of clients as two-way interfaces can assist in proper security design and implementation for this essential IT infrastructure component.

Every client needs the following security elements:

- Antivirus and anti-malware scanners
- Host firewall
- Secured Internet client software such as browser, email, file transfer, and chat
- Password-protected screen saver with auto timeout
- Ability to encrypt network communications
- Ability to encrypt storage devices
- Auditing of all user activity
- An integrity checking system that monitors for unauthorized file changes using hash value comparison
- Configuration integrity management that either alerts when configuration changes are made or locks down configurations that do not allow for alteration

Clients should be subject to NAC procedures to prevent an insecure client from compromising the rest of the network. Use multifactor authentication whenever possible to minimize the risk of an unauthorized person gaining access to the client and hence the entire network. Allow certain necessary functions—such as secure FTP—based on the needs of the client; if there is no legitimate business need for a service, the ability to conduct those actions should be removed.

Another recent shift for some organizations is to allow, even encourage, employees and contractors to **Bring Your Own Device (BYOD)**. The use of an individual's personal cell phone, tablet, laptop, or similar device lowers the cost to the organization, but it also raises many different security questions, ranging from policy to implementation. For instance, who is responsible for the applications that are run on the device and the security of each? Who owns and is responsible for the content and protection of the files created on the device? What if the specific device is incapable of running specific organization-decreed security software, such as antivirus or even VPN software? What if the owner of the device doesn't want to allow a security scan or to run certain software? BYOD exposes an organization to an array of security problems—problems that are growing in importance each day, especially because the population of BYODs is increasing and because the devices are wireless.

To address BYOD security concerns, first consider if BYOD makes sense for your organization. Lowered costs can mean lowered control and access for business purposes. If your organization has legitimate reasons for allowing BYOD, ensure that there are specific policies in place that detail acceptable use and the classification of business data as an asset belonging to the organization. Employ remote-wiping tools to remove data from devices if an employee leaves the organization or is found to have accessed or is in possession of data for which they have no legitimate business need. This is also a handy feature in the event the device is lost or stolen. Ask all employees to sign a release and accept this security measure, and ask the organization's legal counsel to review all corporate policies.

Servers

Servers are the backbone of any IT infrastructure. Whether a file server, database host, email server, proxy server, or authentication server, a server is usually supporting an essential or mission-critical function for the organization. Servers need protection against downtime. This can include redundancy in terms of RAID, duplicate servers, and/or clustering.

Employing duplicate servers means having two identically configured systems running at the same time, but only one is performing services live for the network. The second system is acting as a backup and receives all data changes as those changes occur. In the event of the primary server failure, the secondary server can take over supporting the services for the network. These servers can be located side-by-side in the same facility or in different states connected by a telephone line or the Internet. Some companies rely on a local server and a duplicate server hosted in the cloud to save on operating two physical systems.

Clustering is having two or more identically configured systems running as a collective. All of the systems share in supporting the live service to the network. If any member of the cluster goes offline, the remaining members continue to support the service. Clustering makes it possible for a service to be available consistently while still allowing for scheduled maintenance and occasional individual server downtime. There are additional resources, including auditing costs, when using clustering.

Servers need strong multifactor authentication to ensure that only administrators have the ability to log in from the keyboard. You should sequester servers in a dedicated room or vault to prevent

casual or intentional access by unauthorized users. Lock and monitor the server room at all times.

Routers

Routers are an essential traffic-management device of any network. Network router security is primarily about preventing unauthorized access. First and foremost, a router should be physically inaccessible to any nonadministrator. All router configuration and management interfaces must require strong authentication. Some environments have elected to eliminate local accounts on the routers and rely upon Terminal Access Controller Access-Control System Plus (TACACS+) for router authentication and access. Limit management interface access to a direct console or to a terminal cable connection only, rather than allowing in-network access. Require that all management interface communications and router-to-router communications be encrypted.

If a password is stored in a configuration file for the router, be sure to use an encoding scheme that is not easily cracked or reverse engineered. For example, the Cisco IOS password 7 hashing method uses a weak reversible algorithm and can be easily cracked, using software tools dating back to 1997. Unfortunately, that easily hacked method is still in use today. Be sure to update any configurations to the latest and most secure option available on the device, or consider changing the router to a newer model with better security.

Generally, configure the following on the router:

- Block all directed IP broadcasts.
- Drop all packets from the Internet using an RFC 1918 source address or any other internal address.
- Disable the TCP and UDP small services of echo, chargen, discard, and daytime.

- Enable a warning banner for all attempted connections to the router, especially to the management interface.

Try to limit the use of a router as a filtering device. Keep the router table focused on traffic routing, and implement firewalls to provide content filtering and blocking of suspicious or malicious traffic. Protect any router considered a border router with a firewall.

Switches

Switch security is similar to that of router security. Maintain control over who can physically reach the switch. Limit access to management consoles and require strong authentication to access the management interface. If a password is stored in a configuration file for the switch, be sure to use an encoding scheme not easily cracked or reverse engineered.

Consider deploying switches that support IDS-like features, such as watching for MAC spoofing and ARP flooding. MAC spoofing tricks a switch into thinking a hacker's computer is actually a legitimate host. The hacker steals or spoofs a legitimate MAC address. If a switch is monitoring for MAC addresses that change ports, then MAC spoofing ceases to be a serious threat unless a hacker can physically access the port connection of the MAC they are trying to spoof.

ARP flooding overloads a switch's mapping table so that instead of forwarding packets out the correct port, the switch will default into flooding mode and transmit packets out all ports. This type of attack is part of active sniffing. Active sniffing is an attempt to eavesdrop on switched networks.

Use switch features such as **virtual local area network (VLAN)** support and auditing ports. VLANs are hardware-imposed network segmentation. VLANs isolate traffic. Using VLANs to manage traffic is often cost-effective, as it does not require re-cabling or changing IP addresses; all of the VLAN configuration takes place within the switch.

The audit, mirror, or IDS port on a switch connects an IDS or IPS to monitor all of the traffic traversing the switch. Because a switch

only transmits data only from the port where its intended destination resides, attempting to monitor all traffic on a switched network is not possible without the use of the auditing port. Disable all unused switch ports, which prevents the easiest method of adding rogue devices to a network—namely, plugging into an open port.

Firewalls and Proxies

Firewalls and proxies defend the network, but you need to secure these as well; otherwise, if a hacker can compromise the firewall, then the security it provides is unreliable. Firewalls and proxies, as with routers and switches, need physical access protection. No nonadministrative user should be able to gain direct physical contact with a firewall or proxy. Limit access to management consoles, and require strong authentication to access the management interface. If a password is stored in a configuration file for the firewall or proxy, be sure to use an encoding scheme not easily cracked or reverse engineered. Never leave a default password, change all passwords, and ensure that the same password is not used for multiple devices. Recall that renaming the administrator account also offers an additional control.

A firewall should drop all packets addressed directly to the firewall, as the firewall does not host traditional services accessed in this manner. The same is generally true of a proxy server, but some instances apply where there is either direct communication with the proxy server or the proxy server hosts additional accessible resources.

FYI

Make a final security evaluation of any device, including firewalls, proxies, switches, routers, servers, and clients. You can perform this evaluation using an automated vulnerability scanning tool or a custom manual penetration test. The goal is to find any remaining vulnerabilities in

these devices so each device can be made as secure as possible, as quickly as possible.

If Simple Network Management Protocol (SNMP) is used on the network, require the use of SNMP v3, which allows for encrypted transactions and authentication of SNMP sessions. Then, use custom community names rather than the default public or private communities. When a newer version is introduced, be sure to review the RFC for the protocol to determine if it will work with your existing network or if modifications to hardware or software are necessary prior to changing.

Consider the software or firmware of the router, switch, firewall, or proxy. Install only full final releases, never beta or partial firmware. If you are concerned that the latest release has not been thoroughly tested in the real world, it is acceptable to stick with a previous release, provided there are no published reports of critical security flaws. The enhancements in the latest releases read like a “how to hack manual” to a hacker because all issues fixed in the latest version are still present in the older version. So, do not delay an upgrade if it is deemed the best option for your business.

RISK ASSESSMENT AND MANAGEMENT

How can your organization determine how many firewalls to deploy and where to deploy them? In most cases, this requires a process of **risk assessment** and **risk management**. Risk assessment involves examining values, threat levels, likelihoods, and total cost of compromise versus the value of the resource and the cost of the protection. Many other factors affect risk assessment, but this essential business and security process is as follows:

1. Determine the overall value of the resource or asset. Known as the asset value (AV), this calculation should include both tangible and intangible costs and value.
2. Determine the threats that face that asset. For each threat, calculate the exposure factor (EF), or the amount of potential harm expressed as a percentage. This will create a list of asset-threat pairs with a corresponding EF.
3. Calculate the single loss expectancy (SLE): $SLE = AV \times EF$. This is the amount of potential loss that could be experienced due to a single occurrence of compromise against this asset for a specific threat. This will add an SLE value to each asset-threat pair.
4. For each threat, calculate the potential number of times the threat could be a realized attack within a year's time. This is known as the annualized rate of occurrence (ARO).
5. Calculate the annualized loss expectancy (ALE): $ALE = SLE \times ARO$. This is the amount of potential loss that can be experienced due to any compromise of this asset for a specific threat within a year. This will add an ALE value to each asset-threat pair.

6. Sort the list of asset-threat pairs by the ALE. The highest ALE is the biggest risk for this specific asset/resource.
7. Take the asset-threat with the largest ALE and determine the possible countermeasures that could be used to protect against that threat. This creates asset-threat-countermeasure triplets.
8. For each asset-threat-countermeasure triplet, calculate a new ARO. The countermeasure should reduce the ARO. A perfect countermeasure would reduce the ARO to zero. There are few perfect countermeasures.
9. With the new ARO, calculate a new ALE for each asset-threat-countermeasure triplet.
10. For each asset-threat-countermeasure triplet, make a cost/benefit analysis. The formula for this is: $(\text{Original ALE} - \text{New ALE}) - \text{cost of the countermeasure per year}$.
11. Sort the asset-threat-countermeasure triplets by their potential cost/benefit. The triplet with the greatest cost/benefit is the best choice.

In short, if the cost/benefit analysis for a particular asset-threat suggests a firewall, then installing one is a good business, security, and budget decision. If some other countermeasure delivers greater benefit for less cost, then that countermeasure is the better choice.

Performing detailed risk analysis is a complex task. Many small companies or home network enthusiasts might not want to spend the effort to perform a complex risk assessment. This is understandable, but failure to perform a full risk assessment does not mean that you have no risk or that you have no assets worth protecting.

At a minimum, you need to protect any and all data that is personally identifiable. This includes personal communications like emails; any custom crafted document, picture, video, or other file type; all medical information; all financial information; configuration settings; and passwords and other credentials. Identity theft is one of

the most frequent and devastating crimes of the interconnected world. Worrying only about your neighbors, the random home burglary, or the pickpocket on the street is nothing short of tunnel vision. Today, you must be concerned about all would-be criminals across the globe attempting to gain access to your electronic assets over the Internet.

WHAT ARE ZONES OF RISK?

A **zone of risk** is any segment, subnet, network, or collection of networks that represents a certain level of risk. The higher the risk, the more security you need to protect against that risk. The less risk associated with a zone, the less security is necessary because the threats to that zone pose less chance of harm.

The flip side of zones of risk is **zones of trust**. Highly trusted zones naturally require less security, while zones of low trust require more security.

Each zone of risk needs to be clearly and distinctly isolated from any other risk zone, especially if those zones have different levels of risk. The primary tool used to isolate zones is a firewall.

Most networks have two to four zones of risk. These include the private network (LAN/intranet), DMZ, extranet, and the Internet. See **FIGURE 4-8**. The private network zone has the lowest risk and is the zone of the highest trust. The Internet zone has the highest risk and is the zone with the least trust. A DMZ has less risk than the Internet, but is not as trusted as the private network. A DMZ zone has medium-high risk or medium-low trust. An extranet has more risk than the private network and more trust than the Internet. An extranet zone has medium-low risk or medium-high trust. See **TABLE 4-1**.




FIGURE 4-8 An example of a private network with four risk or trust zones.

**T
A
B
L
E**
4
-
1
R
i
s
k
a
n
d
T
r
u
s
t
L
e
v
e
l
s
o
f
C
o
m
m
o
n

Network Zones

ZONE RISK LEVEL TRUST LEVEL

LAN/intranet	Low	High
Extranet	Medium-low	Medium-high
DMZ	Medium-high	Medium-low
Internet	High	Low

Your organization's written security policy should define where these zones exist and dictate the security requirements for each zone. Such requirements would include traffic management, use of firewalls, use of VPNs to cross the zone divisions, hardening of systems, malicious code scanning, and so on. Firewalls are likely found at risk zone divisions and may be the best cost/benefit security countermeasure in some circumstances.

CHAPTER SUMMARY

Network security implementation relies on a thorough understanding of your organization, its goals, its risks, and the technologies employed within your IT infrastructure. Before you can properly deploy network security, you must first design it. Most network security designs include layers of defense, as well as sufficient capacity for growth.

Network security includes an evaluation of the protocols and methods your organization uses. If the current design is insufficient, replace it with a design that addresses productivity and security. You need to assess the addressing schemes in use—whether public or private, static or dynamic—in light of how each improves or detracts from security.

Other important components of network security design and deployment include controlling communication pathways; hardening systems; and selecting proper equipment, authentication, authorization, accounting, communication encryption, types of hosts, redundancy, and node security specifics.

This chapter also introduced the concept of risk assessment and management to aid in determining the proper number of firewalls and other security

measures. While the terms are important, it is essential to understand how to convert risk into a quantifiable number and address it in a way that properly protects and defends the network.

KEY CONCEPTS AND TERMS

Annualized loss expectancy (ALE)

Annualized rate of occurrence (ARO)

Asset value (AV)

Attack surface

Auditing

Authentication, authorization, and accounting (AAA)

Bring Your Own Device (BYOD)

Brute-force password attack

Cost/benefit analysis

Dictionary password attacks

Dictionary password cracking

Exposure factor (EF)

File Transfer Protocol (FTP)

Hybrid attack

Identity and access management (IAM)

Intangible costs and value

Internet Assigned Numbers Authority (IANA)

Modeling

Multifactor authentication

Network News Transfer Protocol (NNTP)

Piloting

Post Office Protocol (POP)

Rainbow tables

Redundancy

Redundant array of inexpensive disks (RAID)

Regional Internet Registry (RIR)

Risk assessment

Risk management

Router access control lists

Router anti-spoofing

Security through obscurity

Simple Mail Transfer Protocol (SMTP)

Single loss expectancy (SLE)

Single point of failure

Sunk cost

Tangible

Telnet

Uninterruptable power supply (UPS)

Virtual local area network (VLAN)

Zone of risk

Zones of trust

CHAPTER 4 ASSESSMENT

1. Which of the following should be done as part of router configuration?
 - A. Copy and paste the configuration to all routers and firewalls
 - B. Enable a warning banner for all attempted connections
 - C. Require SNMP v 2 or earlier for consistency
 - D. Drop all encrypted packets within the network perimeter
2. Which of the following is *not* a type of attack against password use?
 - A. Brute-force
 - B. Dictionary
 - C. Hybrid
 - D. Recursive
3. Which of the following is part of a defense-in-depth strategy?
 - A. Avoid single points of failure
 - B. Avoid having multiple redundancies
 - C. Avoid removing the default account
 - D. Avoid using devices from different manufacturers
4. Which addressing class is 192.168.32.16?
 - A. Class A
 - B. Class B
 - C. Class C
 - D. Class D
5. Which of the following is an example of redundancy?
 - A. A firewall at each physical perimeter
 - B. Using multifactor authentication
 - C. Encrypting communication outside the network
 - D. An uninterruptable power supply

6. Which of the following best defines security through obscurity?
- A. Changing the names taped to all physical devices
 - B. Changing the logical names of all devices
 - C. Hiding the network in order to secure it
 - D. Expanding the network to hide individual devices
7. When considering multifactor authentication, which of the following is something you have?
- A. An iris scan
 - B. An ID card
 - C. A spoken phrase
 - D. A password
8. Which of the following is commonly referred to as access control?
- A. Accounting
 - B. Auditing
 - C. Authentication
 - D. Authorization
9. Which of the following is true of IPv4 versus IPv6?
- A. IPv4 is more expensive to implement than IPv6.
 - B. IPv4 is more time-consuming when building packets.
 - C. IPv4 is used less and therefore less is known about it.
 - D. IPv4 is plaintext transmission by default.
10. Which of the following best describes a dynamic password token?
- A. A device that shows a random password
 - B. Radio-frequency identification (RFID) chip
 - C. Identification (ID) badge
 - D. Smart card



© Digital_Art/Shutterstock

Securing the Perimeter with Firewalls

CHAPTER 5

Firewall Fundamentals

CHAPTER 6


Firewall Implementation

CHAPTER 7

**Firewall Deployment
Considerations**

CHAPTER 8

Configuring Firewalls



© Digital_Art/Shutterstock

Firewall Fundamentals

CHAPTER 5

TO SOME NETWORK ADMINISTRATORS, A FIREWALL is the key component of their infrastructure's security. To others, a firewall is a hassle and a barrier to accomplishing essential tasks. In most cases, the negative view of firewalls stems from a basic misunderstanding of the nature of firewalls—their purpose and functions.

This chapter clearly defines the fundamentals of firewalls. These include what a firewall is, what a firewall does, how it performs these tasks, why firewalls are necessary, what the various types of firewalls are, and how filtering mechanisms work. Once you understand these fundamentals of firewalls, you will be able to look beyond the unschooled opinions, common mythology, and marketing hype surrounding the devices, and the crucial benefits of effective firewall architecture will become clear. Like any network tool, firewalls are useful in solving a variety of problems and in supporting essential network security.

CHAPTER 5 TOPICS

This chapter covers the following topics and concepts:

- What a firewall is
- Why firewalls are necessary on networks
- How firewalls work
- Types of firewalls
- Options for individual and SOHO (small office/home office) firewalls
- Common uses for host software firewalls, commercial software network firewalls, and hardware and appliance firewalls
- Next-generation firewalls
- Virtual firewalls
- Differences between dual-homed and triple-homed firewalls
- Difference between ingress and egress filtering
- Types of firewall filtering
- Selecting firewalls
- Considerations for buying versus building a firewall

CHAPTER 5 GOALS

When you complete this chapter, you will be able to:

- Define firewalls

- Explain the need for firewalls
- Describe types of firewalls, including network router/interface firewalls, hardware/appliance firewalls, and host software firewalls
- Compare and contrast software and hardware firewalls
- Set up a broadband connection device firewall
- Use Windows Defender Firewall
- Compare and contrast dual- and triple-homed firewalls
- Define the meaning of ingress and egress filtering
- Explain standard filtering methods, including static packet filtering, network address translation (NAT) services, application proxy filtering, circuit proxy filtering, dynamic packet filtering, stateful inspection filtering, and content filtering

WHAT IS A FIREWALL?

A firewall is like a **border sentry**. Think of a firewall as a guard along the border of a country. On a network, a firewall is a filtering device that enforces network security policy and protects the network against external attacks. A **packet** is a unit of information that is routed between one point and another over the Internet or any other network. The packet **header** includes information such as source, type, size, and origin and destination address. As a filtering device, a firewall watches for traffic that fails to comply with the **rules** defined by the firewall administrator.

Firewalls can focus on the packet header, the packet **payload** (the essential data of the packet) or both; or on the content of a **session**, the establishment of a **circuit**, and possibly other assets. Most firewalls focus on only one of these. The most common filtering focus is on the header of the packet, with the payload of a packet a close second.

FYI

The term “firewall” comes from a building and automotive construction concept of a wall that is built to prevent the spread of fire from one area into another. The firewall in a building or car engine compartment is a physical block against the spread of fire.

A firewall is an important part of an organization's comprehensive security policy. A security policy defines the goals, objectives, and procedures of security. Every security policy focusing on network security requires the deployment of a firewall. The firewall's job is to impose all restrictions and boundaries defined in the security policy on all network traffic. A firewall enforces your organization's network security policy, just as a traffic policeman enforces the motor vehicle laws of a town or state. A firewall policy ensures the success of your network's firewall. A company-wide security policy ensures the success of your organization's security infrastructure. Firewalls are in no way a substitute for such a security policy.

FYI

A firewall enforces the network traffic access control security policy. A firewall is the physical embodiment of the security policy. It is the most obvious or direct enforcement of access control on your network's traffic. No other device is as directly involved as your firewall in allowing authorized traffic and denying all else.

Firewall filtering compares each packet received to a set of exception rules. These rules state that content in the packet is either allowed or denied. If the packet matches an allow rule, it continues on to its destination. If the packet matches a deny rule, the packet is dropped. Hence, a deny rule prevents the packet from reaching its

destination. If a packet fails to match any rule, the firewall drops the packet by default.

If properly configured, filtering allows what you want on your network and denies what you do not want. Filtering relies on filtering rules. Each rule identifies a pattern of concern and a response the firewall will make if an incoming element matches the pattern.

Firewalls follow a philosophy or stance of security known as **deny by default/allow by exception**. All the rules on a firewall are exceptions. Some exception rules define what you allow. Some exception rules define what you wish to deny. The final option, sometimes called the final rule, is that anything that did not match one of the exceptions is denied by default.

The filtering rules are the exceptions to the deny-all rule that is the final and absolute rule of a firewall. In fact, if a deny by default rule did not exist, the filtering device would not be a firewall at all. Instead, it would be more like a router or switch, allowing traffic to pass, even if it did not match an expected pattern or rule.

TECHNICAL TIP

Keep in mind the security stance of *deny by default/allow by exception*. A firewall is a filtering device that helps support this stance.

Network security administrators use the term “firewall” to refer to a device or service that allows some traffic but denies other traffic. This is not the same as a building’s firewall, which allows a fire to spread only if the firewall fails. In network security, firewalls allow

traffic through that is considered safe—or at least authorizes traffic without the entire firewall failing. Additionally, if a firewall does fail, it fails into a secured state known as *fail-safe* or *fail-secure*. This means that when the firewall is offline, locked or frozen, or otherwise experiences a problem, the firewall stops all traffic rather than allowing all traffic through.

A gateway is an entrance or exit point to access a controlled space. Gateways are important because high levels of traffic typically pass through a gateway. Thus, positioning a firewall at a network gateway is an aspect of secure design. If the gateway already exists in the network infrastructure, then positioning the firewall at that point is an obvious security improvement. A firewall so positioned intercepts and verifies all traffic crossing that gateway point.

Some other places that firewalls are often deployed include in front of (i.e., on the public Internet side), behind (inside the data center), or in load-balancing systems—or, as they are most often called, **load balancers**. Load balancers are reasonably well-named because that is what these devices do: They take the load coming into a set of servers and ensure that the load is balanced between or among the servers, based on a variety of factors. It is a funnel point for traffic that is an ideal place to enforce policy, and that is done through the firewall.

If a network is still in its design phase, a network designer might make the secure choice to create a gateway in the network's layout. However, in this situation, the concentration point would be known as a **chokepoint** instead of a gateway. If a gateway is an access point to other areas of resources, a chokepoint is a specialized kind of

gateway that focuses on traffic to a single concentrated pathway to simplify the process of filtering.

You can also think of a firewall as a traffic-control device. A significant amount of network security is little more than controlling traffic. Authorized traffic passes through the digital intersection, while unauthorized, unwanted, abnormal, or obviously malicious traffic is blocked. Knowing how to distinguish the differences in the traffic is often difficult, requiring you to rely on well-crafted rules.

The very basic and original form of the modern firewall was a **screening router**. Routers analyze traffic based on destination address. The best-known available route to the destination informs the forwarding decision. However, screening routers added additional rules that could discard traffic based on destination or source address. Once filtering expanded to address protocols and even ports, the screening router became the basic static packet filtering firewall.

In addition to a firewall, a **transparent proxy** (also called an **in-line proxy**) can aid in filtering requests from outside the network. A transparent proxy is most often installed immediately between the firewall, protecting against Internet intrusions, and a web server. A proxy is something that stands in place of something else. In this case, it is a server that stands in place of and behaves initially as if it were the web server. When a request comes in for the web server, it is passed to the proxy if it meets the firewall rules. Upon receipt, the web server proxy either acts on the request or passes the request to the actual web server. This provides a buffer between the web server and the rest of the network.

In addition to filtering, firewalls can also offer routing functions, which are a holdover from this router ancestry. A multi-homed

firewall can grant traffic access to one or another interface or **segment** based on the results of filtering. Traffic for a private local area network (LAN) could traverse one segment, while traffic destined for the demilitarized zone (DMZ) could follow another (see **FIGURE 5-1**). A firewall is an efficient, necessary traffic control device on the highway that is your network.




FIGURE 5-1 A basic multi-homed firewall filtering and routing for two network segments.

A firewall's critical function is to protect your network against external attacks. It is no accident that some threats to computer networks are called viruses; external threats seem to continuously change and evolve. Some attacks are targeted specifically for your network or organization. Other attacks are untargeted and random. Instead, these attacks are directed toward any host that happens to have a specific vulnerability. Malicious code relentlessly infects and compromises unprotected computer systems. Flooding attacks attempt to interrupt timely Internet communications. These are only some of the external threats facing your organization's network.

Whether labeled as a gateway for resource access or a chokepoint to control security, the result is the same. A firewall positioned at a gateway provides filtering services across all traffic. The chokepoint provides filtered access to resources. As a wise poet once said, "Good fences make good neighbors." Similarly, good gateways and well-positioned firewalls make secure networks.

Understanding these definitions and distinctions is only the beginning of understanding firewalls.

What Firewalls Cannot Do

When used properly, firewalls provide reliable and consistent security from external threats. However, firewalls have limitations and are only part of the complete security strategy. By itself, a firewall cannot protect a network against every threat.

A firewall is an essential part of network security, but it is not the whole network security. Don't make the mistake of deploying a firewall while ignoring other security management activities. A firewall is only one piece of the large complex puzzle to address network security. A firewall cannot do all things.

Authentication Server

A firewall is primarily for network traffic filtering. It is not an authentication system. Firewalls are not designed to check logon credentials, compare biometric scans, or even confirm the validity of digital certificates. These are the functions of an authentication service, typically hosted on a domain controller or primary network server.

That said, you might find it necessary for a firewall to allow authentication before granting access to a resource or allowing a session. Some firewalls can have enhancement features that provide for firewall-hosted authentication services. However, in many cases, a better solution would be to have the firewall offload that task to a dedicated authentication server or service, such as 802.1x, public key infrastructure (PKI), or directory services. Most security experts do not recommend using a firewall to authenticate users, or at least not as a replacement for a network's directory service or centralized authentication solution.

Remote Access Server

A firewall is not a remote access server. Connections from remote users do not have an endpoint at the firewall. Instead, the endpoint is a remote access server (RAS) or network access server (NAS). A firewall may function before or after the RAS/NAS to filter remote access traffic. However, that doesn't mean the firewall is the RAS/NAS.

Your firewall should filter all remote traffic, especially because remote traffic is much more likely to be purposefully malicious or accidentally damaging than local traffic. Why? An organization has much more control over who can connect to its network locally than it does when it allows remote connectivity.

Encrypted Packets

Unlike Superman, your firewall does not have X-ray vision into encrypted traffic; in other words, it cannot see the contents of encrypted traffic. A firewall can filter on the header of traffic using **transport mode encryption**, since the original header is in plaintext form. However, a firewall cannot filter the original header of traffic using **tunnel mode encryption**, since the only plaintext component is a temporary tunnel header that only includes information about the endpoints of the tunnel. It would be like trying to guess what's in the boxcars of a train by observing only the locomotive and the caboose.

Position your firewall where it can be most effective. If the security design requires that all traffic content be examined by a firewall, then you need to position the firewall after encryption is removed from the traffic. If the security design requires filtering only on non-encrypted traffic, then positioning the firewall is not as critical.

Firewalls designed for use by web e-commerce sites may have an additional ability to act as the endpoint of a **Secure Sockets Layer (SSL)** or **Transport Layer Security (TLS)** tunnel from an Internet client. This grants the client protection for its data as the information traverses the Internet. This allows a firewall to filter the content of the traffic before the web server receives and processes the information.

Encryption is one method to evade filtering. Users and hackers can employ client-side encryption solutions that encode the data before transmission, or create unauthorized encrypted encapsulation tunnels to prevent firewall filtering. In this situation, your network security policy may stipulate that the firewall needs to block

encrypted transmission initiated by clients, especially if the destination is on the Internet.

Malicious Code Scanner

A firewall is also not a malicious code scanner. Firewalls are traditionally rule-based filtering products. These **rule sets** usually have only a few dozen to at most a few hundred rules. To filter malicious code, the rule list would need millions of entries. In December 2009, the AVG Anti-Virus definition database #270.14.88/2538 was tracking 24,303,360 malicious code infection definitions. By 2018, that number grew into the hundreds of millions. That number of entries is simply impractical to include in a firewall rule set.

Some firewall products include an enhancement or add-on module for malicious code scanning. Such an enhancement is just an add-on component, not a core feature of a firewall. In most cases, it is more efficient and more secure to use separate anti-malware scanners than to add this function to your firewall.

Many firewall rules block traffic with spoofed addresses, uncommon ports, unauthorized protocols, invalid header constructions or values, and so on. Such rules block a significant amount of traffic caused by malicious code, but these rules do not directly block malware from entering or leaving a network.

Intrusion Detection

A firewall is not an intrusion detection system (IDS). An IDS is a type of network burglar or intruder alarm that detects and responds to unauthorized activity inside your network. An IDS performs this task by monitoring all network traffic. Although an IDS can be deployed on a network border or outside the network against the Internet, most operate inside private networks and provide the ability to watch all internal network traffic.

A firewall can detect malicious traffic only when such traffic enters one of the firewall's interfaces. Firewalls usually do not watch over general interior network activity. Firewalls are either border devices for networks and subnets, or they are software products watching over a single host. In either case, they cannot see the same traffic, nor can they perform the same tasks as an IDS. A firewall, therefore, is not interchangeable with a good IDS.

Insider Attacks and Social Engineering

A common misconception is that firewalls protect against insider attacks. A firewall can be a border device, or a firewall can be software on a host. A **border firewall** can filter traffic entering or leaving a network or subnet. A border firewall is unable to see any interior traffic (**FIGURE 5-2**). When an attacker from an inside client attacks a target that is also an internal host, a border firewall is not part of the communication; thus, it can neither detect nor block the attack. A host software firewall can only see the traffic entering or leaving that one host. A host software firewall is unable to see any other interior traffic.




FIGURE 5-2 Border firewalls cannot filter internal communications.

Of course, if traffic does not pass through its interfaces, a firewall cannot filter the traffic. A firewall can filter only what it sees. If malicious or unwanted traffic does not enter an interface of a firewall, the firewall will not be able to filter that traffic. For the best strategy, you should place firewalls on each host, on every border gateway or chokepoint, and between each significant subnet or interior network division.

Another thing firewalls cannot do is protect against social engineering. Social engineering is a category of attacks that focus on the personnel of an organization. These attacks get information from people just by asking for it in clever ways or convincing someone to perform an action that breaches network security. The only real protection against social engineering is worker training and awareness.

■ NOTE

IDSs can detect a plethora of unwanted activities, use several methods of detection, and perform a wide range of responses, both passive and active.

Removable Media

A firewall cannot protect against the significant threat posed by the use of removable media. There are many different types of removable media: USB hard drives, USB thumb drives, CDs, DVDs, Blu-ray discs, HD-DVD discs, other optical discs, flash memory cards, FireWire storage devices, tape media, email attachments, a smartphone connected via USB cable, and more. A firewall is not involved in the use of removable media or in any of the contents these devices may contain. A firewall cannot protect your network against the ongoing threat posed by removable media. Again, the best defense against this threat is good company policy, worker training, and awareness.

Physical Attacks

A firewall, of course, cannot protect against physical incursions or attacks. Physical attacks bypass any and all logical and electronic protection mechanisms. A firewall does not protect against theft of devices, planting of eavesdropping mechanisms, disconnection of cables, connecting a rogue device to an open node, destruction of equipment, dousing electronics with liquids, building fires, or any other form of physical attack. Only effective physical defenses can deter physical attacks. A firewall is not designed or intended to thwart physical attacks.

Misconfigurations and Security Management

Computer equipment and software can do only what each is designed and programmed to do. If an administrator misconfigures a security device, the device does not automatically compensate for that oversight. If a security administrator fails to learn about all of the features and defaults of new equipment, the product cannot secure itself autonomously.

Security requires training, research, careful planning, thoughtful implementation, and ongoing review and maintenance. This process is known as security management—and it takes work. The old expression was never more true than it is today: “Garbage in, garbage out.” What you put into network security is precisely what you will get out of it. So, remember that a firewall cannot compensate for ineptitude or ignorance on the part of administrators.

Firewalls cannot compensate for poor security management. Proactive security management is essential for the success of any security endeavor. Security management is the process of reviewing, testing, tuning, and updating an organization’s security policies and security infrastructure. This is an ongoing effort that requires knowledge, research, and vigilance. The threats and risks facing an organization are constantly evolving to become more persistent and virulent. Your security strategy should be just as rigorous and purposeful in defense. Keeping up-to-date on current threats and trends in network security is a big part of this job. Networking with security professionals, attending relevant conferences, and reading the latest industry literature are ways to keep yourself and your security efforts sharp.

Finally, firewalls are not a perfect solution. They cannot self-adjust to changing conditions or future threats. Firewalls are mostly software (even when operating on dedicated hardware) and are written by fallible humans; therefore, they are subject to bugs and flaws. Fortunately, in spite of all the things a firewall is not, a firewall is a solid filtering solution. It can and should protect the borders of networks and individual hosts as part of a complete security infrastructure.

WHY DO YOU NEED A FIREWALL?

High-speed Internet connectivity has become ubiquitous. Most computer users now have high-speed access at home, at work, and even on the go through mobile devices. Often, these broadband connections are always-on connections. This means computers and networks are online all the time—and exposed to attacks. A drawbridge is easy to cross if it is always down. When a system is always connected, it can be the focus of a concerted attempt to discover its vulnerabilities and breach its security.

How likely is it that a computer will be discovered and attacked over the Internet? It is almost guaranteed. Most systems are detected, scanned, and probed within minutes of obtaining a public **IP address**. It is technical suicide to connect your system to the Internet before installing a firewall, as well as installing the latest vendor patches for the hardware, operating system (OS), and installed applications.

Does this mean that there are armies of hackers just waiting to find new targets to attack? Unfortunately, yes—but not, however, in the way you might think. Malicious programs, such as agents, bots, and zombies, perform most of the scanning and attacking automatically. Groups of malicious code, known as botnets, scour the Internet for new victims constantly.

With a firewall protecting your system from the Internet, hacker scans will be nearly worthless. Responses to probing packets will be filtered. An outsider will learn little about your infrastructure. This, in

turn, means they will discover fewer vulnerabilities, thus wage fewer attacks. Even then, those attacks will be more likely to fail.

Firewalls are not used only for Internet protection. Do not forget that significant levels of threat exist internally. Firewalls stand guard against network segments inside your private network (**FIGURE 5-3**). Think of a firewall as a tool to prevent abuse or misuse of LAN resources. Each major department, subnet, or other relevant distinction within an organization should have firewall protection. This will ensure that accidental compromises, as well as intentional abuses, are minimized internally. All networks protect against both internal threats and external ones.




FIGURE 5-3 An example of a private network using firewalls to separate subnets securely.

Although a firewall has its origins in routing, you still want to use a dedicated router for routing and a dedicated firewall for filtering. In many cases, the combination or multifunction devices offer lots of features, but not necessarily best-of-breed performance and efficiency. These may be money-savers for small networks, but they are often cause for expense for larger networks. If a device fails to perform at commercial levels when necessary, the repercussions can cost more in lost revenue and productivity than deploying the proper safeguards in the first place.

Host firewalls protect a host and also protect a network from the host. A host firewall's job is to filter traffic entering or leaving a single computer system. Attacks and malicious traffic can come from the Internet or from the local network. A host firewall can protect the host against such threats. However, the host itself could be compromised by malicious code or be controlled by a malicious user. In this situation, the host firewall protects the network from the threats coming from the host.

Another way of looking at this is that a host firewall also protects the network from a user in general. A user is the riskiest element in a network infrastructure. Even if every other component—hardware or software—does what it was programmed and configured to do, a user can decide to violate security. Thus, the free will of human users is one of the biggest risks to every computing environment.

Users can be ignorant, make mistakes, be tricked by a Trojan horse, be the target of social engineering attacks, or deliberately perform malicious actions. The security of a host system, including general system hardening, anti-malware scanning, Internet client

restrictions, blocking of nonapproved software, and a host firewall, are all designed to protect the network from the risk of human users.

HOW FIREWALLS WORK AND WHAT FIREWALLS DO

Firewalls work along a communication pathway. This can be the gateway point of a network, a chokepoint within a network, a point-of-zone transition, or on a host. In these locations, the firewall interrupts the traffic flow to inspect packets or sessions. If the traffic is authorized, it continues to its destination. If the traffic is not authorized, it is blocked and dropped.

When one thinks of a firewall, one generally pictures some sort of hardware, when in fact it is more accurate to think of the firewall as a function. There are really two broad types of firewall implementations: **bump-in-the-wire** and **bump-in-the-stack**. The bump-in-the-wire (**FIGURE 5-4**) is a separate hardware firewall implementation, and the bump-in-the-stack (**FIGURE 5-5**) is a firewall that is implemented via software.




FIGURE 5-4 A bump-in-the-wire implementation.




FIGURE 5-5 A bump-in-the-stack implementation.

Firewalls operate on a **bastion host** basis (**FIGURE 5-6**). This is most obvious in the case of hardware firewalls, but it is true of host software firewalls, as well. A bastion host firewall stands guard along the pathway of potential attack, positioned to take the brunt of any attack. A firewall acts as the vanguard, as the front line of defense against any attack. A bastion host can also be called a **sacrificial host**.




FIGURE 5-6 An example of a bastion host firewall deployment.

In the case of a host software firewall, the firewall will attempt to prevent all malicious interactions to and from the host. In the event that the firewall is compromised, it usually disconnects the system from the network. Although this is a form of denial of service (DoS), it is often a preferred fail-safe/fail-secure response over defaulting to an open unrestricted and unfiltered connection. This is especially true if the connection is directly to the Internet.

If the firewall is able to rebuff an attack, then the resources are secure. If the firewall fails due to the attack, it prevents any further communication with the resources behind it. Think of a firewall as a **dead-man switch**. If the firewall fails or goes offline, so does the connection it was filtering.

The most common function of a firewall is to screen or filter traffic. The firewall checks any packet received on its interface against its rule set to determine whether to forward or drop the packet. As you read earlier, firewalls typically function on a deny by default/allow by exception security policy.

The firewall performs most traffic filtering based on information in a packet or segment header. This can include the IP address of the source and destination, as well as the source and destination port. Some firewalls can also filter or block specific protocols or certain uses of protocols. For example, a firewall could block all streaming media protocols and block just the Internet Control Message Protocol (ICMP) type 3 (destination unreachable) and 11 (timeout exceeded).

Firewalls differentiate between networks or subnets. A firewall serves as a clear and distinct boundary between one network area and another. By positioning a firewall between network divisions or

subnets, the network designer and security administrators are using traffic management and control for traffic attempting to cross that intersection. Firewalls serve as boundary devices both on the edge of networks facing the Internet, as well as internally between different divisions within an organization.

Open and unrestricted internal communications might sound like a good idea, but in practice such traffic causes severe degradation of overall infrastructure performance and stability. By preventing—or at least limiting—communications between certain divisions of an organization, traffic efficiency increases. Additionally, traffic control and management decrease risk.

For example, by blocking communications between the programming group's network and the production network, unapproved versions of software cannot be accidentally loaded to the production network. Additionally, blocking general access to the accounting subnet, the research and development subnet, the DMZ, and the extranet from the production network protects against data leakage, spread of malicious code, and other forms of fraud and abuse.

Firewalls can act as a general filter for malicious activity or as a one-way sieve. As a general filter, a firewall will allow all normal benign traffic to pass through. This is the type of firewall used to protect a DMZ or differentiate subnets within a LAN. A general filter firewall works to stop malicious activities.

A general filter firewall can allow communication using any protocol on any port or limit communications to specific protocols and ports (amongst other limitations). If only a few specific resources hosted behind the firewall are to be accessed by external users, the

firewall can allow access to those specific internal systems based on IP address and port but block all other inbound requests.

A sieve firewall will only allow traffic to originate from the private or trusted side. A sieve firewall will allow nonmalicious responses to return from the public or less-trusted side, but it will generally block or prevent any initiation or inbound communication request from the outside. A sieve firewall typically protects a private LAN or an extranet.

Firewalls can support special inbound authorized connections with the LAN. One type of specialized content is virtual private network (VPN) links from telecommuters. This is an important feature for an extranet, as well, since VPN access is often the only method to reach hosted resources through the Internet.

Firewalls can provide port-forwarding services. **Port forwarding** is a form of static reversal of network translation. In traditional or dynamic network address translation (NAT), external entities cannot initiate communications with internal systems; any packet received by the outside interface of the NAT system will not find a matching mapping in the translation table, and will therefore be considered invalid and dropped.

With port forwarding, a translation mapping is coded so that an external IP address and port combination are fixed and redirect traffic to a specific internal system, even if the internal system uses a private IP address. The combination of an IP address and a **port number** is known as a **socket**. In **FIGURE 5-7**, the external port 208.40.235.38:8081 is forwarded to an internal server's port at 192.168.5.74:80. Port forwarding can also be called static NAT,

traffic forwarding, service redirection, **reverse proxy**, and “punching a hole through the firewall.”




FIGURE 5-7 An example of port forwarding across a NAT system.

FYI

The phrase “punching a hole through the firewall” points out a very important concern: namely, that using port forwarding reduces the effective security provided by the

firewall. Without port forwarding or other similar services, a firewall can rebuff attempts to communicate with internal systems. Using port forwarding, you create a pathway across the primary border sentry. Several hacker exploits can take advantage of port forwarding to reach other internal systems on other ports than what was “allowed” by the port forwarding settings. These attacks are usually variations of fragmentation manipulation.

Managing and controlling traffic is a primary concern and function of firewalls. Only authorized communications are allowed; everything else is blocked. The act of determining what to allow and what to block depends on the filtering features of the firewall. These include packet inspection, connection or state management, **stateful inspection**, and others. (A **state** is a logical connection between a client and a resource server.) These filtering concepts are defined later in this chapter.

In addition, firewalls can block or filter outbound traffic. This function can block spoofed traffic or anything with an invalid IP address as the destination or source. Specific protocols or ports can be blocked as well. And, of course, specific domain names or destination addresses can be blocked.

Firewalls can also filter based on content. The firewall can intercept specific content in a packet leaving the network before it reaches the outside. This could result in the packet being discarded, an entire connection being dropped, or the packet being edited to remove the blocked content and replace it with something else.

Content filtering can focus on domain name, URL, filename, file extension, or keywords in the content.

Another capability of firewalls is the ability to filter based on encryption. A firewall can allow encryption without restriction. Or, a firewall can be set to block encryption on some ports and not others. Firewalls can also block encryption from and to all internal IP addresses except for servers, VPN devices, and RAS/NAS.

Firewalls can perform intermediary functions between hosts where network administrators deem direct communication too risky. This is the basic function of a proxy. A proxy firewall positioned along a communication path can hide the identity of one or both endpoints of a communication from the participants. In most cases, proxy services hide the identities of internal systems from external entities.

Firewalls can perform address conservation through the use of an address conversion or translation system. NAT is the most common translation service supported by networks. NAT translates between internal addresses and public external addresses. NAT allows a private network to use RFC 1918 private IP addresses. This, in turn, provides additional security, since Internet hosts cannot address an RFC 1918 system directly.

Because perfect security products do not exist, you must rely on monitoring to watch for attempts (and any successes) to breach or violate security. Security is locking things down to the best of your ability, and then watching for attempts to breach your defenses. Lock, then watch. Firewalls are no exception. If a proposed firewall product is unable to record a log of its actions and network activities, you need to seek out a different firewall product. Remember the time-

tested business adage: “You get what you inspect, not what you expect.”

Firewalls can log events. Firewalls can be set to record any action into a log file. In addition, firewalls can record the content of any malicious traffic into a log file. Additionally, any abnormal network activity, performance levels, and traffic statistics can be listed in a log file. Logging events is an invaluable feature of a firewall.

Through the combination of all these features, functions, and capabilities, firewalls provide consistent, reliable protection for an organization’s computer and electronic resources. But firewalls are not perfect and should not be the sole security component of a network infrastructure. They are only one element—although an essential one—of a network security strategy.

TYPES OF FIREWALLS

Listing the types of firewalls is almost like listing the taxonomy of the animal kingdom in biology. The variations, models, and versions are numerous. In addition, opinions vary about what is and is not a firewall. Many experts begin the discussion of firewall types by dividing the collective into two simple, main groupings: personal and commercial.

A **personal firewall** is designed to provide protection to a single system or a small network, such as a SOHO network. Most personal firewalls offer user-friendly interfaces that may be web-based or graphical in nature (that is, a graphical user interface, or GUI). To take full advantage of their features, most personal firewalls do not require special training or certification.

A **commercial firewall** is designed to provide protection for a medium-to-large business network. Most commercial firewalls are quite complex and often require special training and certification to take full advantage of advanced features. Most commercial firewalls use a Unix-like command line interface (CLI) that, while powerful and efficient, is not intuitive.

These two groupings do not, however, represent the complete collective of firewalls, as several could easily fall into either or both categories. Another common grouping method, therefore, is to classify firewalls as either hardware or software.

A **hardware firewall** is a dedicated hardware device specifically built and hardened to support the functions of the firewall software running on it. A hardware firewall is also known as an **appliance firewall**. It doesn't share dedicated hardware resources with any other

service. A hardware firewall does not require any additional hardware or software for its use. All it needs is one or more network connections and a power source. Although a hardware firewall can protect a single system or an entire network, it can only filter traffic that reaches the network interfaces of its appliance. However, a hardware firewall can be positioned on a network at a chokepoint or gateway to analyze and filter all traffic.

■ NOTE

A software firewall and a hardware firewall are both a form of software, but the hardware firewall has a dedicated appliance as its host, while the software firewall uses a standard client or server as its host.

A **software firewall** is an application installed on a host, which is why it is also known as a host firewall. A software firewall depends upon the host's hardware and operating system. If the host's components are not properly hardened, the software firewall will be less effective, especially if there are other communication pathways or attack points on the host. Software firewalls must compete for resources among all other processes active on the host. At best, a software firewall is only able to protect a single host from malicious network activity. A software firewall is only able to filter traffic that reaches the network interface of its host.

Both software and hardware firewalls can be targets of attack. Exploits can compromise a software component, or physical attacks can harm the host/appliance. A software firewall is often less

expensive than a hardware firewall. A hardware firewall typically offers a wider range of features and capabilities than a software firewall. Both software firewalls and hardware firewalls are options you can use throughout a network infrastructure.

Combining these types of firewalls results in four primary options:

- Personal software firewall
- Commercial software firewall
- Personal hardware firewall
- Commercial hardware firewall

A personal software firewall is a product used on individual home systems, SOHO systems, and even client/server network workstations and servers. Generally, a personal software firewall is free or less expensive than commercial software firewall products.

A commercial software firewall is a product used on client/server network workstations and servers. These can be installed on personal or SOHO systems, but the installation is usually expensive and part of an overall security management or network access control (NAC) system. Most commercial software firewalls can be used in an agent/console infrastructure, where each host's firewall is remotely administered from a master management console.

NOTE

Most wireless access points, both consumer and commercial grade, include some form of firewall to provide filtering services for wireless clients as well as

physical cable connections. Many wireless access points could be accurately labeled as routers and/or switches, especially when some include two to six extra-wired connection ports.

A personal hardware firewall is part of an integrated firewall product, such as a wireless access point (WAP) or a cable/digital subscriber line (DSL) modem. Another variation of the personal hardware firewall is the repurposing of a client or server computer into a home-crafted open-source firewall. One example of this is pfSense by Netgate, a software firewall that provides some routing and VPN functionality.

A commercial hardware firewall is usually a device that handles the complexity of larger organizational networks. A commercial hardware firewall is often very expensive, running in the range of \$10,000 or more. Often these firewalls require detailed training for proper setup and deployment.

FYI

Two common, but different, dichotomies are free versus paid and open-source versus closed-source software. *Free* means you do not need to pay for the firewall to use it. *Paid* means you must pay a purchase price and/or a licensing fee to use it. Many examples of both free and paid firewall products are available. **Open-source** means the original source code is available for viewing and modification. **Closed-source** means the distributed

version is precompiled, and the original source code is undisclosed.

Software firewall products, as well as any other form of IT product, can be free and open-sourced, free and closed-sourced, paid and open-sourced, and paid and closed-sourced. You should not assume that, because something is free, it is open-sourced, or that because something is commercial, it must be purchased.

Also, just because a product is free does not ensure its reliability or trustworthiness. Additionally, being able to review the source code does not warrant the reliability of a product.

For further discussion on free software, visit GNU organization's site at www.gnu.org

Personal versions of software and hardware firewalls might include different add-ons or enhancements than their commercial equivalents. These add-ons or enhancements could include antivirus software protection, password management, registry protection, driver protection, VPN gateways, remote access support, IDSs, IPSs, spam filtering, and more. Usually these add-ons make the firewall products more attractive to the potential buyer. However, most commercial entities would generally avoid integrated firewall solutions in favor of dedicated products to handle distinct security or management functions. An integrated device might offer easier administration, but it represents a single point of failure for multiple

services. Additionally, such bundled solutions are more difficult to troubleshoot due to the complexity of the communications supported.

Another variation of firewall, in fact the original variation, is a screening router. Most appliance routers and many software routers, such as Routing and Remote Access Service (RRAS) in Windows Server, are able to perform firewall-filtering services in addition to routing. Screening routers that perform firewall filtering might provide enough sentry security for your needs. However, if you want more advanced features, a screening router is unlikely to be the best solution for your network.

Regardless of the type of firewall or implementation, it is important to be certain that your firewall implements the security and privacy policies that you want, while not conflicting with other security measures. If, for instance, a remote system that you will access requires use of a certain protocol or procedure that is blocked by your local firewall, then you have a problem that must be solved before you can access the remote system.

INDIVIDUAL AND SOHO FIREWALL OPTIONS

Firewall options for the individual or for those running a small office range from native OS firewalls to special-purpose devices. Most individuals and SOHO users are concerned about security but do not want to spend more than necessary. Obtaining a reasonable level of security for most home or small office environments is actually quite easy and cost effective.

■ NOTE

Installing a handful of security tools doesn't mean that you can throw caution to the wind and randomly explore the entire Internet safely. Instead, by using simple security tools, along with common sense, most threats are either blocked or avoided. Potential exceptions can always affect this rule, of course, such as when a trusted site is compromised and your next visit results in the transfer of malware to your hard drive.

First, an understanding of the common threats facing such an environment suggests that firewall options can be simple and inexpensive. Generally, Internet threats fall into two main categories: passive and active. **Passive threats** are those you must seek out to be harmed. For example, you have to visit a website to be harmed by

malicious code embedded in that site. Likewise, downloading infected content occurs only when you elect to click on a download link. Users can avoid most passive threats by making good and safe choices as to where to go and what to do on the Internet.

To address passive threats, modern web browsers include **pop-up blockers**, **cookie filters** (**FIGURE 5-8**), and malicious site managers to limit exposure. The use of antivirus scanning, an anti-spyware scanner, and anti-spam filters addresses other common threats. Adding a firewall to your **security stance** protects against most of the passive threats on the Internet that might be casually triggered.

 Find in Options

Browser Privacy

Content Blocking



Block third-party content that tracks you around the web.
Control how much of your online activity gets stored and shared between websites. [Learn more](#)

[Manage Exceptions...](#)

☐ **Standard**

Only blocks known trackers in Private Windows.

☒ **Strict**

Blocks all trackers Firefox detects. May cause some sites to break.

 Known trackers in all windows

 Third-party tracking cookies

 Heads up!

Blocking content can cause some websites to break. It's easy to disable blocking for sites you trust. [Learn how](#)

☐ **Custom**

Choose what to block.



FIGURE 5-8 Privacy section of Firefox's Options dialog box. Note the cookie settings.

The other category of Internet threat is active. An **active threat** is one that takes some type of initiative to seek out a target to compromise. This can be a hacker, an intruder, or an automated worm. In any case, an active threat seeks out vulnerable targets. If you do not have reasonable security deployed and an active threat discovers your system, you might be at risk for a compromise.

Fortunately, most individuals and SOHO environments are not significant or primary targets of most hacker activity. There just is not enough value or benefit in spending the time and effort to compromise a moderately secured single computer or small network. If your business is getting lots of traffic and orders, or you happen to be hosting hundreds of systems, then you naturally become a more valuable target to hackers.

In most circumstances, the average home user and most small offices (such as those operating a dozen or fewer computers in a network) can obtain reasonable security protection from very simple and inexpensive firewall options.

The first firewall option to consider is any **native firewall** of the operating system. Most operating systems include a default or native firewall. If yours does, investigate its features and options before automatically tossing it aside. Many of the firewalls found in the latest releases of operating systems are as good as or better than most commercially available third-party host software firewall options.

One of the best-known examples of a native operating system firewall is Windows Firewall. Windows Firewall made its debut in the Windows OS as part of Service Pack 2 for Windows XP. Since its initial release, Windows Firewall appeared as a standard native

feature in Windows Vista, Windows 7, and Windows 10. The feature set and uses of Windows Firewall are discussed later in this chapter.

■ NOTE

Microsoft changed the name of Windows Firewall in Windows 10 to “Windows Defender Firewall”; however, you will see it referred to as “Windows Firewall” in certain instances.

Other operating systems may or may not have a firewall as well recognized as Windows Firewall, but you should still investigate the standard options and feature set provided natively before seeking out alternatives. After all, you would never custom turbocharge a brand-new car until you had driven it and road tested its standard features.

Most Linux distributions do not automatically come with a native software firewall. When Linux does provide a native firewall, it is not likely to be enabled or configured by default. Due to the wide range of variations in Linux distributions, Linux provides a range of firewalls. Linux firewalls include older products, such as ipchains and iptables, as well as newer options such as PF, Netfilter, and Vyatta.

In addition to the free or native OS firewall options, individuals and SOHOs can use firewalls hosted by Internet service provider (ISP) connection devices, routers, or WAPs. An ISP connection device is any hardware connecting a local network—or even a single computer—to a telco’s carrier network to access the Internet. Most ISP connection devices, including cable modems, DSL modems, satellite modems, routers, and wireless modems, include firewall

features. A firewall feature may or may not represent a fully functional firewall. Some are basic filtering tools that block IP addresses, ports, or protocols using a simple blacklist technique.

An ISP connection device may offer wireless connectivity. If the provider does not offer a wireless option or requires you to pay extra for wireless via their equipment, you should consider deploying your own WAP. Most WAPs for consumers and SOHO environments cost under \$100 and offer firewall services, as well as wireless connectivity.

The use of native host OS firewalls and a hardware firewall provided by an ISP connection device or WAP is usually more than sufficient filtering security for a home user or a small office network. However, if you want to explore other options or you have a larger, riskier, or more sensitive environment, consider the other firewall alternatives.

Managing the Firewall on an ISP Connection Device

Most ISP connection devices use a web interface. To initiate access to the management interface, open a web browser and enter the IP address of the device where you would typically enter a URL. In most cases, the device's IP address is the default gateway address, DHCP address, or possibly DNS address of the client's interface directly connected to a physical port on the device.

TECHNICAL TIP

If the ISP connection device provides firewall services, obtain a copy of the most current manufacturer's user manual for the device. This documentation will provide the best initial information on accessing and configuring the device's firewall and is usually available online at the manufacturer's website.

Attempting to open the configuration interface is likely to result in a prompt for authentication credentials. The vendor's user manual should indicate the device defaults for username and password—which is why it is essential that you change these on your devices—but be sure to record the account information and keep it someplace safe. If not, search the Internet using keywords such as “default password” along with the device name, make, and model. If that fails, try username “admin” with a password of “admin” or “password.” It is shocking how often these are correct. In Linux systems, it is common to use “root” as the username and “toor” as the password.

Your ISP has likely locked down access to the connection device. This is sometimes done as a precaution against the uninitiated, who might cause increased technical support hassles. If you discover that your ISP connection device is locked down, try calling your ISP and asking them to grant you access (often by their revealing the credentials to log into the device).

If they refuse to offer this information, then you have four choices:

1. You can accept their refusal as “just how it is” and employ some other device as a hardware firewall.
2. You can change ISPs or service providers to a carrier that will grant you access to configure the device.
3. You can try to replace the carrier’s device with one that you own and fully control.
4. You can seek out ISP and device-specific information on the Internet, which might include bypass or hacking details.

■ NOTE

Cable providers are generally less likely to allow you access into their connection device, while most DSL providers seem willing to grant at least partial access.

However, be very cautious when choosing this last route; it is often unproductive, probably a violation of the ISP terms of service, unethical, and may even be illegal.

If you are unable to find a legitimate path to accessing the configuration interface of your ISP connection device, hacking the device should be your last and final option. Hacking an ISP-provided

device will void your contract and could make you liable for the cost of the device, plus fines and possible legal action. A more ethical solution would be to replace the device with one you own and control, or switch ISPs altogether.

■ NOTE

Be careful about searching for websites that offer information on how to hack or bypass ISP equipment. Often the sites are booby-trapped with malware to trick the unsuspecting visitor. You may be able to download the instructions, but often it is a Trojan file, which includes malware to infect your system. In addition, some ISP equipment may be leased—not sold—and attempting to access internal features may violate the contract, if not the law.

If all else fails, add your own firewall between the ISP connection device and your first networked system. With this configuration, you are still gaining the benefit of an in-line firewall without violating any contracts or flirting with legal troubles. In most cases, you'll want to supplement or disable an ISP connection device's firewall anyway. Thus, if your ISP blocks access to it, just design your security as if the device offered no firewall protection.

Converting a Home Router into a Firewall

Many home routers have sufficiently robust features that you can configure these to function as firewalls. To ensure ease of use (and minimize product returns), these devices are usually shipped with all security features turned off. Because “plug and play” results in immediate connectivity, many users never go back to configure their equipment for security. Thus, hackers and malware are often able to intrude easily into home networks. You can prevent that by understanding what is available in this class of products.

To access a home router, type the IP address of your gateway into your browser. To determine the IP address on a Windows machine, choose Start > Run, type Command, and at the C:\> prompt, type ipconfig. Most routers ship with a default address of 192.168.1.1. If you remember, 192.168.X.X is a nonroutable range of IP addresses, which means that traffic inside that address range will not be shared with addresses outside that range. On a Linux system it is a similar process, except the command is ifconfig on a wired system and iwconfig on a wireless Linux device.

Once logged in, consider changing a number of settings (**FIGURE 5-9**). First, if the router is a wireless device, change the **service set identifier (SSID)** from the default setting. You do not want to be one of eight access points named “ASUS” in your neighborhood. Why is this important? Most PCs will grab the strongest wireless signal from a known SSID. If you have standard SSIDs—such as ASUS—in your wireless access table, you may find yourself on a neighbor’s (or hacker’s) system someday without realizing it. Another good idea is to change the default IP address range to something other than 192.168.1.X. This creates a custom

range for your own network. If you hard-code your internal network range to something like 10.20.30.X, you are less likely to stumble into that situation.



FIGURE 5-9 The initial configuration screen on an ASUS wireless router. Courtesy of Asus.

Most routers enable Dynamic Host Configuration Protocol (DHCP) by default. You could lock down your network by either hard-coding IP addresses into each authorized machine (and turning off DHCP) or setting a strict upper limit on the number of devices permitted to have DHCP leases (most default settings are 100). If you have five devices, limit the number of connections to five. (Do not forget about the Xbox 360 or PS3—these devices need IP addresses too.)

 NOTE

Do not forget to change the default administrator password of your device.

For firewall settings, most routers will have a configuration page to block services or control port access. Determine which ports you need to access the Internet, and then block all of the rest. Note that this setting restricts outbound traffic. So, why block it if it originates from within? Malware, zombies, bots, and other hostile applications usually have to connect to the outside to do any damage or steal your property. If you accidentally download dangerous malware onto your laptop at your favorite coffee bar, and then connect to your home or office network behind the firewall, that malware may have access to all of your peer systems. However, if you are blocking all nonessential outbound traffic, then most malware will not be able to exfiltrate your sensitive information or ask its remote server for instructions on what to do next to your system. Only those

applications that use a common port (like port 80) will be able to get through.

Which outbound ports should you block? First, consult your policy, or determine what programs you are using. For the most part, you'll want to allow the following outbound ports to be open at all times:

- **Port 25**—SMTP (outbound mail)
- **Port 53**—DNS (name resolution)
- **Port 80**—HTTP (Internet access)
- **Port 110**—POP (initiate request for inbound mail)
- **Port 443**—HTTPS (secure Internet)
- **Ports 465 and 995**—SMTP and POP (if using Gmail or other email packages)
- **Ports 1024–1035**—DCOM ports for downloading files (increase number of ports based on number of systems protected; 10 is usually sufficient for a home network)

Beyond port 1035, you may not need to allow outbound traffic, unless you are using cPanel to access an externally hosted website on port 2083, or port 11371 if you are looking up PGP keys. By blocking all other high ports, you will also quickly hear from coworkers or family members who may have been using some of these ports, which are dedicated to specialty software or gaming. At home, you might want to allow (or set time limits); at work, you probably want to block these ports, as gaming is usually not an acceptable use of company resources.

In general, you should not accept any connections that originate from outside your firewall. Your policy may permit exceptions, such

as remote access tools, so be careful about blocking everything. Some home firewalls do not provide a direct way to block specific incoming ports, but most allow you to do port forwarding. A clever way to use this to thwart external attacks is to forward to a nonexistent port. So, for example, if your local network range includes from 10.20.30.40 to 10.20.30.49, you might forward incoming connection requests to port 10.20.30.99—where no device is listening. Ports you might consider forwarding are:

- **Ports 20 and 21**—FTP-data and FTP. Prevent external connections from downloading your files. If your business requires file transfers, safer alternatives are SFTP (Port 22) and FTPS (Port 990).
- **Port 23**—Telnet. Prevent external connections from insecurely logging into your internal systems.
- **Port 53**—DNS. Prevent external entities from poisoning your DNS cache.
- **Port 80**—HTTP (unless you are running a web server from behind your firewall, which is a bad idea).
- **Ports 81 and 82**—Often used as “overflow” for port 80. No valid use, so block these.
- **Ports 137, 138, and 139**—NetBIOS. Often exploited by malware, this provides access into Windows systems.
- **Port 443**—HTTPS (unless you are running a secure web server from behind your firewall, which is still a bad idea).
- **Port 445**—NetBIOS for Windows 2000 and later.
- **Port 3074**—Xbox game port. Do not allow strangers to connect to your Xbox while you’re away. Remember—it “lives” on your internal network.

An excellent way to test your configuration is to go to Steve Gibson's <http://www.grc.com> website and run his free ShieldsUP! port scanning tool (**FIGURE 5-10**). Although the website has an admittedly outdated appearance (at the time of this writing), the tool is legitimate and can be used to scan ports and services, such as Universal Plug and Play (UPnP), often targeted by malicious actors. Most systems end up scoring poorly. If you can achieve 100 percent stealth, you've done a great job.

Shields UP!!

Port Authority Edition – Internet Vulnerability Profiling

by Steve Gibson, Gibson Research Corporation


Universal Plug n'Play (UPnP) Internet Exposure Test

This Internet probe sends up to ten (10) UPnP Simple Service Discovery Protocol (SSDP) M-SEARCH UDP packets, one every half-second, to our visitor's current IPv4 address (**45.19.50.87**) in an attempt to solicit a response from any publicly exposed and listening UPnP SSDP service. The UPnP protocols were **never** designed to be exposed to the public Internet, and **any** Internet-facing equipment which does so should be considered defective, insecure, and unusable. Any such equipment should be disconnected immediately.

Your equipment at IP:

45.19.50.87

Is now being queried:



THE EQUIPMENT AT THE TARGET IP ADDRESS
DID NOT RESPOND TO OUR UPnP PROBES!

(That's good news!)

FIGURE 5-10 ShieldsUP! port scan result confirming a well-configured set of home router firewall rules.

USES FOR HOST SOFTWARE FIREWALLS

The next step up from a native OS firewall (or even an OS without a firewall) is a third-party host software firewall. These options include both open-source and commercial software firewalls for most OSs.

You can use a host software firewall in several situations. The first and most obvious use is simply to protect a client system. Keep in mind that a host software firewall provides protections for both inbound and outbound communications. A host software firewall protects the client from compromises on the network and protects the network from compromises on the client.

Use a host software firewall as an additional layer of protection on a server system. Most server operating systems do not include a host software firewall. Therefore, a dedicated firewall appliance is often deployed on the network. A host software firewall on a server is never a substitute for an appliance firewall. However, it can be a supplement.

■ NOTE

Before installing any third-party software firewall, always double-check for full compatibility with your current OS version and patch level. If the firewall's documentation does not specifically list your OS as fully compatible, do not assume that the device will work properly. Firewall security is not something to leave to chance.

A host software firewall can provide firewall-filtering services in relation to a VPN. Just because the VPN link may be encrypted does not guarantee that the other end of the VPN connection is as secure as you might desire. Using a host software firewall in conjunction with either a software host VPN (such as a transport mode VPN or a remote access VPN) or an appliance VPN adds an additional layer of protection against compromises that could traverse the VPN connection.

NOTE

When installing a third-party software firewall, make sure all native or other firewalls are disabled or uninstalled. Do not attempt to run two software firewalls simultaneously on the same computer system. It is acceptable to run different software firewalls on different systems and to even use one or more appliance hardware firewalls. Just do not attempt to use two software firewalls on one system.

A host software firewall can provide modest protection for small networks. Home networks, gaming networks, and small office networks are sometimes constructed using a primary system connected to the Internet that shares that connection with a small network off a secondary network interface. On a Windows system, the Internet Connection Sharing (ICS) service makes this type of network configuration simple. Use a host software firewall and

provide the secondary network with modest firewall-filtering services.

A host software firewall likely has many other uses. Use host software firewalls in any network configuration with the proper research and testing. The goal is to establish additional layers of security, not to conform to static notions of design and implementation.

Examples of Software Firewall Products

Software firewall products are important options to consider when designing and deploying a security solution for home environments, as well as corporate IT infrastructures. A host of third-party software firewalls are worth considering. Here are just a few of the more widely known options:

- Check Point ZoneAlarm (free and retail)
- Comodo Firewall (free)
- eConceal (retail)
- Jetico Personal Firewall (retail)
- Lavasoft Personal Firewall (retail)
- TinyWall Firewall (free)
- GlassWire Firewall (retail)

In addition to standalone, third-party firewalls, some firewalls come packaged as part of a security suite. These suite-member firewalls are not available as standalone products. However, the collection of security applications might be worthwhile if you do not already have existing solutions for antivirus or anti-malware. Some security suites to consider include:

- AVG Security Suite
- Computer Associates (CA) Internet Security
- F-Secure Internet Security
- McAfee Personal Firewall Plus
- Avast Security Suite
- Microsoft Security Essentials

- Norton Internet Security and Norton 360
- MalwareBytes Internet Security
- Trend Micro Internet Security
- Webroot Internet Security Essentials

Using Windows 10's Host Software Firewall

The native **Windows Defender Firewall** in Windows 10 (**FIGURE 5-11**) is a sufficient security measure for many situations. Before rushing to replace this free security component, take the time to evaluate the benefits of this capable firewall option. Windows Defender Firewall in Windows 10 is a host software firewall. However, it can be used in a variety of situations and network configurations for most home, SOHO, and mobile environments.




FIGURE 5-11 Windows Defender Firewall with Advanced Security configuration dialog box.
Used with permission from Microsoft.

Windows Defender Firewall in Windows 10 includes configuration profiles, so you can create custom firewall configuration settings for work, home, and public connections. This allows strict limitations in public, modest settings at work, and more options available when accessing from home (or whatever your preferences). The benefit is that, once configured, the firewall will adjust its settings based upon the network connection each time you are connected to a known, previously accessed network.

Windows Defender Firewall creates a password-protected homegroup or workgroup that allows file- and printer-sharing between systems authorized by a password. This is an improvement over previous versions of Windows Firewall, which often encouraged users just to turn off the whole firewall rather than properly configure file- and printer-sharing access rules. In addition, this applies not just to Windows systems, but to any devices or computers recognized as media sharing devices (such as an Xbox 360).

Other Windows Defender Firewall improvements in the Windows 10 version include a more granular control and configuration management interface, more extensive logging, and extended ability to be managed from a command line (using the netsh advfirewall firewall command instead of the previous netsh firewall command).

Although it is not revolutionary and still lacks a few features, such as being a true two-way personal firewall with program control, Windows Defender Firewall is a worthwhile host software firewall for most clients in most network situations. That said, you should still explore how this product fits your own computing environment and security needs.

Using a Linux Host Software Firewall

A Linux system can benefit from a host software firewall, or it can support a software firewall for a network. The first idea is just to install a host software firewall for the benefit of the local user. This is the same idea as Windows Defender Firewall in Windows 10. A variety of open-source and commercial host firewall options are available for Linux, including:

- IPFire
- pfSense
- IPtables or Netfilter
- Untangle NG
- UFW
- VyOS
- Smoothwall

If you selected Linux for its low cost of entry, then selecting an equally low-cost host software firewall is often an attractive option. However, paying for commercial host firewall products might offer a greater range of functions or services, along with better technical support.

Using a Linux software firewall as a replacement for a commercial firewall appliance can be a very cost-effective solution. Linux often can repurpose computer hardware that is no longer sufficient to support larger, bulkier, more resource-intensive operating systems, such as Windows. Linux can often extend the useful lifetime of computer hardware by several years. A repurposed computer

system running Linux is a great option for use as a software firewall host.

USES FOR COMMERCIAL SOFTWARE NETWORK FIREWALLS

Commercial software network firewalls install onto your hardware and provide network-level security services. Some commercial software network firewalls install on top of existing operating systems, such as Windows. Others are complete OS replacements, many of which are Linux-based. Some examples of commercial software firewalls include:

- Check Point
- Cisco Systems Next-Generation Firewall Virtual (NGFv) and Adaptive Security Virtual Appliance (ASAv)
- Fortinet FortiGate - VM
- Huawei
- SonicWall

As with any firewall selection, know your environment and understand your security needs before shopping for a firewall solution.

USES FOR HARDWARE/APPLIANCE FIREWALLS

Appliance firewalls, whether called device or hardware firewalls, are common and nearly essential elements of every moderate to large network infrastructure. Appliance firewalls can range from low-end consumer-grade to very expensive high-end, commercial-grade solutions.

A variety of manufacturers and vendors make and service appliance firewalls. Some of the major vendors/brands to consider include:

- Barracuda
- Cisco Meraki
- Fortinet
- Juniper Networks
- Linksys (owned by Cisco)
- NetGear
- SonicWall

When selecting a hardware or appliance firewall, keep a few important points in mind. Some of the most common concerns for businesses include ease of use, secured management interfaces, port-filtering support, stateful inspection filtering, and the ability to be firmware/software upgraded.

Never skimp on throughput. Firewalls often represent bottlenecks to network bandwidth and thus should be selected to maintain wire speed. Be sure a hardware firewall can more than handle the current network speeds and allow for future growth. If you are currently pushing a 1 Gbps network, consider a firewall capable of filtering at 2.5 Gbps wire speed or higher.

For larger networks, centralized and remote management options are often essential. If firewall management requires direct physical contact or if you can configure only a single firewall at a time, you may find these significant hindrances to managing very large networks. An important part of a realistic firewall solution for enterprise networks can be multiple device management, including simultaneous configuration synchronization features.

Consider whether add-ons, upgrades, or extras are available and whether that's important to your decision. Some firewall devices convert to firewall-plus devices or true multifunctional devices. Additional features may include email scanning, message quarantine, attachment stripping, virus scanning, mobile code filtering, anti-spyware, intrusion detection system (IDS)/intrusion prevention system (IPS) features, spam filtering, compliance monitoring, and network access control. Products that support expansion or firewall additions are known as unified threat management tools or may fall under the heading of advanced IDS/IPS.

Whatever the options presented by a vendor, always consider those based on your actual current and future network needs. Just because a product is expensive does not guarantee it will work better than a free, build-it-yourself alternative.

TECHNICAL TIP

It's often possible to replace an appliance or device firewall's OS with a third-party alternative. This can apply even to ISP devices and WAPs. Some of the better-known device firmware replacement options are DD-WRT, Open WRT, and Tomato.

NEXT-GENERATION FIREWALLS

A next-generation firewall (NGFW) is a device that offers additional capabilities beyond traditional firewall functionality. Traditional firewalls may offer stateful packet inspection while a NGFW may provide integrated IDS/IPS functionality. This is considered third-generation firewall technology and can offer several advantages when deployed on a network.

The ability of a single device to perform multiple functions can be cost effective, especially for smaller business networks or situations where there is a lack of expertise in-house to support the technology. One equipment provider is able to aid in the planning, placement, and troubleshooting of the device, although it performs several functions on the network. One device is often easier to manage and support than a series of individual appliances all offering different functionality. Interaction with one appliance interface can be easier for network administrators than learning a variety of interfaces and device languages.

Drawbacks of using a NGFW all-in-one type device include a single point of failure if the device fails. This can be especially problematic if there is no layering of security devices and techniques to provide ample defense in depth. The expense of a multi-purpose device can easily be many times that of a standalone appliance. When something goes wrong, support can be more expensive if not part of an all-inclusive support plan. Finally, the device is more complicated

due to all the various services offered, which makes deployment, maintenance, and log management more complex.

WHAT ARE VIRTUAL FIREWALLS?

The term **virtual firewall** describes a variety of firewall and firewall-like concepts. This can include virtualized software firewalls that provide filtering services for a standard physical network as well as firewalls running between virtualized client and server OSs. In theory, the use of a software firewall as a replacement for a network appliance could work as long as the host OSs network communication is routed through the virtual firewall before leaving the host's network interface controller (NIC).

This is a relatively new and growing area for firewall deployment. Virtualization offers numerous benefits over a traditional, single OS to a single hardware box deployment. Virtualization allows for rapid development, quick prototyping, isolation, traffic management, quick recoveries, testing, and so on.

By virtualizing firewalls along with OSs, you can craft new network architectures that may not exist in the traditional network architectural concepts. For example, with a virtualized firewall, you could route every communication between every virtualized OS through the filtering services. This would be the equivalent of deploying an appliance firewall between every system, but without the hassle, expense, or complexity.

Virtual firewalls are not a panacea. Virtual firewalls will not be useful in every situation, but these are an interesting new option for

deployment to monitor, manage, and filter network traffic over traditional or virtualized network segments.

DUAL-HOMED AND TRIPLE-HOMED FIREWALLS

Firewalls, specifically hardware appliance firewalls, typically have two or more network interfaces. A firewall with two interfaces is known as a **dual-homed firewall**, while a firewall with three interfaces is known as a **triple-homed firewall** or a three-legged firewall.

The benefit of multiple interfaces is that the segments, subnets, or networks connected to each firewall interface are electronically isolated from each other. This prevents unfiltered traffic from leaping from one segment to another in an attempt to bypass firewall filtering.

However, for software firewalls using multiple interfaces, you need to ensure that the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol feature called IP forwarding is disabled. IP forwarding is actually a router rule that allows traffic from one interface to exit another interface without needing to move any further up the protocol stack than where IP resides. In many cases, IP forwarding allows packets to bypass filtering. If the system is to function as a firewall, you should disable this feature.

Software host firewalls are most often single-homed firewalls, since the host has only a single NIC. This is acceptable, since the software host firewall is not providing sentry services between network segments, except between the host and the network.

While firewalls with four or more interfaces are possible, these are rarely deployed. Such a configuration requires significantly more complex filtering and routing rules to operate effectively. It's also a significant single point of failure and a bottleneck to traffic flow between multiple network segments.

■ NOTE

If a system is to function as a router rather than a firewall, IP forwarding might be a desirable function.

IPV4 VERSUS IPV6 FIREWALLS

No real distinction exists between a firewall designed for IPv4 and one designed for IPv6. Many firewalls can already support both versions of IP. If you are planning to migrate to IPv6 or have already started the conversion process, be sure your firewalls support IPv6.

A small issue affects filtering between IPv6 and IPv4 subnets. A protocol translation tool can support interaction between networks using different versions of IP. This translation tool is called Network Address Translation–Protocol Translation (NAT-PT) and was defined in RFC 2766 by the IETF. Be sure your selected firewall supports NAT-PT if you plan to communicate across an IP version barrier.

If you would like to read more on this translation issue or gain further understanding about IPv6, please visit:

- The ARIN IPv6 Wiki at [*https://www.arin.net/vault/knowledge/getipv6_info.html*](https://www.arin.net/vault/knowledge/getipv6_info.html)
- IPv6.com at [*http://www.ipv6.com*](http://www.ipv6.com)

INGRESS AND EGRESS FILTERING

Ingress and egress filtering are common filtering practices to eliminate spoofing. A source address that comes from the opposite side of the firewall than where it is assigned is obviously a spoofed address. For example, this happens when an internal LAN address appears as a source address in a packet on its way into a network from outside. This form of spoof filtering can be part of ingress filtering.

Likewise, the same process can filter for packets leaving a network. If a packet with a source address from the outside, such as an Internet address, is received by a firewall from an interface inside the private LAN, this is also a spoofed address. This form of spoof filtering can be part of egress filtering.

Ingress and egress filtering can expand beyond protection against spoofing and include a variety of investigations on inbound and outbound traffic. This can include blacklist and whitelist filtering, protocol and port blocking, and confirmation of authentication or authorization before communications continue.

Unfortunately, if a packet's spoofed addresses do not violate any of these concerns, the spoofed addresses might not be as easy to detect. For example, if a client spoofs an IP address to look like another client in the same subnet, the rules just described to catch spoofing would fail to notice this spoofed communication. In addition to basic ingress and egress filtering, firewalls can support additional forms of packet examination and investigation.

TYPES OF FILTERING

Filtering is the primary function of a firewall, and most devices can support many different forms of filtering. Through its filtering services, most of the other benefits and capabilities of firewalls apply. Additionally, the terms for the type of filtering and the type of firewall are often used interchangeably. For example, a firewall that supports packet filtering is known as a packet-filtering firewall.

Static Packet Filtering

The most common form of filtering is **static packet filtering**. Static packet filtering uses a static or fixed set of rules to filter network traffic. The rules can focus on source or destination IP address, source or destination port number, IP header protocol field value, ICMP types, fragmentation flags, and IP options. Static packet filtering focuses on header contents and does not examine the payload of packets or segments. Static packet filtering is therefore mainly focused on the Network Layer (Layer 3), but it can also include Transport Layer (Layer 4) elements.

Static packet filtering is fast. Traffic matching a deny rule gets dropped, while traffic matching an allow rule is allowed to continue toward its destination. Static packet filters are invisible or transparent to hosts and users unless the traffic is blocked; then users will notice the actions of the firewall on their communications.

Static packet filtering requires the firewall administrator to define and tune the rule set. Most firewall rule sets, including static packet-filtering rule sets, are a first-match ordered system. Static packet filtering can be problematic when the rule sets get too large. If the rules are in the wrong order or in a chaotic order, the rule set could create loopholes or unintentionally discard authorized traffic.

The complexities of network communications, such as accepting query responses and port shifts, can be difficult to manage with static packet filtering. Static packet filtering may allow the subsequent packets of a fragmented message through, even though the lead packet was dropped. This can result in a DoS on the destination system, which would be waiting for the lead packet that never arrives.

Additionally, static packet filters perform analysis on individual packets, regardless of the relationship or correlation between previous or future packets in a communication stream. This could allow complex multipacket attacks to bypass the firewall if each individual packet is not recognized as malicious independent of other packets in the stream.

Static packet filtering should still be used as a first line of defense in spite of its shortcomings. By using static packet filtering as the first layer of defense, subsequent layers of filtering will have less bulk to address and thus can operate more efficiently. Proper placement and configuration aid in successful packet filtering.

Stateful Inspection and Dynamic Packet Filtering

A state is a session of communication. Stateful inspection addresses the issue of complex malicious traffic. Stateful packet filtering determines whether or not a current packet is part of an existing session, and allow/deny decisions are made based on this determination. Often, state refers to the Transport Layer (Layer 4) protocol TCP's virtual circuits established through the three-way handshake (using the SYN, SYN/ACK, and ACK flagged segments). Stateful inspection systems can also track communications in Layers 5–7. A stateful inspection firewall will keep track of current sessions in a state table stored in memory.

As the firewall encounters each packet, it is analyzed to determine whether or not the packet is part of an existing state. If not, it is likely to be dropped unless it is a packet used to help initiate a new authorized session. This stateful investigation can be considered a dynamic packet filter, as well. With static packet filtering, rules had to be created to allow the outbound requests and the inbound replies. With **dynamic packet filtering**, once a session is established, the firewall watches for packets that do not belong to authorized sessions. Using stateful inspection as dynamic packet filtering allows for simpler rule sets. A rule allows an outbound connection and the firewall's state management automatically allows the return traffic.

Unfortunately, stateful inspection can sometimes be fooled through manipulation of header contents that makes malicious traffic appear to be part of an existing valid session. More advanced stateful inspection filters keep track of the basic endpoints of a session, and additional details about the session, such as the sequencing and

acknowledgment numbers. This reduces the risk but does not fully eliminate it, as a hacker can eavesdrop on a session, learn the sequencing numbers, and predict future valid sequences.

Another issue with stateful inspection is that not all traffic uses states. Specifically, UDP and ICMP are connectionless protocols, so state management will not apply to them. For these protocols, the firewall acts as if a state does exist and keeps track of the source and destination from outbound packets. These are added to the state table with a timeout value. If the timeout occurs before a response is received, the state is removed from the table. A hacker can fool this mechanism in the same way as any stateful protocol like TCP.

■ NOTE

Stateful inspection is a form of dynamic packet filtering. Dynamic packet filtering is the process of automatically creating temporary filters. In most cases, the filters allow inbound responses to previous outbound requests, but on a limited timeout basis. Both forms of state or session awareness, dynamic packet filtering and stateful inspection are usually interchangeable terms.

Network Address Translation

Network address translation (NAT) is not exactly a form of filtering, but it is often included in lists of the filtering services or options provided by firewalls. NAT translates internal addresses into external addresses. NAT can perform this service against IP addresses, as well as port numbers. Any firewall that supports NAT can be a NAT firewall. In most cases, NAT is an additional translation service to the core filtering functions of a firewall. NAT is a common, if not standard, feature of modern firewalls.

Application Proxy

An application proxy, application firewall, or application gateway is an application-specific version of a packet filter. However, unlike a static packet filter that is only able to inspect the header of a packet or segment, an application proxy is able to inspect traffic fully at any layer, including the application payload.

An application proxy, even if given the name firewall or gateway, acts as the go-between or middleman between a client and a server. All communications for the specific application are proxied. This grants the application firewall the ability to inspect application-specific elements of the traffic. Application proxies are application-specific, therefore specific products for email, web, file transfer, database access, Voice over Internet Protocol (VoIP), and other TCP/IP sub-protocols are available.

When an application proxy is deployed, it usually requires that all client software is reconfigured to point communications to the proxy server rather than the actual intended resource server. The application proxy will rebuild the request packet before sending it to the resource server. This can include NAT services, but in most cases, it is just a process of proxying the communication. The application proxy maintains two connections, one with the requesting client and a second with the resource server. Thus, application proxies are not transparent filters, because a client is aware the proxy is in use. The client never establishes a direct connection to the resource server when a proxy is involved.

All other firewalls monitoring a network border must deny access for the application protocols to be managed by the application proxy. This prevents a user from attempting to bypass the application proxy.

Application firewalls can filter on the content of the application payload. This can include IP addresses, domain names, URLs, subprotocols, attachments, keywords, and more. An application proxy can inspect every aspect of an application's communications. This is known as deep packet inspection. Application proxies can also perform caching services to improve performance and reduce connection throughput consumption. The primary limitation of application proxy firewalls is that each unique application will need its own dedicated application proxy. Generic proxy systems are usually ineffective.

Circuit Proxy

A circuit proxy or circuit firewall focuses its filtering on the initial setup process of a session, state, or circuit. This form of filtering can focus on Layers 3–5. It functions similarly to an application proxy, as it acts as a middleman between a client and server. A circuit proxy prevents a direct connection from existing between a client and server to protect the network.

A circuit proxy makes an allow or deny decision on the initiation of the session, state, or circuit. Once a circuit is created, no further filtering takes place. If a client is allowed to initiate communications with a resource server, then the content of the communication is unfiltered and unmonitored (at least by the circuit proxy).

The filtering rules of circuit proxies are similar to those of static packet filtering in that a list of rules of IP addresses, port numbers, domain names, networks, or even resource providers determine which circuits or connections are allowed and which are not. The filter set can be a deny all/but allow exceptions stance, or an allow all/but deny exceptions stance.

Content Filtering

Firewalls can also filter based on content. The firewall can intercept specific content in a packet leaving the network before it reaches the outside. This could result in the packet being discarded, an entire connection dropped, or the packet edited to remove the blocked content and replace it with something else. Content filtering can focus on domain name, URL, filename, file extension, or some other form of keyword. Content filtering is often a feature of application proxy firewalls, stateful inspection firewalls, and dynamic packet-filtering firewalls.

SELECTING THE RIGHT FIREWALL FOR YOUR NEEDS

Choosing the best firewall for your specific environment requires knowledge of firewall options and understanding of the security needs of your network environment. Automatically selecting the most expensive option, the least expensive one, or an open-source solution fails to properly assess the situation. Firewalls are essential elements of security. It would be careless to select a product without due diligence.

■ NOTE

Selecting and purchasing the right firewall for your organization's situation requires careful evaluation. As with any big decision, always resolve important aspects of the process sooner rather than later. For example, first familiarize yourself with the budget for the purchase.

Before you research and evaluate firewalls, you should make sure that your organization has created a written security policy. You would then select the firewall that can fulfill and comply with the policy.

The basic concept of any firewall is the same, so you should evaluate a firewall based on the level of security and features it offers. First and foremost, the firewall should be able to fulfill and

comply with your security policy. Some other characteristics to consider are:

- **Security assurance**—Independent assurance that the relevant firewall technology meets its specifications.
- **Privilege control**—The degree to which the product can impose user access restrictions.
- **Authentication**—The ability to authenticate clients and allow different types of access control for different users.
- **Audit capabilities**—The ability to monitor network traffic, generate logs, and provide statistical reports. Logs may include both authorized traffic and unauthorized attempts and may be able to trigger alarms in response to certain events.

When considering the features, consider the product's ability to meet the needs of the organization. A good firewall product should provide:

- **Flexibility**—The firewall should be open enough to accommodate the security policy of your organization, as well as allow for changes. Security policies can change, and security procedures should be able to change and adapt to those different needs.
- **Performance**—A firewall should be fast enough that users do not notice the screening of packets. The volume of data throughput and transmission speed associated with the product should be reasonable and consistent with the organization's network bandwidth to the Internet.
- **Scalability**—Firewalls should be able to handle additional workload. The additional workload can be due to growth within the organization or due to increased use of the Internet connection.

When considering integrated features, consider the ability of the firewall to meet your network and users' needs. This includes:

- **Ease of use**—The firewall product should ideally have a graphical user interface (GUI), which simplifies the job of installing, configuring, and managing it.
- **Customer support**—Vendors that sell firewalls often provide support. This includes providing prompt access to technical expertise for installation, use, and maintenance. It also includes training.

Knowing the amount of money available for firewall protection helps you quickly eliminate those products clearly outside your price range. Do not automatically discard every product above the budget limit, but realize that anything more than about 15 percent of your maximum allowance is probably out of reach.

Next, consider whether an open-source solution is a viable option for your corporate culture. Some organizations are philosophically opposed to open-source solutions, while others embrace and encourage such endeavors.

Do you want to include the option to build your own or focus on only off-the-shelf, ready-to-deploy options? Making this decision early will help guide you toward a viable choice.

Familiarize yourself with the wire speed of the network, general traffic levels, the types of filtering desired, types or forms of recent attacks, and future growth and expansion plans. Armed with this information, you can begin a serious review for the most appropriate firewall solutions.

Read current firewall reviews both online and in technical publications and trade journals, and consult buyer's guides for expert advice. Browse IT association blogs and discussion forums. Find

firewall options that fall within the parameters you identified based on your understanding of the network and known threats.

Some general guidelines will assist you in this process:

- If speed, flexibility, and simplicity are priorities, then select a packet-filtering firewall or a stateful inspection firewall.
- If real-time applications and high-bandwidth communications are priorities, then select a dedicated application-specific proxy firewall.
- If strong authentication is a priority, then select an application gateway firewall or a dedicated application-specific proxy firewall.
- If detailed logging is a priority, then select an application proxy firewall.
- If customized, unique, or complicated filtering options are a priority, then select a dedicated application-specific proxy firewall or a stateful inspection firewall.
- If stopping internal attacks is a priority, then select a dedicated application-specific proxy firewall or a personal host firewall.

Firewall technology advances quickly. Often, new attacks and exploits crafted by hackers drive these advancements and upgrades. A specific recommendation for a specific model, product line, or even vendor has a finite life. Before selecting a specific firewall product, do research to verify that the vendor is still supporting the firewall product. In the case of a software firewall, ensure that the solution is compatible with and maintained for your OS. You do not want to buy a product that has already been marked for retirement or at risk of being quickly superseded. Check on the revision history and, based on previous time frames, estimate if the current version release is due for a revision.

A single firewall product rarely satisfies every need. However, you should be able to find one or more firewall solutions that address the primary security concerns of your organization.

After purchase, regularly visit the vendor's website and discussion forums (including non-vendor-supported sites) for news and information about your firewall product. Keep current with announcements of updates, problems with updates, newly discovered holes or exploits, alternative configuration ideas, troubleshooting options, and more.

THE DIFFERENCE BETWEEN BUYING AND BUILDING A FIREWALL

Carefully consider the choice of whether to buy a ready-to-deploy firewall versus building it yourself. An off-the-shelf firewall that you simply plug in can be an attractive solution. However, an out-of-the-box product will generally cost more than an equivalent do-it-yourself system. In addition, most preconfigured systems are not as flexible or expandable as a custom-built version.

That's not to say that a do-it-yourself firewall is always the best answer, either. A self-constructed firewall can be less expensive and provide a wider variety of features and options. However, when an in-house, custom firewall experiences a problem, you can usually find little formal technical support beyond a user-community discussion forum.

If you are prepared to fully research and understand a firewall product to install, configure, and manage it yourself, then a do-it-yourself firewall can be a great cost-saving solution. However, if your organization's IT staff is already stretched and overworked, then home-built solutions will likely cause more problems than these solve.

A prebuilt, off-the-shelf firewall is more likely to work right out of the box, requiring only that you plug it in, set the configuration, and connect the network cables. A build-it-yourself firewall may require additional hardware manipulation, installation and configuration of

device drivers, hardening of a host OS, and so on. Ask yourself: Is custom-built worth the hassle?

Do not automatically reject either buy or build firewall options before you research and assess the options. You might be surprised at the low cost and feature richness of a boxed product, or be caught off guard by the complexity, instability, or high cost of a do-it-yourself product. Always look for the best firewall solution for your specific needs and current threats.

■ NOTE

Build-it-yourself firewalls are usually less costly because these usually do not include dedicated hardware and often use open-source (free) source code. However, the tradeoff for the savings to successfully deploy a reliable border sentry solution is that you need to know more about firewalls generally and your specific software and/or hardware.

CHAPTER SUMMARY

An essential element of network security, a firewall is a filtering service used to protect your network and hosts from a variety of threats, both internal and external. Several types of firewalls are available, including screening routers, hardware appliances, and host software products. Each of these firewalls can employ one or more features for ingress and egress filtering. The common filtering features include static packet filtering, stateful inspection or dynamic packet filtering, NAT, application proxy, and circuit proxy.

Firewalls are useful in many different situations. Every network infrastructure can benefit from proper use of a firewall. When making a choice about which firewall to deploy, consider a breadth of options, including the needs of both small and large network environments, host software firewalls, native OS firewalls, third-party OS firewall alternatives, ISP connection device firewalls, commercial firewall options, open-source firewalls, hardware firewalls, next-generation firewalls, and virtual firewalls.

KEY CONCEPTS AND TERMS

802.1x

Active threat

Appliance firewall

Application firewall

Application gateway

Application proxy

Bastion host

Border firewall

Border sentry

Bump-in-the-stack

Bump-in-the-wire

Chokepoint

Circuit

Circuit firewall

Circuit proxy

Closed-source

Commercial firewall

Content filtering

Cookie filters

Dead-man switch

Deny by default/allow by exception

Dual-homed firewall

Dynamic packet filtering

Hardware firewall

Header

Host firewalls

In-line proxy

IP address
Load balancers
Native firewall
Next-generation firewall (NGFW)
Open-source
Packet
Passive threats
Payload
Personal firewall
Pop-up blockers
Port forwarding
Port number
Public key infrastructure (PKI)
Reverse proxy
Rules
Rule sets
Sacrificial host
Screening router
Secure Sockets Layer (SSL)
Security stance
Segment
Service set identifier (SSID)
Session
Socket
Software firewall
State
Stateful inspection
Static packet filtering
Transparent proxy

Transport Layer Security (TLS)

Transport mode encryption

Triple-homed firewall

Tunnel mode encryption

Virtual firewall


Windows Defender Firewall

CHAPTER 5 ASSESSMENT

1. Which of the following is true of firewall rules?
 - A. Rules follow the *allow by default/deny by exception* philosophy.
 - B. No rules on a firewall are exceptions.
 - C. All rules on a firewall are exceptions.
 - D. The final rule is that anything that did not match one of the exceptions is allowed by default.
2. Which of the following attacks is *not* stopped by a border firewall?
 - A. Port scans
 - B. Protocol abuses
 - C. Inside client to internal host attack
 - D. Flooding attacks
3. Which of the following is *not* true of firewalls?
 - A. A firewall is a type of authentication system.
 - B. A firewall can filter traffic.
 - C. A firewall can provide routing functions.
 - D. A firewall is a traffic control device.
4. Which of the following best defines ingress filtering?
 - A. Blocking traffic leaving a network
 - B. Limiting host activities to that host
 - C. Monitoring traffic on its way into the network
 - D. Denying all traffic to specific ports
5. What is another name for dynamic packet filtering?
 - A. Stateful inspection
 - B. Static packet filtering
 - C. Structured packet filtering

- D. Sequential inspection
6. Which of the following is a method of filtering that automatically keeps track of sessions on a limited timeout basis to allow the responses to queries to reach internal systems?
- A. Application filtering
 - B. Deep packet inspection
 - C. Dynamic packet filtering
 - D. Static packet filtering
7. Which of the following is a form of filtering that allows communication, regardless of whether a session was previously established?
- A. Circuit proxy
 - B. Dynamic packet filtering
 - C. Deep packet inspection
 - D. Stateless filtering
8. Which activity differentiates a triple-homed firewall from a dual-homed firewall?
- A. Physical isolation of subnets
 - B. Deployment of the device as an appliance
 - C. Deployment of traffic from the Internet to a DMZ
 - D. Filtering of content including attachment deletion
9. When deploying software firewalls, what is the maximum number that should be operational on a single system at one time?
- A. 1
 - B. 2
 - C. 5
 - D. 10
10. Which of the following is *not* a content-filtering method?
- A. Domain name

- B. Source IP address
- C. Keywords in the packet
- D. Filename



© Digital_Art/Shutterstock

Firewall Implementation

CHAPTER 6

THE LOCAL AREA NETWORK (LAN) ADMINISTRATOR oversees all aspects of Internet and web security administration. This chapter is for those professionals who do not have enough time to dig into the more technical aspects of Internet and web security, but who need reliable options for Internet protection. Here are practical instructions to get a firewall up and running at your organization.

CHAPTER 6 TOPICS

This chapter covers the following topics and concepts:

- How to examine your network and assess its security needs
- How to construct, configure, and manage a firewall
- pfSense firewall requirements and fundamentals

- How to plan and install a firewall implementation with pfSense
- How to configure a pfSense firewall
- Elements of firewall deployment
- How to test a pfSense firewall
- How to troubleshoot a pfSense firewall

CHAPTER 6 GOALS

When you complete this chapter, you will be able to:

- Install a host software firewall
- Assess a network's security needs
- Configure a pfSense firewall
- Test a pfSense firewall
- Troubleshoot a pfSense firewall

EXAMINING YOUR NETWORK AND ITS SECURITY NEEDS

According to the National Security Agency (NSA), attacks to systems connected to the Internet are becoming more complex and dangerous. For instance, hackers can penetrate computer systems, using a variety of techniques to exploit weaknesses hidden in the complex code of many operating systems and applications.

What to Protect and Why

Internet-facing servers are accessible to people anywhere in the world who have Internet access, and these servers are often targets of attacks. They include web servers, email servers, File Transfer Protocol (FTP) servers, and more. If a server has a public IP address, it is a potential target for hacker attack. Firewalls also protect internal clients by providing a layer of protection to limit the ability of attackers to exploit system weaknesses—both hardware and software.

Imagine an organization that sells products over the Internet and processes payment card transactions. Customers insist that the merchant protect their payment card data. To do this, organizations need to address two distinct security areas:

- **Network security**—Computers, hard disks, databases, and other computer equipment attached directly or indirectly to the Internet need protection. Firewalls serve an important role in this aspect of network security.
- **Transaction security**—Web servers must be able to securely complete private transactions with other entities in databases accessible on the Internet. Hypertext Transfer Protocol Secure (HTTPS) is an important tool used to encrypt the transactions. Additionally, firewalls provide protection for the data. For example, it is common to locate a web server behind one firewall. The database server is behind a second firewall that restricts access to the database to only the web server. In this scenario, the area where the web server operates is called a demilitarized zone (DMZ).

Network security and transaction security have four overlapping types of risk:

- **Unauthorized individuals can breach the server's document tree**—Depending on what data resides on the server, this threat

can compromise the confidentiality of documents stored there.

- **Transaction data can be intercepted**—This threat can include personal data, financial data, or payment card information.
- **Information about the server can be accessed**—If attackers learn details about the server, such as the type of server and its operating system (OS), they may be able to identify and exploit its known vulnerabilities.
- **Denial of service (DoS) attacks can occur**—Many different types of DoS attacks are possible, such as when a TCP handshake attack does not conclude and holds the communication channel open. Hundreds of these incomplete sessions in a short time can consume substantial resources on a server and even crash it.

If your organization conducts business on the Internet, you need to take steps to protect those transactions. E-commerce requires a zero tolerance for failure. Protection is not optional; it is required.

Protecting Information and Resources

An organization must protect against attackers trying to access information and resources within the internal network, such as servers and workstations. Servers can host massive amounts of data invaluable to attackers. Database servers may host personally identifiable information (PII) about customers, including their payment card data. Domain Name System (DNS) servers host information such as the IP addresses and names of all systems in the network.

Protecting Clients and Users

Most organizations have a written security policy that outlines specific security requirements. A firewall can be useful in upholding the security policy by providing perimeter security. In other words, the firewall limits the risk of attack from hackers outside the internal network.

However, if you want to guarantee the security of an organization's network and protect its clients' and users' interests, you should do the following, as well:

- **Treat private messages as confidential**—Private messages should be encrypted to ensure that the message remains confidential. Encryption prevents the unauthorized disclosure of the private message.
- **Maintain integrity of information**—Integrity methods verify that data has not changed. For example, hashing is often used to verify the integrity of messages. A **hash** is simply a number calculated by applying an algorithm (a mathematical formula) to the numbers (1s and 0s) that make up a message. No matter how many times you calculate the hash, it will always be the same, as long as the data has not changed. The hash is calculated at the source and sent with the message. The hash is recalculated at the destination against the received message and compared with the hash sent with the message. If both hashes are the same, the integrity of the message is ensured.
- **Use strong authentication and nonrepudiation methods for all transactions**—Digital signatures commonly accompany both authentication and nonrepudiation. For example, a sender can sign a message with a digital signature. The receiver can use this signature to verify the identity of the sender. In other words, the digital signature provides authentication. Additionally, because it was digitally signed, the sender cannot deny sending it. In other words, the digital signature also provides nonrepudiation.

Preserving Privacy

It is more difficult to preserve a user's privacy on the Internet than in the physical world. Usually, you can close your door when seeking privacy or whisper something so that others cannot hear you. But the digital era has changed things. There is no such thing as whispering on the Internet.

When data is sent or received by a user surfing the Internet, that information can be intercepted and collected. Much of this data is stored in databases and can be made available for sale to others. Data mining methods are then used to help build customer preference profiles. From the perspective of an organization selling products, this data is valuable for identifying an individual's buying habits and targeting advertising. However, this data can be easily misused.

The **Electronic Privacy Information Center (EPIC)** was established in 1994, in Washington, DC. Its goal is to alert the public to emerging privacy issues relating to the **National Information Infrastructure (NII)**, such as the Digital Telephony Proposal, medical record privacy, and the sale of consumer data. EPIC's mission is to preserve the right of privacy in the electronic age, as well as to give individuals greater control over personal information and to encourage the development of new technologies that protect privacy rights. EPIC can provide valuable information to IT professionals on emerging cyber laws and other privacy threats.

The **California Consumer Privacy Act (CCPA)**, which went into effect in January 2020, requires companies to protect the consumer information they hold, and it stipulates minimum damages for class action suits in response to data breaches. The requirement to "implement and maintain reasonable security procedures and

practices” often includes keeping sensitive information in a database behind a firewall and restricting access to the database.

Preserving privacy on a corporate network may require more stringent security if you have users accessing resources from all over the world. Administrators are responsible for setting up security policies to guarantee users’ privacy, and each user accessing the company’s network is responsible for browsing the web with ethical regard for others. They also should be aware of information they may reveal while accessing resources from outside the network.

Typically, most network servers (and web servers) log all the hits (connections) received. This log usually includes the IP address and/or the host name of the computer making the request. If the site uses any form of authentication, the server will also log the username. If the user filled out any form during the session, all of the values of any variable from that form are also recorded. The status of the request, the size of data packets transmitted, the user’s email address, and so on can also be logged. Moreover, web servers can make all of this information available to **Common Gateway Interface (CGI)** scripts. Because the majority of web browsers are running on single-user PCs, very likely all of the transactions can be attributed to the individual using the PC.

TECHNICAL TIP

If you want to know more about EPIC, you can check out its webpage at <https://epic.org/>.

Revealing any of this information can harm the user. This is especially true if the user has logged any PII. Attackers can use PII for identity theft. Some PII may allow an attacker to access bank accounts or fraudulently charge payment cards. As a subtler example, if a user researches a job opening opportunity at a site, it may indicate that the user is searching for a new job. Browser lists, hostlists, and caches can also reveal certain patterns about the user. A proxy server, for instance, will track every connection outside the network on the Internet by IP address and the URL requested. If you install a proxy server within an organization, data stored on it needs to be protected so that only authorized individuals have access to it.

PROPER FIREWALL IMPLEMENTATION PROCEDURE

A firewall implementation procedure should prescribe the systematic process to install and configure a firewall properly. Because each firewall product is different and each type of deployment is likely different (for example, a border firewall versus an internal host firewall), you may even find that you need several firewall implementation guides.

The specifics of an implementation guide are different for each organization. Customizing a plan for your specific environment is essential to obtain the best security from the deployed firewall. However, every plan should have certain common elements. Whether you are developing your own firewall implementation policy or revising an existing one, use the following plan components to make yours the best possible policy for your organization.

Every firewall procedure should clearly define the requirements for the firewall. A generic description of a firewall is insufficient. Don't just indicate that you need a software firewall or an appliance firewall. Instead, dictate the specific capabilities, features, and requirements to accomplish the tasks and security goals. You can include a specific example of a vendor's exact make and model. However, products are frequently changed, revised, replaced, and retired by vendors. Include an inventory of features and specifications, even if you focus on a certain off-the-shelf product.

Specify the network design that the firewall will complement. To be effective, network firewalls (as opposed to host firewalls) must be at chokepoints or transition points. The design of the network is an essential element of effective firewall deployment. Define and even map out the network structure that should exist, pinpointing exactly where the firewall is, down to the logical configuration and the physical cables connecting to it. Your network design will change over time, but this can always change in the future when you perform change documentation and guidelines revision.

Write the plan down. Write out step by step, point by point, every action to take from the moment the firewall arrives on site through the point of enabling the filtering of production traffic. Include hardening of the bastion host, applying patches, updating firmware, changing defaults (especially passwords), locking down management interfaces, changing configuration settings, defining the rule sets, configuring interfaces, installing the physical components, backing up the configuration, testing at startup, testing for stress and load, documenting performance and stability, obtaining approval from senior management, and initiating filtering of production traffic.

Include photos in the documentation for quick retrieval of information. This should include the configuration screens and any physical setup, including cable attachment and labeling. The saying “a picture is worth a thousand words” is true when mission-critical functions are down on the network. The faster each is restored, the more likely the business is to survive the network crash or issue.

Write down every aspect of firewall use before deployment. Journal every step performed from start to finish of deployment, and then record into the documentation every action in managing,

administering, monitoring, and troubleshooting the firewall for the remainder of its deployment life. No action or event is too insignificant to record relative to your firewall deployment. This comprehensive firewall documentation is essential to management, troubleshooting, recovery, and incident response.

Be prepared to face a situation where the final design of the firewall may involve political issues from upper management and other departments of the company. For example, the security policy may state that the Network News Transfer Protocol (NNTP) is not essential for the organization. You would then design the firewall to block all NNTP traffic. But even though it may not be essential, upper management may decide to allow the traffic.

You also need to decide how you want to configure the firewall and what strategy you'll use. **FIGURE 6-1** shows a simple firewall solution with one network firewall and all resources placed behind it.




FIGURE 6-1 Simple firewall solution.

FYI

Many security policies have a special section just for the firewall, outlining how to configure and use it. For

example, many organizations recognize that every service and protocol allowed through the firewall represents a risk. By blocking traffic that is not needed to support the organization, these risks are eliminated. A simpler route is to just monitor the activity at the network and web server, but it may subject the network to risks you can easily avoid.

This method is generally not recommended for organizations that host Internet-facing servers such as email servers or web servers. These servers can go behind this single firewall, and you can configure the firewall to direct Internet traffic to the servers using port forwarding. However, this exposes the internal network to potential threats from the Internet. If you instead put your Internet-facing servers directly on the Internet, you expose these to threats of attack from anywhere in the world.

An alternative for organizations that have Internet-facing servers is to use a perimeter network. **FIGURE 6-2** shows one example. The firewall has three connections: one to the Internet, one to the internal network, and one to the perimeter network. This is also known as a bastion host. In this solution, the firewall will direct traffic destined for the Internet-facing servers to the servers on the perimeter network. Unrequested traffic from the Internet will be blocked from entering the internal network.




FIGURE 6-2 Multi-homed firewall used for a perimeter network.

The third solution uses two firewalls to create a demilitarized zone (DMZ) as shown in **FIGURE 6-3**. This provides several

benefits. The two firewalls provide an additional layer of security for the internal network. If the web server requires access to a database, the database server can be placed in the internal network for better protection. Internet users will not be able to access the database server directly.




FIGURE 6-3 Two firewalls used to create a DMZ.

Many organizations use two different brands for their different firewalls. If a flaw or vulnerability appears in one, it is unlikely the second firewall will be vulnerable at the same time. Additionally, an

attacker may be an expert at hacking one brand of firewall, but it is less likely that the attacker has the same level of expertise against two separate brands.

Remember, though, you need to balance the needs of the organization with the security requirements. While it is a true that a DMZ provides more protection than a single firewall, it also costs more.

CONSTRUCTING, CONFIGURING, AND MANAGING A FIREWALL

Many firewall products are available for small businesses and those working from home. SOHO (small office/home office) VPN hardware firewalls are often built on a secure virtual private network (VPN) connection to the company network to transfer email and sensitive files. Consider a SOHO network one in which there are typically 25 or fewer computers and devices on the network.

In many cases, the SOHO firewalls have additional features, such as an antivirus tool, IP filtering, web content filtering, router options (router/firewall combinations), intrusion detection, and denial of service (DoS)/distributed denial of service (DDoS) attack detection. Some of the firewalls come with a Linux or Linux-like OS and use **ipchains** to manage their filter rules.

A SOHO VPN hardware firewall is the best solution when you already have a working network and want to provide remote access. For example, you may already connect your personal computers within your office. If you want to open a new office in another location, you could connect both offices with a SOHO VPN firewall at each office. It will create a secure connection to transfer sensitive data (bank information, customer information, or company-related information) from one office to the other. Another example is a mobile worker, such as a salesperson working in the field. The salesperson could use the VPN's secure connection to access

company files and presentations or customer-related information directly from the company office, even when working in the field.

Recall that a firewall allows you to restrict unauthorized access between the Internet and an internal network. It exists to block unauthorized connections and keep outside attackers from penetrating the internal network. It also prevents inside connections from reaching out to the Internet without authorization. By monitoring inside users, firewalls can prevent the sending of sensitive information, such as personally identifiable information or sensitive corporate data to the Internet.

PFSENSE

You have many choices for firewall applications. This chapter discusses one possible choice: pfSense. The pfSense firewall is a host software firewall. It supports many different network types and is a practical solution for just about any organization. While there are several options for a software firewall, pfSense is a well-known and widely installed option.

Key features of the pfSense firewall include:

- Compatibility with a wide range of hardware and systems
- Flexibility and ease of use
- Comprehensive reporting and logging capabilities
- Web proxy support
- Snort intrusion detection system (IDS) support
- Static and dynamic Domain Name System (DNS) support
- Remote access and VPN support
- Dynamic Host Configuration Protocol (DHCP) and network time server support
- GeoIP blocking
- NAT mapping (inbound/outbound)
- System updates
- IDS/IPS
- Load balancing
- Traffic shaping

- Transparent caching proxy
- Web content filter
- Multiple IP addresses per interface
- Built on a stable, well-supported platform of FreeBSD
- Free documentation and resource library

PFSENSE REQUIREMENTS

As mentioned previously, one of the main advantages of pfSense is that it will run on a variety of hardware. In any case, you will need at least two network interface controllers (NICs) in your pfSense machine, because firewalls traverse two or more network connections. You may need a crossover cable, as well. Many switches today perform auto-sensing; in that case, there is no need to use a crossover cable.

There are many possible configurations for a pfSense installation, but the minimum requirements for pfSense are:

- 64-bit Intel or AMD processor
- Must be able to boot from USB, optical drive, or Memstick
- Two NICs

Installation of pfSense is simple. Visit the Netgate website at <https://docs.netgate.com/pfsense/en/latest/install/installing-pfsense.html> and follow the instructions. Before attempting any download, be sure that the system where the software will be installed meets or exceeds the minimum specifications. You also need a consistent online connection so that the installation process is not interrupted. If something does go wrong with the download, you can delete the application and start again.

PLANNING A FIREWALL IMPLEMENTATION WITH PFSENSE

When planning a firewall, you should use standard hardware, typically from major manufacturers, which will provide you with considerable benefits, including:

- **Flexibility**—You can easily add a redundant array of independent disks (RAID), redundant power supplies, additional random access memory (RAM), and a range of network interface cards.
- **Performance and value for the money**—The average desktop PC will have many times the processing power and memory complement of most mid-range hardware firewall appliances. Office PCs have become a commodity, offering exceptional value for the money.
- **Support**—The major PC manufacturers can support their products virtually anywhere in the world, including rapid, on-site repair or replacement if the customer requires. Alternatively, in most major cities, several third-party computer companies also support such hardware. In contrast, the support available for smaller manufacturers' equipment is often limited to repairs done after the machine was shipped back to the factory.

What kind of firewall is right for an organization? Truthfully, no single correct and definite answer exists. A security policy developed by a Fortune 500 company certainly will not be suitable for a small business owner (and vice versa).

The following are brief scenarios to illustrate different firewall designs to meet different needs.

Firewalling a Small Organization: Packet Filtering or Application-Level Firewall, a Proxy Implementation

Imagine a smaller organization that requires Internet access by employees, but does not host a web server or a VPN server. All users in the organization require access to the Internet, and they need a layer of protection against Internet-based attacks.

A proxy server or a firewall such as pfSense would be sufficient. All users could access the Internet through the pfSense firewall, and it would provide protection against external attacks. Although pfSense is the example examined in depth in this chapter, many other products can provide the same level of service for a small organization.

Firewalling Medium and Large Organizations: Application-Level Firewall and Packet Filtering, a Hybrid System

Employees at medium and large organizations typically access the Internet daily for a wide variety of different business purposes. Additionally, organizations of this size host a web server that has a high volume of access from customers. Mobile workers use a VPN server to access internal resources.

These organizations might decide that a DMZ with two application-level firewalls and packet filtering is adequate for their needs. The web server and VPN server belong in the DMZ to be accessible from the Internet, while the internal network requires additional layers of protection.

Use of a DMZ is attractive to many network managers because it has some security but allows public access outside of the corporate network. It is important to ensure that application-level firewalls are restricting access properly prior to deployment of the DMZ. This is especially important if external people will access high-value information, such as customers accessing inventory available for purchase. Once access is granted, poorly constructed firewall rules may allow access to data that should have been restricted.

Firewalling in a Subnet Architecture

Imagine a larger organization with multiple subnets. There may be a need to limit specific types of traffic within subnets. For example, they may want to ensure that traffic within a specific subnet is encrypted with Internet Protocol Security (IPSec). Additionally, they may want to allow NNTP traffic on one subnet, but not on others. You can configure routers to provide basic packet filtering. In other words, the routers can block or allow traffic based on IP addresses, ports such as port 119 for NNTP, and some protocols such as Internet Control Message Protocol (ICMP) or IPSec.

If the organization wants to protect subnets within the network, basic packet filtering provided by routers might be the most appropriate choice. This model supports each type of client and service within the internal network. No hardware or software modifications or special client software would be necessary. The access through the routers used as packet-filtering firewalls is transparent for both the user and the applications. The existing routers can handle the packet filtering. This solution does not require the purchase of a host or firewall product as an expensive UNIX host, or firewall product would be unnecessary.

INSTALLING THE PFSENSE FIREWALL

One of the best ways to learn about a technology is to use it. Installation of the pfSense firewall can be completed by a person with a basic understanding of computer hardware using the web interface for configuration.

Because pfSense is based on FreeBSD, the list of compatible hardware is the same. The FreeBSD Hardware Notes webpage at <https://www.freebsd.org/releases/11.2R/hardware.html> includes a detailed listing of processors, system boards, and self-contained devices, including cameras and video recording devices. The website is updated as new options become available.

You can install pfSense on a local firewall, on a large-scale security device, or in the cloud. pfSense is the software element of the firewall, which means you must select specific hardware based on known compatibilities. The pfSense support documentation offers guidance and troubleshooting resources.

pfSense software can be downloaded from <https://www.pfsense.org/download>. On the website, you will need to make selections based on the version, architecture (hardware), and installer type (options listed below) to identify the specific download version needed.

To install pfSense, choose between three installer options for the **ISO image** (an .iso file is an exact copy of a file system):

- **Optical disc**—Systems with a CD or DVD drive can use an optical disc.
- **USB**—The BIOS of the system must be set to boot from USB (if not, make the necessary modification or elect a different method).
- **Memstick**—This option uses a serial memory stick (Memstick) via a serial console.

As with most Unix and Linux products, a hash value is provided for verification of the file after download. If the hashes match, the contents are unaltered. If the hashes do not match, the file may have been altered; it should be discarded and the file should be downloaded again.

Once successfully downloaded, the file must be saved to the appropriate media from the list above so that it may be installed on the selected hardware. (Although it is possible to install pfSense via PXE, most installations are on local media.) Specific information about each installation is available on the website.

After booting, pfSense automatically checks the hardware components to determine which hardware is accessible, if it meets the minimum specifications required, and if it is communicating properly. If the hardware does not meet the minimum requirements, a message is displayed; it identifies the issue and whether it can be corrected, such as the need to add additional RAM. The message also notes that pfSense will not install unless the minimum requirements are met. If the hardware is sufficient, pfSense starts and displays a prompt. You must accept the copyright notice or the system will not progress. After you select OK, the installation options display in a menu. The first option, Install, is the one to select.



WARNING

During the installation process, pfSense deletes all data on the memory device. Before you start the installation, ensure that all valuable data is safely backed up.

Next, pfSense prompts you for the filesystem you wish to use (**FIGURE 6-4**). Unix File System (UFS) is the default in the current version, but you may choose the Z File System (ZFS), which is experimental but offered as an option. If you plan to use a partition on the hard drive for storing the application, that partition should be created prior to beginning the installation process for the most seamless installation.



FIGURE 6-4 pfSense partitioning options during the installation.
◆ pfSense

The next prompt is for keymap selection (**FIGURE 6-5**). The standard US keyboard is the default option. After pressing “Select,” the installation process begins, which includes a wipe of all files from the destination media and then installation of pfSense. The installation process typically completes within several minutes.




FIGURE 6-5 pfSense keymap options during the installation.
◆ pfSense

After installation, you are prompted for manual configuration options. Unless you are well versed in those options, you should select “No” and then allow the system to reboot. Remove the drive or disk you saved pfSense to first so that the system will not boot from it the next time the system starts.

After boot-up, you will need to assign interfaces using the system console. Most people use the default installation, with pre-selected options common to most applications of the product, because it is configured for the most common use of the product. The next questions are predicated by your specific use for the product. You are first prompted to assign a virtual local area network (VLAN), if used.

There are additional instructions for using a VLAN on the installation website.

■ NOTE

Netgate provides step-by-step installation troubleshooting guidance on its website:

<https://docs.netgate.com/pfsense/en/latest/install/installation-troubleshooting.html>. Many YouTube videos demonstrate the installation process, as well.

pfSense prompts you for the WAN and LAN interfaces. If you know the information, enter the specifics (such as em0 for the WAN or em1 for the LAN); otherwise, press “Enter” to accept the defaults. pfSense requires one interface; once that is configured, there is no need to enter any others, but you may if they exist. pfSense deems a single network interface assigned to the WAN as “Appliance Mode.” You also have the option of using the pfSense “Auto Assign” procedure if all network cables are unplugged from the system. To use it, simply type “a” and press “Enter.” If the process is not successful, consult the Installation Troubleshooting guide online, or restart using the previous instructions.

CONFIGURING A FIREWALL WITH PFSENSE

Once you have finished installing pfSense, you will need to configure the firewall via a web browser. The interface is uncluttered and logically laid out, and it works well with Mozilla Firefox, Google Chrome, and Microsoft Edge. Once you reboot the computer, you can access it via a browser from a web interface (usually <https://192.168.1.1/24>) to start configuring the firewall. A complete copy of the pfSense default configuration is available on the pfSense website, but it is important to note the following default configuration items:

- All incoming WAN connections are blocked
- All outgoing LAN connections are allowed
- WAN is configured with both an IPv4 and IPv6 DHCP client
- pfSense will function as a DHCP server, by default, on IPv4 and IPv6

That concludes the installation. There are post-installation tasks to be completed with a shell menu on the console. pfSense is available through the interface or via the WAN. The web-based graphical user interface (WebGUI) is the most common way to access and configure the system. Some of the items you can access through the high-level menu are rebooting the system, accessing the shell, pinging the host, filtering log files, and restoring a recent configuration in case it was changed or compromised.

ELEMENTS OF FIREWALL DEPLOYMENT

At this point, you should start thinking about deploying the firewall and about the services you want it to run. You have the ability to run:

- Captive Portal, which is pfSense's authentication front end
- DHCP server
- DNS server
- NTP server
- SNMP daemons
- IGMP proxy daemon
- IDS/IPS deployment
- Virtualization

The specific collection of services should be based on the needs of the network, the existence or absence of other products, and the ability to manage and administer pfSense. Each service has logging options that can be modified to collect more or less information about each instance of the service. When considering what to log, know the storage capacity of the system and the information that will be necessary to investigate an issue or hack to the system. More information on logging is available at the pfSense website.

TESTING AND TROUBLESHOOTING

It is easy to assume that a new device or application will function properly when installed on a network. It may appear to work, but consider whether it is performing the way you assumed it would. Testing helps to ensure that the device or application is functioning properly, and it is critical to firewalls in general.

Create a comprehensive testing plan before installing any firewall. The team that devises the test should consist of network administrators, hardware engineers, and system users to aid in providing the most realistic test environment and data collection. Document the results of the test, and then label the test data so it can be reviewed if there is an issue or if data is passed through the system when it should be blocked.

To test a pfSense firewall, you need test data that closely mimics live data on the network. When possible, companies utilize a sandbox or test environment. It is popular to use virtualization for test environments to save on the cost of equipment.

Checking connectivity is essential. Use common commands such as ping and traceroute. Instructions on how to accomplish these tests and many more are available on the Netgate System Monitoring webpage in the “Status Information” section at <https://docs.netgate.com/pfsense/en/latest/monitoring/index.html>.

The testing process should aid in limiting packet/traffic loss; in the event that traffic is lost or disappears, there is a section on the

website with steps to aid in determining that, as well. To test the system, use a collection of typical and atypical network traffic, if possible. This provides legitimate traffic and offers the best opportunity to observe and test the behavior of the system, as packets are evaluated against the rules. If there is still an issue, troubleshooting may provide insight. Offloading options are available by navigating to the System > Advanced > Networking tab: Hardware Checksum Offloading, Hardware TCP Segmentation Offloading, and Hardware Large Receive Offloading. Pay attention to the wording of each option, because each checkbox identifies that selecting or deselecting it will disable the option.

The Troubleshooting webpage (<https://docs.netgate.com/pfsense/en/latest/troubleshooting/index.html>) provides a convenient list of links to many different pfSense troubleshooting situations. When troubleshooting issues in general, make one change at a time, and then test the change to see if it makes a difference on the network before making any other changes. You should also log each change made, the order in which each change was made, the results of testing the change, and whether the change was maintained or reversed to help prevent unintended consequences while troubleshooting.

One error message you may encounter is “filesystem is full.” This message may also manifest as an inode failure—“no inodes free.” While it may appear intuitive, these messages are not usually due to the lack of space on the firewall drive. Hard drives are prone to failure because they contain mechanical parts that wear out over time. If a hard drive is failing, either message may appear. If the drive is not failing, it can occur if there is a disruption in the operating system

attempting to write information to the hard drive. To identify the cause of the issue, run the `df -hi` command from the shell. Statistics for the system display will include the percentage of the hard drive that is used and the percentage that is still free. The statistics to review are for the actual devices, such as `/dev/md0` and `/dev/md1`. If the `devfs` value is 100%, that is fine and does not represent the space on the physical device.

If the drive has adequate space available, you most likely need to replace the drive. If the drive is full, consider backing up some of the logs to a different location to free up space. Be sure to encrypt the logs in the new location so that unauthorized users cannot access or change the information.

Hardware problems or issues with the FreeBSD kernel, specifically **kernel panics**, could cause the pfSense firewall to unexpectedly reboot. Hardware issues are more common and should be ruled out first. This includes checking the viability of the power supply. Have you encountered issues with the system's power, including surges, a power spike, or even a sustained brownout? Check the operating temperatures of the CPU and RAM on the device where pfSense is installed. Overheating can cause the firewall to stop working and force a reboot. This issue could also stem from a faulty or malfunctioning hard drive. For this issue, perform the `df -hi` command and system check described previously in this section.

From the Troubleshooting webpage, clicking the “Hardware” link in the left pane displays a list of troubleshooting topics, including one for halting and powering off the firewall. This may be required for system or network upgrades, hardware replacement, or troubleshooting network problems. This section also includes

information on how to access various console types, such as if a physical connection is required through a serial console.

A part of every testing and troubleshooting session should include log review. Logs capture information on specific events occurring on a device. Some logs have set defaults that are sufficient for most businesses, while other logs should be reviewed to ensure that what your organization requires is captured. Once an event ends, it is usually too late to recreate that information. Evaluate log details and storage options (such as length of retention, location of storage, and backup of logs) against company policies to ensure compliance. Consider the type of activity where logs would be consulted (testing, troubleshooting, nonrepudiation, and tracking a hacking attempt) and identify whether the information stored will be sufficient. Some of the available log files include firewall logs, system logs, routing logs, gateway logs, filtering logs, copying logs, and routing logs.

CHAPTER SUMMARY

Firewalls provide protection from Internet threats to servers, workstations, and devices, and they can be configured to allow only specific traffic based on what an organization wants to allow. Organizations typically define what kind of traffic is acceptable in a comprehensive, written security policy. IT professionals then identify a firewall solution that can fulfill and comply with the policy within the established budget.

pfSense is an Internet firewall that can protect an organization's network. This chapter demonstrates how to configure and use a firewall such as pfSense using a web-based GUI. pfSense is an open-source firewall, requiring no knowledge of Linux to install or use. pfSense turns any typical PC, or Internet-capable device that meets the specifications, into a network appliance with functionality to route network traffic to and from the Internet, assign IP addresses, and protect the private side of the network from intrusion.

KEY CONCEPTS AND TERMS

California Consumer Privacy Act (CCPA)

Common Gateway Interface (CGI)

Electronic Privacy Information Center (EPIC)

Hash

Ipchains

ISO image

Kernel panics

National Information Infrastructure (NII)


National Security Agency (NSA)

CHAPTER 6 ASSESSMENT

1. Which of the following are the two distinct areas that must be protected with firewalls?
 - A. Access and controls
 - B. File sharing and printing capabilities
 - C. Network and transaction security
 - D. WAN traffic and cellular signals
2. Which agency was created to alert the public to emerging privacy issues related to the National Information Infrastructure?
 - A. CERN
 - B. EPIC
 - C. NSA
 - D. WWW
3. Which of the following is *not* one of the overlapping types of risk in network and transaction security?
 - A. Transaction data can be intercepted
 - B. The server hardware can fail
 - C. Information about the server can be accessed
 - D. Unauthorized individuals can breach the server's document tree
4. Which of the following is the best way to treat private messages as confidential?
 - A. Send messages from personal email accounts
 - B. Send the message to everyone so it is no longer private
 - C. Encrypt the message so it stays private
 - D. Do not send private business messages
5. Which of the following best describes nonrepudiation?
 - A. An action is copied onto several systems so it is redundant

- B. An action is logged so it can be recovered if deleted
 - C. An action is repeated so it is redundant
 - D. An action cannot be denied as occurring
6. Which of the following tools is *not* a troubleshooting tool for firewalls?
- A. ipconfig
 - B. ping
 - C. NAT
 - D. traceroute
7. Which of the following is *not* a type of emerging issue the EPIC would alert the public about?
- A. Alexa personal helper
 - B. The Clipper chip
 - C. Medical record privacy
 - D. Sale of consumer data
8. Which of the following is *not* a function of a firewall?
- A. Block unauthorized traffic from the Internet
 - B. Block unauthorized connections from outside attackers
 - C. Block one device from using too much bandwidth
 - D. Block sensitive information sent to the Internet
9. Which of the following is not an installation method for pfSense?
- A. Streamed across the network
 - B. Optical drive
 - C. USB drive
 - D. Memstick
10. When selecting a firewall solution, which of the following are security concerns to consider?
- A. Audit capabilities

- B. Privilege control
- C. Refresh rate
- D. Authentication



© Digital_Art/Shutterstock

Firewall Deployment Considerations

FIREWALLS CAN BE COMPLEX security solutions. You should plan the deployment of a firewall carefully, whether for a small home office or for a large corporation. Evaluate as many firewall deployment considerations as possible before you begin to ensure that you are making the best choice for your network and circumstances.

Make a clear determination as to which types of traffic you will allow to cross the network border and which types you will block. Evaluate common security strategies, including security through obscurity, principle of least privilege, simplicity, defense in depth, defense diversity, chokepoint, weakest link, fail-safe, and forced universal participation. Integrate each strategy you want to use into the organization's security policy and its firewall deployment.

The actual deployment of a firewall is a fairly straightforward process. This process begins with understanding the network environment, it includes selecting a product that satisfies security needs, and it ends with a more secure network infrastructure. Toward

that end, this chapter examines a few of the remaining concepts that are often part of the decision-making process.

Determine whether features such as reverse proxy and port forwarding are necessary to the infrastructure's network communications. Weigh the benefits of bastion host OSs before using new firewalls. Understand how monitoring and logging, interpreting firewall logs and alerts, and intrusion detection and prevention complement firewalls.

Every organization is different and must evaluate its own business and security needs. Determine which tasks are essential, which are optional, which are personal, and which are malicious. Use firewalls and other controls to support necessary traffic and block everything else. Security administrators are responsible for evaluating needs and solutions, and for preparing a response when security and business needs are incompatible.

CHAPTER 7 TOPICS

This chapter covers the following topics and concepts:

- Common security strategies for firewall deployments
- The purpose of authentication, authorization, and accounting
- How to determine placement of network hardware firewalls
- The benefit and purpose of reverse proxy
- The benefit and use of port forwarding

- Which considerations aid in selecting a bastion host operating system (OS)
- Essentials of monitoring and logging
- How to understand and interpret firewall logs and alerts
- Intrusion detection systems and intrusion prevention systems (IDSs/IPSs)
- How security information and event management (SIEM) aggregate log data for decision-making
- How to evaluate needs and solutions in designing security
- What happens when security gets in the way of doing business

CHAPTER 7 GOALS

When you complete this chapter, you will be able to:

- Describe various firewall security strategies
- Locate firewalls for optimal coverage
- Define the pros and cons of reverse proxy and port forwarding
- Explain the importance of a bastion host
- Design a policy to guide effective firewall monitoring and logging
- Explain the importance of IDS/IPS systems and SIEM
- Assess the business impact of security over availability and performance

COMMON SECURITY STRATEGIES FOR FIREWALL DEPLOYMENTS

A security strategy is a guideline, philosophy, or approach to using security for firewalls, as well as the entire organizational security infrastructure. Many different strategies are available. Some organizations focus on a single strategy, while others combine several strategies into a custom policy.

Security Through Obscurity

Security through obscurity is the idea of gaining protection by using abnormal configurations. This can be either a positive security action or a negative one; it simply depends on the type of obscurity involved. A poor choice for security through obscurity would be to use alternative configurations for standard products.

With a network environment based on standard products—whether OSs, network services, or security solution—just modifying basic configuration settings does not actually provide true security. For example, a service typically operating on one port that is reconfigured to use a different port is not good security. Changing names, addresses, network size, subnetting, or bandwidth consumption, or even spoofing banners, headers, or identity, does not provide true security; however, it may provide some protection against elementary attacks, as the hacker or launched attack may fail if a port match is not found.

Some common mistaken examples of obscurity-based security include:

- Hiding your front door key under a rock in a flower bed
- Connecting to the Internet, assuming that your system is one that won't be noticed
- Keeping an encryption algorithm secret
- Changing the default service port of a network service
- Writing your password on a sticky note and placing it under the keyboard or mousepad

All attempts to provide network security based on hide-and-seek can be overcome or bypassed with basic network scanning

techniques. When a hacker can discover or breach a host, service, or network just by scanning or sniffing network traffic, no real security exists. Hiding might keep colored eggs and candy secure, but as a security solution, it has no place in a modern IT environment.

Security through obscurity is problematic for several reasons. First, if it's the only form of security employed, no real security exists. Second, the obscurity method employed might not be as obscure as you think. Hiding a key under a flowerpot is not as obscure as burying it in an unmarked location in the middle of a field. Third, it often distracts from accurately assessing the true state of security provided by other measures. Fourth, obscurity can instill a false sense of confidence.

True security through obscurity relies on using obscure and nonstandard technologies. If everyone knows that the most common or most popular OS or software product is insecure, it can be a security improvement to use a more obscure product instead. This does not guarantee that other flaws or vulnerabilities won't be present on the alternative system, but usually the exact same exploit or compromise won't be a risk.

However, switching to an alternative technology is a security improvement only if the same flaw does not also exist there. It must provide actual security. Switching technologies, such as changing OSs, does not eliminate the need for security management. It simply changes the product and does not eliminate security concerns in other areas.

Generally, security is obtained through obscurity only when hackers are unable to determine which technologies you are using. The less hackers know about your software, configuration,

deployment, addressing, infrastructure design, patching schemes, and so on, the less successful their attempts to compromise security will be. Obscuring information about the internal IT environment is a primary defense against their hide-and-seek mentality.

STEGANOGRAPHY

Steganography is the art and science of writing hidden messages so that only the sender and intended recipient know a message exists. Steganography is the ultimate security through obscurity technology.

Probably the most well-known technique for steganography in the modern era is the embedding of additional information in a digital image. An image before and after a message has been embedded in it appears unchanged to the naked eye. But if you run the image through the proper program, the secret message appears. However, steganography has been around a lot longer than digital images, and it can take many forms. Writing a message in invisible ink on the page of a book is a form of steganography.

The main advantage of steganography is that messages do not arouse suspicion. If you encrypt an email message and send it, and the message is intercepted, then whoever intercepted the message assumes there is something sensitive or valuable in the message. If, however, you send a copy of a picture of children at the amusement park with a message

embedded in the photo, unless the intercepting party knows the coded message is coming, they will not give it any special attention. Encrypted messages protect the message. Steganographic messages protect the data as well as the intentions of the sender and the receiver.

Least Privilege

The **principle of least privilege** is one of the basic concepts of network security. The idea behind the principle of least privilege is that users should have the minimum level of access to resources needed to complete assigned tasks. Any abilities, access, or privileges beyond that necessary minimum increase the risk of compromise and can lead to accidental or intentional tampering and loss.

In most environments, users need not access every host, every system, every resource, every file, every network service, and every Internet resource. Within a default-deny environment, you block all access to all resources—internal and external—by default. You use the principle of least privilege by adding explicit and specific allow-exceptions only when necessary based on job descriptions.

The drawback to least privilege is the need to control every user's access individually. Each worker will need individually defined resource and activity permissions. This represents a significant increase in user administration. Because of this extra overhead, many organizations only partially use least privilege.

A common practice is to group users into collections of similar job functions or security levels, granting permissions to the group as a whole. This gives users additional privileges beyond what is strictly necessary for their work, but not enough to present a serious risk to the organization's security, stability, or confidentiality. Instead, this small security trade-off may considerably cut down on administrative overhead.

An extension of the principle of least privilege known as **separation of duties** specifically addresses administrative users. A lazy and insecure approach to network security is to grant some high-

end users full administrative privileges across the entire IT infrastructure. This is a recipe for disaster. An administrator who makes a mistake can accidentally harm the entire organization. An administrator who becomes disgruntled or feels wronged by the company can deliberately harm the organization and have the systemwide power to do it. If another user or an external intruder compromised such an administrator's account, a hacker could use the administrator's advanced access to cause significant damage or loss.

A safer approach is to apply the principle of least privilege to all users, including administrators. Under the label of separation of duties, divide the collection of all administrative tasks into small sets of privileges, focusing on a single task, system, service, or issue. Then, assign administrators subtasks within one of these focused areas and permissions only within the scope of their assigned tasks. This compartmentalizes administrative functions and provides a type of blockage against accidents, sabotage, and intrusion.

The result of applying separation of duties is that you eliminate "systemwide" administrators. Instead, each administrator has sufficient privileges only within a limited scope of responsibility. Some areas may have multiple administrators, and some administrators may have powers in multiple areas, but overall the compartmentalization of administrative privileges significantly increases security and decreases risk.

Keep in mind that the principle of least privilege applies to both users and systems. Controlling which resources and communication pathways a service or device can touch will directly reduce the potential for abnormalities, abuse, and compromise.

Simplicity

Keeping things simple is an important part of security. It makes the system easier to understand, manage, and troubleshoot. When systems are too complex, it can be much more difficult to understand and manage the normal traffic and processes and more complicated to troubleshoot. Furthermore, it is harder to verify that a complex system is providing adequate security.

The more complex a solution, the more room for mistakes, bugs, flaws, or oversights by security administrators. The more complex the system, the more likely a hacker can find a vulnerability unseen by the system designers and network managers.

Simple is not always possible, however, especially when it comes to software and network infrastructures. But when you have a choice between a simpler solution and one that is more complex, the simpler option may provide a more realistic and verifiable level of security. Keep in mind that too simple has the same flaws as too complex. Avoid sacrificing security for the sake of simplicity.

Defense in Depth

Installing a security infrastructure requires numerous systems to interact and interlock. The goal of any security solution is to prevent unwanted events while supporting authorized and necessary events. A reliable security infrastructure is typically not a single control, a single defense, or a single countermeasure. Instead, multiple security safeguards are needed to provide complete and exhaustive coverage (**FIGURE 7-1**).




FIGURE 7-1 Defense in depth.

One aspect of infrastructure security design is to think of the defenses as multiple layers or barriers to access secured resources. This is known as layered defenses or defense in depth. When designing security, consider the net result if any component of the security system were to fail. If the failure of a single component results in a compromise or intrusion, the environment has a single layer of protection. When defense in depth is used, a single component failure does not result in compromise or intrusion. Instead, each component has a backup, an alternative, or a supplemental component. Depending on a single security product as the sole component of a security solution is a bad idea. Firewalls are great products, but firewalls alone cannot provide complete security. This is only one component, one piece of a complete security infrastructure. A proper security infrastructure has numerous components interlocked and deployed in layers or levels.

Another aspect of defense in depth is to deploy multiple subnets in series to separate private resources from public. This is known as an **N-tier deployment**. N represents the number of subnets under private control. **FIGURE 7-2** shows a common construction of a three-tiered deployment using a demilitarized zone (DMZ), a database subnet, and the private local area network (LAN).




FIGURE 7-2 An example of an N-tier deployment.

When properly implemented, defense in depth ensures that multiple security controls are involved in every communication, connection, and transaction, regardless of its source or destination. Through the use of multiple or redundant security systems, any attempt to compromise or breach security becomes substantially more difficult, not only for the external intruder, but for an internal saboteur as well.

Diversity of Defense

Diversity of defense is similar to defense in depth, as it supports multiple layers of security. The difference is that in diversity of defense, most layers use a different security mechanism. Thus, the diversity of defense comes from using a collection of diverse security solutions.

Multiple firewalls in a series is defense in depth, but not defense diversity. Adding tools such as an intrusion detection system (IDS), antivirus, strong authentication, virtual private network (VPN) support, and granular access control converts a monolithic defense in depth to a more substantial, multilayered, diverse type of security.

It might be tempting to claim that having multiple firewalls from different vendors (or any security mechanism from different vendors) constitutes diversity of defense. Although using different products from different vendors will reduce some forms of risk, it is not as beneficial as true diversity—which is not just different vendors, but also different security component types. Using a variety of vendors will reduce the likelihood that a single flaw in one product is present in every other product from different vendors.

For example, it is a reasonable idea to use three different vendor antivirus products: one vendor for clients, one vendor for internal servers, and a third vendor for border devices. (However, relying exclusively on antivirus and using no other type of security mechanism is a poor security solution.) Using three different vendors will increase the likelihood of detecting malicious code and decrease the threat from undetected malware. Nevertheless, be aware that using solutions from multiple vendors increases the complexity of managing the network, as there are differences in configurations and

system access. It can be difficult enough to manage single-vendor solutions; when two, three, or more different vendor products of the same product type are involved, the complexity of management increases significantly. Those tasked with supporting the network must have the requisite training on each of the devices and the software in use to ensure that it is configured properly, aids in securing the network, and does not create a false sense of security.

When selecting products—whether by type or vendor—keep in mind the overall infrastructure design, such as parallel versus serial connections, as well as redundancy options. In some circumstances, using two different defenses can result in a larger rather than reduced attack surface. For example, if two different products each have a different vulnerability when deployed in parallel, then both weaknesses are vulnerable simultaneously, giving the hacker the possibility of matching one or the other. Were these items deployed in series, the hacker would have to compromise the first layer of defense successfully before attempting to breach the second. Consider whether the planned diversity is truly a security improvement.

In addition to vendor and technology diversity, consider diversity in the type of control used. The classification of controls includes physical, technical, and administrative. Examples of physical controls include a lock on a facility door (internal and/or external), a fence, and a guard. Technical controls can be hardware or software applied to a network or system to control behavior or traffic, such as a firewall, antivirus software, or an access control list (ACL). Administrative controls are procedural ways to limit a threat or risk, such as the application of permissions, the use of security and firewall policies, and data classification.

There are three control functions: preventative, detective, and corrective. Preventative controls, such as a firewall, aid in the prevention of an attack or access to a system or network. Detective controls offer identification and recognition of unexpected traffic or behavior. An intrusion detection system (IDS) is an example of a detective control. Corrective controls offer actions that aid in returning the network or system to a normal condition, such as rebooting a system or patching a firewall.

When investigating and designing an infrastructure to include diversity of defense elements, keep the following concerns in mind:

- Some companies resell their products, technology, or intellectual property through other vendors. Private-label or relabeled products from company B might be the same product from company A with a different name or packaging.
- Security systems configured by the same security administrator can potentially have the same misconfiguration or design weakness. Consider using a team of security administrators, or at least have several security experts review and check all configurations.
- Many products come from a single original code base or standard. Product version 2.0 from company A and product version 5.0 from company B might both include the same protocol stack code borrowed from an open-source or creative commons–licensed code base. Do not automatically assume that products from different vendors, even with different build versions, represent 100 percent different programming.
- Many systems of the same type will all have the same inherent weaknesses in the underlying technology, design, and security concept.

When implementing diversity, ensure that you are using actual diversity-improving security rather than false or misguided diversity

that reduces security. All security designs and perceptions should withstand the scrutiny of evaluation and testing.

Chokepoint

A chokepoint forces all traffic, communications, and activities through a single pathway or channel. This pathway can be used to control bandwidth consumption, filter content, provide authentication services, or enforce authorization. The purpose of a chokepoint is to ensure that the security device at the control location controls everything. No traffic, user, or data passes unchecked.

Other names for a chokepoint include checkpoint, filter pathway, or bottleneck. A chokepoint can be used to force network traffic along a single pathway monitored or filtered using security devices, including routers, switches, firewalls, IDS/IPS, antivirus, network access control (NAC), and so on. A chokepoint can also control user activity, such as requiring authentication and enforcing access control.

As a security measure, a chokepoint has value only if it is hard to bypass or avoid the bottleneck. If a hacker can interact with a target without going through a chokepoint's filtering system, then the chokepoint is worthless. If the security of a network has chokepoint infrastructure components as a central part of the infrastructure, be sure to evaluate and consider all possible alternative pathways a hacker might employ. For example, a firewall is less effective as a chokepoint if outsiders can access wireless access points (WAPs), physical connection ports, or dial-up modems.

In many cases, a single chokepoint will not be enough. Using multiple chokepoints may be the only effective means of ensuring that your network evaluates and filters all traffic. Some potential chokepoints can include any crossing of a network domain boundary into another. By using chokepoints across the boundaries of a typical

IT infrastructure, you filter most pathways of communications that could host malicious interactions.

Consider potential indirect routes that might bypass “standard” chokepoints. For example, if a VPN link is established between two sites, whether branch offices or business partners, the Internet connection of one site could be an alternate pathway into the second site. This breach is possible if one site uses strong chokepoint security while the other does not. Even though it may be old-fashioned, don’t forget to check for dial-up access to a LAN’s servers or other components. Dial-up access can serve as a backdoor for hackers to circumvent the security protections prepared to block them at the front door.

Weakest Link

Any chain is only as strong as its **weakest link**. A security infrastructure is only as strong as its weakest component. You must know your security infrastructure thoroughly to have an understanding of where the weakest point lies. Once you have identified the weakest link, replace or remove it.

The weakest link security stance is an ongoing process of locating the least secure element of an infrastructure and securing it. Once you have secured the current weakest link, it is no longer the weakest link; therefore, a new weakest link exists. Repeat the cycle by seeking out the next weak point and improving it.

The idea behind this security stance or process is that hackers are performing this exact task as they seek out vulnerabilities to compromise. Ultimately, hackers discover and break the weakest links to gain access and entry into a secured environment. By actively and consistently seeking out vulnerabilities and weak points to secure, you reduce the hacker's potential to find and exploit a weak link. Of course, finding the weakest link is not restricted to examining technical vulnerabilities. Helpful employees with no security training are often the weakest link.

Weakest links are inevitable. But a strong weakest link is always better than a frail weakest link. Using a find-it-then-secure-it mentality helps build security that can withstand most common exploitation and intrusion attempts.

Fail-Safe

Fail-safe, **fail-secure**, **fail-open**, and **fail-closed** are not only design elements of firewalls and other security controls. Fail-safe is also an overarching security stance to drive an organization's security. When any aspect of security fails, the best result of that failure is to fail into a state that supports or maintains essential security protections. Generally, this means to maintain confidentiality and integrity protections. However, most fail-safe solutions will sacrifice availability protection to retain confidentiality and integrity. The fail-safe security stance is not just about using fail-safe security devices, but about designing the overall infrastructure with fail-safe as a core focus. See **FIGURE 7-3**.



FIGURE 7-3 Concept of fail-safe security stance.

Fail-safe does not need to be a standalone security stance. You can integrate fail-safe notions into any security perspective. Keep in mind the ultimate goals and policies of your organization. If availability is the top priority, then fail-safe may not be as viable an option as for environments where you can sacrifice availability to support other security protections.

Forced Universal Participation

It almost goes without saying that for security to be effective, everyone must work within the limitations established by your organization's written security policy. If you make exceptions—for upper level managers, for instance, who see themselves as being above the rules and able to do whatever they want without consequences—then you have only an assumption of security. No true security is present. Security works only when you employ forced **universal participation**.

Every worker, every manager, every senior executive, every temporary worker, every consultant, every vendor, every customer, every business partner, and every outsider must be forced to work within the security policy's limitations. Yes, exceptions often are necessary in the real world, but when exceptions become the norm, security is lost. When it is easy to bypass, avoid, or even ignore security controls, an attacker can use that same pathway to compromise the entire security environment.

Every enterprise that cares enough about security to write a security policy should establish it and enforce it to ensure that everyone follows the same rules. When it is unwritten policy to violate security to accomplish work tasks, an obvious disconnect exists between security and productivity. In this situation, reexamine both security and productivity goals. Security should support productivity, not impede it.

Universal participation is ensuring that everyone abides by security limitations. Potentially, users have many ways to purposefully violate security. This often occurs in environments where workers perceive the security limitations as too restrictive or

when newly imposed rules strongly limit freedoms they enjoyed previously. The less workers believe in and buy into the organization's security policy, the more likely they are to violate rules they perceive as unfair.

Examples of users purposefully avoiding or violating security—that is, not actively supporting and participating in security—include:

- Choosing poor passwords
- Sharing accounts with others
- Using unauthorized personal computer equipment on the company LAN
- Installing an unauthorized wireless access point
- Using unauthorized removable media on company equipment
- Using proxy tools to bypass firewall filtering
- Configuring a dial-out modem connection to establish unfiltered Internet access
- Using Internet based remote-access/remote-control tools to access their workstation from an external system without authorization
- Installing unapproved software

Successful enforcement and implementation of a security policy is dependent upon buy-in from upper management or the owners. To achieve universal participation in the security efforts, either workers must believe that compliance is in their best interests, or you must force compliance through consequences for violations. In most cases, voluntary compliance is better, because workers support the security effort rather than setting up employees to be adversaries pitted against the security infrastructure. It is easy to become disgruntled about security when some employees appear entitled or have a different set

of rules than others. Consistent enforcement is important so all understand the consequences and adhere equally to the policy.

AUTHENTICATION, AUTHORIZATION, AND ACCOUNTING

Security is supported and enforced by authentication, authorization, and accounting (AAA), a framework for controlling access to computer and network resources, enforcing policies, and more.

Authentication is the process of verifying the digital identity of an entity—for example, an employee, a customer, a server, an application, a website, a smart meter, a mobile phone, or a security guard. Authentication is also commonly called logging on.

Authentication can take place at the firewall or be passed through the firewall to take place beyond the firewall. Authentication at the firewall takes place when the allow/deny decision is dependent upon the identity of the host or user. Authentication beyond the firewall happens when the allow/deny decision is not dependent upon the identity of the host or user; instead, the firewall filters for traffic characteristics and leaves authentication up to the internal host receiving the connection.

When authentication takes place at the firewall:

- The firewall serves as the gatekeeper, allowing communication beyond the firewall (or not). Whenever different types or levels of port, protocol, or circuit access are granted based on individual user identity or membership in specific groups, authentication at the firewall takes place.

- When authentication takes place at the firewall, the firewall can perform the logon process or it can forward the task to a dedicated authentication service. Native firewall authentication methods can be basic, such as supporting a username and password online, or complex, such as supporting credit card payments or the use of smart cards. However, most firewalls have a limited number of authentication options.

Generally, authentication on or at a firewall supports only one—or at best a few—options. At-firewall authentication often requires defining credentials on the firewall separately and distinctly from any identity on the network. Thus, most at-firewall authentication solutions do not integrate into a network's single sign-on system. For this reason, large organizations rarely employ at-firewall authentication.

When a firewall hands off the authentication task to a dedicated authentication service that performs the logon process on behalf of the firewall (**FIGURE 7-4**), a much wider assortment of authentication options becomes available. A common use of this mechanism is IEEE 802.1x. IEEE 802.1x is known as **port-based network access (admission) control (PNAC)**. In this term, port is not referring to a port number in OSI Layer 4; instead, it refers to a more generic reference of a connection point to a network infrastructure.




FIGURE 7-4 A firewall handing off authentication tasks via 802.1x to a dedicated authentication server.

IEEE 802.1x is commonly found on many network devices, including firewalls, switches, routers, and wireless access points. It also has the nickname of “portal authentication.” An 802.1x device establishes the initial electronic connection or virtual circuit between the requesting host and the device, but before communication to a host beyond the portal device takes place, authentication must occur. 802.1x can support RADIUS, TACACS, Kerberos, certificates, passwords, smart cards, biometrics, special access codes, and even credit cards as means of authentication.

Handoff authentication, such as when 802.1x is used, is a fully scalable architecture. You can use it to manage at-firewall authentication for any number of users or hosts. The only limitation to 802.1x is the scalability of the authentication server that takes over the responsibility of managing logons on behalf of the firewall.

When authentication takes place beyond the firewall, the firewall makes allow/deny decisions based on traffic factors rather than user- or host-identification factors. In this configuration, the firewall is not involved in the authentication process. Instead, the host beyond the firewall that receives the communication handles any necessary authentication.

For example, a firewall can protect a DMZ. Most DMZ firewalls do not perform at-firewall authentication. Instead, the web servers hosted in the DMZ can provide both anonymous content and authenticated content to visitors. Visitors log on using any means of authentication supported by the web server, the protocols in use, and the visiting client. Commonly, this is forms-based authentication secured through a Secure Sockets Layer (SSL) or Transport Layer Security (TLS) encrypted session.

When any form of authentication takes place, whether at or through the firewall, you should always encrypt all authentication-related communications. Nonencrypted or plaintext authentication traffic is subject to eavesdropping and, thus, theft of credentials. Encrypted authentication traffic is resistant to eavesdropping as well as replay attacks, and in some circumstances even man-in-the-middle attacks.

Many different authentication protocols, both encrypted and plaintext, can be employed by at-firewall or beyond-firewall

authentication. Whatever the solution, evaluate all authentication options in light of the native security and inherent risks. If you are using authentication encryption, realize that this does not ensure that subsequent data exchanges are automatically encrypted. Select solutions that provide protection for authentication traffic, as well as data traffic.

You can avoid authentication completely with anonymous connectivity. In an anonymous connection, the user and/or host identity is unknown, and access is granted or denied based exclusively on the traffic content and the proper formation of resource requests. For example, public website access can be anonymous, such as at Cybersecurity at Ready.gov (<https://www.ready.gov/cybersecurity>), which supports mostly anonymous connections to online resources.

Authorization takes place after authentication. Authorization is the process of defining which resources can be accessed by an electronic entity and what level or type of access is granted. Authorization is also commonly called access management or privilege management. Authentication determines, within established parameters, who you are. Once that is known, it is possible to determine whether you are allowed to do what is requested. Authorization, like authentication, can take place either at the firewall or beyond the firewall. When authorization takes place at the firewall, the firewall makes a decision on allowing subsequent communications across the firewall. This is sometimes known as controlling the circuit. If the authenticated user does have access to a resource, then the firewall allows and supports the connection. If the

authenticated user does not have access to a resource, the firewall will block or drop the connection.

Usually, the firewall is involved only in authorization of the connection or circuit. Any content within the subsequent communication might be subject to content filtering, but that is a distinct and separate service and function of the firewall. This form of authorization control is the same as a circuit proxy.

The destination host controls what is and is not accessible to users when authorization takes place beyond a firewall, just as with beyond-firewall authentication. The firewall is involved only in monitoring for abuses or potentially filtering content. When the firewall performs content filtering on a communication not managed by the at-firewall authorization, this is the same as an application proxy.

While authentication and authorization are commonly used together to control users' access based on their identity, they can be employed independently. For example, a firewall performing at-firewall authentication does not need to also provide separate authorization. If users authenticate, then they have access to whatever is beyond the firewall. If users fail to authenticate, then they do not have access. In this example, the authentication process is a form of authorization.

Similarly, a firewall can perform at-firewall authorization without authentication. This is the basic or normal function of firewall rules. Authentication is assumed when addresses and ports in traffic headers "identify" the communications; the rule sets themselves provide authorization restrictions. If a port is closed, access is denied. If a source address is blocked, access is denied.

Authentication and authorization are important components of a complete network security structure. Use these services on a firewall or beyond the firewall, as necessary. Additional important services related to firewall security are monitoring and logging of traffic and events that occur on or across a firewall.

The third A in AAA is accounting. It is different from financial accounting in that it logs what a user has done on the system. In other words, it keeps track of who did what to which assets via one or more log files.

PLACEMENT OF NETWORK HARDWARE FIREWALLS

The placement of your network's firewalls depends on several key but subjective factors. No single correct answer or deployment strategy is perfect for every application. However, several good general guidelines can be of help in planning firewall placement.

- **Understand the structure of your network**—Where are natural or organizational divisions on the network? Do these correspond to distinct subnets or geographic/physical locations? In many situations, placing a firewall between each division or physical location (even different floors of a building) can be beneficial. Do you have a DMZ or extranet? You will need to isolate these from the private LAN and from the Internet.
- **Know the traffic patterns of your network**—What are the common vectors and pathways of communication, both between internal systems and between internal and external systems? Consider positioning firewalls to filter all traffic. This may involve the creation of chokepoints.
- **Protect Internet access points**—You should firewall-protect all Internet access points or gateways. In fact, any external gateway or portal should have a firewall.
- **Protect remote access and wireless connections**—You should position a firewall between each remote access server (RAS) or network access server (NAS) and the private LAN. Assume that all remote and wireless connections are potentially malicious.
- **Determine hacker pathways**—Based on previous compromises, what are the most likely access pathways an internal or external

hacker might use to breach your network? Place firewalls along these pathways.

- **Assess the entire environment**—Does any other unique aspect or feature of your IT infrastructure warrant consideration for firewall protection?

Placing a firewall everywhere mentioned in these guidelines may not be practical, cost effective, or secure. Most commercial-grade firewalls are costly. Even if there is no purchase cost (due to promotions or open source), the ongoing maintenance costs of firewalls add up significantly over time. With firewalls deployed on every host and on every segment, the overhead of firewall management would far exceed the total IT budget of most businesses. Before any deployment, conduct a risk assessment to determine whether a firewall is the best choice of countermeasure.

A second drawback to ubiquitous deployment of firewalls is over-dependence. Firewalls are only part of a complete security solution, not a whole solution by themselves. By deploying firewalls on every host and every segment, other essential security endeavors might be overlooked or deemed unnecessary. This would be a disastrous security stance. Firewalls cannot compensate for poor security policy, poor security management, or a lack of proper system hardening.

BENEFIT AND PURPOSE OF REVERSE PROXY

Reverse proxy is a firewall service that allows external users access to internally hosted web resources. This service takes the traditional proxy function and inverts it. Instead of hiding the identity of the client reaching out to the Internet, reverse proxy hides the identity of the web server accessed by the Internet (or external) client.

External users direct their queries to a public Internet IP address and a default or assigned port number as when accessing any service. However, the IP address is the address of a proxy server, not the actual web server hosting the requested resource. The proxy server then performs a function similar to network address translation (NAT) on the request to convert the destination address to the internal (likely private) address of the resource host. See **FIGURE 7-5**.




FIGURE 7-5 Reverse proxy.

You can deploy reverse proxy to support load balancing or load distribution across multiple internal resource hosts. This allows clients to use a single public address to access a cluster of internal web servers.

Other common reasons to deploy reverse proxy include:

- **Reverse caching**—Reverse caching takes static content and caches and sends it by the proxy, rather than requiring that each request for the same content be served by the web server.
- **Security**—Using a reverse proxy adds an additional layer of protection and control between Internet-based users and internally hosted servers. Proxying allows the real identities of the internal servers to remain unknown or at least obfuscated.
- **Encryption**—The proxy server can serve as the endpoint for Secure Sockets Layer (SSL) or Transport Layer Security (TLS) encryption tunnels. This can allow a firewall or IDS to monitor

and filter the contents of the traffic before it reaches the web servers. Proxy-based encryption can also benefit from hardware and software acceleration to maintain high-performance communications.

Reverse proxy does not have to function solely on the private network's borders. It can also operate internally between subnets or departments, especially within IT infrastructures of very large corporations. Reverse proxy is also useful in an extranet or a DMZ. Indirect access by reverse proxy can often be a significant security benefit when you deploy and manage it properly.

USE AND BENEFIT OF PORT FORWARDING

Port forwarding is a firewall, proxy, and routing service that can receive a resource request on an interface at one port, then forward the request to another address on the same or different port. Port forwarding is used in reverse proxy, but only for web traffic like TCP port 80 or port 443. Port forwarding can support any service on any port.

Port forwarding is a variation or enhancement of NAT. This is why port forwarding is sometimes called port address translation (PAT). When a request reaches an outward-facing interface to a specific port, the request goes to an internal host. The routing is controlled by a static NAT mapping that defines which internal IP address and port will receive communications sent to an external IP address and port.

Port forwarding does not support caching, encryption endpoint, or load balancing. Only a single internal machine can use a forwarded port at a time. The internal receiving (destination) host will perceive the source of the communications as the port-forwarding device because the translation service will rewrite the source address as its own. Thus, the destination system will not know the real originator of a communication.

You can find port-forwarding services on almost any service or device that supports NAT. This includes most firewalls (hardware and software), WAPs, and Internet connection devices (such as DSL and

cable modems). Port forwarding is also an essential element in the Internet Connection Sharing (ICS) service of Windows that allows multiple systems to share a single Internet connection through the primary connected computer.

CONSIDERATIONS FOR SELECTING A BASTION HOST OS

A **bastion host OS** is a system designed, built, and deployed specifically to serve as a front-line defense for a network. A bastion host is the first (or nearly so) host accessed by external entities on their way to access DMZ, extranet, or private network resources. The bastion host withstands the brunt of any attack attempt to provide protection for hosts behind it.

Bastions were the highly fortified sections of medieval castles designed to assist with defense. The bastion was located along the castle perimeter wherever access attempts were likely, such as around the main gateway or entrance, any side or back entrances, or natural landscape pathways of approach. The bastion areas usually had thicker, stronger walls; room for additional warriors; and special defensive features. These special defensive features might have included slits for shooting arrows, holes for ramming spears or pushing away ladders, and pots of boiling oil to pour on attackers.

On modern computer networks, a bastion is a fortified computer device—possibly a host, firewall, or router—placed in the line of fire between privately owned and controlled networks and the public Internet. The purpose of a bastion host is to provide frontline filtering and defense against typical attacks. Most firewalls—especially appliance, hardware, or device firewalls—operate as bastion hosts.

Software firewalls can also be bastion hosts if the host is properly locked down and hardened.

Knowing that firewalls are commonly deployed as bastion hosts raises the question of what type of OS you should use as the host OS on a bastion host firewall. Two main categories or divisions of bastion host OSs are **proprietary OSs** and **general-purpose OSs**.

Proprietary OSs are operating systems built exclusively to run on a bastion host device. Most appliance firewalls employ a proprietary OS. This includes commercial firewall devices as well as many ISP connection devices and wireless access points. These proprietary bastion host OSs support the functions or services critical to security (or their other primary purposes) and little else. An example of a proprietary bastion host OS is Cisco IOS.

A firewall device's bastion host OS supports only firewall functions. An ISP connection device's bastion host OS supports firewall services and connectivity services. A wireless access point's bastion host OS supports firewall services, wired connectivity services, wireless connectivity services, and possibly other services, such as 802.1x.

General-purpose OSs include Windows, Linux, Mac OS, UNIX, and others. These are OSs that support a wide variety of purposes and functions, including serving as client or server host OSs. When used as a bastion host OS, each must be hardened and locked down. Otherwise, an insecure host OS can render the security provided by a firewall worthless. If an attacker can crash the firewall host or bypass the firewall filtering, then the firewall is not providing effective security.

It is possible to reasonably harden a general-purpose OS for use as a bastion host OS, but this takes specific and detailed modifications. Most software firewall products include a guide to perform OS hardening to improve the security of the bastion host and reduce the vulnerabilities and backdoors that might permit firewall compromise.

Some firewalls include a prehardened version of a general-purpose OS as the bastion host OS. This is the case for many Linux-based firewalls, such as pfSense.

The benefit of using a general-purpose OS as a bastion host OS is that such OSs are widely available. Some—mostly Linux variations—are free. Also, you can leverage your current knowledge and skill with a general-purpose OS when you are using it as a bastion host OS. Keep in mind, however, that general purpose OSs are much more widely attacked, and the risk of new exploits is high.

The benefits of a proprietary OS as a bastion host OS are fewer known attacks and less risk of future exploits. However, proprietary OSs might be significantly different from other OSs in the environment and may require system administrators to learn a completely new system to properly and securely administer the firewall. Also, most proprietary OSs are officially released only on the bastion host hardware. This makes it more difficult to obtain and more expensive in most cases (at least in comparison to free and open-source alternatives).

MONITORING AND LOGGING

Monitoring and **logging** are the watching and recording of events and traffic that take place on or across a firewall. In theory, you can record every packet traversing a firewall; however, this is usually impractical because the storage space required would be immense. Certain products can maintain a sliding window of recorded traffic, encompassing hours of communications using multiple terabytes of storage. However, these devices are for advanced security management and beyond the scope of this chapter.

Firewall monitoring and logging are about recording events that take place on or across a firewall (**FIGURE 7-6**). You can use various approaches to logging. You can log everything, or you can log only events of seeming importance. Generally, you will fine-tune logging through real-world experience over time. Most security professionals recommend capturing everything initially. Later, you can tune logging and reduce it once you have identified specific types of information, traffic, or events as having no forensic, evidentiary, or historical value.


System ▾ Interfaces ▾ Firewall ▾ Services ▾ VPN ▾ Status ▾ Diagnostics ▾ Help ▾

Status / System Logs / System / General

System Firewall DHCP Captive Portal Auth IPsec PPP VPN Load Balancer OpenVPN NTP Settings

General Gateways Routing DNS Resolver Wireless

Last 50 General Log Entries. (Maximum 50)

Time	Process	PID	Message
Apr 23 10:04:37	check_reload_status		Syncing firewall
Apr 23 10:04:37	php-fpm	340	Local User Database: Successfully edited user admin
Apr 23 09:57:50	php-fpm	339	/index.php: Successful login for user 'admin' from: 172.21.0.1 (Local Database)
Apr 23 09:22:36	php-fpm	339	/rc.openvpn: OpenVPN: One or more OpenVPN tunnel endpoints may have changed its IP. Reloading endpoints that may use WANGW2.
Apr 23 09:22:36	php-fpm	339	/rc.openvpn: Gateway, none 'available' for inet6, use the first one configured. '
Apr 23 09:22:35	check_reload_status		Reloading filter

FIGURE 7-6 An example of pfSense’s firewall logging display interface.
© pfSense

Firewall logging ensures that the defined filters or rules are sufficient and functioning as you expect. By watching traffic coming through a firewall's filtering scheme, you can discover traffic that the firewall should have blocked, but for which an effective rule was not defined (or properly defined).

■ NOTE

One example of a device able to record all traffic within a sliding period is the Solera DS series of network forensic appliances (owned by Symantec). For more information on this type of advanced network monitoring product, visit the Symantec website at

<https://www.symantec.com/products/network-forensics-security-analytics>.

Generally, record every connection allowed by the firewall, as well as every packet dropped or blocked. This stance of “log and monitor everything” ensures that nothing goes unnoticed and that future investigations of log file contents will reveal complete data. You can always throw out extra and unnecessary data later, but you cannot recover important data not recorded during the actual event.

Firewall logging can take place on the firewall or offload to a dedicated logging system. In either case, the firewall collects the event and/or packet you want to record. The difference is whether the information goes into a log file stored locally on the firewall or moves across a network connection to a dedicated logging server.

One well-known and widely used centralized logging system is syslog. **Syslog** is a standard for forwarding log messages from a client device (the firewall in this situation) to a centralized server. Linux and Unix systems have supported syslog capabilities for decades. Windows systems prior to Windows 7 required third-party products to interact with syslog services. Most network devices, including firewalls, switches, proxies, and so on, natively support syslog (meaning that no third-party media is required for use)

Firewall logging:

- Creates a historical record of activity for traffic and trend analysis, as well as growth prediction
- Tracks usage levels and times for load balancing, accounting, or even back-charging users
- Discovers new methods or techniques of attack, especially those based on network packet manipulation
- Detects intrusions or attempts to breach security
- Creates legally admissible evidence for use in prosecution

A firewall log of all events produces a significant amount of recorded information. Always track available remaining storage capacity to ensure that sufficient space remains to store new logged events. As a current storage device reaches capacity, back up and move the records into long-term archival storage. In a future investigation, you can never know how far back you will need to look to discover evidence of a security breach. Your organization's written retention policy should define how long you will retain logs.

In addition to juggling the storage requirements, you will need to be able to analyze logged data. Log file analysis is not usually an exciting proposition for most people. Long and tedious, it can be

mind-bendingly boring. Fortunately, a growing field of automated analysis tools can reduce the time and burden of log file investigation. Firsthand knowledge, understanding, and experience are always the best tools for interpreting and understanding firewall log contents.

Upon infiltrating your network, one of the last steps most experienced hackers take is to remove all evidence of their activities. They normally do this by deleting log files or surgically editing log files with specialized tools. For example, SolarWinds Event Log Consolidator is a tool that edits Windows log files and removes only those entries matching a keyword, such as an IP address or username. To combat this potential threat, use a write-once read-many (WORM) storage device to save all log files. Employing WORM means logs can be written once and then accessed for read-only access many times. Intruders cannot edit data written to a WORM device.

If you make the decision not to record everything taking place across the firewall, here is a short list of essential things to log:

- All connection attempts rejected by the firewall
- All traffic directed at the firewall
- Any access to internal resources by external hosts
- Any access to external resources by internal hosts
- Any modification or disabling of firewall rules and filters
- All firewall reboots, storage device shortages, or other abnormal operational events

Every firewall management interface is different. Thus, the methods and mechanisms you employ to configure your specific logging and

monitoring will be different, as well. Always be sure to consult the firewall's documentation about the type of logging and what types of events are logged before deploying a firewall.

Define firewall logging in the firewall security policy included in your organization's network security policy. The logging and monitoring features of a firewall are as essential as the rule and filter mechanism you use to allow or deny traffic. Configuring and recording logs is important, but understanding how to interpret firewall logs is more important.

UNDERSTANDING AND INTERPRETING FIREWALL LOGS AND ALERTS

Enabling firewall logging is often a fairly simple task. Recording traffic and events into a file is not really a complex task. But deciding which events warrant the triggering of an alert and what the contents of a log file actually mean is not always so straightforward.

An alert is the automated notification to an administrator when a specific event affects the firewall. Some software host firewalls are preconfigured with a wide number of pop-up messages triggered each time one of the all-too-common activities occurs. For example, each time a new application attempts to communicate with the network, a firewall pop-up alert might notify the user of this event and prompt an action (FIGURE 7-7).



FIGURE 7-7 An example of a Windows Defender Firewall pop-up alert.
© Microsoft Corporation

Appliance firewalls and those generally designed to protect a network rather than a single host are not as commonly configured with preset default alerts. Instead, the firewall administrator makes choices about which events actually need alerts. Some common events that usually warrant alerts include:

- A firewall reboot
- A connection attempt with the firewall host
- Communications to or from a pre-identified IP address or subnet
- Detection of some or all attack attempts
- Detection of any successful intrusion or security breach

Ultimately, the organization's security stance and the level of risk determine which events require alerts, as opposed to simply recording each in a log file. A common recommendation is to alert only on those events likely to require an administrator's immediate response or action. Everything—or nearly everything—should go into a log, but not everything warrants an alert.

With or without alerts, a firewall's logging system creates a massive amount of raw data that is not useful. You can perform analysis of a log file either in real-time or in periodic batches. You can hire a large staff of professionals whose primary task is to perform real-time log file analysis, or you can use an automated analysis or intrusion detection/intrusion prevention (IDS/IPS) tool. In most cases, automated log file analysis tools both reduce the strain on manpower and leverage the power of high-speed computing against the massive volumes of information collected by network firewalls.

Through careful analysis of firewall logs, you can detect several different types of unwanted or malicious activities, including worms,

Trojan horses, remote control tools, port scans, and directed intrusion attempts. When you perform manual log analysis, the first hurdle is to gain access to the log file content.

Many firewalls record the log files in binary or obfuscated form so that they are difficult for external or unauthorized users to access. Such log files require administrators to access the log contents through a log reading interface. Some third-party products function as alternative management interfaces and can also interpret the log file format for viewing.

Once you obtain access to the log contents, either directly or through an interpretation interface, the next hurdle is to ensure that you are viewing the correct log. Several different types of log files can operate on your system, including traffic logs, audit logs, alert logs, console logs, and even a host device's OS logs for software host firewalls. Another nuance is that the firewall device, a management host, or a centralized logging server might store its own logs. Finally, several dated versions of a log might exist. Be sure to open or access the most recent log—or whatever date range is of relevance to a specific investigation.

Of course, the key to interpreting and understanding the contents of a firewall log lies in understanding how network communications function. Every communication session involves two hosts, each with an IP address and port number. Some forms of communications have less information, such as blocked connection attempts, connectionless traffic, and Layer 3 protocols such as ICMP and IGMP.

By examining each established session to evaluate the source and destination IP addresses and the port numbers in use, you can quickly determine whether the transaction was valid or invalid. If the

addresses are known—or at least not known to be malicious—and the ports are for acceptable services, then the communication is most likely benign. However, if an address is a known malicious host or an invalid address (such as an unassigned internal address) or you discover a known malware port, these are signs of unwanted and potentially malicious communications (or at least attempts to communicate).

The Internet Assigned Numbers Authority (IANA) and some security vendors maintain lists of well-known port numbers and the associated applications. They often list known malicious software and their commonly used port numbers. Search for “malicious ports” or “Trojan ports” to find current lists.

To examine and interpret firewall logs, you might:

- Identify and discard all packets involved in benign sessions between valid hosts and valid ports.
- Identify and discard all valid UDP transmissions to and from valid hosts and ports.
- Identify and discard all valid ICMP (and other authorized protocols) to and from valid hosts.
- Identify any communication from known malicious addresses. (This is a symptom of communications from known malicious hosts.)
- Identify any communication from unauthorized or unassigned internal addresses. (This is a symptom of rogue internal hosts.)
- Identify any communication to an unknown or unauthorized port on an internal host. (This is a symptom of rogue internal services.)
- Identify any patterns of serial communication attempts across a sequential or randomized set of addresses and/or ports. (This is a

symptom of scanning or probing.)

- Identify significant volumes of nonsession traffic of the same or similar construction to one or a small group of internal addresses. (This is a symptom of flooding.)
- Identify any packet of invalid size or invalid header construction. (This is a symptom of scanning, probing, or an exploitation attempt.)
- Identify any packet with configured source routing. (This is a symptom of attempted intrusion or traffic shaping.)
- Identify packets destined to nonresource-sharing systems, such as firewalls or routers. (This is a symptom of probing, network scanning, or an exploitation attempt.)

This is just the start of a firewall log file analysis procedure. For each and every threat, exploit, attack, or concern that an organization wishes to identify, add a step to search log contents for the appropriate symptoms. Remember that many symptoms are shared across a wide variety of malicious activities.

Keep in mind that hackers can perform malicious activities within seemingly benign traffic. This is especially true when the payload of traffic is the malicious content. Firewalls can detect malicious or abnormal characteristics based on header information and packet construction oddities, but many types of malicious payloads are not firewall filterable. For example, password guessing against a logon prompt is typically not a filterable activity. An IDS or IPS could potentially detect online password guessing, but a firewall cannot.

This brings up an important issue discussed previously, but it bears repeating: Firewalls are not perfect. Firewalls are only part of a security infrastructure; these are not the totality of security. Even when you use them properly, firewalls have limitations. For example,

firewalls are of no help in protecting against an insider attack. Firewall limitations are discussed later in this chapter.

The best tool for interpreting the contents of a firewall log is a thorough understanding of the normal and expected traffic that occurs across your network. By knowing what is normal and expected, you are much better prepared to detect and identify the abnormal and unexpected. Keep complete configuration documentation and actively seek direct, hands-on knowledge and experience to improve your ability as a firewall administrator to examine, interpret, and take action on the contents of log files.

Do not become complacent or assume that previous administrators did their due diligence. Hackers can be patient. One tactic shared by the Federal Bureau of Investigation (FBI) is the persistence of some hackers once they gain access to a network and stay dormant, sending only an occasional packet stream but posing no threat. This allows the account and traffic to appear normal and legitimate, which would allow for the future leaking of information when the hacker identifies something of value.

INTRUSION DETECTION SYSTEMS AND INTRUSION PREVENTION SYSTEMS

A controlled network border sentry device like a firewall filters any traffic attempting to cross. Not all traffic that needs monitoring crosses a network border guarded by a firewall. That's where an IDS is valuable. An IDS or an IPS monitors internal hosts or networks, watching for symptoms of compromise or intrusion. Effectively, an IDS is a form of burglar alarm that detects when an attack is occurring within the network.

An IDS serves as a companion mechanism to a firewall. Once an IDS detects an intruder, it can send commands or requests to the firewall to break a connection, block an IP address, or block a port or protocol. You must configure the firewall to receive these commands and authorize the IDS to send each. However, not all IDSs and firewalls are compatible in this manner.

Eventually, security professionals developed the IPS concept. The IPS strives to detect the attempt to attack or intrusions before it can be successful. Once detected, the IPS can respond, preventing the success of the attempt, rather than waiting until after a successful breach to respond. IPSs do not replace IDSs. Instead, these are often used as an initial layer of proactive defense, relegating the IDS to a reactive measure against those events that the IPS misses or that internal personnel perform.

IDS and IPS are important components of a complete network security solution. However, these devices are not without issues. IDS/IPS solutions can create a false sense of security under certain conditions. Two commonly known conditions include unknown zero-day attacks and **false positives**.

■ NOTE

Next-generation intrusion prevention systems (NGIPSs) are available that include features such as application control and user control. An NGIPS extends the traditional functionality of an IPS by monitoring a much longer list of network protocols to detect more attacks.

When zero-day attacks threaten the network, an IDS or IPS might not have the mechanism to detect the attack. Thus, the lack of an alarm could cause administrators to assume that no attacks are occurring. In most cases, the lack of an alarm does mean that nothing malicious is happening. So, this assumption is not completely unwarranted. But if an IDS never triggers an alarm, you probably have a poor detection system rather than the complete lack of compromise attempts or attacks.

When an IDS fails to detect an attack, whether it is a zero-day or a known attack that was simply overlooked, this is called a **false negative**. A false negative is one of the most problematic issues with an IDS, as it can lull the organization into a false sense of security. When alarms are not going off, it is common to assume that no malicious events are taking place. If that is a false assumption, then

real attacks are occurring and the security staff is unaware. This is the worst type of security breach.

False positives may create a false sense of security for the opposite reason: namely, too many alarms from benign occurrences. Remember the story of the little boy who cried wolf? After the initial alarms turn out to be benign activity, the urgency of responding to alarms diminishes. After additional false positives, an administrator might not investigate the alarms. Eventually, you ignore the alarms altogether. Once this situation arises, you treat alarms for malicious events as false positives, once again reinforcing a false sense of security.

IDS and IPS can use rule-based detection mechanisms similar to those of a firewall, and can also employ detection mechanisms borrowed from antivirus technology. For example, some IDSs and IPSs have a database of signatures or patterns of known malicious activity. This is known as **signature-based detection, database-based detection, or knowledge-based detection**. Any attack within the database can be detected in live activity by the IDS.

Another form of detection is anomaly based. **Anomaly-based detection** systems look for abnormalities. To do that, you have to first define “normal.” Normal can be defined by rules or filters that prescribe all of the valid packet constructions and header contents. Then, anything that fails to match the definition of normal is an anomaly. Since mistakes, errors, and poor network application programming can occur, an anomaly is not necessarily malicious or an intentional attack.

Yet another detection mechanism is behavioral based. **Behavioral-based detection** looks for differences from normal based

on a recording of real-world traffic that establishes a baseline. A normal baseline can be recorded over hours or days. Once recorded, it is the yardstick by which to measure all future activities. Any future event not similar to behaviors in the baseline set represents a possible violation.

As with anomaly detection, behavioral-based detection has limitations. For example, all forms of benign behavior may not have occurred during the behavior recording time. Or, malicious events might have taken place during the recording period. In either case, future situations can generate false positives (mistaking benign behavior as malicious) or false negatives (mistaking malicious behavior as benign). The record should include data that is date- and behavior-specific, such as a system update that occurs on the second Tuesday of each month, to aid in reducing false positives. Because a hacker may already be resident on the network, do not ignore traffic simply because it is part of the baseline. Investigate normal traffic initially to be sure.

Some IDSs and IPSs can import and interpret firewall logs as another source of information in their efforts to detect unwanted activity. Thus, in addition to reviewing log files manually or with separate investigation tools, they can also benefit IDSs and IPSs and improve scope and effectiveness.

IDS and IPS are not substitutes for firewalls, but they are potential supplemental or companion solutions that can expand the effective detection methods already available. When considering an IDS and/or IPS to improve the filtering already provided by a firewall, carefully consider deployment options and thoroughly plan out the improved infrastructure. Compare the plan to a typical IT

infrastructure and consider where the placement of an IDS and/or IPS makes the most sense for your organization.

SECURITY EVENT AND INFORMATION MANAGEMENT

The combination of securing a network using security event management (SEM) and security information management (SIM) is **security information and event management (SIEM)**. The combination of these techniques allows for relevant information collected from multiple systems and processes to be compared and aggregated so it may be used in decision-making. Statistical measures and analytics aid in changing the facts about events into something logical and easier to comprehend. The relevant information pertaining to an event that is stored for review or later access is called a log. What is logged is usually determined by default or by customization of each device. SIEM logs are usually much more comprehensive than standalone system logs. This allows reviewers to identify whether an event is an anomaly or something requiring security action. It is a more holistic approach to network event and log management.

Logs should be written using the WORM principle. The collection of disparate log information from systems on the network is **aggregation**. Identifying how the events captured in the logs relate using statistical measures is **correlation**.

In addition to enhanced reporting options, SIEM systems offer automated alerts with trigger thresholds that can be customized for the specific needs of an organization. Log management and review is monotonous and often overwhelming if many devices on a network

capture event activity to logs. This is because logging occurs in real time, and keeping up with events as each happens is difficult.

Another benefit of SIEM products is time synchronization, which manages the disparity between system clocks—such as those in different time zones of a global network—so that events can be assessed based on when each occurred, not on the recorded time based on location. When tracking events that occurred up to and after a breach, knowing the exact time and its relationship to recorded time on other systems is essential to successful tracking and prosecution.

A feature many managers appreciate is event deduplication services. SIEM can be set to eliminate or silence duplicate events from multiple devices when pertaining to one item, packet, or occurrence. That is, the same event logged by three different devices would not become escalated unless there was an associated issue with those three devices. One entry would capture the essence of all three events so as not to make any event appear more significant within the scheme of network traffic.

Device and information management can occur in-band or out-of-band. In-band management is the most common method and requires using common protocols to collect information about the network. These protocols could include Secure Shell (SSH), Simple Network Management Protocol (SNMP), or Telnet. The traffic is only internal, so the use of Telnet is not a security risk, as no outside access is gained through the protocol. A serious issue with in-band network management is that if the network is down for any reason, there may be no way to capture or review data and logs. Often, this type is supplemented with out-of-band capabilities; this does not

happen but in all installations, which leaves some networks vulnerable.

Out-of-band management (OBM) is when a device, commonly a server, has access to the specific devices controlled through a management port in each. This is often accomplished using a networked device with a public IP address that accepts only traffic from whitelisted network devices. When there is a routing or traffic congestion issue on a network, this is a benefit if the appliance configuration still allows some network access. OBM techniques are often deployed on mission-critical networks so there is an alternate independent path to those devices.

EVALUATING NEEDS AND SOLUTIONS IN DESIGNING SECURITY

Any organization larger than a single person will generate multiple opinions about what should and should not be allowed across the firewall. Almost everyone would agree on allowing—at some level—the necessary protocols, applications, and services essential for mission-critical (or at least operationally necessary) business tasks to cross the firewall-secured network boundaries.

Opinions may differ about what constitutes a necessary business task, but if it is a necessary task, the organization's security solution should make the task possible. Restrictions, limitations, filtering, and logging of even necessary business communications must occur. The point, however, is making a determination about what should and should not cross a firewall.

Generally, four main types of communications take place within a business environment: business-essential, business-wanted, personal, and malicious. Business-essential communications must take place or the business will suffer. Blocking critical transactions will directly cause problems for the business by impeding the production and distribution of products, inhibiting customer service, reducing profits, and so on. The security infrastructure must support business-essential communications.

Business-wanted communications are not essential to the core function or purpose of the organization. The failure of these business

communications might be inconvenient or reduce quantity or quality, but the essential business functions will go on. Business-wanted communications help the business do what it does better, faster, cheaper, cleaner, more efficiently, or in a flashier way. But when those wants are not available for whatever reason, the business's core activities still function sufficiently.

Personal communications are those transactions between individuals inside the company and with external entities not directly related to business tasks. If you eliminated personal communications, all business functions would continue unhindered. However, that's in a perfect world. In the real world, stifling human communications can lead to low job satisfaction, feelings of isolation, and even worker disaffection. This, in turn, causes a reduction in productivity and quality.

Most organizations must strike a balance between allowing only business communications and allowing all communications. Obviously, allowing every communication is a bad idea from a security standpoint, as well as for productivity. Often, a modest level of website access, filtered email access, and potentially other nonthreatening services as forms of personal communications can travel on and over company equipment during work hours with a reasonable amount of security.

But no clear or distinct boundaries define what level of personal communication will maintain worker morale and job satisfaction. It is often fairly obvious when restrictions are too tight, because job performance suffers or breaches of security increase. Part of the solution to this problem is educating users or workers about what is acceptable and unacceptable behavior on the company's network.

This includes defining what types of communications to allow or block, whether for business or personal use.

Malicious traffic, the fourth common type of network traffic, should always be blocked. This is a widely held security stance and the whole point of using security in the first place. However, it is not easy to identify all malicious traffic, and it can sometimes appear to be one of the other three common forms of communications. Improving security over time is essential to maximizing protections against malicious communications, while at the same time supporting necessary, desired, or benign personal communications.

Every organization must define its own parameters and justifications for each type of transaction allowed across any security perimeter, whether network traffic or physical access. Whatever the determination, clearly spell out the designations for valid versus invalid transactions in your organization's written security policy, and then implement them in the security infrastructure.

Once you have defined necessary transactions and prescribed the desired security, use these concepts in actual security-enforcing software and hardware solutions. This is where your understanding of the goals of security, the organization's mission, and the budget is critical. There is always a product that costs more, but more expensive does not always mean better security. In fact, in many cases, cheap or free solutions can offer equivalent or better security than the most expensive product available, if it is properly configured and implemented.

Using security is an essential part of the risk assessment and management process. Part of evaluating security is understanding the threats, attacks, and risks, as well as the available countermeasures.

One aspect of understanding countermeasures is performing a cost/benefit analysis to obtain the best security solution for the invested dollar. Don't just purchase the most expensive solution; that is never a reliable measure of security quality.

Also, do not automatically purchase the product your cost/benefit analysis says is the best option. You must evaluate each proposed solution in light of all other elements of the security infrastructure, as well as the size of your budget. It is usually not the right choice to spend 50 percent or more of the budget on a single security component. Instead, devise a spending plan or ratio tool to guide the procurement of security solutions.

One possible method is to allocate 10 percent of the budget to each of the 10 major sections or divisions of security addressed. This assumes that you can divide the overall security of your organization into 10 compartments—you might need 30 sections, or maybe you can get by with only four. Whatever the number of divisions you use, grant each an equal or weighted portion of the budget based on risk. When a solution costs more than is available within a certain section, you can then shift funds between sections as needed.

The purpose of this system is to make security administrators and designers more aware of the cost of security and to prevent overspending in one area and underspending in another. You should spend security funds somewhat evenly to secure the overall organization, rather than over-securing one area and neglecting another. A hacker or intruder who encounters a highly fortified defense is likely to go looking for a less secure backdoor.

Ultimately, you should treat your security budget just like any other budget. Categorize and prioritize expenses. Outline spending on

paper before writing the first check or funds transfer, or swiping the corporate credit card. Allocating funds based on needs and importance, then adjusting those allocations based on market conditions or changing threats, is all part of proper security management (as well as security budget management).

WHAT HAPPENS WHEN SECURITY GETS IN THE WAY OF DOING BUSINESS?

Businesses exist to earn a profit. Profits come from performing tasks, such as producing products or providing services. When the security infrastructure interferes with essential business tasks, something has to change. Security should prevent compromise while supporting and allowing business functions. When security changes or when business tasks change, the security policy and the business should adjust.

Absent this adjustment, security can interfere with business operations. Some organizations will never face this situation, but every business must have a response plan in place in case it does. When business tasks are at odds with security protections, what should you do? A few short-term options are available; none of them are optimal. One short-term option is disabling security so the business tasks go forward. Whether the security is disabled for a short period or permanently, you provide the opportunity for compromise, intrusion, or sabotage. Using this shortsighted solution might not immediately cause harm to the infrastructure, but it establishes a pattern and mindset that security is not important.

Even if no attack or breach occurs during this period of lesser or no security, the real damage is the change in mentality toward security: a shift from viewing security as essential and mandatory to optional and nonessential. Once an organization believes that security can be turned off when inconvenient, then security suffers, and it is

only a matter of time before a catastrophic compromise occurs. Security depends on vigilance. Once the commitment to vigilance is lost, so is real security.

Another short-term option is not to perform the task that security is blocking. If the task is not essential, it might be possible to move forward as an organization without performing that specific task again. However, if it is an essential or critical task, failing to support it will negatively affect the organization.

There is another way. The preferred and secure long-term response is to reevaluate the business task in the light of the security infrastructure. Design a new security solution or modify how the task is accomplished. Support both security and business functions. These are both essential to the long-term stability and vitality of an organization.

CHAPTER SUMMARY

Never deploy firewalls hastily. Instead, use careful investigation and planning in the design and use of firewalls. Consider the numerous firewall deployment issues regarding your business needs and network situation.

Some important concerns include deciding:

- Which security strategies to integrate
- How AAA enforces security
- Whether you need reverse proxy
- Whether to configure port forwarding
- Which bastion host OS to use
- How to monitor firewalls and handle log data
- How to deal with IDS/IPS false positives and false negatives
- How the firewall will interact or interfere with business processes

KEY CONCEPTS AND TERMS

Aggregation

Alert

Anomaly-based detection

Bastion host OS

Behavioral-based detection

Centralized logging system

Correlation

Database-based detection

Deduplication

Diversity of defense

Fail-closed

Fail-open

Fail-safe

Fail-secure

False negative

False positives

General-purpose OSs

In-band management

Knowledge-based detection

Logging

Monitoring

N-tier deployment

Out-of-band management (OBM)

Port-based network access (admission) control (PNAC)

Principle of least privilege

Proprietary OSs

Reverse caching

**Security information and event management
(SIEM)**

Security through obscurity

Separation of duties

Signature-based detection

Steganography

Syslog

Time synchronization

Universal participation

Weakest link


Write-once read-many (WORM)

CHAPTER 7 ASSESSMENT

1. When is a reverse proxy useful?
 - A. To grant outside users access to internal email servers
 - B. To support internal users accessing the public Internet
 - C. To allow internal users to access external web servers
 - D. To offer external entities access to an internal web server
2. Which of the following is an event found in a firewall log file that is a symptom of a rogue host operating within the private network?
 - A. Packets from an unassigned internal address
 - B. Packets to an unknown port on an internal host
 - C. Packets in a serial grouping, attempting to access a series of ports
 - D. Packets that are all exactly the same, directed toward a single target
3. Which security strategy is based on locking the environment down so users can perform their assigned tasks, but little else?
 - A. Creating chokepoints
 - B. Diversity of defense
 - C. Principle of least privilege
 - D. Separation of duties
4. Which of the following statements is true regarding a reverse proxy?
 - A. The reverse proxy server can act as the endpoint for a TLS tunnel.
 - B. A reverse proxy cannot be used in conjunction with secured websites.
 - C. A reverse proxy can be used with tunnel mode IPSec VPNs.
 - D. A reverse proxy cannot support simultaneous SSL tunnels.

5. Which of the following is the most important feature of a bastion host OS?
 - A. Leveraging existing OS administrative knowledge
 - B. Inherent ease of use
 - C. Resistance to attacks and compromise attempts
 - D. Unlimited remote administration
6. When considering deployment of an IDS or IPS, what is the biggest problem?
 - A. Failing anomaly detection
 - B. False positives
 - C. False negatives
 - D. Failing to operate at wire speed
7. Which of the following is a highly recommended method or technique for keeping firewall logs secure and uncorrupted?
 - A. Storing the logs in binary form
 - B. Using high-capacity hard drives
 - C. Using timestamps
 - D. Using WORM drives
8. Which standard allows a firewall to hand off authentication to a dedicated service hosted on a different system?
 - A. IEEE 802.1x
 - B. IEEE 802.3
 - C. IEEE 802.11
 - D. IEEE x801
9. Which security stance focuses on the use of firewalls as its primary means of controlling communications?
 - A. Chokepoint
 - B. Universal participation
 - C. Fail-safe
 - D. Weakest link

- 10.** Which term describes the deployment of multiple subnets in a series to separate private resources from public?
- A. Diverse
 - B. Separation
 - C. N-tier
 - D. DMZ



© Digital_Art/Shutterstock

Configuring Firewalls

CHAPTER 8

A FIREWALL IS A KEY COMPONENT of a complete security infrastructure. Although it is common, a firewall is not a simple security measure. You must understand the complexities of a firewall to design a security policy, build the firewall, and configure it properly.

This chapter discusses the foundations of firewall rules, authentication, and authorization. You will also learn about firewall limitations, how to improve firewall performance, how encryption relates to firewalls, how to evaluate various firewall enhancements, and how to think about security concerns related to firewall management interfaces.

CHAPTER 8 TOPICS

This chapter covers the following topics and concepts:

- What constitutes good firewall rules

- How to compose a firewall policy defining what to allow and what to block
- What guidelines to follow for ordering firewall rules
- How to know what to allow and what to block when using a firewall
- Essential elements of a firewall policy
- Limitations of firewalls
- How to improve firewall performance
- The downside of encryption with firewalls
- Firewall enhancements
- How to use firewall management interfaces

CHAPTER 8 GOALS

When you complete this chapter, you will be able to:

- Construct examples of common firewall rules
- Explain the importance of firewall rule ordering
- Describe a properly written firewall policy
- Evaluate the limitations and weaknesses of firewalls
- Describe methods to manage firewall performance
- Define encryption concerns related to firewalls
- Evaluate the benefits and drawbacks of firewall enhancements

- Demonstrate how to access and use firewall management interfaces

FIREWALL RULES

Firewalls filter traffic using rules or filters. All firewalls, of whatever type—static packet filtering, stateful inspection, application proxy, circuit proxy, or content filtering—use rules to filter traffic. Rules are instructions that evaluate and take action on traffic traversing the network.

Firewall rules are used to control what traffic enters or leaves a secured network area. Depending on where the firewall is positioned, this mechanism can protect the private network from the public Internet or filter traffic between internal subnets or departments. The security administrator or the dedicated firewall administrator configures firewall rules in accordance with the organization's security policy.

Generally, two main philosophies or security stances govern the use of rules: default deny or default allow. Default deny is also known as *deny by default* or *deny all*. Default allow is also known as *allow by default* or *allow all*. These two rules define the foundation for rules governing traffic crossing the firewall.

A default-deny stance assumes that all traffic is potentially malicious—or at least unwanted or unauthorized. Thus, everything is prohibited by default. Then, as benign, desired, and authorized traffic is identified, an exception rule grants it access to the network. This is known as *deny by default/allow by exception*. This method of rule creation allows security administrators to focus on what is wanted or needed, rather than watching out for all possible forms of unwanted

or malicious activity. This is the whitelist-like action in a firewall whereby only those exceptions you define are allowed through.

A default-allow stance assumes that most traffic is benign. Thus, everything is allowed by default. Then, as malicious, unwanted, or unauthorized traffic is identified, an exception rule blocks it. This is known as *allow by default/deny by exception*. This method of rule creation forces the security administrator to be on constant watch for new and different forms of unwanted, unauthorized, and malicious activity. This is the blacklist-like action in a firewall, whereby everything is permitted with a few noted exceptions.

Most security experts agree that deny by default/allow by exception is the more secure stance to adopt. This stance automatically prevents most malicious communications by default, while the opposite allow-by-default stance allows most malicious communications by default. In most situations, fewer exceptions are required for a default-deny solution than for a default-allow system.

Most firewalls, both hardware and software varieties, usually come preconfigured with rules above and beyond the default rule. Firewalls are normally factory-configured in a deny-by-default stance with some common exception rules thrown in for end user/customer convenience (**FIGURE 8-1**). Often, these rules allow the most common forms of communication to occur across a new firewall without requiring you to configure the firewall fully upon initial installation. Some of the more common factory default rules allow for web, email, instant messaging, and file transfer through common Internet services on default ports.



FIGURE 8-1 pfSense’s default rule configuration interface, showing a default deny as the last rule.
© pfSense

Although this can be a convenience, it is never in your best interest to rely upon a third party's assumptions about your environment or guesses as to what types of traffic you do or do not want to cross your firewalled boundaries. Always take the time to review any factory-installed rules before deploying the firewall in the production environment. Delete or disable any rule that you do not need or want. With regard to any rules that you do want or need, you should double-check to make sure the rules are in line with your security policy. If you need rules that are not present by default, add each carefully and double-check your work as you go.

Your organization will need to decide which firewall rules to define; this is an essential part of its security policy. If the appropriate sections of the organization's security policy related to firewalls do not predefine which rules to implement on a new firewall, perform the following procedure:

1. Inventory all essential business processes and communications that will cross the firewall/checkpoint.
2. Determine the protocols, ports, and Internet Protocol (IP) addresses of valid traffic for both internal and external hosts.
3. Write out the rules or use a firewall rule designer/simulator.
4. Test the rules in a safe, nonproduction environment.
5. Obtain written approval for the rule sets from a change approval board or administration.
6. Document the rules into a security policy procedure amendment and submit the amendment to the security policy management team or administrator for inclusion in the official document.

Ultimately, this is the basic process for creating any new element of security. The goal is to have a written security policy for every security component. If there is no current policy or procedure

defining the steps to take for the deployment of a new security element, then you must write, test, and get approval for a new policy or procedure. Once a procedure exists, use it to judge successful deployment.

The exact rules to add to a new firewall depend upon the business processes that are unique to every organization. However, some common types of rules are found on most firewalls. These include:

- Access to insecure Internet websites (HTTP)
- Access to secure Internet websites (HTTP over SSL or TLS)
- Access to other Internet website protocols (SQL, Java, and so on)
- Inbound Internet email
- Outbound Internet email

If other Internet services are essential to a business task, rules allowing or enabling such communications would be needed. Try to keep the number of rules to a minimum. Grant access only to traffic that is essential. Each rule added to a firewall increases its attack surface—and therefore its security vulnerability.

Inbound and Outbound Communications

Keep in mind that most network communications are two-way transactions, exchanges, or sessions. Often, an internal client makes an initiation request to an external server. The server responds and sets up a communication channel, often called a virtual circuit (at the Transport Layer based on TCP) or, more generically, a session. Traffic can traverse the session in either direction. Thus, most rules should allow inbound responses to initial outbound requests.

Depending upon the firewall, a single rule can sometimes define both the outbound and inbound communication parameters. Other firewalls need two separate rules: one rule for the initial outbound request and a second rule to handle the resultant inbound response. In either case, each desired session-based communication service must be supported or allowed through proper two-way rule construction.

In addition to rules that allow traffic based on internal user initiation, consider rules that manage inbound communications for externally initiated communications. When a firewall is to allow an external host to request access to internal resources, you need to create an inbound rule or ingress filter. Define rules that allow inbound initiations only if external entities need to access services and resources hosted internally.

Otherwise, allow the firewall to serve its primary purpose and prevent outsiders from gaining easy access to internal systems. Most firewalls deny by default all access to internal resources from external entities. This is done either by ignoring all requests received on a specific interface (such as on dual-homed or triple-homed firewalls) or by having a specific rule (or a set of rules) that filter out all traffic originating from noninternal IP addresses.

Access Control Lists

A firewall rule can also be called a filter or an **access control list (ACL)**. These terms are often used interchangeably in documentation, in books, and in the firewall's configuration or **management interfaces**. Don't be confused by this. A rule is a written expression of an item of concern (protocol, port, service, application, user, IP address) and one or more actions to take when the item of concern appears in traffic. A filter is the same thing as a rule, but the point or purpose of using a filter is to emphasize the intention to block or deny unwanted items of concern. Thus, the use of the term ACL emphasizes the intention to grant or deny traffic on an access control/authentication basis. An ACL focuses on controlling a specific user or client's access to a protocol or port.

Rules allow traffic to pass unhindered or block or deny traffic to prevent it from reaching its intended destination. A rule can also include a logging element, so the traffic is logged whether allowed or denied.

TECHNICAL TIP

Use your firewall to ensure that only properly originated communications are authorized. Unless you are running internal services (such as a web server), external entities should never be able to initiate a connection.

TECHNICAL TIP

Most session-based communications will use a well-known default port for the destination port where the service or resource resides. These often use a randomly selected higher order port (any port above 1023) for the client source port. Most firewalls, especially those capable of stateful inspection, will automatically handle and adjust for the random source port when establishing a session. If this is not the case with a specific firewall, you may need to create a custom rule or consider replacing the firewall with a better product.

COMPOSING FIREWALL RULES

Setting up or defining rules for a home, portable host, or SOHO (small office/home office) environment can be fairly easy if the number of different types of access is minimal. But, as the number of controls, limitations, restrictions, and exceptions increases, it can become more complex to define rules properly. Obviously, larger networks with more advanced infrastructures and communications require more intricate rule sets.

The basics of defining or crafting rules are the same—or at least similar—across all firewall products. However, some firewalls support editing or writing the rules directly, while others employ a graphical interface or use a design wizard to accomplish this task. Software host firewalls, especially those designed for end users, typically use a graphical wizard-based system, while hardware firewalls typically expose the raw rule itself.

Most rules have six main elements or components (**FIGURE 8-2**):

Edit Firewall Rule

Action

Pass ▼

Choose what to do with packets that match the criteria specified below.

Hint: the difference between block and reject is that with reject, a packet (TCP RST or ICMP port unreachable for UDP) is returned to the sender, whereas with block the packet is dropped silently. In either case, the original packet is discarded.

Disabled☐ Disable this rule

Set this option to disable this rule without removing it from the list.

Interface

WAN ▼

Choose the interface from which packets must come to match this rule.

Address

IPv4 ▼

Family

Select the Internet Protocol version this rule applies to.

Protocol

TCP ▼

Choose which IP protocol this rule should match.

FIGURE 8-2 pfSense's rule configuration interface.
© pfSense

1. **Base protocol**—Set as TCP; UDP; ICMP; and potentially other Layer 2, 3, or 4 protocols.
2. **Source Address**—Set as a specific IP address, a subnet, a range of addresses, or ANY for all possible addresses.
3. **Source port**—Set as a specific port, a set of ports, a range of ports, ports less or greater than, or ANY for all possible ports.
4. **Target address**—Set as a specific IP address, a subnet, a range of addresses, or ANY for all possible addresses.
5. **Target port**—Set as a specific port, a set of ports, a range of ports, or ANY for all possible ports.
6. **Action**—The two standard actions are Allow and Deny. Some firewalls include Log and Alert actions that allow additional actions in a single rule.

FYI

Some firewall rule systems include application protocol designations in the rule set. This is usually a human-friendly naming convention pointing out the Application Layer protocol rather than an element of the filter used to select traffic. However, some firewall interfaces will substitute a common protocol name for the destination port. For example, port 80 would be displayed as HTTP and port 25 as SMTP. However, these are usually for human viewing; the port number is still used in the actual protocol-level filtering activity.

Depending upon the firewall, TCP and UDP rules (as well as ICMP; IGMP; ARP; and other OSI Layer 2, 3, or 4 protocols) may be defined in the same place or within different interfaces, pages, or rule sets. When all rules are defined in the same place, use an additional rule element designating TCP, UDP, ICMP, and so on. All of the following examples of rules assume TCP-based communications, unless otherwise noted.

The common rule structure is:

[Protocol]	[Source Address]	[Source Port]	[Target Address]	[Target Port]	[Action]

When defining outbound rules, the source address and port are often set as ANY unless the rule is to apply to a specific system(s) or port(s). For example, a rule that allows any internal client on the 192.168.42.0/24 subnet to access any insecure Internet website would be:

TCP	192.168.42.0/24	ANY	ANY	80	Allow

A rule to allow access to any secured Internet website would change the destination port 80 to 443:

TCP	192.168.42.0/24	ANY	ANY	443	Allow

A rule to block internal clients from accessing Internet-based FTP sites would be:

TCP	192.168.42.0/24	ANY	ANY	21	Deny

FYI

Keep in mind that firewall vendors may have crafted unique management interfaces that display firewall rules differently from the basic line form demonstrated here. For example, some firewalls have a configuration page for Allow rules and a separate page for Deny rules. In that example, the rules on each page would not include an Allow or Deny action, as the page where each is defined implies the action.

When defining inbound rules, the source address and port address are often ANY, unless the rule is to apply to a specific system(s) or port(s). For example, a rule that allows any external host client to access an insecure website hosted on the 192.168.42.98 would be:

TCP	ANY	ANY	192.168.42.98	80	Allow

An inbound rule to deny external hosts access to a telnet server hosted on the same internal system as the web server would be:

TCP	ANY	ANY	192.168.42.98	23	Deny

Generally, inbound rules are needed only when an internal resource is specifically hosted for the purposes of external access. This type of rule is commonly found on demilitarized zone (DMZ) and extranet firewalls, but should not be used on intranet border firewalls. The only common inbound rule that might be required is one that allows responses to previous internal client requests to resources outside of the network. A stateful inspection firewall usually allows response traffic automatically, but static filtering firewalls might require a specific respond rule set, such as:

TCP	ANY	ANY	192.168.42.0/24	>1023	Allow

When a UDP rule set needs to be defined, you must either switch to a different rule set interface on the firewall, or change the protocol designation from TCP to UDP. Assuming a firewall interface with a leading protocol designation for each rule, a rule set allowing user queries with an internal DNS server at 192.168.42.104 by external entities, but blocking attempts to perform DNS zone transfers, could be:

TCP	ANY	ANY	192.168.42.104	53	Deny
UDP	ANY	ANY	192.168.42.104	53	Allow

To write a firewall rule controlling ICMP, change the protocol designation to ICMP—or change interfaces—and define an ICMP type. The port designations are dropped completely, as these are not

used by ICMP at OSI Layer 2. Multiple ICMP rules can be defined for various subtypes to allow some ICMP traffic and a deny-all catchall rule for everything else. For example, to block only inbound ICMP Echo Request packets, use the following filter:

ICMP	ANY	192.168.42.0/24	Type 8	Deny	Allow

■ NOTE

Internet Control Message Protocol (ICMP) is a communication protocol used in most networks to deliver error messages. Responding with these messages is a security problem because it enables hackers to learn about the version of software and patching levels within the network.

TECHNICAL TIP

The ICMP protocol's header defines the purpose of the protocol's control message through the use of numbered types. There are dozens of defined types. The most commonly used are Type 8 Echo Request, Type 0 Echo Reply, Type 3 Destination Unreachable, and Type 11 Time Exceeded. Many types are further specified with the use of a numbered code. For example, Type 3 Code 3 is destination port unreachable.

For a more exhaustive presentation of the ICMP types and codes, read RFC 2939 at

<https://www.iana.org/assignments/icmp-parameters/icmp-parameters.xhtml>.

When defining firewall rules, here are a few basic guidelines:

- Keep the rule set as simple as possible.
- Document every rule.
- Use a change control mechanism to track rule modifications.
- Always confirm the default deny before using changed/updated rule sets.

A simple rule set is a set with as few rules as possible. Fewer rules mean fewer complications. Fewer rules mean less chance for a loophole. Fewer rules are easier to test. Fewer rules are harder to attack and compromise.

Every rule enforced by a firewall should be written in the security policy. Only rules in the written security policy should be enforced by a firewall. Whenever you need to change a firewall rule, first update the security policy, and then reconfigure the firewall to match the settings defined in the security policy. Always fully document the rules actively filtering traffic.

Do not just document that a rule exists and what the basic rule structure is; also include a description of the intention or purpose of the rule. Sometimes, a rule written in a way that appears straightforward may not achieve the intended purpose. When other security administrators review the documentation and examine the rules, they may see the discrepancy and correct it. Get in the habit of

making notes about intentions, thoughts, and concerns. Often, these stray comments will lead to solutions or assist in future troubleshooting.

A change control mechanism tracks and monitors the changes to a system. It does this by recording every change, modification, or adjustment. Often, a change management solution requires manual change logging; thus, it must be enforced and followed by a written component of the security policy. The change control process produces historical documentation of the state of security components. Change documentation is often essential in troubleshooting and repair. Often, it has been the last change made to a system or rule set that has caused it to break or malfunction. Knowing exactly what was done—and why—makes it faster and easier to get the system back to its last secure state and allow for correct programming of the rule.

Most firewalls operate on a default-deny stance, but it may be possible to modify this stance. Stance modification may be a general setting in the overall configuration of the firewall, or it might just be the last rule in the rule set. If it is just a rule, be sure that the final rule is always default deny and that the default-deny rule is never edited or modified. A default-deny rule should look like or enforce filtering as follows:

TCP	ANY	ANY	ANY	ANY	Deny

Because a firewall is a first-match-apply rule-based system, rules are ordered on the firewall. Traffic is compared with each rule, starting from the top rule. If a rule matches the traffic, the rule's

action is applied. Once a matching rule's action is applied to the traffic, no further rule matching is attempted. If the traffic does not match any rule except the final rule, then it will be denied by default. All rules previous to the default deny can be either exceptional Allow actions or specific Deny actions. No matter what the structure of the rules or the number of rules in the set, always ensure that there is a default rule. As rule sets get more complicated, and with the potential for a rule conflict, the default rule helps set precedence (from top down).

Here is an example of a basic TCP rule set:

TCP	ANY	ANY	192.168.42.0/24	>1023	Allow
TCP	192.168.42.1	ANY	ANY	ANY	Deny
TCP	ANY	ANY	192.168.42.1	ANY	Deny
TCP	192.168.42.0/24	ANY	ANY	ANY	Allow
TCP	ANY	ANY	192.168.42.55	25	Allow
TCP	ANY	ANY	192.168.42.98	80	Allow
TCP	ANY	ANY	ANY	ANY	Deny

All of the firewall rule examples are using the default common port for common services as defined by [***https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml***](https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml).

This collection of seven rules performs the following filtering functions:

1. Allow response to TCP connections to internal hosts.
2. Prevent the firewall (192.168.42.1) from directly connecting to anything.
3. Prevent external hosts from directly accessing the firewall.
4. Allow internal hosts to access external resources.
5. Allow external hosts to send email inbound to the email server at 192.168.42.55.
6. Allow external hosts to access an internal web server at 192.168.42.98.
7. Apply a default-deny rule to all traffic not matching a previous exception.

ORDERING FIREWALL RULES

Building a rule set is not difficult, but it is necessary to focus on the details, as some rules can be complex. The most important aspect of a firewall rule set is its order. Getting rules out of order causes unexpected and unwanted consequences. This can include traffic you want to block and other unwanted traffic crossing the checkpoint. Rule-set ordering is critical to the successful operation of firewall security.

When the wrong rule is positioned first, this creates a potential loophole. A **loophole** is a flaw in the logic of filtering that will allow an unwanted action to occur. A firewall can perform only the operations for which it is programmed, and the specifics of and the order of the rules are a form of programming. Take the time to evaluate, test, and verify that you have defined all firewall rules correctly and placed each in the best order.

The first and most basic rule-set-ordering convention is that the universal Deny rule should be the last and final rule. The use of deny by default or default denial rests on the premise that the last rule is the catchall rule to block all traffic not allowed access due to a previous rule-based exception.

Another common guideline to rule-set ordering is to place critical Deny exceptions first or early in the rule set. When specific internal or external IP addresses or ports, or even entire protocols, are to be absolutely blocked, you may need a Deny exception rather than relying upon the default-deny final rule. Some of the previous Allow

exceptions might inadvertently permit communications due to universal application (with the use of ANY). By using a preemptive specific enforced denial before any of the Allow exceptions, you eliminate the possibility of accidentally allowing a known malicious or unwanted communication.

Whenever possible, use fewer rules rather than more rules. Even with proper ordering, the more rules you have, the greater the likelihood of configuring something incorrectly or creating a loophole. One issue that causes more rules rather than fewer is infrastructure design specifically related to addressing. A need for more rules arises if a range of IP addresses is allowed access, but within that range, some addresses are refused access. For example, compare two scenarios.

First, a network has a host address range of 192.168.42.140–190. All hosts except for 188, 189, and 190 are allowed access to a certain port. A single rule allowing hosts 140–187 is all that is necessary because the default-deny rule takes care of blocking the remaining nonincluded hosts.

Second, a network has a host address range of 192.168.42.140–190. All hosts except for 165, 171, and 188 are allowed access to a certain port. You need multiple rules to use this configuration. One or more rules must define Deny exceptions for 165, 171, and 188, followed by the Allow rule of the 140–190 range. If the firewall allows only a single address or a range of addresses per rule rather than allowing a list of nonsequential addresses, then three Deny rules would be necessary in this scenario.

In this example, network design and addressing can be used to make firewall rule-set construction either larger and more complex or

shorter and more distinct and compact. The latter is preferred for administrative purposes, as well as security and efficiency. If the process of creating rules requires a significant number of special exceptions to modify or adjust ranges of addresses or ports, consider reconfiguring the network rather than using a too complex or too long rule set. When designing or writing firewall rules, especially when writing pairs or sets of rules, consider using a single rule or a simpler rule set if the network's addressing scheme, infrastructure design, or subnet layout is adjusted.

As another guideline to ordering rule sets, consider placing rules related to more common traffic earlier in the set rather than later. Comparing traffic to the rule sets takes time; each check of each rule takes some finite amount of time. The fewer rules it needs to check before granting an Allow, the less delay to the traffic stream. Prioritize in the rule-set list the more commonly used forms of traffic, whether by IP address, port, or protocol. Put the less commonly used forms of traffic further down in the rule-set list.

As rule sets get larger, each becomes more complex. Often, the complexity stems from having explicit Allow rules with additional Deny specifications. This results in rules that overlap. Overlapping is acceptable when you understand it and use it on purpose. When overlapping occurs accidentally, it can result in undesired loopholes. Again, the solution is to keep the rule set as simple as possible, document every rule with the intentions, and test the rule thoroughly before deployment.

Ultimately, rule sets are about enforcing security relevant to the organization. The rule set should reflect the guidelines prescribed in your written security policy, specifically the firewall policy. The goal

of designing, writing, and ordering rules for a firewall should be to focus on obtaining the necessary security. Elegance and speed are dividends, but not as essential as blocking the bad and allowing the good traffic. Never lose focus on the primary goal: filtering traffic in accordance with your security policy.

WHAT SHOULD YOU ALLOW AND WHAT SHOULD YOU BLOCK?

The most commonly asked question when installing a firewall is, “What should be allowed and what should be blocked?” This question assumes a distinct and definitive answer. The answer, however, is subjective and variable. A single security stance or filtering configuration that works for every situation, every organization, and every system doesn’t exist. Network security managers must investigate the needs and threats to the business and make informed decisions about what traffic to allow and what traffic to block.

The first stage in making this determination is to perform a complete inventory of all needed or desired communications. This should include every transaction between internal systems, as well as any interaction with external systems. As part of this inventory, indicate the protocol in use, the port(s) in use, and the likely source and destination addresses.

A possible inventory of network communications could look like **TABLE 8-1.**

**T
A
B
L
E**

**8
-
1
A**

P
a
r
t
i
a
l
L
i
s
t
o
f
C
o
m
m
u
n
i
c
a
t
i
o
n

S
O
C
C
U
R
I
N
G
O
N
A
N
E
T
W
O
R
K

TRANSPORT LAYER

PROTOCOL	PORT	PROTOCOL	SOURCE	DESTINATION
TCP	8080	HTTP	Internal Clients	192.168.42.101 (Internal website)
TCP	8081	HTTPS	Internal Clients	192.168.42.101 (Internal website)
TCP	80	HTTP	Internal Clients	Numerous Internet websites
TCP	443	HTTPS	Internal Clients	Numerous Internet websites
TCP	110	POP	Internal Clients	192.168.42.115 (Internal email server)
TCP	25	SMTP	Internal Clients	192.168.42.115 (Internal email server)
TCP	25	SMTP	192.168.42.115	External email server (ISP based)
TCP	22	SSH	Internal Clients	Numerous internal servers

This is not an exhaustive table, but it shows the basic idea of an inventory of network communications. Notice that some communications are internal only, while others cross the network boundary. Based on a complete and exhaustive inventory of your network's communications, you can establish a potential rule base.

From the inventory, block from crossing the network border all communications that should only be internal. In most cases, the default-deny rule of a rule set will block those communications from exiting the network. However, when a port or protocol is in common, it is possible that an Allow rule for a valid external communication would also allow external entities access to an internal-only communication.

In Table 8-1, an example of this issue is the email protocol Simple Mail Transfer Protocol (SMTP) operating over port 25. By defining an allow-exception rule, you could create a loophole that threatens the internal communication on this same port. Reduce or eliminate this risk by using a very specific rule or by adding an additional Deny exception. A very specific rule in this example would allow external communication over port 25 only if it originated from the internal email server at 192.168.42.115. With this rule defined, if any other internal host attempted an external SMTP communication on port 25, it would not be allowed because the IP address would not match this specific rule. The communication would instead be blocked by the default deny.

If an allow-exception rule for external communications also applies to internal communications, you can add a specific deny-exception rule just prior to the Allow rule. The Deny exception can specifically block internal communications from crossing the network

border, but still allow the subsequent Allow exception to grant access to the valid transactions.

After you complete a thorough inventory of network communications, you should determine which communications are mission-critical, important, optional, for personal recreational use, or actually malicious. Block any traffic you deem malicious or just unwanted, whether through an explicit Deny or via the catchall default deny.

Evaluate all other forms or types of traffic against several factors, including policy. If it is not in support of the business, is it really needed? Are communications for personal entertainment or nonbusiness functions permitted by policy or considered a waste of resources? Are any communications duplicated or redundant? Is this a waste of resources, an expansion of the attack surface, a valid redundancy, or avoidance of a single point of failure? Are any of the communications no longer necessary? If so, should you disable or remove that communication type?

TRAFFIC INVENTORY

A common—and free—network protocol analyzer (or sniffer) for conducting a traffic inventory is Wireshark (www.wireshark.org). Wireshark can be used in the absence of a firewall, with a firewall set to allow all traffic, or in the presence of a firewall to inventory all traffic on the network. Traffic can then be partitioned into that which needs to cross the firewall and traffic that should not cross the firewall. Use of a sniffer tool in

general, or Wireshark specifically, for traffic inventory is done to highlight all traffic and make sure that the firewall rules are accurate. During the traffic inventory, you can uncover a lot of previously hidden information about the network. You can learn which applications and users are the top consumers of bandwidth. You can find out which protocols are actually in use on the network. You can identify malicious uses that were previously unknown. Some of the more advanced firewalls also have this functionality built-in.

Base your evaluation of the many other possible questions or issues on your organization's policy statements, including mission and goals. Take the time to understand the communication needs of your network, then design the filtering rule set accordingly. Remember that the firewall is often the front sentry or the first line of defense against malicious communications originating from outside networks. Leaning toward caution and locking things down is a more effective security stance than leaving communications open and hoping detection mechanisms will notice and identify malicious events.

Although some exceptions to these suggestions will arise from your organization's business model or strategies, the following items are commonly considered communications to block in most, if not all, circumstances:

- All Internet Control Message Protocol (ICMP) traffic originating from the Internet
- Any traffic directed specifically to the firewall

- Any traffic to known closed ports
- Any traffic to ports of known malware, such as 31337, used by Back Orifice
- Inbound Transmission Control Protocol (TCP) 53 to block external Domain Name System (DNS) zone transfer requests
- Inbound User Datagram Protocol (UDP) 53 to block external DNS user queries
- Any traffic from IP addresses on a blacklist
- Any traffic from internal IP addresses that are not assigned

As an organization manages its security over time, especially as it reacts to intrusion attempts, experiences scanning, and becomes the target of hacking and flooding, it must create additional communication filtering rules to respond to real-world conditions. Hackers are constantly evolving their skills and techniques. Thus, the filtering rules that work well today may be insufficient in the near future.

Networks can be very large and very complex. Be fully aware of the communications taking place across your network to know what to allow and what to block. In addition to this knowledge, having established a security strategy through a written security policy will direct and guide the deployment of firewall rules, as well as other aspects of your security infrastructure.

ESSENTIAL ELEMENTS OF A FIREWALL POLICY

A firewall policy is a security policy that focuses on the deployment of firewalls within the organization's IT infrastructure. The first step in deploying a firewall is constructing a firewall policy. Once you have established a firewall policy, deploy subsequently installed firewalls in full compliance with the policy.

A written firewall policy provides several benefits and serves several purposes, such as a:

- Guide for installation
- Guide for configuration
- Tool to assist in troubleshooting
- Guideline to detect changes and differences
- Mechanism to ensure consistent filtering across all firewalls

A firewall policy should address several specific issues and contain specific configuration details. One of the first important elements of a firewall policy is defining and designating security zones. Every network has different zones of risk and zones of trust. Often, zones are differentiated along the same lines as a typical IT infrastructure. Create clear descriptions of what each subnet does, its level of risk, and its level of trust. Based on this information, you can formulate a basic firewall deployment strategy.

A firewall policy should define specifically what type of firewall you need. It may even prescribe the specific vendor, make, and model

of firewalls at each zone transition or interface. Firewall types include static packet filtering, proxy, and stateful inspection.

For each prescribed firewall, define a complete firewall rule set. Define each rule in full detail, not just the configuration settings to implement on the firewall's interface. Support each rule with justifications and reasons why you defined or selected it. Give a clear indication of the order of the rules.

A firewall policy should also prescribe host software firewalls for deployment on clients and servers. This should include configuration settings and rule sets, as well. Design and configure host firewalls to complement appliance firewall security.

The firewall policy should also address:

- What to log
- Logging location and redundancies
- Where to store log files
- How often to review log files
- What add-ons or enhancements to use
- Who is responsible for firewall administration
- How to access firewall configuration interfaces
- Where to physically and logically locate the firewall
- What level of physical access control is necessary
- What form of backup or redundancy is present
- How to manage encryption and where to use or disallow it
- How to deploy and configure IDS/IPS to interact with firewalls

As with every security policy, the firewall policy must be as exhaustive as possible. It should address and define every aspect of

firewall design, deployment, implementation, management, tuning, repairing, recovery, troubleshooting, and monitoring. A firewall policy should be the first and last authority on all things firewall within the organization.

The first draft of a firewall policy will probably not be exhaustive or accurate. Review and improve your firewall policy periodically, just as you do with the rest of your organization's security policy. Each review period should assess whether the firewall is meeting the security needs and how to improve it so it continues to provide sufficient security. Work all improvements into the written security policy. Then adjust the deployed infrastructure to comply with the written security policy.

LIMITATIONS OF FIREWALLS

Firewalls are essential components of security. Systems ranging from standalone home systems to global IT infrastructures all need firewalls as part of the overall security solution. However, firewalls are not perfect solutions. In fact, when it comes to firewalls, you need to account for several well-known limitations in design and management.

Firewalls are ultimately software code written by humans. In either a host firewall or an appliance firewall, the logic and controlling mechanisms of the firewall are software. Software is simply code that is designed and written by people. Whenever people are involved, the possibility exists that they will make mistakes or oversights. Fortunately, most firewall products are rigorously and thoroughly tested as a counterbalance to this threat. But do not assume that the firewall software is inherently perfect.

While it doesn't happen often, administrators and hackers have discovered software coding bugs or flaws in real-time use. Hackers are constantly using scanning, testing, and probing tools to discover exploitable weaknesses. Once a vulnerability appears, a hacker will create an exploit that takes advantage of the flaw. Some exploits have caused firewalls to freeze or crash, while others have given hackers the ability to read or adjust filtering rules. The initial discovering of such a flaw is termed a zero-day exploit.

Some exploits are based on buffer overflows. A buffer overflow is a condition whereby a memory buffer exceeds its capacity and the

extra content “overflows” into adjacent memory. Hackers often use this as an attack against poor programming techniques or poor software quality control. Hackers can inject more data into a memory buffer than it can hold, which may result in the additional data overflowing into the next area of memory. If the overflow extends to the next memory segment designated for code execution, a skilled attacker can insert code that will execute with the same privileges as the current program. Improperly formatted overflow data may also result in a system crash.

■ NOTE

Most patching documentation released by software vendors notes the purpose of the patch and the issue it fixes. These release notes read like a “how-to” for a hacker to learn of existing vulnerabilities in a system without the specific patch applied. Do not delay in testing and implementing patches.



Once the firewall vendors become aware of an exploit in the field, they quickly develop and release a patch. By having an effective security management and patch management system, you can quickly test and install these new updates for the firewalls protecting your production environment. If the release of a patch or the application of a patch is delayed, the window of opportunity for hackers to compromise your environment remains open.

Firewalls are sometimes vulnerable to fragmentation attacks. Fragmentation attacks are an abuse of the fragmentation-offset

feature of IP packets. Fragmentation may occur in a network where many different network links join to construct a global infrastructure. Some network segments support smaller datagrams (another term for packet or frame) than others, so larger datagrams fragment into the smaller size. When the fragmented elements of the original datagram reassemble, reassembly of the fragments can cause several potentially malicious reconstructions, such as overlapping and an overrun.

Overlapping can cause full or partial overwriting of datagram components, creating new datagrams out of parts of previous datagrams (**FIGURE 8-3**). An overrun can cause excessively large datagram construction. Other fragmentation attacks are designed to cause DoS or attempt to confuse IDS detection or firewall filtering.

Normal Fragmentation



Overlapping Header and Data Sections




FIGURE 8-3 Fragmentation and overlapping.

Protections against fragmentation attacks include using modern IDS and firewall filtering features, as well as performing sender fragmentation. Sender fragmentation queries the network route to determine the smallest maximum transmission unit (MTU) or datagram size. Then the sender prefragments the data to ensure that no en-route fragmentation will occur.

Firewalking is another known firewall limitation. Firewalking is a technique to learn the configuration of a firewall from the outside. The technique uses a valid IP address of an internal host. Then, from an external system, a hacker attempts to establish a communication session with the internal host over a multitude of different ports. Effectively, this is a form of port scanning, but with the addition of the known internal host, the hacker can learn not only which ports are open, but also which ports actually allow communications with an internal system. Effectively, firewalking discovers the rules or filters on a basic packet-filtering firewall. Modern stateful inspection firewalls are usually not as vulnerable to firewalking due to the ability to detect this mechanistic activity.

Internal code planting is another known firewall limitation. Firewalls are often deployed as border sentries. These are intended to protect internal systems from communications that originate from external entities. Unfortunately, some security administrators use only inbound firewall filtering, leaving outbound traffic uncontrolled and unfiltered.

In this situation, if a hacker can plant code internally, trick a user into downloading or running code, or an employee could bring in code intentionally or accidentally in a file or download, the

possibility exists that outbound connections to malicious external entities could occur. Many hacker tools are based on this technique. A hacker creates a server of sorts that is hosted on an external host. The automatic connection client utility executes on an internal host and establishes the connection, and then the external host is able to send data or commands back to the internal client. Often, this form of attack results in a hacker gaining modest to complete remote control over the compromised internal host.

Denial of service (DoS) is another common problem that reveals a limitation of a firewall. A DoS attack, specifically a flooding- or traffic-based DoS, sends massive amounts of data to a target system. If that system has a firewall, the firewall can detect and discard the potential DoS traffic. The firewall's filtering service can usually prevent the DoS traffic from breaching the network's perimeter and affecting internal systems.

However, since the firewall has to collect, analyze, and respond to every packet received on its interfaces, a well-organized DoS can consume all available bandwidth of the connecting segment to the firewall, as well as consume all of the processing capabilities of the firewall. This, in turn, prevents any legitimate traffic from reaching the network. Thus, even with a firewall protecting the internal network, a DoS flooding attack can still successfully disconnect or interfere with external communications.

The vulnerability to DoS flooding is the one limitation or weakness of firewalls that you cannot fix, improve, or repair by either upgrading the firewall or applying a patch. Upgrading to a stateful inspection firewall addresses fragmentation, firewalking, and even internal release of code. Patching will address programming bugs and

buffer overflows. Upstream filtering is one possible countermeasure to prevent flooding attacks from reaching an Internet-facing firewall, but this technology requires cooperation from Internet service providers (ISPs) and other third parties.

Knowing these limitations, as well as recognizing the likelihood of other possible limitations, brings the security administrator back to basics. Security management is mandatory to maintain any semblance of security in any environment. Be familiar with authorized traffic and network usage. Keep systems current with patches, use a hardened configuration, stay knowledgeable about new exploits, and monitor the environment for successful and attempted compromise. These are the essential, long-term strategies for maintaining security.

IMPROVING PERFORMANCE

Using a firewall is mainly about establishing and maintaining border security, but security managers must also look into other areas of importance and concern beyond just security. One of these additional nonsecurity areas is performance. Because the goal of security is to protect the production environment, installing security measures that severely hinder capabilities necessary to accomplish work is counterproductive.

Network communication speed is one of the most common concerns with the use of firewalls. When selected and installed properly, a firewall should function at wire speed. Functioning at wire speed means the firewall does not introduce any delay or latency in communications because it operates at the same speed as the network. If there is a 1,000 megabits per second (Mbps) network, then a 100 Mbps-capable firewall is too slow.

Two primary criteria can help you to improve firewall performance above and beyond purchasing a firewall suitable for the network's known transmission speed. The two criteria are caching and load balancing. By considering these criteria in the network design and business throughput planning, security can be implemented.

Caching is the holding of often-accessed content in storage or memory on the firewall. Then, when a future request for the same content is received, the cached copy goes into service rather than the original external source. This does require that the content be

relatively static and that a staleness value be used to trigger a content refresh. Caching is a technique to seemingly improve performance by borrowing from proxy servers. In fact, adding caching to a firewall effectively transforms it into a proxy server for whatever service you configure the caching to supplement.

Unfortunately, caching is beneficial only on two main services: web and file transfer. If the potential or actual bottlenecks to performance across a firewall are not due to these two types of communication, then caching is not a good solution. For many organizations, the main external communication is web transfer (employees viewing content on the web), which fits the caching scenario.

Load balancing is the distribution of the firewall filtering workload across multiple parallel firewalls. You can use this as shown in **FIGURE 8-4** with a phalanx of two or more firewalls deployed in parallel between routers that perform load balancing traffic management on both inbound and outbound communications. Load balancing allows for deeper packet inspection while maintaining wire speed performance using parallel-distributed firewall processing.




FIGURE 8-4 Firewalls deployed in a load-balancing configuration.

TECHNICAL TIP

Router load balancing traffic management can use a variety of scheduling algorithms, such as **round-robin** or **fair queuing**. Round-robin simply hands out tasks in a nonpriority sequence. For example, with three firewalls and five transactions, transaction one is sent to firewall one, transaction two to firewall two, transaction three to

firewall three, transaction four to firewall one, and transaction five to firewall two. Fair queuing operates by sending the next transaction to the firewall with the least current workload.

The use of load balancing often has additional benefits, such as redundancy and fault tolerance. Both of these improve the availability of the firewall-filtering service for network communications. Load balancing uses a device or software to better distribute traffic or tasks among member devices or network paths. If the load balancing is among devices, traffic would be queued to go to the first device, then the next communication stream to the next device, and so on. If the load balancing is network traffic, it would be a combination of software and hardware that responds to traffic it sends to learn of the time it takes packets to reach destinations at various locations on the network.

Two main classifications of load balancing are based on the Open Systems Interconnection (OSI) Reference Model Network Layer where the balancing occurs. Layer 4 (Transport Layer) directs the traffic encountered based on Network Layer and Transport Layer protocols. Layer 7 (Application Layer) determines routes based on application and header information.

High availability is a term often associated with load balancing and network availability. This measure is the percentage of time when a service or device is available to users on the network. A device may be available 95 percent of the time, and down the remaining time only to allow for hardware updates and patches. Load balancing aids in user satisfaction, as the traffic that is distributed to its destination

faster may be responded to and any response sent to the user more quickly.

Having a fail-safe or fail-secure plan to disconnect communications if the firewall is compromised is essential. However, a more secure and productivity-friendly solution is to maintain continued communication with filtering by having redundant pathways with duplicate firewalls.

THE DOWNSIDE OF ENCRYPTION WITH FIREWALLS

Encryption is a common security mechanism you can employ to protect data in storage and in transit. When used properly, encryption provides confidentiality protection. Encryption prevents unauthorized third parties from gaining access to secured content. With the risk of eavesdropping and data theft constantly increasing, the use of encryption is more than a good idea, it is an essential element of storage and transaction security.

However, encryption does have a downside relative to firewalls. A firewall is typically not the intended destination or direct communication partner of a communication, especially an encrypted communication. Thus, any encrypted data cannot be filtered by a firewall. However, just because packets are encrypted does not automatically mean that nothing about the packet remains in plaintext.

Transaction or communication encryption takes two main forms: tunnel mode and transport mode. Tunnel mode encrypts the entire original payload and header, while transport mode encrypts only the payload. In tunnel mode, a temporary header goes with the encrypted packet to guide its path across the VPN tunnel. In transport mode, the original header remains in plaintext.

A firewall can still view and filter on the contents of the tunnel mode header or the transport mode header. However, the tunnel mode header contains details only about the endpoints of the tunnel, not the

endpoints of the actual communication itself. So, tunnel mode header filtering is not very useful as a filter against malicious traffic.

Filtering on the transport mode header is a viable option, because the header in the packet in transport mode is not encrypted. Thus, any header-only filtering rules could still apply to transport mode—encrypted communications. However, any filtering that required an examination of the payload will be rendered null.

When building or designing firewall filtering rules, you must make a choice about how to handle encrypted content. Consider both the valid and invalid reasons for content encryption. The organization can choose to support or allow encryption of specific types over specific protocols or ports, but disallow and prevent encrypted communications elsewhere. Encryption for web communications and email exchanges is often acceptable, while other transactions with the Internet might not be encrypted.

When designing the firewall rules, the management of encrypted traffic can range from full allowance to full denial. Whether to allow encryption over a specific port but not another, and whether to allow encryption all the time, for only certain users, or for no one are common issues your organization needs to address and include in its security policy. No ultimate right or wrong decision about whether to allow encryption across a network border is possible. Decide based on the security stance of your organization, the risk presented by both encrypted and plaintext transactions, and the types of communications essential for your particular business tasks.

A growing trend in secure deployments, especially for DMZ or other hosted resources designated for public access, is to support encrypted communications but have the encrypted tunnel end at or on

a firewall. This allows you to examine the contents of the communications before the packets reach the actual destination. This is becoming a common infrastructure design component for public facing websites. **FIGURE 8-5** shows an example of a configuration where SSL or TLS links from external Internet clients end at the firewall, and then traverse the last segment in the clear to reach the destination web server.




FIGURE 8-5 A firewall deployed as an SSL/TLS endpoint to allow for web traffic filtering.

Another concern is that when encrypted traffic crosses a firewall, the firewall is usually unable to perform NAT. This is dependent on the encryption system in use, but it is worth investigating before deployment. If the encryption mechanism hashes the header, even if it remains in plaintext, the modification to the header by NAT will cause the hash verification process to fail, thus rendering the transmission void. With an encrypted header, NAT will be unable to modify the addresses. Be sure to select encryption protocols that are NAT-compatible if this is a necessary function of the network's infrastructure.

FIREWALL ENHANCEMENTS

Firewalls emerged from basic router traffic control systems into full-featured content filters. However, throughout this evolution, firewalls have remained focused on detecting unwanted, unknown, or malicious traffic and blocking it. The filtering mechanisms have typically been identification-based. Whether IP address, port, protocol service, MAC address, content keyword, or even user authentication, the Allow and Deny decisions relied directly on easily definable elements.

Modern firewalls now offer a variety of enhancements, improvements, or add-on features that some organizations might find attractive. You should consider the value of these enhancements in light of the ability of the device to continue to perform the essential services of firewall filtering. Do not let features that are not essential to your use of the firewall distract you from obtaining and deploying a reliable filtering device. Be sure to research any newly offered enhancement thoroughly before making a purchase. Unnecessary features not only inflate the cost, but some can cause security problems or introduce vulnerabilities into the network.

A common or popular firewall enhancement is malware scanning. Adding an antivirus, anti-spyware, anti-Trojan, anti-whatever scanner to a firewall is not a significant stretch of the core firewall's capabilities, especially for application proxies and stateful inspection firewalls. The main concerns should be whether such an enhancement can maintain wire speed performance and whether the firewall's anti-

malware abilities are on par with existing standalone malware protections. If you are considering a firewall malware scanning enhancement, pay attention to the details. Important considerations are the detection engine and the mechanism of periodically updating virus signature definitions.

Firewalls can also offer IDS and IPS features. In fact, the merging of intrusion detection and prevention is a logical combination. The ability of a single device that can filter traffic, as well as watch for and defend against intrusions, is attractive. But the question remains: Does the combination device perform both tasks at the same level of excellence as common independent products? If not, then the combination is more of a detriment than a benefit. Also, does the firewall maintain the ability to operate at wire speed?

Some firewalls are equipped to function as a VPN endpoint. When designing a firewall and VPN architecture, this might be an attractive option, but as with every firewall enhancement, you should evaluate its features and performance against that of standalone solutions.

Unified threat management (UTM) is the deployment of a firewall as an all-encompassing primary gateway security solution. The idea behind UTM is that you can use a single device to perform firewall filtering, IPS, antivirus scanning, anti-spam filtering, VPN endpoint hosting, content filtering, load balancing, detailed logging, and potentially other security services, performance enhancements, or extended capabilities. UTM has its advantages, mainly in the ability to deploy a single product and manage multiple security services from a single interface.

Any given UTM can be a jack-of-all-trades product and a master of none. With many of the larger firewall and security product vendors developing product lines supporting UTM, however, it is possible to deploy a reliable single-device solution. But even a reliable UTM product remains a single point of failure. UTMs are obvious improvements to environments that use simple or outdated firewalls and those that lack independent coverage in the non-firewall security categories. An all-in-one UTM device can quickly and effectively improve your organization's security.

As the trend toward virtualized networks, hosts, and applications continues, the need for security within virtualized networking environments increases as well. Just because the virtual host exists in memory alone does not imply that it is immune to hacking or exploitation. In fact, the business may be at even greater risk. Include firewalls in the construction of virtualized networks.

You can still use a hardware firewall when traffic between virtual hosts crosses a physical network segment. However, when virtual host communications occur within memory alone, a virtualized firewall is necessary. A virtualized firewall is really the same as a virtualized host; it is a software construct of a hardware environment that hosts the OS so it can function in memory rather than on actual physical devices. You can install a software firewall into a virtual host to act as a firewall for virtual network connections. When designing and using a virtualized network, make the effort to include virtual firewalls as part of the infrastructure.

MANAGEMENT INTERFACES

You must configure firewalls for proper performance and the security functions you expect of them. Every firewall has a management interface used to configure its functions and features. Through the management interface, you can control and define all aspects of a firewall.

A firewall's management interface may be command-line based (**FIGURE 8-6**) or graphical based (**FIGURE 8-7**). A graphical user interface (GUI) might be a standalone client application you must install onto a host, or it may be a web interface hosted on the firewall's own internal mini-web server.

```
Starting syslog...done.
Starting CRON... done.
pfSense 2.4.4-RELEASE (Patch 3) amd64 Wed May 15 18:53:44 EDT 2019
Bootup complete

FreeBSD/amd64 (pfSense.localdomain) (ttyv0)

VMware Virtual Machine - Netgate Device ID: 2eb1708fbe01888f85c3

*** Welcome to pfSense 2.4.4-RELEASE-p3 (amd64) on pfSense ***

WAN (wan)      -> em0      -> v4: 202.20.1.2/24
LAN (lan)      -> em1      -> v4: 10.0.1.1/24

0) Logout (SSH only)          9) pfTop
1) Assign Interfaces          10) Filter Logs
2) Set interface(s) IP address 11) Restart webConfigurator
3) Reset webConfigurator password 12) PHP shell + pfSense tools
4) Reset to factory defaults  13) Update from console
5) Reboot system              14) Enable Secure Shell (sshd)
6) Halt system                 15) Restore recent configuration
7) Ping host                   16) Restart PHP-FPM
8) Shell

Enter an option: █
```

FIGURE 8-6 An example of accessing a firewall (pfSense) via command line (SSH).
© pfSense



COMMUNITY EDITION



Status / Dashboard

+ 0

System Information



Name	pfSense.securelabsdemand.com
User	admin@172.21.0.1 (Local Database)
System	VMware Virtual Machine Netgate Device ID: 929ab454fd013f6b8c07
BIOS	Vendor: Phoenix Technologies LTD Version: 6.00 Release Date: Tue Apr 5 2016
Version	2.4.4-RELEASE-p3 (amd64) built on Wed May 15 18:53:44 EDT 2019 FreeBSD 11.2-RELEASE-p10
CPU Type	Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz

FIGURE 8-7 An example of a GUI firewall interface (pfSense).
© pfSense

A firewall's management interface may offer encrypted access by default. If the firewall does not offer encryption of the management interface, then you should replace it with a firewall that does. Always enable the encryption so that all future access to the management interface is protected from eavesdropping, interception, and session hijacking. Encryption of the session accessing a firewall's management interface is the most important and critical aspect of management interface configuration.

Secure a firewall's management interface both physically and logically. This implies that direct physical contact with the firewall device, or firewall host, will be limited and restricted to authorized administrative personnel only. Prohibit any physical contact or proximity to external or unauthorized entities.

Secure all logical access to the firewall's management interface through a strong authentication process. If multifactor authentication is supported, use it. If not, require long and complex passwords. Ensure that administrators do not reuse the same password on any other system. Administrators, as well as all users, should never reuse an old password on any system once it expires. Passwords are like milk—they last only so long, even with proper care and storage, and once expired, it should never be reused.

TECHNICAL TIP

The first step in deploying a new firewall is to change the default password of the administrator account. Better yet, modify every predefined user account and every default access code of any type. Use something unique, original, and difficult to predict. Be sure to properly document and

securely store this information so it can be retrieved when needed but protected from unauthorized use.

Limit the methods to access a firewall's management interface. Manufacturers often enable all possible avenues of access to ensure ease of installation upon initial deployment. However, you should disable all methods of accessing the management interface, except for the sole secure mechanism selected for use by your organization. Potentially available avenues of access include telnet, encrypted telnet, web, encrypted web, SSH, and other proprietary options. Select one of the secure options and disable the rest. Test access via the disabled methods to ensure that the rules written and the exclusions made actually are working properly.

Also, consider the communication pathways for accessing the management interface. Usually, the physical cable connection port—often labeled as CON for console, TER for terminal, or ADM for administration—cannot be disabled through software. This port is usually a dedicated port used for direct physical connection management and not for any other purpose. A crossover cable or a special console cable (often RS-232 based) is necessary to connect a computer directly to the firewall to gain access to the management interface.

However, most in-network or over-the-network logical access pathways can be disabled. Be sure to disable access to the management interface on every physical network interface or port that faces outward. In other words, the less-secure NIC ports should not support access to the management interface.

If more than one NIC port exists on the secure side, then limit access to only one of those ports. If the firewall is a multifunctional device that also supports wireless connectivity, disable management interface access over a wireless connection. The resulting configuration should allow management interface access only via the physical CON connection and one logical pathway. Disable all other methods of access.

Configure the firewall to record every attempted and successful connection to the management interface into a log file. It might also be a good idea to record all subsequent configuration changes made via the management interface. This log can serve as a component of change documentation and will assist with future troubleshooting, investigations, and reconfiguration.

Finally, periodically verify manually that all settings remain as configured. This is done by viewing the configuration or management interfaces and comparing the live settings to those defined in your documentation. Whenever you discover a discrepancy, investigate and repair it at once.

CHAPTER SUMMARY

Firewalls are essential to maintaining security, such as supporting valid communications and blocking malicious traffic. However, firewalls are often more than just simple filtering tools. The standard and enhanced features of firewall products require knowledge and skill if these are to be securely deployed.

This chapter discussed building firewall rules, ordering rule sets, understanding what to block, dealing with the limitations of firewalls, maintaining high-speed network performance with a firewall, managing encryption across a firewall, evaluating firewall enhancements, and handling firewall management interfaces.

KEY CONCEPTS AND TERMS

Access control list (ACL)

Default allow

Default deny

Fair queuing

Filters

Firewalking

High availability

Internet Control Message Protocol (ICMP)

Load balancing

Loophole

Management interfaces

Round-robin

Unified threat management (UTM)

Wire speed

CHAPTER 8 ASSESSMENT

1. Which of the following firewall rule guidelines is most important?
 - A. Include every possible address and port in a rule within the set to ensure that an explicit callout exists for every type of communication.
 - B. Place explicit Deny rules for individual systems before explicit Allow rules for ranges that include those individual systems.
 - C. Place universal Allow rules for individual systems before universal Deny rules for systems in that range.
 - D. Include all specific denials for known malicious remote control tools after explicit Allow rules.
2. What form of encryption allows a firewall to filter based on the original source and destination address? (Assume that the firewall is located along the path between session endpoints.)
 - A. Tunnel mode
 - B. Transport mode
 - C. Traffic mode
 - D. Transaction mode
3. Which of the following is a default-deny rule?
 - A. TCP ANY 192.168.42.0/24 ANY ANY Deny
 - B. TCP ANY ANY 192.168.42.0/24 ANY Deny
 - C. TCP ANY ANY ANY ANY Deny
 - D. DENY TCP ANY ANY ANY ANY
4. Which of the following is a potential weakness of a firewall that cannot be fixed with the application of a patch?
 - A. Fragmentation
 - B. Programming bug

- C. Buffer overflow vulnerability
 - D. DoS from external sources
5. Which type of communication session can be improved using caching on a firewall?
- A. Email
 - B. Instant messaging
 - C. Remote access
 - D. Web
6. What is always the most important element within a firewall rule set?
- A. Using specific addresses instead of ANY
 - B. Listing Deny exceptions after Allow exceptions
 - C. Listing inbound exceptions before outbound exceptions
 - D. Listing the final rule of default deny
7. Which of the following is the primary factor when composing firewall rules?
- A. Bandwidth
 - B. Business tasks
 - C. Traffic levels
 - D. User preferences
8. Which of the following is *not* satisfied with a firewall policy?
- A. Assisting in troubleshooting
 - B. Detecting changes in deployed settings
 - C. Ensuring consistent filtering across the infrastructure
 - D. Network load balancing
9. Which of the following is a firewall rule that prevents internal users from accessing public FTP sites?
- A. TCP 192.168.42.0/24 ANY ANY 21 Deny
 - B. TCP ANY ANY ANY FTP Deny

- C. TCP 21 192.168.42.0/24 ANY ANY Deny
- D. TCP ANY ANY 192.168.42.0/24 21 Deny

- 10.** When constructing a rule set, where should you place the default-deny rule?
- A. First, before any other rule
 - B. After any explicit Allow rules
 - C. After any explicit Deny rules
 - D. Last, after all other rules



© Digital_Art/Shutterstock

Establishing Remote Access with VPNs

CHAPTER 9

VPN Fundamentals

CHAPTER 10


VPN Management

CHAPTER 11

VPN Technologies

CHAPTER 12

VPN Implementation



© Digital_Art/Shutterstock

VPN

Fundamentals

VIRTUAL PRIVATE NETWORKS (VPNs) allow organizations to transmit private and sensitive data securely over public intermediary networks. The Internet can serve as an inexpensive long-distance carrier for wide area network (WAN) connections established using VPNs. VPNs use native operating system (OS) features or third-party software, as well as hardware devices, including edge routers, firewalls, and VPN appliances. Remote access for mobile users, linked multi-office intranets, and secured access into an extranet are common VPN architecture solutions.

The use of VPNs on networks ranging from SOHO (small office/home office) networks to corporate global networks has dramatically expanded the ways people connect and do their work. VPNs allow remote employees, business partners, traveling workers, suppliers, distributors, and others to benefit from secured connectivity while geographically distant. With a basic understanding of VPNs, you can begin designing and planning your own VPN solution.

CHAPTER 9 TOPICS

This chapter covers the following topics and concepts:

- VPNs and their uses
- Benefits of a VPN and why you would deploy one
- Limitations of a VPN
- Effective VPN policies
- VPN deployment models and architectures
- Difference between tunnel mode and transport mode
- Relationship between encryption and VPNs
- VPN authentication and how is it implemented
- VPN authorization and why is it necessary

CHAPTER 9 GOALS

When you complete this chapter, you will be able to:

- Define VPNs
- Explain the business and personal uses of VPNs
- Describe the pros and cons of VPNs
- Explain deployment models and architectures of VPNs, including an edge router, a corporate firewall, a VPN appliance, a remote access server, a site-to-site VPN and supporting devices, and a host-to-host VPN and supporting devices

- Differentiate between a transport-mode VPN and a tunnel-mode VPN
- Describe the importance of encryption, authentication, and authorization to VPNs

WHAT IS A VIRTUAL PRIVATE NETWORK?

A virtual private network (VPN) is a mechanism to establish a secure remote access connection across an **intermediary network**, often the Internet. VPNs allow remote access, remote control, and highly secured communications within a private network. VPNs employ encryption and authentication to provide confidentiality, integrity, and privacy protection for network communications.

The term VPN has its origins in the telecommunications world. A telephone VPN created a PBX-like system for businesses without the need for deployment of true **private branch exchange (PBX)** hardware. Instead, the system used a public telephone service and the PBX services at the telco's central offices. This service/product sold under the name Centrex (a combination of *central* and *exchange*) in the 1960s through the 1980s.

After the proliferation of computer networks and Internet connectivity, the term VPN evolved to refer to tunneling connections across network links. Early computer VPNs focused on the tunneling or encapsulation processes and rarely included encryption services. Today, VPNs are almost always secured using encryption. However, you should never assume that anything is totally secure, especially connections over **public networks**. Always confirm that a product performs encryption properly before you depend on it for sensitive operations.

A VPN creates or simulates a network connection over an intermediary network. But what makes a VPN private? Several possible mechanisms exist:

- The primary organization owns all of the network infrastructure components, including switches, routers, and cables. A true private VPN occurs when a single organization owns all of the hardware supporting its VPN. However, few organizations actually own all of the connections between their locations, so this is usually impractical or prohibitively expensive. This wholly owned and operated system constitutes a **trusted VPN**.
- A dedicated set of **channels** is used across leased telco connections. This method provides physical isolation even on third-party equipment; hence, privacy is maintained. This type of system is more practical, but it is still expensive. This can also be called a trusted VPN, since you must be able to trust the owner of the hosting infrastructure to protect network communications against eavesdropping.
- Encryption ensures privacy even over public networks, such as the Internet. This method is the most reliable, as the other two options are still at risk to eavesdropping. Additionally, encryption to provide privacy is not only practical, it is the least expensive option as well. This system can be called a **secured VPN**.
- A **hybrid VPN** establishes a secure VPN over trusted VPN connections. A trusted VPN allows an organization to know and control the pathway of its transmissions. However, a trusted VPN does not protect against eavesdropping or alteration. A secure VPN protects the confidentiality and integrity of data, but it does not control or ensure the transmission path. When you combine these two VPN techniques, you create a potentially more secure and practical solution. Two possible layouts of hybrid VPNs are shown in **FIGURES 9-1** and **9-2**.




FIGURE 9-1 A hybrid VPN consisting of a secure VPN across an intermediary trusted VPN.




FIGURE 9-2 A hybrid VPN consisting of a secure VPN segment within a trusted VPN.

VPNs are often associated with remote access or remote control. However, these associations need clarification to have value. Remote control is the ability to use a local computer system to remotely take control of another computer. In a way, this process is the application of the thin-client concept on a fully capable modern workstation to simulate working against a mainframe or to virtualize your physical presence. This application is generally the same as a VPN, which creates a remote network connection rather than a remote control session.

With a remote control connection in place, the local monitor, keyboard, and mouse control a remote system. This process looks and feels as though you are physically present at the keyboard of the remote system, which could be located in another city or even on the other side of the world. Every action you perform locally acts as if you were physically present at that remote computer virtually via the remote control connection. The only limitations are the speed of the intermediary network link and the inability to physically insert or remove media and peripherals.

You might think of remote control as a form of software-based thin-client or terminal client. In fact, many thin-client and terminal client products sell as remote control solutions. Many modern operating systems include remote control features, such as Remote Desktop, found in most versions of Windows. Once enabled, a Remote Desktop connection remotely controls another Windows system from across the network.

Remote access is different from remote control. A remote access link enables access to network resources using a WAN link to connect

to the geographically distant network. In effect, remote access creates a local network link for a system not physically near the network. Over a remote access connection, a client system can technically perform all of the same tasks as a locally connected client. Network administrators can impose restrictions on which resources and services a remote access VPN client can use.

Remote access and VPNs were originally supported over dial-up telephone links using modems. Today, remote access encompasses a variety of connection types, including Integrated Services Digital Network (ISDN), digital subscriber line (DSL), cable modem, satellite, mobile broadband, and more. Due to the wide availability of high-speed Internet connections, VPNs and other remote access solutions have become popular for both personal and business purposes.

In many cases, a remote access connection is created from a remote client back to a primary network. If the remote client needs to connect directly to the local area network (LAN), such as over a dial-up connection, a remote access server (RAS) will host a modem to accept the connection. If the remote client can use the Internet to access the LAN, then a local Internet connection is necessary. Once a normal LAN connection or Internet connection runs from the client, the VPN link is possible. Once the connection is established, the remote client can interact with the network as if it were locally connected.

In **FIGURE 9-3**, a LAN and a remote client have a connection to the Internet. These two connections are independent of each other. The LAN's connection is usually a permanent or **dedicated connection** supporting both inbound and outbound activities with the

Internet. The remote client can have a dedicated or **nondedicated connection** to the Internet. In the latter case, the connection is established before a VPN can be created. Once both endpoints have a connection to the intermediary network, in this case the Internet, the VPN can be created. In **FIGURE 9-4**, a new network connection is established from the remote client to the LAN across the intermediary network. This new network connection is the VPN.




FIGURE 9-3 A corporate LAN and a remote client with Internet connections.




FIGURE 9-4 A VPN established between a remote client and a LAN over the Internet.

VPNs can operate over standard Internet connections or dedicated business communication circuits, such as Multi-Protocol Label Switching (MPLS), Asynchronous Transfer Mode (ATM), and Frame Relay. However, the additional expense of a dedicated, isolated, and even secured business circuit is not necessary with a VPN.

A VPN can operate securely over the Internet and still provide high levels of security through encryption. This allows inexpensive insecure links to replace expensive business-leased lines without sacrificing security.

VPNs are one of the most efficient and cost-effective means to provide secure remote connectivity. VPNs take advantage of cheap long-distance connections when established over the Internet, since both endpoints need only a local Internet link. The Internet serves as a “free” long-distance carrier. Because the speed of the VPN depends

on the speed of the local Internet link, you still might want an **optical carrier (OC)** line (such as an OC-1 at 51.84 Mbps up to OC-768 at 40 Gbps) for high-speed connectivity.

Connections from a LAN to an intermediary network can support VPN traffic only, or they can allow a combination of both VPN and normal non-VPN traffic. The latter configuration is less secure, but it offers flexibility as to whether all communications must be VPN-secured or not. An Internet connection reserved solely for VPN use, therefore, is not necessary. An OC-1 line is more than capable of supporting one or more VPN links—in addition to numerous non-VPN Internet sessions—with no difficulty or latency. Latency is a delay in the transfer of data that is not desired or planned.

Setting up VPNs can require extensive knowledge and expertise on the part of the IT or security administrator. For example, some important concerns of secure VPNs include:

- All VPN traffic must be authenticated and encrypted. A VPN without authentication is not private, and a VPN without encryption is insecure.
- All VPN endpoints must abide by the same security parameters and algorithms. Each VPN tunnel must have corresponding **encryption key** sets to securely exchange encrypted content. Additionally, the same security policy should govern all endpoints.
- Proper encryption protocols must ensure that no external third party can affect the security of the VPN. Weak encryption makes a “secure” VPN vulnerable to attacks.

When you use a trusted VPN, you need to consider these concerns:

- A trusted VPN is based on the provider’s ability to limit and control access to the VPN’s content. Only the trusted VPN

provider should be able to modify the channels or pathway of the VPN.

- Only the trusted VPN provider can add, remove, or change data in the trusted channel. Violating this violates the trust the client places in the provider.
- The addressing and routing performed within the trusted VPN must be defined before the VPN goes online. These services are usually predefined in the SLA (service level agreement), but they may be dynamically modified for each VPN connection.

Even hybrid VPNs have an important focus for concern, namely that the segments of the VPN that are trusted versus secured need clear definition. Mistaking security for trust—or vice versa—can have devastating results.

VPNs use tunneling or encapsulation protocols. Tunneling protocols encase the original network protocol so that it can traverse the intermediary network. In many cases, the tunneling protocol employs encryption so that the original data traverses the intermediary network securely. The protocols that create VPNs include IPsec (Internet Protocol Security), PPTP (Point-to-Point Tunneling Protocol), L2TP (Layer 2 Tunneling Protocol), SSH (Secure Shell), SSL (Secure Sockets Layer), OpenVPN, and TLS (Transport Layer Security). The dominant forms of secure VPNs use IPsec or SSL/TLS as the tunneling/encapsulation protocol. Most VPNs use software that operates on top of the OS of a host. However, some VPN appliances can support VPN connectivity without adding any software to the host. A **host VPN** software product allows a single host access to VPN services, while a **VPN appliance** allows an entire network to access VPN services.

VPNs simplify many business networking problems by providing an easy and efficient means to securely connect headquarters, remote offices, traveling workers, and telecommuters.

WHAT ARE THE BENEFITS OF DEPLOYING A VPN?

The reasons to deploy and use VPNs vary greatly among organizations. Cost is always a significant factor in any business decision. Budgets are never unlimited, so organizations must consider their options amid limited funds to accomplish their missions and goals. One common goal is high productivity. Granting workers the ability to access and use resources in a timely and efficient manner assists in the completion of work. When those resources are computer files or network services, employees no longer need to be in the same building as those resources. Remote access to resources, therefore, is becoming more common than ever.

Secure remote access is essential. As the proliferation of access and connectivity spreads from work to home to portable/mobile devices, access to the Internet and private LANs is becoming ubiquitous. Companies must use security controls on resource access or suffer the consequences of insecure access methods. With the removal of physical limitations for access comes the loss of control over where and how workers connect into the private LAN.

Workers may connect to the company LAN from mobile phones, through Internet cafés, over hotel networks, and at other random Wi-Fi hotspots. In this age of Bring Your Own Device (BYOD), many use personally owned laptop computers and mobile devices rather than officially issued company systems. All of these are outside the control of the company's IT and security department. The only option

is to limit LAN connections to those that can be secured. Thus, VPNs have become a necessity in this mobile and interconnected world.


VPNs support remote access from a wide variety of complex devices, reduce risk caused by insecure access locations, and enable interaction with all LAN resources. Furthermore, flexibility, scalability, ease of administration, reliability, and more make VPNs an obvious choice in the face of modern connectivity risks and challenges. **Scalability** is the ease of which an organization can quickly increase capacity and use or shrink capacity and use of a device, system, or network.

■ NOTE


Does every worker and every organization need VPNs and remote access? No. For many work situations, VPNs are not the correct solution. These include any form of work that requires special tools, physical access to equipment, or close supervision by managers.

Remote access, mobile connectivity, and secured communications are solid reasons to deploy and use a VPN. But are these the only positive aspects of a VPN? The most often touted benefit of VPNs is cost savings. VPNs are a great way to save on long-distance charges for telecommuters and traveling workers. These also create huge savings for businesses that would need only local Internet links for a VPN, rather than a dedicated **leased line** between each location (**FIGURE 9-5**). The farther away each business location is and the more locations a company has, the more cost savings a VPN can

generate. A leased line is a semi-private connection subscription whereby the traffic on the line is limited to those that have subscribed (often one to a few businesses in a geographic area). Leased lines can potentially transmit the data of business competitors, so eavesdropping on unencrypted data packets is a concern, as is the high cost of such an exclusive means of communication.



Dedicated Lines Mesh Network



VPN Network

FIGURE 9-5 A corporate network using dedicated leased lines versus using VPNs over the Internet.

Additionally, to truly compare the connectivity that a VPN offers to dedicated leased lines, you need a full mesh of leased lines. A full mesh requires a line between each business location. This allows for direct communication between one site and another. Since a VPN across the Internet would provide the equivalent site-to-site communication capabilities, only a mesh network of **dedicated leased lines** can truly compare. In this leased line situation, your organization is the only entity using the communication pathway; it is dedicated for your business. This solution is obviously very expensive compared to a VPN's significant cost savings.

As corporations seek to reduce IT infrastructure costs, a common technique is to allow employees to telecommute. **Telecommuting** allows workers to access corporate resources whether the employee works from home, while traveling, or while on site with a customer. In the past, telecommuting clearly implied the use of dial-up connections to connect with the company LAN; now it refers to any geographically remote employee connecting to the network to do his or her job.

With the proliferation of high-speed broadband connections and Wi-Fi, telecommuting has become more plausible and realistic. Through the use of VPNs, telecommuting enables a true remote office, rather than just a file exchange and communication system. VPNs make telecommuting not only possible, but also practical and secure. Because of VPNs, expanding the workforce is no longer a geographically limited proposition.

Extranets are often deployed as businesses establish new partnerships or seek more interaction with suppliers, distributors, and other external entities. Extranets are border networks, similar to a

demilitarized zone (DMZ), where resources are hosted for access by external entities. However, unlike a DMZ, an extranet is not open for public use. Only a limited and specific set of users is allowed to connect into an extranet. Often, this limitation means that a specific VPN configuration is necessary to access the extranet's resources. With VPNs, extranets are both possible and practical. Extranets can suffer from extensive hits or visitor traffic to the site. Consider the process of checking the delivery status of a package you ordered by logging into the carrier or delivery service's website. That is using an extranet. You have an ability to see only the status of your specific package within the system, and the means to gain that access is with a specific tracking number.

VPNs allow system administrators to manage and control a network remotely. Because of VPNs, employees can work from anywhere, friends can create WAN links to support multiplayer games, and technical support can repair client systems remotely. A VPN is the solution whenever a network connection is needed between two systems or two networks but it is not feasible to install a direct cable connection.

Often, the real benefits of a VPN are not from the VPN itself, but from all of the new possibilities for work, research, learning, and play that are now feasible. These benefits include:

- Reduced equipment costs
- Unlimited geographic connectivity
- Increased flexibility and versatility of worker location
- Improved privacy and confidentiality due to strong encryption
- Verified transmission integrity

- Fully scalable global infrastructure and architecture
- Rapid deployment options
- Flexible integration with existing networks and technologies
- Faster return on investment (ROI) than traditional WAN infrastructures
- Reduced dependence on long-distance carrier solutions
- Reduced support burden on Internet service provider (ISP)

Individuals and organizations that use and integrate VPNs in new and unique ways are sure to reap additional benefits. History shows us that as new means of communication are created, these often change and are used in ways that were unpredictable at the beginning of the installation. However, VPNs are not perfect, and some challenging issues limit their use.

WHAT ARE THE LIMITATIONS OF A VPN?

Although VPNs offer many benefits, you need to evaluate the very real and distinct limitations before you put a VPN in place.

A VPN connection offers flexible secure communication options, but it does not ensure quality of service. A VPN link is dependent upon the stability, throughput, and availability of the ISP connection, as well as the intervening network connections between endpoints. VPNs over the Internet can easily suffer from latency, fragmentation, **traffic congestion**, and dropped packets. This also results in a lack of dedicated bandwidth between business sites because of the volatility of the Internet.

Fragmentation occurs when a packet's size exceeds the size allowed on a segment of the network or Internet. The too-large packet is broken or fragmented into smaller packets that meet the size limit. Traffic congestion on a network is similar to rush hour traffic on most highways. During peak times, more traffic (packets or cars) attempts to join the artery (bandwidth on a wired or wireless device, or physical highway), which causes all traffic to significantly slow down because capacity is greater than the highway (network or physical) can support. This slowdown affects both network traffic and employee productivity. A malware or denial of service (DoS) attack against the network shares some common behaviors with traffic congestion, which can be dangerous if an administrator mistakes an attack for congestion.

TECHNICAL TIP

TECHNICAL TIP

Encrypted traffic does not compress. Compression reduces the size of a data set, removing redundancies or repeated sections within the data set. Properly encrypted data produces ciphertext that does not contain redundancies or recognizable patterns. If ciphertext did have these characteristics, it would not be as secure. Thus, without these redundancies, it is not possible to compress encrypted data.

Although VPNs are excellent solutions over nearly every broadband connection option, a VPN can be difficult to maintain over dial-up. In more rural areas, dial-up connections are still in use and may be the only alternative to wireless communication. VPN traffic is encrypted, and encrypted traffic does not compress. Most dial-up modem connections rely on **compression**—mainly hardware compression—to improve connection speed. When compression is not possible, a significant and noticeable speed reduction occurs. Additionally, VPN tunnel management can impose a significant increase in management overhead because of changes in protocol headers, potential authentication latency, and a prolonged connection establishment negotiation.

Another area of concern is the minor risk or potential of data exposure while in transit over the Internet. This is only a real concern if the VPN does not use strong encryption or configures the encryption improperly. Proper security management will eliminate this as a serious concern.

Vulnerabilities exist at VPN endpoints. With a VPN, side attacks against the encrypted link are nearly eliminated. However, data entering or leaving the VPN is at risk. An end-user computer could be infected by malicious code that can traverse the VPN link into the company LAN. Also, private and confidential data from the company LAN can be copied across the VPN link to the end-user computer. On this computer, that data is less secure and subject to a wider range of threats.

You should also consider the increased difficulty in providing technical support remotely. This is especially true when the remote connection is not functioning. In addition, it is more difficult to keep remote systems in compliance with security settings, conduct training, allow supervisory oversight, enable HR management, and monitor user activities.

An even larger concern is granting open or blanket unrestricted network-resource access to those connecting via VPN. You must enforce stronger authentication and authorization limitations on VPN users, especially on VPN telecommuters. Remote users should have access only to those resources necessary for their current tasks. Unlimited access to network resources can quickly result in exploitation and confidential data leakage if the remote user or the remote computer is compromised.

If you understand these limitations and address each properly, you can help to avoid catastrophic mistakes when correctly installing and productively using VPNs. One of the primary tools to accomplish this is the VPN policy.

WHAT ARE EFFECTIVE VPN POLICIES?

Effective VPN policies are those that clearly define security restrictions imposed on VPNs and align with the overall IT mission and goals of your organization. VPNs can offer numerous exciting possibilities for mobility and interconnection. However, VPNs can also be a risk to the confidentiality and stability of your organization's infrastructure.

Like all security policies, your VPN policy should start with a thorough risk assessment and analysis. Without fully understanding the assets, processes, threats, and risks of VPNs, you cannot effectively use or manage the device or traffic.

Developing your VPN security policy is not a simple or straightforward task. You need to plan for time and effort to address a wide variety of issues and concerns. Some of the aspects of design and planning of a VPN policy include (but are not limited to):

- Considering the benefits and drawbacks of software and **hardware VPN** solutions
- Imposing stringent multifactor authentication (multiple elements you know or that validate who you are) on all VPN connections
- Implementing strong access control (authorization) restrictions on all VPN connections
- Defining how the VPN will be managed, through what interfaces, and by whom
- Exploring the complexities of patch management over VPN

- Defining the mechanisms of providing remote technical support for VPN telecommuters
- Enabling detailed auditing on all activities occurring across or through a VPN
- Defining distinct qualifications on granting user access to telecommuting VPNs
- Prescribing the user training requirements for all VPN activities

VPN DEPLOYMENT MODELS AND ARCHITECTURE

One of the first decisions you face when deploying a VPN is determining which device will serve as the termination point of the secured tunnel. You have several options, but often the decision rests on where in the network infrastructure you want to position the tunnel endpoint. The features the VPN device provides may be a factor in the decision.

Such factors include deciding which devices have sufficient processing power to maintain wire speed, even with heavy traffic and complex encryption. Another concern is whether network address translation (NAT) is present, as this can impose problems for tunnel-mode encryption.

VPN Deployment Models

A hardware VPN uses a dedicated device to perform VPN functions for the network. The three primary VPN device models are **edge router**, **corporate firewall**, and VPN appliance.

Edge Router

With edge routers as the VPN termination point, the VPN link exists only over the public intermediary networks, not within the private LAN(s). This does require that the edge router support VPN connectivity. Edge router VPN termination ensures that a firewall can filter the traffic exiting the VPN to move into the LAN. All traffic, regardless of transportation means, must comply with the firewall's filtering rules. If the VPN terminates inside the firewall, then traffic from the VPN could potentially violate security because it was not firewall-inspected.

VPN termination on edge routers is best suited for controlled access into the DMZ. Such a configuration grants business partners easy access to the DMZ without exposing traffic to the Internet or granting unnecessary access to the private LAN.

Corporate Firewall

Terminating the VPN at the corporate firewall is possible if the firewall supports VPN services. Not all firewalls provide this service, so this depends on your choice of make and model of the firewall product. With a firewall-to-firewall VPN across the public network, users from one network LAN can access resources in another network LAN without additional complexities. Primarily, this configuration treats the VPN link between the firewall endpoints as just another route in the LAN (actually, WAN). The benefit here is that users do not need to reauthenticate or abide by additional firewall restrictions when the VPN terminates at the corporate firewall.

Corporate firewall termination of the VPN means that the traffic entering or leaving the VPN does not pass through the filtering restrictions of the firewall. Instead, the firewall just serves as a tethering point for the VPN tunnel endpoint. Any traffic not associated with the VPN, however, is subject to firewall investigation.

This configuration has a potential limitation. As the number of VPN links or the traffic load of the VPNs increases, the resulting increase in cryptographic computations could interfere with the firewall's wire speed filtering performance. To prevent this, perform a trend analysis to monitor for this condition and improve firewall performance before a bottleneck occurs.

VPN Appliance

A third device option is a dedicated VPN appliance. Unlike an edge router or firewall termination point, a dedicated VPN appliance specifically handles the load of a VPN, instead of VPN support as an add-on service. You can position a VPN appliance outside the corporate firewall, in a similar location to an edge router, so that all VPN traffic passes through firewall filters.

A VPN appliance can reside inside corporate firewalls to prevent firewall filtration. This deployment is similar to the corporate firewall concept, at least in terms of not filtering VPN traffic. This deployment method also ensures that no external entity can interfere with the endpoints of the VPN tunnel. These techniques are useful when corporate firewalls already exist that do not support the VPN technology or architecture you want. So, rather than replacing the firewalls, you install an additional dedicated VPN appliance.

VPN Architectures

In addition to the selection of the VPN device, you have several architectural decisions to make. These focus on the intended purpose, function, or use of the VPN, such as remote access, host-to-host, site-to-site, and extranet access.

Remote Access

The first VPN architecture is basic remote access (**FIGURE 9-6**).

Remote access VPN is also known as **host-to-site VPN** because it supports single-host VPN connections into a LAN site. This design grants individual telecommuters or traveling workers easy access into the private LAN. A single LAN can support several remote users with any of the VPN endpoint concepts: edge router, corporate firewall, or VPN appliance.




FIGURE 9-6 A VPN used to connect a single LAN with remote mobile users.

Site-to-Site

A second VPN architecture is site-to-site (**FIGURE 9-7**). **Site-to-site VPNs** are also known as **LAN-to-LAN VPNs** or **WAN VPN** connections between LANs. Regardless of the name, a site-to-site VPN supports secure connections between LANs over intermediary public networks. When installed properly, a site-to-site VPN can be an inexpensive mechanism to create a single **distributed LAN** (also known as a WAN) for a multilocation organization. A site-to-site VPN uses any of the VPN endpoint concepts: edge router, corporate firewall, or VPN appliance.




FIGURE 9-7 A VPN used to connect multiple LANs.

A slight variation on the site-to-site VPN is the site-to-site VPN with remote mobile users. This is a combination of the site-to-site

VPN concept with remote access VPNs (**FIGURE 9-8**). The benefit of this architecture is the ability to add new employees with secure connections to the LAN without needing additional corporate facilities to provide on-site workspaces.




FIGURE 9-8 A VPN used to connect multiple LANs with remote mobile users.

Host-to-Host

A third VPN architecture is host-to-host. Host-to-host VPNs are also known as client-to-server, remote-to-office, or remote-to-home VPNs. A **host-to-host VPN** is a direct VPN connection between one host and another. This mechanism operates over a public network or within a private network. Over a public network, a host-to-host VPN provides a connection that is secure from the public. Over a private network, it provides an additional level of security for mission-critical or highly sensitive transactions.

Host-to-host VPNs labeled as **client-to-server VPNs** create secure client interaction with the services of a resource host. This is similar to, but not exactly the same as, a secure web link between a web browser and a web server. SSL can be used for application protocol security, as it is with secure web sessions, or as a VPN protocol. As a VPN protocol, SSL operates at the Network Layer of the Open Systems Interconnection (OSI) Reference Model; as a web session security tool, it operates at the top of the Transport Layer.

A **remote-to-office VPN** is a direct link between a portable or home system and an office workstation. This VPN link allows a user to work from home or while traveling without sacrificing access to resources, services, or applications that might only be installed (or licensed) for use on the office workstation computer. This VPN can establish a remote control session as easily as a remote access session.

A **remote-to-home VPN** connects a work computer or a portable computer back to a home system, effectively the opposite configuration of the remote-to-office. The remote-to-home VPN

grants you access to a home computer while you are away from the house.

The host-to-host VPN variations usually depend on software VPN solutions native to the operating system or third-party applications installed on the host. However, some VPN appliances are adaptable to these simpler host-to-host connectivity architectures.

Extranet Access

A fourth VPN architecture is extranet access. With a VPN tunnel endpoint positioned at or inside the perimeter of an extranet, this option serves as a pathway for business partners, distributors, suppliers, and so forth to gain access to corporate resources without exposing their traffic to the Internet or granting them unnecessary access to the private LAN.

A VPN linked into the extranet, as opposed to the DMZ, provides greater security to the remote entities. A VPN link to the DMZ exposes the remote entities to any threats found in the DMZ. Since the DMZ is publicly accessible, it is risky. An extranet VPN grants the remote entity secure communications without significant risks at the VPN termination point.

VPNs commonly serve as a choke point to control which external entities have access to the extranet. Only those granted specific access, assigned user accounts, and provided configuration details are able to configure and establish a VPN link with an extranet.

TUNNEL VERSUS TRANSPORT MODE

VPNs can use two main types of encapsulation encryption. These are known as tunnel-mode and transport-mode encryption.

Tunnel-mode encryption protects the entire original IP packet's header and payload. The encrypted packet becomes the payload of a new IP packet with a new IP header. This form of encryption ensures that the identities of the original endpoints of the communication are kept confidential while the traffic traverses the secured link. Tunnel-mode encryption is commonly used by VPNs linking network sites together or providing secure remote access.

Transport-mode encryption protects only the original IP packet's payload. The encrypted payload retains its original IP header. This form of encryption protects only the payload, not the identities of the endpoints. Transport-mode encryption helps VPNs link individual computers together.

THE RELATIONSHIP BETWEEN ENCRYPTION AND VPNS

Encryption and a secure VPN are virtually inseparable. A secure VPN exists only because the traffic is encrypted, but some trusted VPNs may or may not use encryption. To fully understand and appreciate the operations of VPNs, you need a reasonable understanding of encryption.

Encryption is just one aspect of the larger topic of cryptography. **Cryptography** is the art and science of changing information so it is not easily recognized or understood by unauthorized third parties. Cryptography occurs through a complementary and reversible set of processes: encryption and decryption. Encryption is the process of converting original usable data, called plaintext, into an unusable chaotic form, called **ciphertext**. **Decryption** is the process of converting ciphertext back into plaintext. A secure communication product must provide both encryption and decryption.

Modern cryptography is based on algorithms. An **algorithm** is a set of rules and procedures, usually mathematical in nature, that define how the encryption and decryption processes operate. Algorithms are often very complex. Many algorithms are publicly known and anyone can investigate and analyze the strengths and weaknesses of an algorithm.

TECHNICAL TIP

Less than 10 years ago, 64 bits was considered strong encryption. Today, 128 bits is the smallest symmetric key length that can be considered strong; in many circumstances, 256 bits is expected. Is a key of twice the size of the previous key length stronger? Actually, while a 128-bit key is twice as long as a 64-bit key in that it has twice as many digits, a 128-bit key creates a keyspace much more than twice the size of a 64-bit key, and a 256-bit key is more than twice the size of the 128-bit key. Every time an additional bit is added to a key length binary number, it doubles the size of the possible keyspace. A 128-bit key creates a keyspace that is doubled 64 times that of a 64-bit key. That is 2^{64} or 1.8×10^{19} times as large as that of a 64-bit keyspace.

Encryption algorithms use a key. The encryption key is a unique and secret number that controls the encryption and decryption processes performed by the algorithm. A key is a very large binary number measured or defined in terms of its bit length. The bit length of the key is the number of binary digits that comprise the key. For example, a key of 128 bits is 128 binary digits long.

The bit length of an algorithm's key defines that algorithm's **keyspace**. The keyspace is the range of keys that are valid for use for that specific algorithm. Any key created using the specific number of binary digits of the key length is part of that algorithm's keyspace. This keyspace includes every number between a key of all zeros to a key of all ones (remember, this is binary code with only ones and zeros as digits). This means that 2^{128} number of keys are available

to an algorithm with a key length of 128 bits. 2^{128} in decimal form is more than 3.4×10^{38} .

Three main types of algorithms operate within the realm of cryptography: symmetric, asymmetric, and hashing.

Symmetric Cryptography

Symmetric cryptography is based on algorithms that use a single, shared secret key. A common way to remember the type of key used by symmetric is to remember a synonym for symmetric: same. The same key must encrypt and decrypt data, and that key must be shared with all communication partners of the same session.

Symmetric cryptography is very fast in comparison to asymmetric cryptography, typically 10,000 times as fast with similar key length and message size. Generally, the longer the length of a symmetric key, the stronger the encryption produced. Most algorithms have one or only a few key length options, so picking an algorithm with a longer key often ensures stronger encryption. Keys shorter than 128 bits are considered weak, keys of 128 to 256 bits are considered strong, and keys longer than 256 bits are considered very strong. That's why when you select a symmetric cryptography product, you should only use solutions employing 128-bit or longer keys.

Selecting a strong algorithm is more involved than just key length; you also need to consider other issues, such as the use of random numbers, use of the avalanche effect, and resistance to reverse engineering attacks. Random numbers are numbers in a sequence that is not easily assumed or derived as a pattern. The **avalanche effect** is associated with hashing, where a small change elicits a large effect in the output hash. Most weak algorithms are quickly broken once released to the public, so most consumers do not need to focus as much on the details of the algorithm as to look for algorithms that support longer keys. The general rule (which does gloss over some of the details) is that longer keys are better. For more information on cryptography, selection of algorithms, and the internal

workings of algorithms, consult *Schneier's Cryptography Classics Library* or *Introduction to Modern Cryptography*.

■ **NOTE**

Hypothetical communication partners Bob and Alice are commonly used to describe and discuss the functions of cryptography. This is because we often note communication from points A to B, which is why Alice and Bob were selected for the examples.

Symmetric cryptography works as follows (**FIGURE 9-9**):




FIGURE 9-9 The symmetric cryptography process.

1. Bob and Alice exchange a symmetric key (K1). (This process is normally performed by an asymmetric process or out-of-band communication.)
2. Bob encrypts a message with the shared symmetric key.
3. Bob transmits the message to Alice.
4. Alice decrypts the message using the shared symmetric key.

■ NOTE

Out-of-band communication means communication through a different channel or means. It can also refer to a communication stream that occurred separate in time from the existing communication. For successful key exchange, this process often occurs just prior to the rest of the communication using the key.

The security service provided by symmetric cryptography is confidentiality. Symmetric encryption, with a reasonably long key, prevents unauthorized third parties from accessing or viewing the contents of communications or stored data files. Symmetric cryptography protects files on storage devices, as well as data in transit. In fact, due to its strength and efficiency, symmetric cryptography is the preferred method to secure data in storage or in transit of any size.

Some examples of symmetric cryptography algorithms or systems include DES, 3DES, AES, CAST, RC4, RC5, RC6, RC7, IDEA, Twofish, and Blowfish.

VPN solutions employ symmetric cryptography to protect communications in transit. This protection prevents unauthorized eavesdropping or midstream modifications.

Asymmetric Cryptography

Asymmetric cryptography is based on algorithms that use either **key pairs** or some other special mathematical mechanism.

Asymmetric cryptography that uses key pairs is commonly known as **public key cryptography**. All public key cryptography is asymmetric, but some asymmetric algorithms are not public key algorithms. A common way to remember the types of keys or non-keys that are used by asymmetric cryptography is to remember a synonym for asymmetric: different. Different keys are used for different purposes, different keys are used by different members of the communication session, and some systems use something different from keys altogether.

TECHNICAL TIP

Keep in mind that an average 2,048-bit asymmetric key is not just 8 times as long as a very strong 256 symmetric key; in fact, it has a keyspace that is doubled 1,792 times that of the keyspace of a 256-bit key, an astronomically large number: $2^{1,792}$ or 2.79×10^{539} times that of a 256-bit keyspace. Numbers of this size are very hard to manage, much less run through complex mathematical computations.

Asymmetric cryptography is very slow in comparison to symmetric cryptography. This is due to the complexity of the math used by these algorithms, as well as the length of asymmetric keys

(when there are keys). Public key cryptography uses keys that are 1,024 to 8,192 bits long (and sometimes longer).

Asymmetric cryptography also uses math based on a concept known as one-way functions. A **one-way function** is a mathematical operation performed in one direction relatively easily, but impossible or nearly so to reverse. This type of math was formalized into cryptographic algorithms in the late 1970s. Just because it seems old, do not discount the strength of this type of function.

Asymmetric cryptography is slow in comparison to symmetric due to its extremely long keys and its extraordinarily complex mathematical functions. Thus, asymmetric is generally not suitable for encrypting bulk data for storage or transmission; instead, this type of encrypting is best suited for **identity proofing** and **key exchange**.

The protections provided by asymmetric cryptography are authenticity and nonrepudiation. **Authenticity** is a term used to convey the combination of authentication and access control. Based on usage, especially with public key cryptography, asymmetric solutions can prove the identity of the source of a message (authentication) or control the destination or receiver of a message (access control). Additionally, when used properly, asymmetric cryptography may provide proof that a sender sent a message and prevent that sender from being able to credibly deny sending it, which is known as nonrepudiation.

Public key cryptography, a subset of asymmetric cryptography, is based on key pairs. Each participant in a communication or community has a discrete key pair set. The key pair set consists of a **private key** and a **public key**. The private key is kept secure and private at all times. The public key is put out for open public access

and use. The key pairs work together as opposites. The encryption or encoding that one of the keys performs can only be undone by the opposite key of the pair.

One mechanism afforded by asymmetric cryptography is a **digital envelope**. It works as follows (**FIGURE 9-10**):




FIGURE 9-10 The asymmetric cryptography process of a digital envelope.

1. Bob obtains Alice's public key.
2. Bob encodes a message with Alice's public key.
3. Bob sends the message to Alice.
4. Alice decodes the message using her private key.

The encoded message is readable only by Alice, who has the corresponding private key. A digital envelope ensures that the message is accessible only by the intended recipient, namely the owner of the corresponding private key. The mechanism of digital envelope exchanges symmetric keys between communication partners. In this mechanism, one side of the conversation performs key generation, and then exchanges that key securely. This is the mechanism VPNs use to exchange keys securely.

Another mechanism afforded by asymmetric cryptography is a digital signature. A **digital signature** proves the identity of the sender. It works as follows (**FIGURE 9-11**):




FIGURE 9-11 The asymmetric cryptography process of a digital signature.

1. Alice uses her private key to encode a message.
2. Alice sends the encoded message to Bob.
3. Bob accesses Alice's public key.
4. Bob uses Alice's public key to unlock the encoding.

This process proves to Bob that the message came from Alice. This is authenticity. A digital signature does not prevent others from viewing the message or confirming the source of the message. A digital signature is not encryption; it does not provide confidentiality protection. VPNs use digital signatures, but these usually employ the version that uses hashing as part of the process (see the next example).

This mechanism of digital signatures improves through the addition of hashing. (Hashing is described in more detail in the next section.) A digital signature using hashing is a bit more complex, and uses the following procedure (**FIGURE 9-12**):




FIGURE 9-12 The asymmetric encryption process of digital signature with hashing.

1. Alice computes the hash value of a message.
2. Alice encodes the hash value with her private key. This encoded hash is the digital signature.
3. Alice adds the encoded hash to the message.
4. Alice transmits the message.
5. Bob strips off Alice's digital signature.
6. Bob hashes the original message to compute its hash.
7. Bob obtains Alice's public key.
8. Bob decodes Alice's digital signature to reveal Alice's original hash.
9. Bob compares Alice's original hash to the hash Bob calculated.

If the hashes are the same, Bob accepts the message as retaining its integrity. This has also proved that the message was from Alice (authenticity). Furthermore, this process ensures that Alice cannot deny having sent it (nonrepudiation). VPNs commonly employ this form of digital signature to prove the identity of the endpoints and confirm the integrity of transmissions.

Some examples of asymmetric cryptography algorithms or systems that are not public key cryptography are Diffie-Hellmann, El Gamal, and elliptical curve cryptography (ECC). The most common example of an asymmetric cryptography algorithm or system that is public key cryptography is Rivest-Shamir-Adleman (RSA). VPN solutions employ asymmetric cryptography to secure symmetric key exchange (via digital envelopes), authenticate VPN link endpoints or users (via digital signatures), and, when used in conjunction with hashing, verify source and integrity of transmitted messages (also digital signatures).

Hashing

Hashing is the cryptographic function that takes the input of a file or message and creates a fixed length output. The input can be of any size, but the output is a fixed length based on the hashing algorithm used. Hashing does not modify or alter the original data in any way; it simply uses the original data to generate a hash value as a new data item. Common hash value output lengths include 128 bits (such as those produced by MD5), 160 bits (such as those produced by SHA-1), and even 512 bits or more (such as those produced by various members of the SHA-2 family). The output of a **hash algorithm** can be called a **hash**, a **hash value**, a message digest, a message authenticating code, a fingerprint, a digital value, or a checksum.

Hashing checks integrity. Hashing computes a hash value upon storage or transmission of a file or message, then computes another hash value at the end of the storage period or upon receipt of the transmission. If the before-and-after hash values are the same, then the message has retained its integrity. If the before-and-after hash values are different, then the message has lost its integrity and is no longer trustworthy or usable.

Hashing verifies integrity based on the complex mathematical functions used in the hashing algorithm. A central component or feature of these functions is the avalanche effect. This effect ensures that small changes in the input data produce large changes in the computed hash value. A single binary digit change in a file should produce a clearly recognizable difference in the resultant hash value.

Hashes are a one-way function. Hashes are fixed length values, regardless of the size of the source data. The hash will always be a certain number of characters, although the values of those characters

depend on the data that was hashed. A hash value cannot be reversed directly back into the original data file from which it was calculated. Thus, if an attack obtains a hash value, the data is not at significant risk of being extracted from the hash.

Some examples of hashing algorithms or systems include MD5, HAVAL, SHA-1, and the SHA-2 variants (SHA-256, SHA-384, SHA-512, etc.).

VPNs employ hashing, often as part of digital signatures, as a method to confirm that transmitted data has retained its integrity. Any mismatching of hash values causes the system to discard the received data and request a retransmission. This ensures that either end of a VPN connection accepts only fully valid and true data.

Establishing VPN Connections with Cryptography

Now that you know about symmetric cryptography, asymmetric cryptography, and hashing, let's take a look at the overall process VPNs use to provide secure communications. A basic VPN setup process occurs as follows:

1. Bob computes the hash from a message (**FIGURE 9-13**).
2. Bob encodes the hash using his private key to create a digital signature.
3. Bob generates a random symmetric key (K1).
4. Bob uses K1 to encrypt the message.
5. Bob obtains Alice's public key.
6. Bob encodes K1 with Alice's public key to create a digital envelope.
7. Bob sends the digital signature, the encrypted message, and the digital envelope to Alice.
8. Alice uses her private key to decode the digital envelope to reveal K1 (**FIGURE 9-14**).
9. Alice uses K1 to decrypt the message.
10. Alice computes the hash of the message.
11. Alice obtains Bob's public key.
12. Alice decodes the digital signature with Bob's public key to obtain Bob's hash value.




FIGURE 9-13 The overall cryptographic process of establishing a VPN connection (1/3).




FIGURE 9-14 The overall cryptographic process of establishing a VPN connection (2/3).

Alice compares the pre- and post-hash values. If the hash values are the same, then the message has retained its integrity. Thus, the correct public key decoded the digital signature. Bob's identity is verified (authenticity and nonrepudiation).

This exchange enables the symmetric key to secure subsequent messages from Bob to Alice or from Alice to Bob (**FIGURE 9-15**). Once a VPN has been established between Bob and Alice, an exchanged session key (K1) exists at both endpoints. Using this shared session key, Bob can encrypt messages (envelope #1) to send to Alice, and Alice can encrypt messages (envelope #2) to send to Bob.




FIGURE 9-15 The overall cryptographic process of establishing a VPN connection (3/3).

Variations on this basic VPN session establishment are possible. One variation would be to encrypt both the message and the digital signature with K1. Another variation would have Bob digitally sign K1 with his private key before using Alice's public key to encode it into a digital envelope.

This initial symmetric key exchange establishes the encryption that will protect all subsequent messages for the current communication session. When this session ends and a new VPN session begins, a new symmetric key will initiate the process. However, as long as the current session lasts, the same symmetric key applies.

Reusing a symmetric key poses a slight risk, even for multiple messages in the same session or conversation. To minimize or reduce this risk, some VPNs use **rekeying** processes. Rekeying discards the current in-use key and generates and exchanges a new symmetric key. Rekeying comes in several types:

- **Time rekeying**—Rekeying triggers at a specific time.
- **Idle rekeying (lag or delay)**—Rekeying triggers when a specific amount of idle, lag, or delay in the conversation occurs.
- **Volume rekeying**—Rekeying triggers when a specific amount of traffic is transmitted.
- **Random rekeying**—Rekeying triggers at random time intervals.
- **Election rekeying**—Either member of a VPN session can elect to rekey at any time.

If a VPN solution performs mid-session rekeying, you may or may not have an administrative configuration option to manage the

rekeying processes. Rekeying may be embedded in the encryption algorithms. **One-time pad** encryption systems are also possible. A one-time pad system uses a unique and random symmetric key for each segment of a communication. This is a more complex system, but it offers greater security. Attempting to crack a multi-random-key encryption scheme is one of the most difficult attacks against encryption known.

Digital Certificates

You can increase the reliability of authenticity and nonrepudiation by using **digital certificates** instead of plain public and private keys. A digital certificate is a public key and private key pair digitally signed by a **trusted third party**. This third party is a **certificate authority (CA)**. The CA first verifies the identity of the person or company, then crafts and issues the digital certificates.

■ NOTE

Technically, computers provide pseudo one-time pads, as they are currently unable to produce true random numbers. Instead, a **pseudo random number generator (PRNG)** is used. A PRNG uses a complex algorithm and the timing chip to produce seemingly random numbers. The result is very good, but not truly random.

In addition to improved reliability, certificates also resolve a scalability issue with public key cryptography systems. Without certificates, each system has to manage an ever-expanding library of public keys. With a third-party certificate-based system, each system no longer needs to retain public keys. Instead, keys are exchanged at the beginning of each session with the assistance of the CA.

Each host stores only the trusted public keys of the CA. The CA's public keys are stored in the trusted roots list (TRL). Any certificate issued by a trusted CA is accepted as valid, providing the certificate has not expired and has not been revoked.

The process is similar to the overall cryptographic process to establish a VPN connection. The primary difference is that instead of a sender's generic private key, the sender's digital certificate encodes the message. Then, the recipient uses the CA's public key to start the unpacking process. The CA's public key decodes the CA's private key encoding around the sender's public key. (This is the sender's digital certificate.) This confirms the identity of the sender through the CA's issued digital certificates. Then, the recipient uses the sender's public key to decode the sender's private key encoding on the message.

The use of digital certificates adds a few additional steps to the overall process, but this improves the identity verification of the participants in the secured communication. Once a communication session ends, the recipient can discard the sender's public key. Each time a new session starts, the sender's public key is re-exchanged via the digital certificate process. This reduces the burden of public key management and makes any secure communication service, including VPNs, much more scalable.

WHAT IS VPN AUTHENTICATION?

Authentication is the process of confirming or proving the identity of a user and is of significant importance for VPNs. Because VPNs allow external entities to connect to and interact with a private network (or system), verifying the identity before granting access is paramount. VPN authentication takes place on two levels: connectivity and user.

When a VPN link starts, the hardware and software components at each endpoint must authenticate to establish the connection. Once the communication link begins, the user performs a separate and distinct authentication process. If either of these identity proofs fail, the VPN is severed.

The actual mechanism of authentication used in either case can vary. Options include username and password, smart cards, token devices, digital certificates, and even biometrics. As a rule of thumb, you should avoid username- and password-only solutions and instead use multifactor authentication options. Password-only authentication is notoriously exploitable and is not reliably secure enough for a VPN.

The VPN services, whether software or hardware, may support authentication directly or indirectly when offloaded to dedicated authentication servers. Offloading can point to any number of widely used authentication, authorization, and accounting (AAA) services or

technologies, including RADIUS, TACACS, 802.1x, LDAP, and Active Directory.

When selecting an authentication solution for a VPN, consider the strengths and weaknesses of the authentication factors independently of their VPN integration. If an authentication factor has weaknesses on its own, these are not alleviated when used with a VPN.

VPN authentication should be scalable and support interoperability among potential connecting hosts. Do not use authentication that cannot support more than a few dozen users or hosts. Likewise, using a form of authentication that is available on only one platform or operating system will impose limitations difficult and expensive to resolve later.

WHAT IS VPN AUTHORIZATION?

Authorization is controlling what users are allowed and not allowed to do. Authorization is also known as access control. You must establish clearly defined policies as to what activities will and will not be supported for VPN connections.

An initial concern is defining the users allowed to establish a VPN connection. By considering VPN a resource rather than just a method of connectivity, clearly not all users should have access to all resources. If VPN connectivity is neither essential to nor conducive of a worker's assigned tasks and responsibilities, they should not be allowed to open or use VPN connections.

Authentication can help enforce this access control issue. Whenever a user account that is not authorized to use VPN attempts to authenticate across a VPN, that attempt should fail. In fact, automatically locking such an account is a reasonable security response. Users should have a clear understanding of whether or not they have been granted the privilege of VPN use. Any access of a non-VPN-approved user across a VPN, therefore, is either a sign of a policy-violating employee or an outside intrusion attempt.

In addition to locking down use of the VPN itself, you should restrict access to resources across a VPN, even for authorized users. A VPN does establish a network connection indistinguishable except for speed from a local direct-wired connection, but that does not mean that the same level of authorization is necessary, required, or even recommended.

A stronger application of the principle of least privilege is needed with VPN connectivity. Not all of the tasks and resources employees use when on site are necessarily required for activities they perform remotely. A user should have one level or sphere of access when on site and a different, smaller sphere of access when connecting through a VPN. This, of course, will depend upon the assignments of the worker and the sensitivity of the resources needed to carry out those assignments, but thoroughly consider this issue before granting default or wholesale access.

Mission-critical resources, processes, and information may be at risk if exposed to users connecting over a VPN. Because a remote user's computer is potentially less secure than an organization's workstations on site, the additional risk may be significant. Does a telecommuter or traveling worker actually need access to a specific resource? If not, block access.

Should a remote VPN user be able to transfer sensitive, private, or valuable data to the remote host? This is the risk of information leakage. Once data leaves the LAN, the corporate security infrastructure loses its ability to fully control and protect the information. On a remote host, once the VPN disconnects, the LAN's security protections disappear. The only remaining security left is on the host and practiced by the user.

Remote access is inappropriate or just impractical for some resources. For example, consider preventing printer access by remote VPN users. Do traveling workers need to print documents they are unable to physically handle? Yes, in some situations, remote printing makes sense. But should you allow every local activity over a VPN just because it is possible?

Another authorization concern is whether or not to allow Internet communication for VPN users. If local LAN users are able to access the Internet, then technically so can VPN users. But should they? Is this an additional risk or just an additional bandwidth consumption burden?

If VPN users are unable to access the Internet from the LAN through the VPN, can they access the Internet simultaneously over the same ISP link as their VPN link to the LAN? This configuration is known as a **split tunnel**. Many organizations see this as a significant risk. This scenario makes it possible for an Internet attack to breach the remote host and then use the VPN to access the private LAN. This would be a nearly unrestricted pathway between the Internet and the LAN. Most organizations employ VPN connection solutions that prevent simultaneous local VPN and Internet connectivity on remote hosts.

Even with these VPN authentication issues, the overall process of enforcing access control is essentially the same for both local and remote hosts. In most cases, access control is defined on the resource itself, regardless of where the users are located. Keep this in mind as you craft policies defining the parameters for the use and management of VPN.

CHAPTER SUMMARY

VPNs are mechanisms to remotely connect LANs and mobile hosts. VPNs can save organizations money and offer worker flexibility. VPNs secure communications over public intermediary networks. VPNs should be used whenever infrastructure costs are high or you need mobile access and flexibility. VPNs cannot guarantee quality of service over the Internet or protect against endpoint device/host vulnerabilities. You should deploy VPNs based on thoroughly researched VPN policies.

VPNs offer many advantages, including a wide range of implementation choices. Various devices and software products support VPN connections. VPNs' endpoints can be inside, on, or outside a firewall. VPNs can support remote access for mobile hosts, create links between individual systems, and support channels between networks. VPNs can employ either tunnel-mode or transport-mode encryption. VPNs can use the same authentication and authorization techniques that you deploy elsewhere for general network security.

KEY CONCEPTS AND TERMS

Algorithm

Asymmetric cryptography

Authenticity

Avalanche effect

Certificate authority (CA)

Channels

Ciphertext

Client-to-server VPNs

Compression

Corporate firewall

Cryptography

Decryption

Dedicated connection

Dedicated leased lines

Digital certificates

Digital envelope

Digital signature

Distributed LAN

Edge router

Encryption key

Extranet VPN

Hardware VPN

Hash

Hash algorithm

Hash value

Hashing

Host-to-host VPN

Host-to-site VPN

Host VPN

Hybrid VPN

Identity proofing

Intermediary network

Key exchange

Key pairs

Keyspace

LAN-to-LAN VPNs

Leased line

Nondedicated connection

One-time pad

One-way function

Optical carrier (OC)

Out-of-band communication

Private branch exchange (PBX)

Private key

Pseudo random number generator (PRNG)

Public key

Public key cryptography

Public networks

Rekeying

Remote access VPN

Remote-to-home VPN

Remote-to-office VPN

Scalability

Secured VPN

Site-to-site VPNs

Software VPN

Split tunnel

Symmetric cryptography

Telecommuting

Traffic congestion

Trusted third party

Trusted VPN

VPN appliance


WAN VPN

CHAPTER 9 ASSESSMENT

1. Which feature in tunnel-mode encryption is not supported in transport-mode encryption?
 - A. A footer is added to contain the hash value.
 - B. The payload is encrypted.
 - C. The source address is encrypted, but not the destination address.
 - D. The header is encrypted.
2. What are the two most important characteristics of VPN authentication?
 - A. Single factor and replayable
 - B. Scalable and interoperable
 - C. Transparent and efficient
 - D. Replayable and scalable
3. What form of cryptography encrypts the bulk of data transmitted between VPN endpoints?
 - A. Symmetric
 - B. Hashing
 - C. Public key
 - D. Asymmetric
4. Of the following VPNs, which prevents filtering of VPN traffic?
 - A. Edge router
 - B. Extranet VPN
 - C. Corporate firewall
 - D. Host-to-site VPN
5. When designing the authorization for VPNs and VPN users, what should you use as the primary security guideline?
 - A. Distributed trust

- B. Multifactor
 - C. Principle of least privilege
 - D. Grant by default, deny by exception
6. Which of the following should be completed prior to building a VPN policy?
- A. Risk assessment
 - B. Patch management
 - C. Financial audit
 - D. Awareness training
7. Which VPN access control issue can be enforced through VPN authentication?
- A. Controlling access to printers
 - B. Restricting access to the Internet
 - C. Limiting access to files
 - D. Blocking unauthorized VPN users
8. Which form of VPN deployment requires additional authentication for accessing resources across the VPN?
- A. Corporate firewall
 - B. Site-to-site VPN
 - C. Host-to-site VPN
 - D. Edge router
9. Which of the following is the primary difference between a VPN and a LAN connection?
- A. Authentication factors
 - B. Resource access
 - C. Speed
 - D. Security
10. Which of the following is a limitation of deploying a VPN?
- A. Intermediary networks are insecure.

- B. VPNs cannot run on native operating systems.
- C. It is not possible to have a trusted VPN.
- D. Vulnerabilities exist at endpoints.



© Digital_Art/Shutterstock

VPN Management

VIRTUAL PRIVATE NETWORK (VPN) MANAGEMENT is a critical component in your organization's computer security. VPNs extend the internal network beyond the perimeter secured by firewalls and other security technologies. Proper management of a VPN requires an understanding not only of VPN technologies, but also of the business requirements driving VPN implementation.

Using the wrong type of VPN—or a VPN that does not meet your organization's requirements—can create security problems where none previously existed. Before selecting a VPN product or technology for your organization, create a detailed requirements document. This should factor in the security requirements for the VPN and business requirements. Although security is important, keep in mind that security exists to support your organization's goals, not the other way around.

As soon as you have a set of prioritized requirements, start looking for a solution that meets as many of these requirements as possible. You may not be able to find a single product that meets all

requirements. When that occurs, choose the solution that best fits the highest priority requirements. In some cases, using multiple solutions to meet all of the requirements may be the right solution.

Research is crucial to selecting the appropriate VPN. Do not rely solely on websites, technical magazine reviews, or an attractive sales pitch. Depending on your requirements, you may want to explore robust public domain solutions rather than an off-the-shelf commercial product.

Keep in mind that VPNs are a security technology, but they may also degrade your security perimeter. Putting a VPN in place is just the beginning of the project; you'll need to address a number of additional factors for a successful VPN deployment.

CHAPTER 10 TOPICS

This chapter covers the following topics and concepts:

- VPN management best practices
- How to develop a VPN policy
- How to develop a VPN deployment plan
- Common VPN threats and exploits
- Trade-offs are between commercial and open-source VPNs
- Differences between personal VPNs and network VPNs
- How to balance anonymity and privacy

- How to protect VPN security to support availability
- The importance of VPN user training
- How to troubleshoot VPN issues

CHAPTER 10 GOALS

When you complete this chapter, you will be able to:

- Describe VPN best practices
- Write a VPN policy
- Describe the issues involved with deployment, placement, and implementation of a VPN
- Appraise the threats and attacks against VPNs
- Contrast the needs and features of personal and enterprise or network VPNs
- Compare anonymity and privacy
- Compose an introductory VPN training program for users
- Formulate a procedure for troubleshooting VPNs

VPN MANAGEMENT BEST PRACTICES

Before you begin a VPN deployment, you should familiarize yourself with the recommendations, guidelines, and procedures that will allow you to manage your VPN securely, efficiently, and as cost effectively as possible. These techniques and recommendations are known collectively as “best practices” and are generally developed and published or shared by experts in the field. VPNs have been around since the late 1990s, which means that a significant amount of real-world experience is available in the industry for your reference.

A best practice is generally not a tool, but rather the collected wisdom of fellow security practitioners sharing what they have learned. One of the great things about working in the information security field is that numerous experts are generally willing to share their experience. Be warned, however: Security experts tend to have strong opinions, so any time you are reviewing a best practice, be sure to keep the needs and requirements of your own environment in mind. A process that works great in a 25,000-employee global manufacturing company may not work nearly as well in a 50-employee medical records processing company.

You may be able to adapt a process from another company to fit your environment. The key to getting the most out of best practices is to consider each, keeping your specific environment and requirements in mind. Keep what works, modify what you can adapt, and ignore what doesn't make sense.

Build in Redundancy

Provide redundancy, because everything, including VPNs, can break. If your organization will rely on your VPN for remote access, encrypting and securing data, or providing a business partner access to your extranet, you will find out quickly how critical your VPN is—on the day that it breaks. Most commercial VPN products offer a failover or load-balancing capability so that in the event one device fails, the other will pick up the traffic.

Often, VPNs are provided via a cloud implementation. In this case, be sure that a level of redundancy is built in and that the redundancy is clearly spelled out, and guaranteed, in the contract or **service level agreement (SLA)**; it will otherwise be out of your direct control. Although passing the burden for security in a cloud implementation may seem efficient, if the service provider does not actually provide the level of service and expertise needed, it can be disastrous. Corporate legal counsel should review all SLAs with the IT management and administrative staff to ensure that what is promised is the level of support required.

Alternatively, keep a spare VPN product on your shelf, configured and ready to go live in the event of a failure. Generally, waiting for tech support or ordering a spare part can take more time than your organization is willing to wait for restored service. This level of redundancy and preparedness is often referred to in the corporate disaster recovery plan. You will learn more about this later in the chapter.

Choose the Right VPN Product for Your Environment

Choosing the right VPN product is critical to the long-term success of your VPN deployment. Take your time, document your requirements, carefully evaluate the capabilities of each VPN product you review, check with peers if available, and review appropriate industry literature. A security magazine review of VPN products is a valuable tool for starting your search—or even narrowing the field once you have your requirements documented—but do not select your solution based solely on which product won last year’s Editor’s Choice award from your favorite industry magazine. The editors and reviewers at most industry magazines, websites, and blogs are technically capable and spend a lot of time looking at products, but they do not have to support the product in your environment. Ultimately, you will be responsible for meeting your organization’s requirements and maintaining support.

Avoid purchasing products based on promises rather than proven capabilities. **Slideware**, also sometimes known as vaporware, is any product that appears in a vendor’s PowerPoint presentation but is not yet available for sale. When possible, road-test a product in your environment before purchasing. It is always a good idea to see what a product can do firsthand.

Finally, when looking for a product, consider using resellers and consultants. A good reseller can do some of the legwork for you by narrowing the search to a smaller pool of products. Although some resellers work much like a standard department store, selling whatever is on the shelves, many will go the extra mile to ensure that you get a product that will work in your environment. This saves

them the challenge of trying to support a poor product after they have sold it; if you like the solution, resellers have the opportunity to sell you additional products in the future.

Develop a VPN Policy

A VPN policy (often referred to as a remote access policy) documents your organization's rules for using the VPN. You will read about creating a VPN policy at length later in this chapter. Recognize, however, that proper policy framework is a key best practice when dealing with security technologies, especially those offering remote access to your computing environment.

Prohibit Split Tunneling

Split tunneling is a configuration setting that allows simultaneous access to both an untrusted network (like a home network) and a secured VPN network connection. This may not sound like a bad idea at first—after all, why wouldn't you want someone connected to the VPN to access the Internet or their home network (or another network) at the same time? The reason split tunneling is a bad idea is that it potentially opens a door into your network that you cannot control. This is known as **hairpinning**, because malicious code can enter from the nonsecure network, make a hairpin (or sharp) turn, and enter your secure network with little or no trouble because it is entering from a secure and verified endpoint.

If the client machine is compromised by a virus that permits remote control of the system by an attacker, and that client machine connects to the VPN, the attacker will have access to your internal network from anywhere on the Internet. If you prohibit split tunneling, however, then even if the attacker can compromise the client, the external connections terminate as soon as the VPN connects, ensuring that your network is secure (even if the client is not).

Ensure Client Security

For the most part, VPN technology is both mature and secure. VPNs are subject to denial of service (DoS) attacks, but VPN servers are rarely hacked. In October 2019, the private VPN provider NordVPN confirmed rumors that it had suffered a hack. The vulnerability was due to the expiration of the internal private key, leaving the connection susceptible to spoofing by hackers spinning up their own servers to pretend to be NordVPN servers. Nevertheless, VPNs remain a viable vector for someone who wants to attack your network. The target of these attacks is typically the weakest link of the VPN chain—the client system.

One of the challenges of working with remote access VPNs is making sure that the client at the other end of the connection is secure. Remote access VPNs permit access to a secure network from a remote location across an untrustworthy network. If the client system is not secure, you run the risk of compromising your secure network. A typical VPN client runs an operating system (OS) that needs patching at least once a month; runs applications that may need patching almost as often; and is vulnerable to viruses, spyware, and other attacks. Be sure to install antivirus software, anti-malware software, and a software firewall on every client that will connect to your network through the VPN.

One school of thought is: If it does not belong to the company, it should not connect to the company's network. Fortunately, you have some control over your organization's computers. You can require antivirus, anti-malware, intrusion detection systems (IDSs), and firewall software on any computers your company owns. For systems not owned by the organization, you can require that the equipment

meet specific security standards, including running an antivirus test prior to allowing access. This is a complicated process and may be beyond the scope of some system administrators. As a result, if you do not prohibit noncompany systems from connecting to the network, you cannot control whether those systems are secure; that is, you cannot ensure that your network remains secure. It is very easy for an end user to load a VPN client on a home PC and connect to your network. Although some VPN solutions offer techniques to prevent uncontrolled systems from successfully connecting, it is important that you lay the groundwork by prohibiting the practice.

This is a problem that must be dealt with if the organization has a Bring Your Own Device (BYOD) policy and a virtual private network. How should BYO devices be admitted to the VPN, and what are the policy and technical guidelines for connection? How much security is the organization willing to give up to gain the financial and operational benefits of BYOD? In many cases, cloud-based providers will offer support for a wider variety of client OS types and versions than a traditional VPN product. Maybe it is time to consider a cloud-based solution. These are all factors your organization must consider when considering your VPN and the security, or insecurity, of your information.

Practice Vulnerability Management

Effective **vulnerability management** can help manage your remote clients. These are the technology and business processes used to identify, track, and mitigate known weaknesses on hosts or applications within a computing environment.

Remember, everything is vulnerable to attack. UNIX, Windows, routers, network printers—and even your VPN solution—will have vulnerabilities. Examples of vulnerabilities include software coding errors, improper configurations, and poor password choices. The danger with a VPN is that it expands your network from systems you can closely control to include systems in a home office, a branch office, a hotel, a coffee shop, or even a business partner's network.

Vulnerability management is a combination of tools and processes that allow you to reduce risk in your computing environment, including VPN-connected systems and networks. Use tools that periodically test your environment, including the VPN systems, for missing patches, configuration issues, known exploits, and other vulnerabilities. This will ensure that your remote systems as well as your local systems are secure. Scan often and address issues when found.

TECHNICAL TIP

When addressing vulnerabilities in your environment, be efficient. If you are using a tool that ranks vulnerabilities from Level 1 to Level 5, where Level 1 is informational and Level 5 is critical, you might be tempted to address the Level 5s first, and then come back and address Level

4s next, then Level 3s, and so on. But this would be similar to having a keyboard with several broken keys. You wouldn't fix the "E" key first, then come back later to fix the "U" key because it's less critical—you would fix them all at once. Take a similar approach with vulnerabilities. Determine what level of vulnerabilities you want to address in your environment (all Level 3 and above, for example); then tackle them, system by system. This approach may leave some Level 5s in the environment a little longer, but it will take a lot less time to secure the entire environment, since you will not need to touch the systems multiple times. It can be more efficient, as the lower priority areas may be under attack.

Use Multifactor Authentication

Your VPN is only as secure as your authentication method. One of the easiest ways to compromise a VPN is by compromising the authentication credentials. All it takes is one user with a password of “password” to open a direct connection to your network. A best practice is to use multifactor authentication for VPN access. This is a method of proving identity using several different authentication factors. Authentication factors are something you know, something you have, or something you are. Examples include a smart card (something you have) with a PIN (something you know); a biometric device (something you are) coupled with a password (something you know); or a proximity card (something you have) that activates a fingerprint reader (something you are). The most common form of multifactor authentication is **two-factor authentication**, which uses two authenticators (most commonly the user ID/password and a texted number or code to be entered). Each additional factor added increases the security exponentially but also increases cost and complexity.

Document Your Implementation Plan

If you fail to plan, you plan to fail. Document your implementation plan. You cannot simply find an open rack in the data center, run a cable, plug in a device, and keep your fingers crossed that it will work. Document your implementation and support plans. You will learn more about this later in the chapter.

Monitor VPN Availability

Monitoring the availability of the VPN can save time and aggravation. The typical (and worst) method to discover issues with your VPN is when users start calling for help. This can be particularly challenging when the callers include members of your organization's senior management team. A corollary of Murphy's Law may be, "The likelihood that a senior management team member will try to access the network over the VPN is directly proportional to the likelihood that the VPN will fail." It is far better to alert senior management to temporary problems than for them to tell you that something is not working. Because VPNs are network components, you can generally use the same monitoring equipment that you use to monitor routers, switches, and other network gear.

■ NOTE

Vendors can be very helpful when your equipment is down. Many technical and information security professionals believe they can fix anything, and sometimes they will work on a problem for long hours before finally reaching out to tech support. Although tech support can be problematic, especially when wading through the first level of support, the vendor can be invaluable when the problem is in the VPN equipment or software.

Perform Regular Reviews, Backups, and Updates

Once you have your VPN deployed, review usage regularly. If you notice there are employees who do not use the VPN, you may want to remove their access. If you see employees who have multiple concurrent connections, you may have a security issue and should investigate further. Suspend access for anyone who will be unavailable or not allowed to access the VPN due to corporate travel outside the country or vacation.

Back up your VPN configuration regularly. This is a good practice for any network equipment, but in the event that your VPN hardware fails and needs replacement, you will need to be able to restore your last known working configuration quickly. Rebuilding a VPN configuration from the default settings can be a long and challenging task—not to mention making post-incident review meetings an unpleasant experience.

Patch regularly. Vendors typically release patches and updates to VPN code throughout the life of the product. These patches address security issues, fix bugs, or provide additional functionality. In an ideal environment, you will have a development VPN that you can use to test patches and updates. In most environments, however, you will not have the luxury of a development VPN and will have to test when you implement in production. In either circumstance, work closely with your vendor to make sure that you receive prompt notice of patches and updates. Establish an operational process and maintenance window to apply patches and updates in a timely fashion.

Your VPN solution may end up being a critical component of the organization's business-continuity planning and disaster-recovery planning. In the event of an incident that prevents employees from getting to their work location, a VPN that provides work-from-home support is a key component of many recovery plans. Events such as earthquakes, snowstorms, tornadoes, floods, pandemics, and other natural disasters can make working remotely a viable alternative to standard operations.

This collection of VPN management best practices is meant to serve as the starting point for your successful VPN deployment. It is not meant to be comprehensive, but instead to offer some common practices and processes to help you with your VPN deployment. Depending on your VPN solution—as well as your environment, business requirements, and experiences—you may use every one of these suggestions or only a few of them. The key is to ensure that you are doing what works in your environment. If you need help, do not be afraid to reach out to your peers for advice and suggestions. Over time, you will develop your own set of best practices as you gain more experience as a security practitioner, and soon others may be asking you for advice.

DEVELOPING A VPN POLICY

If you are implementing or supporting a VPN solution, use a VPN policy to ensure that your users understand the requirements for computing on the VPN. A VPN policy is sometimes called a remote access policy, a term used when dial-up lines and modems were the primary means to access the network remotely.

Keep in mind that a VPN policy should be a part of your overall policy framework, not a standalone. If you try to develop your VPN policy in isolation from the overall policy framework, you may find that you are duplicating information, or potentially writing VPN policy that conflicts with other aspects of your overall security policy framework. For example, if you put a requirement in your VPN policy that user passwords must be 15 characters long but the password policy states that passwords have to be eight characters long, you will confuse end users.

The components of a solid VPN policy include:

- **Introduction**—State the policy by name and tell how it fits into the organization’s policy framework.
- **Purpose**—Describe the issues the policy addresses and how the policy should be used. Include references to any applicable governance, risk, or compliance issues, as well as any specific legal or regulatory requirements supported by the document.
- **Scope/binding nature statement**—Describe the systems, networks, or people covered by the policy, and outline the penalties for not following the policy. The phrase “disciplinary action up to and including termination” is common in security policies.

- **Definitions/acronyms**—Define technical terms or acronyms used in the policy. Readers of the policy may not be familiar with the terms or jargon used.
- **Document**—Include the document creator, creation date, version, document status (e.g., draft, template, policy, and guidelines), as well as any version tracking information.
- **Policy**—This is the actual policy language. Be very clear in this section, leaving as little open to interpretation as possible.
- **Optional elements**
 - **Summary**—If your policy is long, you may want to summarize it in a bulleted list at either the beginning or end of the policy. This provides employees a quick method to check for policy statements.
 - **Roles and responsibilities**—If your document is lengthy, or you need to document who does what under the policy, include roles and responsibilities. For example, a policy dealing with infrastructure might include roles for the system manager, system architect, end user, developer, or other key people within the organization.

Some specific topics to include in your VPN policy are:

- Restrict remote access to the organization's VPN solution.
- Prohibit split tunneling.
- Define which classes of employee can access the network by VPN. This could include regular employees, vendors, contractors, and temps, or it could be restricted to only home office workers, depending on business requirements.
- Define which types of VPN connections will be permitted.
- Define authentication methods permitted.
- Prohibit sharing of VPN credentials.

- List the configuration requirements for remote hosts, including current virus protection, anti-malware, host-based intrusion detection system (HIDS), and a personal firewall. Some VPN solutions include the ability to check for these types of configurations.
- Prohibit the use of noncompany equipment or, if personal systems may connect to the VPN, define the minimum standards for those connections.
- Define required encryption levels for VPN connections.
- If you will be using your VPN for network-to-network connections, define the approval process and criteria for establishing a network-to-network connection.

Have your policy reviewed and approved by your communications, legal, and human resources departments before release. Document the appropriate approvals in the document status portion of the policy, then communicate the policy to the employees. Posting the policy to an information security or security policy intranet website is a common practice. Once it is available on the intranet, you can use standard communications methods to make employees aware of the policy requirements. These methods can include email, a structured awareness program, inclusion in new-hire training, or even web-based or in-person policy training. The method you select will be based on your organization's size, locations, and requirements. The key with any type of awareness training is to structure your communications to the correct audience. A group of engineers will require a very different introduction to a technical policy than a sales team.

These are just some of the factors in developing a VPN policy for your organization. Although these should form the basis for the

development of your organization's policy, be sure to cover all applicable requirements. One sure way to alienate your employees is to release a policy that makes no sense to them or that you have to revise too soon to cover things you didn't think of in the first version. Take your time, consider all the requirements, and you will end up with a usable VPN policy.

DEVELOPING A VPN DEPLOYMENT PLAN

Now that you understand best practices and VPN policy, it's time to look at how to select, place, and use a VPN.

Consider these criteria during your search:

- What types of VPN connections does the solution support (user access, site-to-site, or both)?
- Which encryption protocols are supported?
- How many VPN connections are supported?
- How well does the VPN perform compared with a high-speed network?
- How well does the VPN interoperate with your existing network infrastructure?
- What are the support options available from the vendor?
- How easy or difficult is the VPN to set up?
- What are the management capabilities available with the VPN?
- What additional features does the VPN offer?
- Does the VPN support failover or high availability configurations?
- How scalable is the VPN?

Once you have gathered this information, you can start to narrow your selection based on how well these features meet your criteria.



Before selecting and purchasing a VPN solution, consider a number of factors. One that is too frequently ignored is the budget for the solution. Carefully consider your criteria and evaluate your options, always keeping in mind that you need to live within your budget. When budgeting, be sure to look not only at acquisition costs, but also at the ongoing support and maintenance costs for at least the first three years.

Next, consider where you will deploy the VPN on your network. Three common architectures typically accompany VPN solution deployments. The three deployment types are bypass, internally connected, and DMZ-based.

Bypass Deployment

A bypass architecture (**FIGURE 10-1**) deploys the VPN so that traffic to the VPN and from the VPN to the internal network is not firewalled in any way. This is referred to as **bypass deployment**.




FIGURE 10-1 A bypass VPN implementation.

This was a common architecture when VPNs were first introduced to the market. The logic behind this deployment architecture is that security is adequate without additional firewalling because the VPN accepts only encrypted connections on a specific port. The assumption is that because traffic is encrypted, it does not require additional protection. Even if you did place a firewall on the Internet-facing VPN connection, the firewall would not be able to analyze the encrypted VPN traffic. A bypass architecture also considers anyone connected to the VPN as a trusted host. Finally, passing traffic across a firewall always causes concerns about performance impacts, and justifiably so.

Two significant issues arise with this implementation model. First, VPNs are like any network device; each can have a variety of vulnerabilities. As a result, VPNs are still vulnerable to an attack against the device itself. Second, the uses for VPN have expanded greatly since the technology first appeared, and it is not uncommon in today's environment to leverage a VPN to provide untrustworthy hosts access to portions of the network. These could be customers accessing an order management portal, vendors supporting systems on the internal network, or suppliers who are transferring invoices to your internal systems. Although some controls (usually routing table-related) are available in most VPN solutions, terminating a VPN directly on your internal network is a dangerous practice. As a result, the circumstances where a bypass VPN architecture is still a viable solution are limited, unless your risk tolerance is fairly high.

Internally Connected Deployment

An internally connected architecture (**FIGURE 10-2**) deploys the VPN so that traffic to the VPN and from the VPN to the internal network is not firewalled in any way.




FIGURE 10-2 An internally connected VPN.

An internally connected VPN architecture recognizes that the VPN is vulnerable to attack if placed directly on the Internet, so it places the Internet-facing VPN connection behind a firewall. Selecting the appropriate firewall for this solution is critical to a successful implementation. The VPN then connects any remote users or site-to-site connections directly to the internal network.

This is an improvement over the bypass architecture, but it still does not address the potential security issues associated with untrustworthy VPN connections. This architecture is not recommended, although it will work if the only hosts connecting are trusted hosts.

DMZ-Based Implementation

A demilitarized zone (DMZ)–based implementation (**FIGURE 10-3**) addresses the main shortcomings of the previous two architectures.




FIGURE 10-3 A VPN implemented in a DMZ.

This architecture features a firewall in front of the VPN to protect it from Internet-based attacks, as well as a firewall behind to protect the internal network. The firewall on the inside can be configured to protect important infrastructure like finance department or research department servers, restrict business partners to only the systems to which they need access, or limit vendor access to only those systems supported.

The largest negative in this design is the cost of deploying multiple firewalls for the implementation. However, many companies already have DMZs set up in this configuration, so it may be just a matter of leveraging the existing infrastructure for your needs.

After selecting the ideal VPN and determining where you are going to place it in your infrastructure, it is time to work up your deployment plan.

Before you look at the components of a successful deployment plan, keep in mind that your environment is unique to your organization. Some elements you will read about will not apply to your deployment. These are, rather, just some common elements found in most typical VPN deployment plans. Just as with the other topics discussed here, review these while thinking of the requirements of your organization.

To deploy your VPN successfully, you need to:

■ **NOTE**

You can use a project management application to plan your VPN rollout, or use something as simple as a word

processing or spreadsheet application. The purpose of this plan is to allow you to structure your deployment more formally than a plan written on the back of a cocktail napkin.

- Plan the physical location of the VPN. This is commonly rack space in your data center.
- Ensure that your selected location meets power and cooling requirements. Get information on power and cooling requirements from your VPN's technical specifications.
- Plan your Internet Protocol (IP) addressing for the external and internal network connections on the VPN, as well as a pool(s) of addresses assigned to clients when they connect to the VPN. Plan for peak usage when assigning these pools to the VPN. Although an average number of users is an excellent use benchmark, peak use determines the maximum number of IP addresses to ensure that all users can connect. If this is a site-to-site connection, your IP addressing plan will not be as complex, but it will still be important when setting up the tunnel.
- If you are using firewalls as discussed in the previous section, plan the rules you will need on the firewalls to permit the VPN to work. Generally, IPSec VPNs will require UDP port 500 for the IKE packets and TCP port 443 for the IPSec traffic. SSL VPNs use port 443 exclusively. Most VPN rule sets permit ICMP packets from the Internet. If you are experiencing issues with the VPN, it is frequently useful during troubleshooting to determine whether the user can reach the VPN server using tools like ping or tracert.
- Configure the VPN server by setting up the IP address pools, assigning IP addresses to the interfaces, establishing your banner message, and disabling split tunneling. It is also a good idea to have the vendor review your configuration before going live to ensure that you have not missed anything in your planning. In

some cases, you might want to have the vendor install the VPN for you while you concentrate on client rollouts, policy communications, and other related tasks.

- Set up the authentication mechanism. Ideally, this will be a token-based authentication solution, but it could be RADIUS-based authentication or, in some smaller environments, user accounts set up on the VPN.
- Follow your organization's change management policies when deploying the VPN. If your organization does not have a formal change management process, be sure to inform management of your planning schedule to avoid surprises. The last thing you want is to bring the VPN online during the same weekend the accounting department is doing a major upgrade to your organization's financial systems. If there is a problem with their upgrade, you can count on at least one person blaming it on your VPN rollout.
- Once the VPN is on the network, create a pilot group to test the deployment. Address any issues before you roll the deployment out to a larger user pool. Minimize any issues that employees might face with the new solution. The best way to sabotage a new solution is to deploy it with lots of problems. You may find that some employees have a low tolerance for issues.
- Develop your operations manual, which will document the configuration, operations procedures, change planning processes, and change history.
- Develop your user documentation. Include how to install the client (if a manual install is required), how to log on, the number to call for support, frequently asked questions, and any other information you think will make the user's experience with the VPN easier.
- Develop your support processes. Who gets the first call when there's an issue? What is the escalation path if the issue cannot be resolved?

- Communicate the rollout plan to management and affected employees. Let them know the timing, benefits of the new solution, and anything else that you believe will ensure a successful deployment.
- Install the VPN client for any remote access users. If this is a site-to-site deployment, configure the tunnel connection.
- If using a clientless VPN, configure the software.
- Distribute authentication credentials, tokens, and so on.
- Train your users.
- Go live—and enjoy a successful VPN rollout.

You now should be able to select, plan, and successfully deploy a VPN for your organization.

VPN THREATS AND EXPLOITS

Consider the threats and possible attacks against your VPN and how to mitigate each one. A VPN can be a critical component of your information security infrastructure. A properly implemented VPN addresses a number of common attacks against your infrastructure, including eavesdropping, man-in-the-middle, replay, and others. However, while your VPN can mitigate attacks, it also opens up an entirely new set of security issues.

Remember, VPNs are essentially network devices and subject to many of the same security issues you would find in a router or a switch. If you are running a software VPN, then your VPN can have many of the issues you find on your network's servers.

A hardware VPN solution can suffer from a number of security vulnerabilities, including:

- Weak default password
- Insecure default configuration or misconfiguration by the installer

One of the most common and easily exploited vulnerabilities on any hardware network device is the default password. Vendors set an initial password on their equipment. Usually, all it takes to discover this password is a quick search on the Internet; often, it is not particularly hard to guess. Try the vendor name, “admin,” or “password,” and odds are good that you will be able to log on. It seems like a sensible mechanism—the vendor does not want to

distribute its equipment without a password, so it sets a standard password that is easy for the installer to remember.

The potential problem occurs when the installer forgets to change that password. It is not uncommon for an installer to leave the password in place for the duration of the installation. Typing in “admin” after each reboot is much easier than typing in “\$\$Th!s!sAS3cur3P@ssw0rd!!” every time you make a change requiring a reboot. The problem occurs when either the installer forgets to change the password when the work is done, or an attack occurs while the installer is in the middle of the installation.

Imagine that the installer has a few additional settings to change, but it’s 6:00 p.m. on a Friday and he’s ready for the weekend. He decides to come in early Monday to finish the configuration. For the entire weekend, the VPN that you are counting on to provide secure access to your network is on the Internet with the password “admin.” The best way to address this type of vulnerability is through stringent system configuration procedures and strong awareness training for your support staff and contractors. Disciplinary action when an installer fails to follow the instructions is also remarkably effective at getting the message across.

The second common issue with hardware VPNs is a device that is installed in the default configuration. Very few network devices come out of the box in a fully secure configuration. For example, an important configuration setting is disabling split tunneling. A default VPN configuration might not disable that setting; instead, it may require modifying the configuration to disable the feature. If you are completing your first install, it is tempting to just change enough of

the configuration to get the VPN up and running and avoid any unfamiliar settings.

A related issue occurs when an inexperienced installer modifies the configuration without understanding the impact of the changes. For example, an installer with a limited understanding of encryption protocols might decide that using the Data Encryption Standard (DES) algorithm for VPN encryption would be “good enough” and would improve performance over a higher encryption algorithm like 3DES. The installer probably read somewhere that the longer the key length used for encryption, the higher the impact on performance. This is often true, but longer key lengths usually mean more secure communications. Installer-induced security threats are some of the hardest to track down. Because they were caused by someone with a limited understanding of what they were doing, the installer’s ability to help you track down the issues is equally limited.

To mitigate the risk of these issues, make sure that you train your installer (or yourself) before installing the VPN. If you do not have the time or justification for training on the product, then engage a vendor or systems integrator to assist with the install or even complete it for you. Once the installation goes online, it pays to have an expert perform a vulnerability or penetration test against your VPN. It never hurts to have an expert check your work.

A software VPN solution can suffer from the following:

- OS vulnerabilities
- OS misconfiguration
- Application conflicts
- Instability

- Viruses and malware

Although software VPN solutions are not typical in corporate environments, they are common in smaller organizations, academic environments, and other areas where a hardware VPN does not meet business requirements or exceeds the budget. The main threats to software VPN solutions arise in the OS. Operating systems like UNIX or Microsoft Windows consist of highly complex coding. Each is designed to support a number of business-related tasks and are more highly complex than the OS found on hardware-based VPNs. As a result, the software-related threats can be much more complex.

Operating systems contain millions of lines of software code, and as with anything written by people, they can contain mistakes. Those mistakes are what attackers seek to exploit when attacking an OS. Exploits may include buffer overflows, privilege escalation, or any number of other issues. As a result, if you run your VPN on a general-purpose OS, your VPN becomes susceptible to the same software vulnerabilities as your OS. In addition, software VPNs have many more possible OS configuration errors than hardware VPNs.

Some organizations run VPN software on a server that supports multiple applications. For example, the VPN server might also be a SQL server, a web server, or a file server. If this is the case in your environment, be aware of potential vulnerabilities created by application conflicts. This could be a matter of two applications that contend for resources on the servers causing issues, but it could also be the chance that one of the applications opens a new vulnerability in the VPN software. Think about a VPN server that is also a web server. If the web server is configured incorrectly, you could expose the VPN configuration files to an attacker through the web server

interface. In that case, you could configure your VPN server securely but still have an attacker pull a copy of your configuration files out of a web directory, bypassing your security.

Another potential threat associated with software VPNs is instability. Many information security professionals are reluctant to run software VPNs due to the challenges of the operating system crashing and taking the VPN with it. Many security professionals wouldn't ever want to run a VPN that could be taken out by a "blue screen of death"—the nickname for the Microsoft server crash screen.

Finally, OSs are vulnerable to viruses and other forms of malware. These would include Trojan horses, rootkits, or any of the other destructive malware currently active on the Internet. Any of these infections could compromise the security of your VPN, as well as your internal network.

If you are planning to run a software VPN, be sure to run it on a dedicated server, double-check both the OS configuration and the VPN configuration, and keep the OS fully patched and up to date. When your VPN is fully configured, run a vulnerability management tool against it, or have a professional come in and conduct a penetration test. Moreover, be sure to install and maintain current antivirus and anti-malware software on the server. It is a good idea to install a firewall application and an intrusion detection/prevention (IDS/IPS) application to ensure that your VPN remains secure.

Vulnerabilities common to both hardware and software VPN implementations include:

- DoS attacks
- Missing patches

- Backdoor attacks
- Unpublished vulnerability in the code
- Weak client security
- Weak authentication
- Hairpinning
- Credential sharing

CREDENTIAL SHARING

Credential sharing is sharing a username and password, or other security credentials, with others. A person may share credentials with another authorized user such as a boss, secretary, or coworker, or a family member. A person may also share credentials for reasons such as unofficial outsourcing all or a part of their job, such as data entry or coding, unbeknown to the employer. The practice does not necessarily result in compromising the system, but often it does. It is a violation of the confidentiality of the data and defeats the authentication process.

In a DoS attack, the attacker is trying to prevent others from accessing the VPN service. Attackers use specially crafted packets designed to crash the VPN or, more likely, direct a flood of traffic at the VPN in an attempt to overload it. This is not a very common attack because of the large amounts of traffic required to overload most companies' network infrastructures. Generally, you will not see

a VPN targeted with a DoS attack; VPN attacks are typically designed to allow the attacker access to the internal network. DoS is more common against popular websites like Twitter or Facebook, where there is significant publicity, or against online merchants that may be susceptible to blackmail. When your website is a significant source of revenue, a DoS attack can be very expensive.

Any hardware or software platform that can run VPN software is vulnerable; that is the nature of any computer technology. When these vulnerabilities are discovered, the vendor will typically develop a patch or an update to address the issue. If you do not keep your VPN current, you leave your network open to these issues, and attackers will try to exploit known vulnerabilities with your VPN.

A backdoor account attack is pretty rare, but can happen in some instances. An example of the traditional backdoor attack was featured in the classic movie *War Games*, where the system developer left a user account and password on a secure system in case he needed to get back into the system. This account was exploited by an attacker (in the movie, a kid just trying to play online games)—nearly causing the end of the world. Needless to say, the issues in your environment will not be quite as dramatic, but these are still of concern. Code running on VPNs and OSs is closely scrutinized, so system developers are less likely to be a threat, but what about a system administrator who thinks he is about to be fired? Or a vendor who comes in to support your VPN, creates an account for support, and forgets to delete it before leaving? The best way to mitigate this type of attack is through scheduled auditing of all accounts on your VPN, as well as strong, documented procedures for how and when accounts

are created and deleted. This particularly applies to accounts with elevated privileges or on systems accessible from the Internet.

One of the trickiest—and fortunately rarest—security threats is an unpublished vulnerability in the VPN code. Researchers or potential attackers can discover an unpublished vulnerability. If it is unpublished, it means the vendor is not yet aware of it, and thus has not yet developed a patch or an update. This is known as a zero-day attack. These are probably the most difficult threats you will face because there is not a lot you can do beyond following security best practices, creating a layered security environment, and monitoring the behavior of your environment for anomalies. Should you encounter this issue, be sure to follow your incident response process and work with the vendor to identify the issue and create a patch.

Another vulnerability is not directly on the VPN, but in the devices connecting to the VPN. VPNs offer significant security advantages, but if the client at the other end of the connection is not secure, then your network is at risk. To mitigate this risk, use a standard client configuration, which includes antivirus, anti-malware, firewall, and maybe even intrusion-detection software. Also make sure that the only clients connecting to the VPN are company-owned. Permitting personal assets to connect to your network opens a number of additional issues, because you are not able to supervise the configuration of those systems. All it takes is one employee whose child clicks on the video link that a friend sent (generally with a comment like “check out this video I took of you”), and you could end up with a virus on your network when the employee connects to the VPN using the same system.

A final vulnerability to consider is the challenge of weak authentication. Your VPN is really only as secure as your authentication mechanism. If you are authenticating via user ID and password, you run the risk of someone guessing or stealing those credentials and accessing your network without permission. To mitigate this risk, rely on either token-based or biometric authentication methods instead of—or in addition to—a user ID and password. If you must use a user ID and password for authentication, be sure to properly train users in the use of strong passwords.

A familiarity with the various threats, attacks, and mitigations is critical to the long-term support of your VPN and your total network. Be aware that no list of threats and attacks can be all-inclusive because attackers are constantly developing new methods for compromising networks. In the absence of specific threats and attacks, rely on good security practices to keep things safe.

COMMERCIAL VERSUS OPEN-SOURCE VPNS

When looking at VPN solutions, decide if you want to include the option to leverage open-source VPNs. If so, be aware of some of the trade-offs you will encounter when leveraging an open-source VPN product.

A commercial solution offers the following benefits:

- Ease of installation and management
- Available management tools
- Access to vendor support
- Available hardware maintenance

The benefits of an open-source solution include the following:

- Low cost
- Flexibility
- Ability to run on existing hardware
- Access to Internet-based support

Both commercial and open-source solutions pose some challenges. Commercial solutions are typically easier to implement, but that ease of installation comes at a cost. Commercial solutions are more expensive and usually less flexible than open-source solutions. Open-source solutions like OpenVPN (<https://openvpn.net/>) offer significant flexibility, but they generally require significantly higher skill in installation and support of the product. You rarely have access

to vendor support, as open-source solutions usually rely on the knowledge of the user and the development community. When you have an issue, you can normally post a request on a discussion board, and the community will assist you in resolving your issue.

When weighing the pros and cons of open-source VPN solutions, take into account a few other factors. Many organizations have policies about the use of open-source software. These policies derive from the company's tolerance for risk, the open-source licensing agreements, and intellectual property implications of open source. If your organization embraces open source, an open-source VPN is probably worth considering.

The best VPN application still comes down to identifying your requirements and finding the best solution for your organization. If you are on a low budget or have access to solid technical resources, an open-source solution could be the best choice for your needs.

DIFFERENCES BETWEEN PERSONAL AND ENTERPRISE VPNS

Much of the information in this chapter has dealt with VPNs deployed in support of an enterprise or at least a mid-sized or large network. Another class of VPN applications is the personal or individual VPN, sometimes referred to as the SOHO (small office/home office) VPN solution.

Many home routers include support for VPN connections, making open-source VPNs popular in this space. OpenVPN (<https://openvpn.net/>) offers an open-source VPN client that connects to many standards-based VPN gateways, such as SoftEther (www.softether.org/). The value of these personal VPNs is twofold. First, implementing a personal VPN provides you secure access to your home network. Even better, implementing a personal VPN provides you with some valuable experience with VPNs that you can then transfer to your organization's enterprise VPN.

BALANCING ANONYMITY AND PRIVACY

A key concept to understand about VPN technologies is the difference between anonymity and privacy. **Anonymity** is the ability of a network or system user to remain unknown. Tools that support anonymous access include Tor (www.torproject.org), an open-source software and network to surf the web anonymously. A hardware solution that leverages the Tor network for anonymity is VyprVPN (www.anonabox.com/partners/vyprvpn). The problem with solutions for anonymity is that they do not always protect your privacy. In many cases, the traffic carried anonymously by these applications is unencrypted, which means an attacker can read it with access to the right part of the network. If you are passing user IDs and passwords, credit card numbers, or other information that you would like to keep private, these are not the correct solutions for you.

Privacy, on the other hand, is keeping information about a network or system user from disclosure. If you want to protect your information, a VPN connection does an excellent job of this by encrypting the data that it carries. Leverage VPNs whenever you are connected to an untrustworthy network and want to send sensitive information. Wireless networks are a particularly ripe environment for attackers trying to get private information from a public network. VPNs do not offer any anonymity, however, because you are always able to track the endpoints of a VPN connection, and that information is needed to maintain the VPN connection.

PROTECTING VPN SECURITY TO SUPPORT AVAILABILITY

One of the design decisions you made when selecting your VPN was whether you needed a highly available solution. Once your users expect to access the network from a hotel, from a customer location, from a coffee shop, or from home when they have a sick child, you will find they have little tolerance for outages. A downed VPN means that none of your remote workers can do much work, and in some organizations that can mean that a majority of the company could be off the network until the VPN is back up and running.

TECHNICAL TIP

When considering VPN availability, do not overlook the little things. Put your VPNs in separate racks, connected to separate power supplies. Nothing is more embarrassing than explaining to your manager that the VPN was down for two hours because someone accidentally switched off a power strip.

The most common method for implementing a highly available VPN is really quite simple. You buy two VPN hardware units (or implement two open-source VPNs) and then configure each as a highly available pair using the vendor's high-availability mechanisms. Cisco offers the **Hot Standby Router Protocol**

(HSRP), a long-time proprietary standard, which allows configuration of a pair of Cisco VPNs so that in the event the primary VPN fails, the backup takes over seamlessly. If configured correctly, end users do not notice the cutover. A more industry-standard protocol that offers similar functionality is the **Virtual Router Redundancy Protocol (VRRP)**. You also have the option to use third-party solutions. Many of today's load balancers offer the ability to load-balance VPNs. When one VPN fails, the load balancer automatically directs all traffic to the remaining gateway. An example of load balancing over multiple routers is the proprietary Cisco Gateway Load Balancing Protocol (GLBP).

In addition to ensuring that your VPN is highly available, also ensure that your Internet circuits have similar redundancy. If you have a single Internet connection and it goes down, VPN users will not be able to access the network. In the event of circuit outages, the user will typically need to reconnect to the VPN. This is still better than not being able to connect at all.

When you are considering a highly available VPN solution, be sure to consider the acquisition costs and the ongoing maintenance costs over the following three to five years. Continued maintenance and updates can add significant costs over the life of a product. It is important to understand the full cost of a solution rather than just the purchase price.

THE IMPORTANCE OF USER TRAINING

Once your VPN is up and running, you must train your users on how to use it. One of the most common challenges with IT infrastructure deployments, especially with security infrastructure like a VPN, is for the IT team to assume that, because they understand how the solution works, their end users understand it as well. That is almost always a bad assumption. It can turn a successful rollout into an unsuccessful one.

To develop good user training, start by setting appropriate training program goals. One of the first mistakes security practitioners make when developing training is jumping past any planning and starting with taking screen shots, typing up documentation, and scheduling meetings. However, much like a successful VPN deployment, the first part of successful training is planning. Determine the following before designing your user training:

- Who is your target audience for the training? Before you start writing your training, understand who will be learning.
- What is the technical awareness level of your audience? Are you training tech-savvy IT professionals, or salespeople who just want to turn on their PCs and start working?
- Where is your audience? Is your audience in one central location, or spread out regionally—or even nationally? This will affect the choice of media you use to train users.
- What level of training are you trying to deliver? Will users need to install the client software? Do you want to include basic

troubleshooting in the training? Are you trying to sell features, or do you just want to train them on how to log on? Should you include additional security issues in this training?

Once you gather the appropriate information, determine the best mechanism for delivering the training. For example, you may distribute some written instructions, create a training video, have a conference call, or hold a web conference. For high-profile users or those who require significant support, hold classroom training. If you decide to go with live training, plan to do additional information security training. After all, you have their attention for the period of the training; you should maximize the benefit not only to your users, but also to your company. Remember that well-trained employees tend to be the most secure computer network users.

Other things to keep in mind as you plan user training: Make sure your training plan includes how you will train new employees. Will your training post to an intranet website that new hires can watch? Will you offer new-hire training at a regular interval to ensure that everyone understands how to use the VPN?

Finally, do not forget your support staff. If a user with VPN issues is supposed to call the help desk, make sure the help desk knows how to help.

VPN TROUBLESHOOTING

To successfully troubleshoot and resolve VPN issues, you must be organized, methodical, and prepared. Like anything else on your network, the VPN will experience an issue at some point. The preparations you undertake during installation and maintenance will prove to be critical components of your troubleshooting process.

Before you start troubleshooting your VPN issue, have access to the following information:

- A network diagram showing the placement of the VPN and other key network components like firewalls, routers, and switches
- A copy of the current VPN configuration
- Any error, system, or alert logs
- An operations guide
- Maintenance logs and change management records

Now that you have your basic information, you can start troubleshooting. There is no one way to troubleshoot an issue, particularly a VPN issue. What you can do is determine what works best for you, and stick with that process. Some people like to start at the far end and work back to the local VPN server. Others will start at the network level and work their way up to the application. The actual process is not as critical as having an established procedure and sticking with it. Here is one way to tackle this process:

TECHNICAL TIP

When troubleshooting issues with a VPN, or any network equipment, be sure to review all recent changes to the environment, not just changes to the VPN. An update to a DNS server, the replacement of a router, or changes to IP addressing can affect your VPN.

1. **Identify the symptoms**—Frequently, when dealing with a user-reported issue, the reported item will bear little resemblance to the actual problem. Users tend to generalize with statements like “No one can get on the VPN” or “The intranet is down from the VPN.” You need more specifics before you start correcting an issue. Ideally, you will receive an automated alert from your monitoring system, which generally permits you to track down issues much more quickly than a help desk ticket.
2. **Determine the scope of the problem**—Know whether your problem is isolated to a single user or system, a workgroup, or the entire network. Do you have one user who cannot connect, 100 users, or the entire company? Not only will this help you get a handle on the urgency of the issue, but it also gives you valuable clues to what the problem might be. If it is one user, the odds are pretty good that it is a client issue, or possibly a network issue on the user’s end, and the help desk can probably assist with the issue. If a large group—but not everyone—is affected, you can look for commonalities. Do all the users belong to the same VPN group? Do they all connect from the same part of the country (indicating that it could be an ISP issue)? If you have multiple VPNs, are the problems all on one VPN?
3. **Look for changes**—Before you assume something is broken, see if anything has changed. If the issue is with a single user, did the user install a new application like Skype or Tor that could interfere with the VPN? If the VPN is unreachable from the Internet, was someone working on your Internet connection? Did you upgrade any software over the weekend? The

importance of formal change management is critical to successfully maintaining a production computing environment. When you are troubleshooting, determine what has changed. Be aware that sometimes the toughest part of chasing down change-related issues in a complex environment is figuring out which of the 20 or 30 changes performed during a weekend change window is causing your problem.

4. **Call the vendor**—It is never a bad idea to call the vendor to see if it is a known issue. If you do not have phone support with the vendor, visit its website and check its knowledge base. The Internet is an invaluable tool for helping with troubleshooting.
5. **Try the most likely solution**—As the resident expert, once you have gathered all the information and reviewed the likely issues, try a fix. Record these fixes in your operations guide and in your change control process.
6. **Test it**—You made a change; test to see if it worked. If not, back out of the change you made and repeat the previous step with the next most likely fix. Repeat the process until the problem is solved. Do not keep trying fix after fix without backing out previous fixes; you could end up making so many changes that you will not be able to easily get back to a stable configuration.
7. **Check to see if you broke anything else**—This is a critical step. Changes to the environment—even if you are trying to fix something—can have unpredictable effects. The last thing you want is to assume the problem is solved, just to get a call later saying no one can get to the Internet.
8. **Document, document, document**—This is probably the most important part of the process. If you do not write processes and procedures down, it is as if those procedures were not followed. When the next expert comes along to troubleshoot and your changes are not documented, they will start with incorrect information or guesswork. If you encounter the same problem 12 months later, you will be glad you can refer to your previous process to speed up the second recovery. Create a section in your operations guide just for this kind of documentation.

9. **Remember to make a backup**—While included in the documentation, be sure to capture any changes to configurations in a backup. A network upgrade or fix for another problem could cause the issue to reoccur. It is much faster to reinstall than to perform a second level of troubleshooting once one item is corrected.

Now that you have a general troubleshooting process, consider some other things when troubleshooting a situation:

- **Do not panic**—During a major outage, people tend to panic. They will start sending emails to large distribution lists, sending alerts, escalating issues, and doing a variety of other things that are not particularly helpful to you while you are troubleshooting an issue. Keep your cool, work on the problem at hand, and keep management informed of the progress you make.
- **Do not overcommit or make promises you cannot keep**—If you do not know what the issue is, do not tell your manager or users that the network will be back up in an hour.
- **Focus on the problem at hand**—Getting sidetracked when troubleshooting a complex issue is easy. Be sure that the problem at hand is the one you are working to address. If you identify other potential issues unrelated to the current outage, document each and plan to correct under your change control process when the outage is over. But do not go off on a tangent and delay addressing the immediate issue.


■ NOTE

When you are troubleshooting a VPN issue, management will often require frequent status updates. The worst place for you to be when the VPN is down is spending all your time on conference calls fielding questions like, “When will it be fixed?” Management needs to be educated to the

concept that things will not be fixed until you get off the phone and back to working on the problem.

You can do a few quick and easy tests if you are having a major VPN outage and you just want to be sure that your VPN is still on the network. Generally, your VPN will have two interfaces. One will be accessible from the Internet, and the other will be the connection to your internal network (possibly through a firewall). In the event of a major outage, start by verifying that both of those ports are functioning.

To verify the internal port is available, you can use the ping command from the Windows command prompt or from a UNIX command line. See **FIGURE 10-4** for an example of a successful ping response.

A screenshot of a Windows Command Prompt window. The title bar reads "Command Prompt". The window content shows the following text:

```
Microsoft Windows [Version 10.0.18362.657]
(c) 2019 Microsoft Corporation. All rights reserved.

C:\Users\kim>ping 2607:f8b0:4000:806::200e

Pinging 2607:f8b0:4000:806::200e with 32 bytes of data:
Reply from 2607:f8b0:4000:806::200e: time=36ms
Reply from 2607:f8b0:4000:806::200e: time=31ms
Reply from 2607:f8b0:4000:806::200e: time=30ms
Reply from 2607:f8b0:4000:806::200e: time=30ms

Ping statistics for 2607:f8b0:4000:806::200e:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 30ms, Maximum = 36ms, Average = 31ms

C:\Users\kim>
```

FIGURE 10-4 Troubleshooting a VPN issue using PING.
Used with permission from Microsoft.

Once you have validated the internal port, also validate the external port. This is frequently a challenge because in most organizations, a direct connection to outside the network is not readily available. In some larger organizations, you may have a digital subscriber line (DSL) Internet connection dedicated for

testing, but if you do not have that type of access, a number of Internet sites will let you run network diagnostics from their websites. See **FIGURE 10-5** for a sample Internet tracert command from *www.network-tools.com*. If you do an Internet search on “network tools,” you can find a number of sites that offer a similar capability.

Tool Traceroute ▼

☐ Convert Base-10 to IP

2607:f8b0:4000:806::200e

Go!

Traceroute Check for: '**2607:f8b0:4000:806::200e**'

traceroute to 2607:f8b0:4000:806::200e (2607:f8b0:4000:806::200e), 10 hops max, 80 byte packets

1 2600:3c00::8678:acff:fe1c:ec41 (2600:3c00::8678:acff:fe1c:ec41) 0.965 ms 2600:3c00::e6c7:22ff:fe10:9cc1 (2600:3c00::e6c7:22ff:fe10:9cc1) 0.740 ms 2600:3c00::8678:acff:fe1c:ec41 (2600:3c00::8678:acff:fe1c:ec41) 0.969 ms

2 2600:3c00:2222:17::1 (2600:3c00:2222:17::1) 0.710 ms 2600:3c00:2222:15::1 (2600:3c00:2222:15::1) 0.500 ms 2600:3c00:2222:17::1 (2600:3c00:2222:17::1) 0.675 ms

3 2600:3c00:2222:10::1 (2600:3c00:2222:10::1) 0.479 ms 2001:678:34c:5e::2 (2001:678:34c:5e::2) 1.066 ms 0.938 ms

4 2001:678:34c:5e::2 (2001:678:34c:5e::2) 1.078 ms 1.015 ms 2001:4860:0:11e3::1 (2001:4860:0:11e3::1) 0.987 ms

5 2001:4860:0:11e3::1 (2001:4860:0:11e3::1) 1.021 ms 0.958 ms 2001:4860:0:1::1a95 (2001:4860:0:1::1a95) 1.879 ms

6 dfw28s22-in-x0e.1e100.net (2607:f8b0:4000:806::200e) 0.841 ms 2001:4860:0:1::1a95 (2001:4860:0:1::1a95) 1.562 ms dfw28s22-in-x0e.1e100.net (2607:f8b0:4000:806::200e) 0.905 ms

FIGURE 10-5 Using Internet-based traceroute tools to determine whether your VPN is accessible from the Internet.
Courtesy of Network-Tools.com.

You now have the tools to create your own troubleshooting process for your VPN, using some or all of the steps and tips discussed. You just need to determine what works best in your environment. Be methodical, do not panic, keep focused on the issue, and you will be very successful in quickly addressing VPN issues in your environment.

NOTE

Be careful using any third-party website. A site that appears to offer assistance may do so but may also log information from your systems and network, or download malware to your system while conducting the analysis.

CHAPTER SUMMARY

VPN management is a complex activity that requires attention to detail and good documentation. A variety of best practices may assist you in your VPN deployment and ongoing support. Selecting and deploying a VPN solution that meets your business requirements is essential. This includes deciding between open-source and commercial VPN solutions when selecting your product. Understanding the threats, attacks, and mitigations and ensuring that your VPN is available are key components of success. In addition, you must understand the difference between privacy and anonymity; know the difference between personal/individual and enterprise/network solutions; and, finally, be able to train your users well. Once you have mastered all these facets of successful VPN rollout and support, you will be ready to move on to other security topics.

KEY CONCEPTS AND TERMS

Anonymity

Bypass deployment

Hairpinning

Hot Standby Router Protocol (HSRP)

Privacy

Service level agreement (SLA)

Slideware

Two-factor authentication

Virtual Router Redundancy Protocol (VRRP)


Vulnerability management

CHAPTER 10 ASSESSMENT

1. Which of the following is *not* one of the three most common VPN deployment architectures?
 - A. Bypass
 - B. DMZ
 - C. Internally connected
 - D. Modified
2. Which of the following is *not* a threat common to software and hardware VPNs?
 - A. Backdoor attack
 - B. Denial of service
 - C. Missing patches
 - D. Split tunneling
3. Which of the following is one of the most common and easily exploited vulnerabilities on any hardware network device?
 - A. Weak default password
 - B. Accessible power switch
 - C. Default tunneling
 - D. No multifactor authentication
4. Which of the following is *not* a threat common to software and hardware VPNs?
 - A. Weak authentication
 - B. Weak client security
 - C. Weak encryption key
 - D. Weak user name
5. Which of the following is a benefit of a commercial VPN solution over open-source solutions?
 - A. Less cost

- B. Product support
 - C. More flexibility
 - D. Access to the code
6. Which of the following is *not* a component of a VPN policy?
- A. Definitions
 - B. Introduction
 - C. Scope
 - D. Troubleshooting
7. Which of the following is *not* a VPN best practice?
- A. Back up the configuration
 - B. Permit split tunneling
 - C. Employ vulnerability management
 - D. Select based on best reviews
8. Which of the following is the term for malicious code entering the network and making a sharp turn into the secure network?
- A. Boomerang
 - B. Curving
 - C. Hairpinning
 - D. U-turning
9. What is the term for a VPN deployment in which traffic between the VPN and the internal network is not firewalled?
- A. Alternative deployment
 - B. Bypass deployment
 - C. Differential deployment
 - D. Internally connected deployment
10. Which of the following is *not* part of multifactor authentication?
- A. Something you are
 - B. Something you have
 - C. Something you know

D. Something you wear



© Digital_Art/Shutterstock

VPN Technologies

VIRTUAL PRIVATE NETWORK (VPN) is a general industry term that actually covers many different technologies.

Anyone who works in a large company and needs to access a corporate network remotely will probably confirm using a VPN to connect to the network for email, intranet, and other corporate applications. Ask anyone how that VPN works, however, and you will almost always get a blank stare. End users are typically interested only in the result of a solution, not in how it works. For you, the information security practitioner, however, just how a VPN works is as important as the benefits a VPN provides your organization.

VPNs are deployed in a number of different ways, leveraging a variety of technologies, platforms, and protocols. Determining which VPN is the right fit for your organization requires successfully gathering and interpreting your business requirements. Once you have documented those requirements, it is up to you as the security practitioner to understand the various options and capabilities to fit the VPN technology to the appropriate business requirements.

A variety of technical factors affect the selection and installation of a VPN solution. Some VPNs are available as software installed on a workstation or a server. Other VPNs are software components of other devices, like a router or a firewall. Finally, dedicated VPN hardware appliances provide secure remote connectivity.

A variety of underlying protocols can provide different functions, features, and levels of encryption. When a vendor starts talking about L2TP, IPv6, SSL and SSH, or IPSec, you will need to understand the jargon and make the right technology decision for your organization.

Finally, other infrastructure considerations when working with VPN technologies include how networks function—including network address translation (NAT) and Internet Protocol (IP) version—and the use of virtualization, which can affect how you deploy, maintain, and troubleshoot a VPN.

CHAPTER 11 TOPICS

This chapter covers the following topics and concepts:

- The differences between software and hardware solutions
- The differences between Layer 2 and Layer 3 VPNs
- Internet Protocol Security (IPSec)
- Layer 2 Tunneling Protocol (L2TP)
- Secure Sockets Layer (SSL) and Transport Layer Security (TLS)
- Secure Shell (SSH) protocol

- How to establish performance and stability for VPNs
- How to use VPNs with network address translation (NAT)
- Different types of virtualization

CHAPTER 11 GOALS

When you complete this chapter, you will be able to:

- Contrast hardware and software VPN solutions
- Describe VPN protocols, uses, features, and problems
- Explain the problem of using VPNs with NAT
- Evaluate hardware VPN devices

DIFFERENCES BETWEEN SOFTWARE AND HARDWARE SOLUTIONS

A key topic in discussing VPN technologies is the differences between software and hardware solutions. The following definitions help reveal those differences:

- **Software VPN**—Software-based VPNs are sold either as part of a server operating system (OS), as part of an appliance operating system, or as a third-party add-on software solution.
- **Hardware VPN**—A hardware VPN is a standalone device, dedicated to managing VPN functions such as authentication, encapsulation, encryption, and filtering.

The functionality of software and hardware VPN solutions is essentially the same, but providing a secure remote connection on demand highlights some important differences.

■ NOTE

In the early days of VPNs, VPN software commonly ran on a Windows or UNIX server. Today, server-based implementations are mostly in small environments because of the poor scalability and reliability of the OS. Most software VPNs act as a component of a firewall or router, not as an add-on to a server.

Software VPNs

When evaluating whether or not to deploy a software VPN, consider the two types of software VPNs:

- **Operating system–based VPN**—An operating system–based VPN solution is an application that runs on a Microsoft Windows or UNIX OS. These are generally used in smaller companies, as this device tends to be less scalable and less stable than other VPN solutions. Operating system–based solutions are generally less expensive, especially when running on a shared server, which may host other applications. Many of these solutions are open source, which further reduces the cost, although it can increase the complexity of the solution.

One potential security issue with this solution is that the server communicates with the public network (generally the Internet) for the connections to occur. Any time you connect a server and the associated OS to a public network, you increase the security risk, because you are exposing it to a much larger pool of attackers. You can mitigate this risk by limiting the ports open to the server, and, if you are using only site-to-site VPN connections, you may be able to restrict the connections to specified IP addresses. You can also add a firewall to the VPN server, which adds to the overhead on the server, but it will provide additional protection to offset the increased exposure.

- **Module-based VPN**—A module-based VPN runs as a component on a larger system. These are sometimes included as part of the overall feature set, or in other cases may require the purchase of additional licenses to use the VPN. An example of this would be the VPN capability included with many firewalls. Many routers also offer this type of capability, which permits the easy encryption of wide area network (WAN) links for security-conscious companies. The benefit of this type of VPN solution is reduced complexity of the environment, since you have fewer discrete devices to manage. VPN modules are also typically less expensive than hardware VPN solutions. Some vendors also offer

hardware accelerators for improving overall performance of this solution.

Hardware VPNs

Hardware VPNs are dedicated appliance-based solutions, generally based on a router-type platform. Hardware VPNs are the most common type of VPN deployed in corporations today. Although hardware VPNs can be complex to deploy, these devices are typically more scalable than software counterparts, and they can be easily deployed in a redundant manner. Hardware VPNs can increase the complexity of an environment, because you are deploying additional equipment. The good news is that you can usually manage this additional hardware with the same types of network management tools you use to manage the routers and switches in an environment.

Hardware VPNs can create some security issues, largely related to potential vulnerabilities in the VPN software code on the appliance. A number of security alerts related to VPN vulnerabilities have appeared in recent years. Fortunately, you can manage this issue fairly easily by keeping current on your vendor's security alerts and by upgrading VPN code in a timely fashion. If your organization does not have the skills or ability to test the solution thoroughly, it might be preferable to run the $N-1$ version of code, where N is the current version of code, unless a known issue with that previous version of the code has been published.

Ultimately, the requirements of your business will drive your selection of a software or hardware VPN. The good news is that many options are available for every size network.

NOTE

You must make a risk-based decision on whether to accept the possibility of undiscovered bugs in a new version of hardware or to live with the known bugs in a previous version. If an upgrade is issued primarily for fixing security flaws, you may be better off with the new version, surprises and all.

DIFFERENCES BETWEEN LAYER 2 AND LAYER 3 VPNS

One distinguishing component of current VPNs is the use of a variety of transport protocols to establish connections. This is helpful not only because the protocols have different capabilities, encryption strengths, and authentication mechanisms, but because each can run at different layers of the Open Systems Interconnection (OSI) Reference Model. The OSI model is the standard seven-layer conceptual tool that describes protocols and their corresponding functions. Each layer communicates with its peer layer on the other end of a communication session. Although the OSI model is helpful in discussing protocols, most protocols are not in full compliance with it.

In the case of VPNs, the protocols used by the vast majority of solutions work at Layers 2 and 3 of the OSI model.

Layer 2 of the OSI model is the Data Link Layer. The Data Link Layer is the protocol layer that transfers data between adjacent network nodes. In the case of a VPN, this protocol transfers data from one VPN endpoint to the other. Layer 2 Transport Protocol (L2TP) is an example of a protocol that communicates using Layer 2. Layer 3 of the OSI model is the Network Layer. The Network Layer is responsible for end-to-end packet delivery and includes the ability to route packets through intermediate hosts. IPsec is an example of a VPN protocol that communicates at Layer 3.

TECHNICAL TIP

SSL/TLS and SSH are protocols that operate at Layer 7 of the OSI model, the Application Layer.

TECHNICAL TIP

An RFC is a **Request for Comments**. These “requests” are published by the **Internet Engineering Task Force (IETF)**. The IETF is the standards body for Internet-related engineering specifications. The IETF uses RFCs as a mechanism to define Internet-related standards. (See www.ietf.org for more information.)

INTERNET PROTOCOL SECURITY (IPSEC)

Internet Protocol Security (IPSec) is a standards-based protocol suite designed specifically for securing Internet Protocol (IP) communications. IPSec authenticates and encrypts each IP packet in an IP data stream. In addition, IPSec has protocols that can establish mutual authentication and cryptographic key negotiation during a session. IPSec operates at the Network Layer (Layer 3) of the OSI model.

The IPSec standard utilizes three major components:

1. Authentication Header (AH)
2. Encapsulating Security Payload (ESP)
3. Internet Key Exchange (IKE)

Authentication Header (AH) provides integrity protection for packet headers and data, as well as user authentication. It can optionally provide replay protection and access protection. AH cannot encrypt any portion of a packet. In the initial version of IPSec, the ESP protocol could provide only encryption, not authentication, so AH and ESP were often used together to provide both confidentiality and integrity protection for communications. Because authentication capabilities were added to ESP in the second version of IPSec, AH has become less significant; in fact, some IPSec software no longer supports AH. However, AH is still of value because AH can authenticate portions of packets that ESP cannot. Also, many existing IPSec implementations use AH.

Encapsulating Security Payload (ESP) is the second core IPsec security protocol. In the initial version of IPsec, ESP provided only encryption for packet payload data. Integrity protection was provided by the AH protocol if needed. In the second version of IPsec, ESP became more flexible. It can perform authentication to provide integrity protection, although not for the outermost IP header.

TECHNICAL TIP

Why does the layer matter? Because IPsec operates at Layer 3 of the OSI model, it can encrypt any traffic in Layers 4 through 7 of the OSI model. That means IPsec can be used to encrypt any application traffic.

NOTE

IPsec supports the Data Encryption Standard (DES), a 56-bit encryption protocol, and 3DES, which effectively yields a 168-bit encryption protocol. (In 3DES, pronounced “triple DES,” data is encrypted three times using DES.)

Therefore, in all but the oldest IPsec implementations, ESP can provide encryption only, encryption and integrity protection, or integrity protection only. This chapter mainly addresses the features and characteristics of the second version of ESP.

Internet Key Exchange (IKE) negotiates, creates, and manages security associations. Security association (SA) is a generic term for a

set of values that define the IPSec features and protections applied to a connection. You can also create SAs manually, using values agreed upon in advance by both parties, but because these SAs cannot be updated, this method does not scale for real-life large-scale VPNs. In IPSec, IKE provides a secure mechanism for establishing IPSec-protected connections.

IPSec supports two different modes:

- **Transport mode (host-to-host)**—In transport mode, only the data packet payload is encapsulated, while the packet header is left intact. In this mode, the destination host decapsulates the packet. This is the mode used in a Windows environment to secure a client-to-server connection using IPSec.
- **Tunnel mode (gateway-to-gateway or gateway-to-host)**—In tunnel mode, the IP packet is entirely encapsulated and given a new header. The host/gateway specified in the new IP header decapsulates the packet. This is the mode used to secure traffic for a remote access VPN connection from the remote host to the VPN concentrator on the internal network. (A **VPN concentrator** decrypts information coming from a remote user over the Internet, and it encrypts information sent back over the Internet to the remote user.)

IPSec is a little different from some of the other protocols. It provides high-quality, interoperable, cryptographically based security for IPv4 and IPv6. As a result, the protocol supports a comprehensive set of security services, including:

- Access control
- Connectionless data integrity checking
- Data origin authentication
- Replay detection and rejection

- Confidentiality using encryption
- Traffic flow confidentiality

IPSec-based VPNs have been the dominant VPN platform for many years, although in recent years SSL-based VPNs have been making significant inroads. IPSec is popular for three reasons:

- It supports all OS platforms.
- It provides secure, node-on-the-network (where the computer is treated as part of the local, secured network) connectivity.
- It offers a standards-based solution, permitting easier interoperability between different devices and vendors.

For more information on IPSec, see

<http://tools.ietf.org/html/rfc4301>. While RFC 4301 provides the main IPSec standard definitions, separate RFCs exist for some IPSec features. The full set of RFCs defines the entire standard.

LAYER 2 TUNNELING PROTOCOL (L2TP)

Layer 2 Tunneling Protocol (L2TP) is an older protocol largely replaced by IPSec and SSL/TLS-based VPNs in production environments; however, you will still encounter references to the protocol in current VPN literature. It may still be in use in some older environments, where backward compatibility could still be an issue. L2TP was used extensively in early VPN solutions, but it lost its popularity as other protocols proved to be more usable as industry standards developed.

L2TP is a combination of the best features of **Point-to-Point Tunneling Protocol (PPTP)**, a Microsoft proprietary protocol, and the **Layer 2 Forwarding (L2F) Protocol**, which was an early competing protocol for PPTP, developed by Cisco Systems. Like PPTP, L2TP was an extension of **Point-to-Point Protocol (PPP)** to allow PPP to be tunneled through an IP network. L2TP support was first included in Microsoft server products with the release of Windows 2000 Server and continues in Windows Server 2019.

One of the challenges with L2TP is that it only provides a mechanism for creating tunnels through an IP network. It does not provide a mechanism for encrypting the data being tunneled. As a result, L2TP was typically used in conjunction with IPSec protocol's ESP for encryption. For this reason, you may sometimes see L2TP referred to as L2TP/IPSec.

RFC 3193 defines the use of IPSec to secure an L2TP implementation. For more information on L2TP, see [***http://tools.ietf.org/html/rfc2661***](http://tools.ietf.org/html/rfc2661). Although RFC 2661 provides the main L2TP standard definitions, separate RFCs for some other implementations and features exist. The full set of RFCs defines the entire standard.

Remember that you may never encounter L2TP in a production environment, but you should know the basic aspects of this protocol when looking at tunneling and VPN protocols in general.

SECURE SOCKETS LAYER (SSL)/TRANSPORT LAYER SECURITY (TLS)

One of the key VPN protocols today is SSL/TLS, which is the main alternative for a VPN solution if you do not want to leverage an IPSec solution. However, before you consider this protocol in conjunction with VPNs, it is important to understand the origin of this protocol.

If you have ever surfed the World Wide Web (WWW), you have used the Hypertext Transfer Protocol (HTTP) to connect to a website. One of the drawbacks of HTTP is that it does not include the ability to encrypt or otherwise protect the data stream between the client and server. This was not an issue until the early 1990s, when the need to protect against eavesdropping on communications became critical to the ultimate success of the World Wide Web. Several technologies have addressed this need, but one solution has rapidly become the industry standard: Secure Sockets Layer (SSL).

HTTPS is the web protocol that utilizes SSL to encrypt HTTP, and it is used worldwide today for secure web communications. (See **FIGURE 11-1.**)

Krebs on Security

https://krebsonsecurity.com

Advertisement

Two-Factor Authentication Just Got Better.

START MY TRIAL

Subscribe to RSS

Follow me on Twitter

Join me on Facebook

KrebsonSecurity

In-depth security news and investigation



ADVERTISING/SPEAKINGABOUT THE AUTHOR

A Little Sunshine / The Coming Storm / Web Fraud 2.0 — 11 comments

17 Pay Up, Or We'll Make Google Ban Your Ads

FEB 20

A new email-based extortion scheme apparently is making the rounds, targeting Web site owners serving banner ads through Google's AdSense program. In this scam, the fraudsters demand bitcoin in exchange for a promise not to flood the publisher's ads with so much bot and junk traffic that Google's automated anti-fraud systems suspend the user's AdSense account for suspicious traffic.

Important: Regarding your AdSense account!

AdSense Syndicate

From: AdSense Syndicate <adsense.syndicate@protonmail.com> [P]
Subject: Important: Regarding your AdSense account

Advertisement

Two-Factor Authentication Just Got Better.

START MY TRIAL



FIGURE 11-1 A secure browser session using SSL. Note the “https” in the address bar.
Courtesy of Krebs on Security.

TECHNICAL TIP

SSL supports 128-bit and 256-bit encryption, while TLS supports the Advanced Encryption Standard (AES) with 128-, 192-, and 256-bit keys.

SSL was originally proposed as a standard by Netscape. Version 1.0 had serious security flaws, which were corrected in versions 2.0 and 3.0. As this protocol has become more widely used, it has been formalized in the IETF standard known as Transport Layer Security (TLS). The SSL/TLS protocol suite provides a method for secure client/server communications across a network. SSL/TLS prevents eavesdropping and tampering with data in transit. SSL/TLS also provides endpoint authentication and communications confidentiality through encryption.

TECHNICAL TIP

SSL/TLS requires a large overhead for processing the RSA algorithm. By utilizing SSL/TLS acceleration and load balancing, the overhead can be better managed with additional benefits, including a centralized control point for private keys and certificates, which streamlines administration and makes policy application easier. With acceleration, the identity of a client can be validated

before access is allowed to web resources, and traffic can be reencrypted at the load balancer for additional security.

TECHNICAL TIP

Although the IETF standard refers to the protocol as TLS, the industry still uses the acronym SSL when referring to the protocol used to secure browser communications.

In typical end user/browser usage, SSL/TLS authentication is one-way. The server is authenticated when the client compares the information entered to gain access to the server to information on the SSL certificate stored on the server. The client knows the server's identity, but not vice versa; the client remains unauthenticated or anonymous. See **FIGURE 11-2** for an example of the information in a server certificate.




FIGURE 11-2 A certificate used to authenticate a server in an HTTPS connection.

Used with permission from Microsoft.

SSL/TLS can also perform bidirectional authentication by using client-based certificates. This is particularly helpful when using this protocol to access a protected network, as it adds a layer of authentication to the access.

FYI

SSL/TLS is supposed to authenticate servers, but it is not as robust as originally planned. Phishing attacks (attacks that direct a user's browser to a counterfeit website) are a major problem for security professionals. Users typically do not check to see if a site is encrypted or if the certificate information matches their destination website. Newer versions of browsers check encryption status for the user and warn of a mismatch between a certificate and the domain presenting it.

SSL/TLS and VPNs

As you learned in the section on IPsec, a VPN creates a secure tunnel through a public network like the Internet. Although SSL VPNs still leverage the concept of tunneling, tunnel creation is different than IPsec. An SSL VPN establishes connectivity using the SSL protocol. IPsec works at Layer 3 of the OSI model, while SSH functions at Layers 4 and 5. SSL VPNs can also encapsulate information at Layers 6 and 7, which makes SSL VPNs some of the most flexible available.

One additional function of an SSL VPN is that it usually connects using a web browser, whereas an IPsec VPN generally requires client software on the remote system.

SSL VPNs create predominantly remote access VPN connections, where a client connects to applications on an internal network. This is different from a site-to-site connection, where two gateways connect disparate private networks across the Internet.

Some of the SSL/TLS VPNs benefits versus IPsec VPNs include:

- **Less expense**—Because an SSL VPN is typically clientless, you do not have the costs of rolling out, supporting, and updating client software.
- **Platform independence**—The access to an SSL VPN comes through the standard SSL interface, which is a component of virtually every web browser on any OS that runs a browser. Although the OS support for VPN clients is good for common OSs, a significant lag can exist in client development following the release of a new OS or the release of a new client version. You may see a Windows client 90 to 180 days earlier than the Apple Mac or Linux version.
- **Client flexibility**—As a general rule, IPsec clients are installed only on corporate systems. Due to the additional configuration

flexibility, SSL VPNs allow access from a variety of clients, including corporate systems, SOHO (small office/home office) systems, and even a kiosk system, such as in a library or an Internet café. This wider accessibility tends to increase satisfaction with the technology.

- **No problems with NAT**—Historically, NAT caused issues with IPsec VPNs. As a result, virtually all IPsec vendors have created workarounds. With an SSL VPN, you do not have these issues because SSL works at a higher OSI layer than IPsec and, as a result, is not affected by NAT.
- **Granular access control**—This is a benefit or a drawback, depending on your environment. SSL VPNs require a greater granularity of access than a typical IPsec VPN, due to the creation of a tunnel from the host to the internal network. SSL VPNs require explicit definition of each resource accessed. The upside is that, unless you have defined it explicitly, an SSL VPN user cannot access the resource. Although this barrier has security benefits, in a complex environment, this could add significant overhead to your VPN support.
- **Fewer firewall rules required**—To access an IPsec gateway across a firewall, you need to open several ports to support the individual protocols for authentication and the tunnel. With an SSL VPN, you only need to open port 443, which is easy because of the prevalence of the HTTPS protocol.

For more information on TLS, see <http://tools.ietf.org/html/rfc5246>.

SECURE SHELL (SSH) PROTOCOL

The **Secure Shell (SSH)** protocol is a method for secure remote login and other secure network services over a public network such as the Internet. SSH services a number of applications across multiple platforms including UNIX, Windows, Mac, and Linux.

You can use SSH:

- For login to a shell on a remote host, replacing the need for telnet and rlogin (see **FIGURE 11-3** for an example of an application that uses SSH for this application)
- For executing a single command on a remote host, replacing the need for rsh
- For file transfers to a remote host
- In combination with rsync to back up, copy, and mirror files securely
- In conjunction with the OpenSSH server and client to create a full VPN connection




FIGURE 11-3 PuTTY is an application, providing terminal emulation, that leverages SSH to provide secure connections to hosts.

Courtesy of PuTTY.

The SSH protocol consists of three major components:

- TLS provides server authentication, confidentiality, and integrity with **perfect forward secrecy** (use of an independent session key

for every web session a user starts in order to protect data)

- User Authentication Protocol (UAP), which authenticates the client to the server
- Connection Protocol, which multiplexes the encrypted tunnel into several logical channels

For more information on SSH, see <http://tools.ietf.org/html/rfc4251>.

ESTABLISHING PERFORMANCE AND STABILITY FOR VPNS

Now that you have an in-depth understanding of the many options available to you for VPNs, as well as other secure access protocols, it is time to discuss some of the challenges you might encounter when supporting your VPN. For your VPN rollout to be successful, consider two factors: performance and stability.

Performance

The performance of a VPN on the network affects many forms of communication differently. The type of VPN selected directly affects its performance. Some critical factors can affect the performance of your VPN:

■ NOTE

Do not forget time-of-day considerations when investigating load-related performance issues. You will generally find your peak times for VPN use will be first thing in the morning, as employees get started for the day; after lunch, when everyone comes back from the break; and at the end of the business day. If performance issues arise, be sure to note the times of day to help correlate the data and identify root causes or at least common factors.

- **VPN type**—The performance characteristics of a VPN supporting remote clients can be very different from the performance characteristics of a VPN supporting site-to-site connections, or even a mixture of remote clients and site-to-site connections.
- **Protocol**—The performance characteristics associated with an IPsec VPN can be very different from what you may find with an SSL VPN implementation. How you apply IPsec and SSL/TLS in a VPN solution can affect your VPN's performance. Validating the performance specifications of the solution before you roll it out into production should address any performance issues associated with the protocol startup.

- **Load**—The number of remote access or site-to-site VPNs will affect the overall performance of your VPN rollout. The challenge in addressing this issue, particularly in an environment supporting a large pool of remote clients, is that the performance issues will tend to crop up during peak use. To appropriately diagnose these issues, you will need to understand the network and usage fluctuations to identify times of peak usage. Many of the current reporting tools available for VPNs tend to show averages over time, which can hide peaks and valleys in the use numbers. Be sure you fully understand these performance reports.
- **Client configuration**—In a remote VPN connection, much of the performance is actually related to the client's capabilities. If the remote client is running on old hardware with limited memory and an underpowered processor, the overhead associated with encrypting the traffic will affect performance of the VPN connection. Another factor contributing to overhead is other programs and processes running on the remote PC. If the user is running a memory-intensive application such a photo editing suite, you may find that this resource impact reduces the performance of the VPN.
- **Bandwidth**—The bandwidth available to your VPN can have a significant impact on its performance and can vary widely among the remote hosts and gateways. If your VPN is supporting site-to-site VPNs connecting two locations, the bandwidth allocated at either (or both) ends of the connection may affect performance. You may find, for example, that an unreliable DSL connection at a remote client location creates unacceptable delays for the user.
- **Topology**—Depending on the location of your VPN endpoints, the topology may affect performance. For example, if your VPN connection has to traverse a firewall or a proxy server, you may find reduced performance depending on how well those devices handle the VPN traffic.
- **Encryption level**—The higher the encryption level, the stronger the security, but also the greater the impact on the memory and processor of the endpoint devices. That being said, you should

always run the highest encryption available. If you suspect encryption is causing performance issues, you can look into a dedicated processor for handling encryption or upgrading the processing capabilities of the central processor.

- **Traffic**—An issue related to bandwidth is traffic loads. Let's say the sales department likes to stream video baseball games across the computer network on Wednesday afternoon, and you have VPN performance issues during that time. Increasing bandwidth may fix the problem but does not really address the root cause, which is a traffic spike rather than too little bandwidth. To diagnose a performance issue related to traffic, devise a way to look at the traffic on your network. Another facet of this issue is what the traffic load looks like across the VPN. Do your remote users store their documents on servers in the core network for backups and security? If you are running your VPN with split tunneling disabled, are remote users doing web browsing through the VPN connection? Optimizing traffic within the VPN and on the outside network traffic can go a long way to ensuring that you do not encounter performance issues.
- **Client version**—Sometimes, the software version installed on the client can affect performance. Managing older versions on remote devices can be very difficult depending on how your organization manages those devices. Keep your clients up to date, and you should be able to avoid any performance issues related to client versions.

Stability

To ensure a successful VPN deployment, the implementation must be stable. Stability refers to the ability of the VPN solution to stay “up” and available with consistent performance (no dropped connections or drag). Some factors that can affect VPN stability include:

- **Configuration**—Ultimately, how you configure the VPN will have a major impact on your VPN deployment. Check your internal VPN configuration, and factor stability into your initial design. If access to the network through the VPN is mission-critical, you should ensure use of configuration that includes some level of high availability or failover.
- **Location**—Consider placement of the VPN in the network. If the VPN connection has to traverse three firewalls, multiple local routers, and a proxy server, you may find that the connections are not as stable as needed.
- **Software version**—The version of VPN software (or in the case of a hardware VPN, the concentrator code) can have a significant impact on the stability of the rollout. Be sure to keep the VPN software updated. Careful testing of all software version updates is essential prior to release into the production environment.
- **Underlying OS**—The OS on which you run your VPN can definitely affect the stability of your VPN implementation. A VPN running on an old Windows OS might have issues with the dreaded blue screen of death. A hardware-based VPN could run into challenges if there are firmware or OS issues, although those problems are typically less common than with an OS-based solution. The number of lines of software code needed to run a hardware VPN are quite a bit less than the current OS coding. The leaner, more precise the software, the less risk you run of issues with the OS.

While these lists contain many of the most common sources of performance issues you may encounter, be sure to reference your

troubleshooting processes and procedures when diagnosing VPN performance and stability issues.

USING VPNS WITH NETWORK ADDRESS TRANSLATION (NAT)

VPNs and **network address translation (NAT)** have historically suffered from some conflicts when used together. NAT is an Internet standard that allows you to use one set of IP addresses on your internal LAN and a second set of IP addresses for the Internet connection. A device (usually a router or firewall) stands between the two connections and provides NAT services, managing the translation of internal addresses to external addresses. This allows companies to use large numbers of unregistered internal addresses while needing only a fraction of that number of addresses on the Internet, thus conserving the addresses. This is similar to a company that may have hundreds of phones in a building but pays for only a small number of connections to the phone switch, as it is unlikely that every employee would pick up the phone at the exact same time.

NAT was created as a workaround to IP addressing issues. Because the Internet relies on the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol for communications, it also relies on the IPv4 addressing that is an integral part of the TCP/IP protocol suite. The explosive growth of the Internet threatened to exhaust the pool of IPv4 IP addresses. Without unique addresses, the Internet would be unable to successfully route TCP/IP traffic. This was clearly unacceptable because the Internet was fueling the explosive growth of many businesses. As a result, NAT

was proposed and adopted widely as a way to conserve critical IPv4 addresses.

In the early days of the Internet, when IP addressing was created, developers believed the 32-bit addressing scheme (known as IPv4) to be adequate for any potential network growth. Theoretically 4,294,967,296 unique addresses were available using 32-bit addressing and, even discounting the reserved ranges, more than 3 billion addresses were possible. At the time, there were enough addresses to provide one for every person on the planet.

Unfortunately, the designers of the addressing scheme dramatically underestimated the explosive growth of the Internet, as well as the popularity of TCP/IP in business and home networks. There are no longer enough addresses to go around. IPv6 contains an addressing scheme that allows for a dramatically larger pool of addresses, but because of the cost of switching to IPv6, the move has been slow for many networks on the Internet today. This is due in large part to the use of NAT. Many companies will not move to IPv6 until forced to do so. For some it is due to installed hardware and software not supported in the newer standard. For others it is the lack of skilled employees to make the necessary changes, as well as perform training and testing. Presently, IPv6 networks account for around 40 percent of business networks, with that number to continue to increase. For more information on NAT, see RFC 3022 at <http://tools.ietf.org/html/rfc3022>.

Two main types of NAT are available:

- **Static NAT**—This version of NAT maps an unregistered IP address on the private network to a registered IP address on the public network on a one-to-one basis. This is used when the translated device must be accessible from the public network. For

example, a web server on your private network might have an unregistered address of 10.10.10.10, but a NAT address of 12.2.2.123. A user trying to connect to that website can enter 12.2.2.123, and the router or firewall at the other end will translate that address to 10.10.10.10 when the packet reaches it. A static address is used for devices that require public access, such as to a website.

- **Dynamic NAT**—This version of NAT maps an unregistered IP address to a registered IP address from a group of registered IP addresses. This is more commonly used when large pools of devices on the internal network need to access the Internet and do not have a requirement for a static address. The workstation's address is translated to the next available registered address as soon as it initiates a connection to the public network.

The critical thing to remember about NAT is that due to limitations in the IPsec standard, IPsec has issues traversing a translated network. VPN vendors have addressed this issue, but the workaround can make troubleshooting more challenging. If possible, run your IPsec VPNs on untranslated addresses or deploy an SSL VPN. Because SSL runs at a higher level in the OSI model, it is not affected by NAT. NAT traversal is a general term for techniques that establish and maintain TCP/IP network and/or UDP connections traversing NAT gateways.

For IPsec to work through NAT, configure the firewall to permit the following protocols and ports:

- **Internet Key Exchange (IKE)**—User Datagram Protocol (UDP) port 500
- **Encapsulating Security Payload (ESP)**—IP protocol number 50
- **Authentication Header (AH)**—IP protocol number 51

TYPES OF VIRTUALIZATION

There are several types of virtualization available for SOHO and business installations. Operating system virtualization is the emulation of an OS environment hosted on another OS. A virtual machine exists logically, but it does not have an associated dedicated physical device. Thus, a single physical machine can host multiple virtual machines. Virtualization of storage allows a storage manager to abstract the underlying physical storage technologies and manage storage at a logical level. Network virtualization allows you to run multiple logical switches within a single physical switch chassis. Some SSL VPNs support VPN virtualization.

All of these technologies can potentially affect your VPN environment.

Desktop Virtualization

Desktop virtualization (sometimes called client virtualization) is a concept that separates the personal computer environment from the physical machine by using a client/server model of computing. This model is reminiscent of the thin-client client/server architectures of years ago, where the OS was run centrally. This generation of the model offers significantly more business benefits. A virtualized desktop is hosted on a remote central server instead of on the local hardware of the remote client. This allows users to work from their remote client, while all of the programs, applications, processes, and data used are stored and maintained on the corporate server.

This has an immediate impact on your VPN in ensuring the compatibility of the VPN client with the virtualized desktop. The possibility always exists that the virtualization of the client might cause issues with the client or complicate troubleshooting any issues with the VPN. The best approach to working in a virtualized desktop environment is to test the VPN applications carefully before the virtual implementation moves into production.

Another issue associated with the VPN is how you publish the virtual desktop to a remote client that is not connected to the network. One model is to make the VPN connection first, then publish the virtual desktop across the VPN connection. This works well with SSL VPNs due to the lack of dedicated client software. With an IPSec VPN deployment, a working VPN client on the remote desktop is needed before publishing the virtual desktop. This adds to the complexity of your VPN environment. The other model is to publish the virtual desktops natively from a DMZ environment across the

Internet and then connect back to the network via VPN, publishing the VPN client (if using IPSec) as part of the virtual desktop image.

Some limitations of desktop virtualization include:

- Additional security risks associated with the complex network and desktop image environments used when working with remote virtual desktops.
- Loss of employee privacy as all storage and processing is centralized.
- Maintaining VPN clients in the virtual images.
- Increased downtime in the event of network failures.
- Outages affecting a greater number of employees.

The concept of virtualization is now a standard among enterprise customers. For SSL VPNs, the need for virtualization is natural. Enterprises often provide different remote access VPN presences to different user groups, such as partners and different departments of employees. The following section covers some of the basic capabilities to consider for a “virtualized” SSL VPN deployment.

SSL VPN Virtualization

Virtualized SSL VPNs provide a unique virtual VPN configuration for each individual user group. Much as in other types of virtualization, a virtualized SSL VPN allows the separation of the physical and logical sides of the VPN.

Some benefits of a virtualized SSL VPN environment include:

- **Greater flexibility**—A virtualized SSL VPN provides you with the ability to create custom authentication methods and VPN group policies for different user groups.
- **Delegation of management**—Virtual VPNs allow the delegation of management roles for each virtual instance. If one virtual VPN instance allows a business partner to connect to an internal application, the management of that virtual VPN can be assigned to someone supporting that business partner.
- **Added security when working in a multigroup environment**—Virtual VPNs provide a total logical separation of the VPNs' instances in terms of system resources, routing tables, user databases, and policy management interfaces.

The use of virtualized VPNs is especially attractive to companies that provide VPNs as a service. The ability to create completely separate environments and resource pools is critical. Setting up dedicated hardware environments for each customer does not allow a service provider to take advantage of the economies of scale offered by a single implementation.

CHAPTER SUMMARY

There are different types of VPNs and related protocols. These include IPSec, Layer 2 Tunneling Protocol (L2TP), Secure Sockets Layer (SSL)/Transport Layer Security (TLS), and Secure Shell (SSH). There are several advantages and disadvantages of hardware and software solutions. After reading this chapter, you should be able to determine which solution would be appropriate in your environment.

You learned how some of the challenges associated with establishing performance and stability for VPNs can affect the performance and stability of the VPN. You also learned about the issues network address translation (NAT) can present when rolling out a VPN, as well as the impact of both client and VPN virtualization on your VPN and network.

KEY CONCEPTS AND TERMS

Authentication Header (AH)

Desktop virtualization

Encapsulating Security Payload (ESP)

Internet Engineering Task Force (IETF)

Internet Key Exchange (IKE)

Layer 2 Forwarding (L2F) Protocol

Layer 2 Tunneling Protocol (L2TP)

Network address translation (NAT)

Perfect forward secrecy

Point-to-Point Protocol (PPP)

Point-to-Point Tunneling Protocol (PPTP)

Request for comments

Root cause

Secure Shell (SSH)


VPN concentrator

CHAPTER 11 ASSESSMENT

1. What are the two modes supported by IPSec? (Multiple answers are correct.)
 - A. Transition
 - B. Tunnel
 - C. Transport
 - D. Encrypted
2. Pick two benefits of SSL/TLS over the use of IPSec VPNs. (Multiple answers are correct.)
 - A. Many more deployments in corporate settings
 - B. Fewer firewall rules required
 - C. Granular access control
 - D. Support typically available immediately following the release of a new OS
3. Which of the following is an IPSec protocol that negotiates, creates, and manages security associations?
 - A. Authentication Header
 - B. Encapsulating Security Payload
 - C. Internet Key Exchange
 - D. Transportation Authentication
4. Which of the following is *not* a part of IPv6 IPSec cryptography?
 - A. Confidentiality
 - B. Data origin authentication
 - C. Data integrity
 - D. Translation services
5. Which of the following is *not* a benefit of virtualized SSL VPN environments?

- A. Added security for multigroup environments
 - B. Delegation of management
 - C. Greater flexibility
 - D. Redundant hardware installation
6. Which of the following is a true statement regarding IPSec?
- A. It supports Windows and Linux environments, but not Apple Mac OSs.
 - B. It provides secure node-on-network connectivity.
 - C. It is set to 56-bit encryption for speed.
 - D. It has replaced the need for SSL-based VPNs.
7. Which of the following is *not* a major component of the SSH protocol?
- A. Connection Protocol
 - B. Datagram Protection Protocol
 - C. Transport Layer Protocol
 - D. User Authentication Protocol
8. Which of the following is used by IPSec and provides integrity protection for packet headers and data, as well as user authentication?
- A. Authentication Header (AH)
 - B. Encapsulating Security Payload (ESP)
 - C. Internet Key exchange (IKE)
 - D. Secure Shell (SSH)
9. Which of the following are the two main types of NAT?
- A. Dynamic and static
 - B. Inbound and outbound
 - C. Encrypted and filtered
 - D. Hidden and shielded
10. Which of the following is *not* a benefit of SSL/TLS over the use of IPSec VPNs?

- A. Client flexibility
- B. Guaranteed uptime
- C. Less expense
- D. Platform independence



© Digital_Art/Shutterstock

VPN Implementation

VIRTUAL PRIVATE NETWORKS (VPNs) are a key technology component in today's computer security environment. Recall that the Internet is open to virtually any user, while an intranet is open only to individuals within the organization. A third kind of network is the extranet. Extranets lie somewhere between the Internet and an intranet.

Companies use extranets to connect with suppliers and customers, two constituencies that are essential to business processes. Extranets often use VPN technology to ensure confidentiality and integrity of information. VPNs are becoming an increasingly important component of any successful business's information technology plan. VPNs can substantially enhance the technological architecture of any organization's network through the use of efficient and flexible collaborative technology.

In this chapter, you will learn how to plan a complete VPN installation at your organization, from operating systems (OSs) and VPN appliances to remote desktops.

CHAPTER 12 TOPICS

This chapter covers the following topics and concepts:

- OS-based VPNs
- VPN appliances
- Remote desktop protocol (RDP)
- How to use remote control tools
- How to perform remote access
- Remote Desktop Services (RDS)
- Microsoft DirectAccess
- DMZ, extranet, and intranet VPN solutions
- Internet café VPNs
- Online remote VPN options
- The Tor application
- Planning a VPN implementation
- VPN implementation best practices

CHAPTER 12 GOALS

When you complete this chapter, you will be able to:

- Plan for a remote control VPN using Remote Desktop
- Evaluate hardware VPN devices
- Experiment with Tor
- Set up an Internet café VPN client

- Assess online remote control products, such as GoToMyPC and LogMeIn
- Configure an IPSec VPN

OPERATING SYSTEM–BASED VPNS

An OS–based VPN is convenient because remote servers may be referred to by their assigned Internet Protocol (IP) addresses rather than use network address translation (NAT). This avoids problems inherent in connecting to servers behind a many-to-one NAT configuration. There are several ways to install a VPN using computers running commercial OSs. Configuration of a VPN connection from a client computer can be accomplished using a variety of OSs, including Windows, Linux, and UNIX.

A VPN is a hardware and software solution for remote workers, providing users with a data-encrypted gateway through a firewall and into a corporate network. VPNs were once practical only for large businesses. Today, however, most businesses—large and small—can afford the technology. VPNs are becoming increasingly popular in the small to mid-sized business market. VPNs are ideal for companies with telecommuters, satellite offices, or employees who travel and need to connect to the corporate network via the Internet. If used properly, VPNs block hackers from attempting to access your network to steal sensitive data.

As a data-encrypted tunnel over the Internet, a VPN can offer a robust and secure Internet connection for the organization. It can also be an inexpensive alternative to a dedicated leased line. How do you know if your organization needs a VPN? Here are some key factors you should consider before deciding to use a VPN:

- Does the organization send or receive sensitive data? For most businesses, the answer is probably yes. Most companies have customer information and records, financial records, and proprietary information within internal networks that merit protection.
- Does your organization employ telecommuters, traveling employees, or other remote workers? If so, a VPN can provide two main advantages: It can offer secure network access to employees away from the office, traveling or working off-site, and it can extend the corporate network to them, enabling them to remain productive outside your office.
- Does your company already use Secure Sockets Layer (SSL)/Transport Layer Security (TLS)—encrypted Internet pages? Some companies using Microsoft Exchange servers for email, for example, may already have the encryption protection necessary for remote workers—at least for accessing their email (via Outlook on the web). In this case, the VPN is a built-in feature of the OS. Businesses without sensitive information can use OS-based VPNs and web-based alternatives to a VPN for authentication and encryption, although these may be less secure.
- Does your organization have more than a few employees? A hardware VPN may be more difficult to manage for a company with fewer than five employees. Software- or cloud-based alternatives might work better for such an environment if such expertise is not available in-house.

Suppose you have considered these issues and concluded that your organization does need a VPN. In this case, here are six further important factors to consider:

1. Consider the difference between a VPN based on customer premise equipment (CPE) and one based on an OS. A CPE solution represents the majority of VPNs on the market and is commonly referred to as a VPN appliance. This solution is easy to set up, manage, and maintain. Windows Server 2019 Network

Access is an example of an operating system–based VPN. If you have a server running Windows Server 2019, you can install the Network Policy and Access Services (NPAS) role and configure the server as a VPN server. This requires some expertise with Windows Server 2019 and can be a little more challenging than a CPE solution. However, the OS-based VPN can be cheaper and easier to manage than a CPE.

2. Should you install the VPN yourself or use a managed service? Any competent IT staff can probably install leading commercial products from vendors such as Cisco or SonicWall. Although the do-it-yourself (DIY) approach provides more control over setup and usage, installing a VPN incorrectly can inadvertently open a security hole in your organization’s network. In addition, the administration and management of a VPN in-house can sometimes be complicated. Telecommunications companies including Qwest, Verizon, and BellSouth, as well as several Internet service providers, offer managed security solutions that can save you time and money.
3. Do you have a firewall? A VPN cannot replace a firewall. Some administrators tend to use a VPN instead of a firewall, which is not a smart choice. The purposes of a VPN are to create an encrypted tunnel or gateway through your network’s firewall and to keep out hackers. The VPN encrypts the pieces of data, but the firewall still protects the internal network from outside threats. A VPN without a firewall is not secure.
4. Do you have an OS-based VPN? Regardless of the strategy you end up using, make sure you have an IPSec (Internet Protocol Security)–compliant OS. IPSec is a VPN-supporting technology included in Windows XP through Windows 10 and Windows Server 2008 R2 through Windows Server 2019. Used with compatible VPNs, IPSec guarantees the authenticity, integrity, and confidentiality of network traffic. Interoperability with a VPN may be an issue with Macs or some variants of UNIX or Linux. If you decide to buy a VPN, make sure it is compatible with the organization’s OS.

5. Do you have a wireless local area network (WLAN)? The VPN should operate securely with it. A VPN can enhance the capabilities of a WLAN, but improperly layering a VPN on a wireless network can result in security holes. One method places the WLAN outside the firewall, hosting the VPN behind the firewall to ensure security. Otherwise, wireless network traffic can access systems behind the firewall, canceling the benefits of the VPN. Many organizations use layered firewalls so that the wireless network is protected from the outside while restricting access to the inside. In essence, the WLAN operates in what is called a demilitarized zone (DMZ).
6. Can your organization tolerate a potential decrease in network performance? A VPN may cause a performance lag for internal users accessing the Internet. This happens when 10 to 15 percent of the Internet bandwidth serves as security overhead. VPNs are great for setting up a secure connection, but they can take a measurable toll on connection speed. The tradeoff is that VPNs are worthwhile investments for providing a secure connection for remote and traveling workers.

VPN APPLIANCES

One of the easiest and most cost-effective ways to provide secure access to a network is to purchase an inexpensive VPN appliance and set it up, which will take about an hour of your time. VPN appliances can make secure remote access easy.

When considering the purchase of a VPN appliance, ensure that you have the required complementary hardware in place. First, the VPN appliance must have access to the Internet. Remote users will use the public IP address assigned to the appliance to connect to it. Second, the VPN server must have access to the internal network. It will use internal routing to connect remote users from the Internet to the internal network. Of course, resources in the internal network must be on and available for the VPN users to access them.

VPN USAGE IN ORGANIZATIONS

VPNs serve an organization's computer network in two primary ways. They give remote users access to internal networks or they connect two separate offices. These are the host-to-gateway model and the gateway-to-gateway models:

- **Host-to-gateway VPN**—In a **host-to-gateway VPN**, the mobile user takes specific actions to connect to the VPN. For example, the mobile user would first connect to the Internet from a remote location outside the organization. Once

connected to the Internet, the user could then initiate the VPN to tunnel through the Internet. The VPN appliance or server then acts as a gateway for the user to access resources on the internal network.

- **Gateway-to-gateway VPN**—A **gateway-to-gateway VPN** is used to connect two offices in different locations. For example, an organization could have a main office in Virginia Beach and a remote office in Miami. VPN appliances or servers can operate in both locations with an always-on VPN connection between them. Now users in Miami can connect to resources in Virginia Beach using this gateway-to-gateway model. In this model, users in the remote office do not need to take any additional steps to connect. The gateway-to-gateway model is also called a site-to-site model.

Not long ago, VPN appliances were expensive and required client licenses for each computer, in addition to the appliance itself. VPN technology was too expensive for all but the largest companies. But new products make it possible to install a VPN appliance on virtually any size network for budget-minded organizations and SOHO (small office/home office) networks. For example, Buffalo Technology's 125 High-Speed Mode wireless secure remote gateway is a VPN gateway/firewall router and a wireless access point rolled into one neat package. Another great product is the Linksys WRV54 Wireless-G VPN broadband router, a similar product that provides robust protection for your network. You should know that some VPN

appliance products on the market are designed for home installations. Although these products are very easy to install, they allow only a very limited number of accounts and some of them provide relatively slow access.

Configuring a Typical VPN Appliance

Most VPN appliances are designed for simple and quick installations, with plenty of wizards and an automated setup that makes it easy even for non-computer-savvy people. All you typically need to do is to plug the appliance into your network between your Internet service provider (ISP) connection and your internal network. If your network does not have a router or hub, this device can serve that purpose as well. Once you turn on the VPN appliance, you can use any computer on the network to log on to a web page, complete your configuration, and add user access accounts.

VPN appliances are a secure technology, but you need to take basic security measures to preserve the security of your network and remote connections. When you are configuring user account access on the VPN gateway system, for instance, always change the default settings and never use the default passwords. Also, you should give each VPN user an individual access account. In practice, that means if an employee leaves the company, you do not have to change the access passwords for everyone—you just turn off the associated account.

Client-Side Configuration

Once you have configured your appliance, you will need to configure the software on the computers (clients) connected to the network. The systems designed for small installations assume that you will use Microsoft or Mac VPN client software. Some variants of Linux and UNIX may have built-in client VPN software.

Adding a VPN appliance to your office network gives you a remote access solution that lets you and your staff be more productive from anywhere in the world. Not a bad return on a few hundred dollars and an hour of your time.

REMOTE DESKTOP PROTOCOL

Remote Desktop Protocol (RDP) is a proprietary protocol developed by Microsoft that provides a user with a graphical interface to another computer. The protocol is an extension of the ITU-T T.128 application-sharing protocol. Clients exist for most versions of Microsoft Windows, Windows Mobile, Linux, UNIX, Mac OS X, and other modern operating systems. By default, RDP uses Transmission Control Protocol (TCP) port 3389.

Remote Desktop Connection (RDC) is a built-in application that uses RDP. When RDC is enabled, you can connect to another computer, log on, and perform almost any action as if you are sitting in front of the remote computer. You can do this from a desktop PC to another desktop PC using the same operating system. In large organizations, administrators commonly use this to remotely manage servers from their desktop PCs.

Remote Desktop Connection must be enabled on the remote computer. In Windows 10, you do this by entering “remote desktop settings” in the Windows search bar and selecting Remote Desktop settings. On the Remote Desktop screen, set the Enable Remote Desktop setting to On and then confirm the change (see **FIGURE 12-1**). This opens port 3389 on the remote computer. If the connection goes through a firewall, port 3389 must be open on the firewall. You can launch RDC differently in different Windows OSs. However, one method that works in all current Windows versions is to type MSTSC

at the command line. The initials represent Microsoft Terminal Services Connection.




FIGURE 12-1 Enabling Remote Desktop on a Windows client.
Used with permission from Microsoft.

The client can run other OSs, such as Mac OS, Linux, or UNIX, as long as the terminal services protocol is supported. When connecting to Windows Server 2019, Windows 10, or newer systems, you will need to weaken security to support the non-Microsoft clients.

GoToMyPC is another remote desktop technology that allows you to remotely access your computer from any other Internet-connected computer in the world with almost any OS through a secure, private connection. The application is ideal for organizations that need remote desktop access for up to 50 computers. It is an easy and secure remote-access solution that enables convenient access to email, files, programs, and network resources from home or the road.

NOTE

Microsoft changed the name of Terminal Services to Remote Desktop Services in Windows Server 2008 R2. However, the MSTSC command still works. It is also worth noting that Terminal Services and Remote Desktop Services have much broader usage than just connecting to remote desktops. For example, you can use these services when a Microsoft server is configured as VPN server.

USING REMOTE CONTROL TOOLS

As companies continue to expand networks and use remote offices and telecommuting, there is a need to manage devices from virtually any location. Microsoft has offered a built-in remote control solution called Remote Assistance for modern OSs since Windows XP. It allows help desk professionals or other IT administrators to remotely control a user's system, while the user is watching.

For example, a user may not know how to configure an application. The user calls the help desk for assistance. Instead of trying to talk the user through the steps, the help desk pro can show the user how to do it. The help desk pro can use Remote Assistance to take control of the user's desktop. While connected, the pro will have control of the user's desktop as long as the user allows it. The user is able to disconnect the pro at any time.

Although the built-in Remote Assistance is great as a free tool, it does not meet the needs of every organization. Several third-party tools are available that can provide additional features. These products include TeamViewer, LogMeIn, and RemotePC, among others.

USING REMOTE ACCESS

VPNs allow remote users to connect to a private network over a public network. The private network is the organization's internal network. The public network is often the Internet, but it is also possible for an organization to use leased lines from a telecommunications company to create the VPN connection.

Remote users can be:

- Salespeople on the road
- Field technicians
- Consultants working in customer work sites
- Anyone authorized who needs to have access to internal company resources while away

Because data transmits over a public network, it needs protection.

VPNs use tunneling protocols to establish secure connections. These tunneling protocols include different types of encryption to protect the data.

The Technology for Remote Use

Several protocols support VPNs. These include:

- **Point-to-Point Tunneling Protocol (PPTP)**—This protocol supports Microsoft’s remote access servers and has known issues. It uses Microsoft Point-to-Point Encryption (MPPE). Although PPTP is still used for some remote access solutions, IPsec and SSL/TLS-based solutions are replacing it.
- **Layer 2 Tunneling Protocol (L2TP)**—Cisco and Microsoft collaborated to create this by combining strengths from Cisco’s Layer 2 Forwarding (L2F) protocol and Microsoft’s PPTP. It uses IPsec for encryption. A significant weakness is that IPsec cannot go through a network address translation (NAT) server because NAT breaks IPsec.
- **SSL/TLS-based tunneling protocols**—Due to the limitations of IPsec with NAT, newer tunneling protocols use SSL/TLS for encryption. For example, Microsoft can use Secure Socket Tunneling Protocol (SSTP). VPN appliances can use SSL/TLS-based tunneling protocols. SSL/TLS requires public key infrastructure (PKI) support to obtain and use a certificate.
- **Internet Key Exchange v2—Internet Key Exchange v2 (IKEv2)** is an IPsec-based VPN protocol that uses NAT traversal (NAT-T). NAT-T allows IPsec traffic to pass through a NAT server. IKEv2 provides significant improvements over IKE and has been adopted by several companies, including Microsoft (in Windows Server 2008 R2 and forward), Cisco, and OpenVPN Inc. **OpenVPN** is an open-source solution presented later in this chapter. IKEv2 requires public key infrastructure (PKI) support to obtain and use a certificate.

Each method has advantages, depending on the access requirements of your users and your organization’s IT processes. Although many solutions only offer either IPsec or SSL/TLS, some vendors, including Microsoft and Cisco, offer multiple technologies integrated

on a single platform with unified management. Offering both IPSec and SSL/TLS technologies can enable organizations to customize the remote access VPN without any additional hardware or management complexity.

SSL/TLS-based VPNs enable remote access connectivity from almost any Internet-enabled location using a web browser and its native SSL/TLS encryption. It does not require pre-installation of any special-purpose client software. This makes remote access SSL/TLS VPNs capable of “anywhere” connectivity from company-managed desktops and non-company-managed systems, such as employees’ PCs, contractor or business partner desktops, and Internet kiosks. Any software required for application access across the SSL/TLS VPN connection is dynamically downloaded on an as-needed basis, thereby minimizing desktop software maintenance.

IPSec-based VPNs are the deployment-proven remote access technology used by organizations. IPSec VPN connections use pre-installed VPN client software on the user systems, focusing it primarily on company-managed desktops. IPSec-based remote access offers versatility and customizability through modification of the VPN client software. Using APIs in IPSec client software, organizations can control the appearance and function of the VPN client for use in applications such as unattended kiosks, integration with other desktop applications, and other special deployments.

Choosing Between IPSec and SSL/TLS Remote Access VPNs

Both IPSec and SSL/TLS VPN technologies offer access to virtually any network application or resource. SSL/TLS VPNs offer additional features such as easy connectivity from non– company-managed desktops, little or no desktop software maintenance, and user-customized web portals upon login. The primary drawback with IPSec is that it cannot traverse a NAT server. If you are deploying a VPN server and want the connection to go through a NAT server, SSL/TLS is a sound solution.

It is possible to use NAT traversal (NAT-T) to allow IPSec traffic to pass through a NAT server, but be aware of some issues with it. For example, Microsoft has specifically recommended that NAT-T not be used, though IT professionals still recommend NAT-T with non-Microsoft hosts.

REMOTE DESKTOP SERVICES

Remote Desktop Services (RDS) is a built-in Microsoft Server product with multiple uses. It works in two main modes: Remote Desktop Session Host (RD Session Host) and Remote Desktop Virtualization Host (RD Virtualization Host). With RD Session Host, the server typically provides pools of desktop sessions and applications for everyday user workloads. In RD Virtualization Host mode, the server typically provides personal desktop sessions in virtual machines. Both modes allow an administrator to provide multiple virtual desktops, each of which can be used just like a user's primary workstation.

There are other components to RDS as well, including RD RemoteApp and RD Web Access.

RD RemoteApp

One of the challenges with running applications on remote servers is the odd appearance to users, and they have trouble adapting. RD RemoteApp is a Microsoft solution that runs on a Microsoft RDS server but appears to end users as if it were actually running on their systems, regardless of what OS they are using. Users do not need to open an RDS session; instead, they launch the application from their Start menu or a shortcut on their computer. The application appears in a window on the users' computers just as if it were running on the local computer.

RD Web Access

RD Web Access is an RDS server role that allows desktops and RD RemoteApp applications to launch from a web browser. In addition to the usual RDS infrastructure, RD Web Access needs Internet Information Services (IIS) to work properly.

An administrator can integrate virtual desktops and RD RemoteApp applications into web pages, which appear to launch from a web server. In other words, the clients use a web browser to access a website. From within this website, they can click on a link for the desktop or RD RemoteApp application. RD Web Access can be configured in an internal intranet or accessible to users from the Internet.

RD Web Access allows remote clients to connect to internal resources without the need for a VPN. Depending on what your remote clients need, this may be a suitable substitute.

RDS and Hosted Services

Over the past few years, many software publishers have experimented with offering hosted services. The basic idea behind hosted services architecture is that an organization does not have to purchase licenses for software applications or have the hassles of installing or maintaining those applications. Instead, an ISP or a software vendor leases the applications to the organization. The application actually runs on the service provider's servers, and users interact with the application over the Internet.

This arrangement has some drawbacks, however. For instance, a hosted service takes an application's configuration out of an organization's direct control. It is not uncommon to hear about network administrators who were put out of a job because the companies they worked for decided to outsource all of their applications to a hosting provider. Another compelling argument against the use of hosted services has to do with service availability. If your Internet connection fails, no one can access the hosted applications. Of course, Internet service is more reliable in some areas than others.

Remote Desktop Services for hosted applications has many benefits. The primary one is that the service provider takes care of all of the application maintenance for you. Many of these benefits are things that you just do not get if you install the applications locally on each individual workstation or if you outsource your applications to a hosting provider. Microsoft products can provide hosted applications using RD RemoteApp and RD Web Access.

MICROSOFT DIRECTACCESS

DirectAccess is a Microsoft solution that can be used as an alternative to a traditional Internet Engineering Task Force (IETF) VPN. It allows remote clients to connect to internal servers without initiating a VPN connection. As long as clients have Internet connectivity, the system should be able to access internal resources using DirectAccess.

Microsoft introduced DirectAccess in Windows 7 and Windows Server 2008 R2 products, as well as continued support in Windows 10 and Windows Server 2019. Once it is configured on the clients and servers, it is relatively invisible to the clients. Client computers connect to the DirectAccess computer, which acts as a gateway to internal resources. Only resources configured for access with DirectAccess are available from clients. In other words, you could have 10 servers in the internal network, but choose to make only a few accessible.

For example, you could configure a Microsoft Exchange server (used for email) with DirectAccess. When a DirectAccess-enabled Windows 10 client connects to the Internet, it would automatically connect to a DirectAccess server. When the user starts Microsoft Outlook, DirectAccess automatically makes the connection to the internal Microsoft Exchange server. In other words, users can be on the road and still use their email client just as if they were in the office. The same process works for any servers that an administrator wants to make accessible on the Internet.

A significant added benefit of DirectAccess is that administrators can execute control over the remote clients. For example, in a Microsoft environment, Group Policy can ensure that a system meets the minimum-security settings. Although this is normally not possible for systems that are disconnected from the internal network, DirectAccess with Unified Access Gateway (UAG) allows an administrator to apply Group Policy to these remote computers.

It is possible to use Network Access Protocol with DirectAccess. You can create policies to ensure that the remote system has specific security measures in place. For example, you can ensure that the system is up to date with current security patches and that it has up-to-date antivirus software installed and enabled.

DMZ, EXTRANET, AND INTRANET VPN SOLUTIONS

A demilitarized zone (DMZ) is a physical or logical subnetwork that contains and exposes an organization's external services to a larger untrusted network, usually the Internet. The purpose of a DMZ is to add an additional layer of security to an organization's LAN. An external attacker can gain access to equipment in the DMZ, but not parts of the network behind the firewall.

For example, public-facing servers that need accessibility from the Internet are placed in the DMZ. This could include web servers, email servers, File Transfer Protocol (FTP) servers, and more. Organizations often place VPN servers for remote access in the DMZ.

In a network, the hosts most vulnerable to attack are those that provide services such as email, web, and FTP servers to users outside of the local area network. Because of the increased potential of compromise, these hosts are placed into a subnetwork to protect the rest of the network if an intruder were to succeed.

Hosts in the DMZ have limited connectivity to specific hosts in the internal network, but connectivity to the external network is allowed. For example, a web server in the DMZ may be able to connect to a database server in the internal network, but not to any other hosts in the internal network. This allows hosts in the DMZ to provide services to both the internal and external network, while an

intervening firewall controls the traffic between the DMZ servers and the internal network clients.

Intranet VPNs

An intranet is an internal network. Although users within the intranet can access the Internet using different resources such as a proxy server, access to the internal network is severely restricted. Because traffic in the intranet is primarily from internal clients, the intranet is a trusted zone and needs fewer security measures.

An intranet VPN is a VPN that connects two or more internal networks. Earlier in this chapter, you learned about the concept of gateway-to-gateway VPNs. A gateway-to-gateway VPN provides connectivity between two locations, such as a main office and a branch office. This is also known as an intranet VPN.

It is important to realize that even though the VPN may be called an intranet VPN, it will still traverse a wide area network (WAN) link. Most organizations will rent access to this WAN link; it is very rare that a company has exclusive access to it. In other words, the WAN link will be accessible to users outside the organization. The same level of security measures used in a DMZ VPN should also secure an intranet VPN.

Extranet VPNs

Extranet VPNs link customers, suppliers, partners, or communities of interest to a corporate intranet over a shared infrastructure. For example, an organization may hire a consulting company to evaluate different processes within the organization and provide recommendations for improvement. The organization could create an extranet VPN to allow the consultants access to some internal resources.

Extranets are commonly configured to connect via the Internet, but can use leased lines or even dedicated connections. Extranets differ from intranets in that remote users outside of the enterprise are allowed access to the resources inside the network.

FIGURE 12-2 illustrates an extranet VPN topology. Using digital certificates, clients establish a secure tunnel over the Internet to the enterprise. A certification authority (CA) issues a digital certificate to each client for device authentication. The CA server checks the identity of remote users and then authorizes remote users to access information relevant to their functions.




FIGURE 12-2 An extranet VPN topology.

INTERNET CAFÉ VPNS

An Internet café is a public location that sells Internet access, often by the minute. The café will often sell typical café items such as coffee and sandwiches. However, with the explosion of wireless in recent years, many eateries provide free wireless Internet access to bring in customers. For example, many coffee shop and restaurant locations provide free Wi-Fi.

The challenge when using an Internet café or even an open wireless connection at an eatery is security. Others in the local area may be able to view data in the connection unless it is encrypted. Because the majority of Internet traffic is not encrypted, attackers may be able to gain valuable information by capturing another user's Internet use.

The owner of the Internet café can capture any data that passes through with a free packet sniffer such as Wireshark. Additionally, many free wireless sniffers are available that an attacker can use over a shared wireless connection to capture all of the traffic. Any system on the wireless network has access to the traffic on that network.

An alternative is to use an Internet café VPN connection. As soon as you connect with the Internet café or the wireless connection, you would connect to the Internet café VPN. This would be hosted at your organization. It will encrypt all the traffic and prevent any sniffing attacks.

For example, ExpressVPN is a product your organization can purchase and use as an Internet café VPN connection. Once you set it

up, you can direct your users to connect to it for all Internet access, including from their smartphones.

 NOTE

You can get more details about Express-VPN at
<https://www.expressvpn.com/>.

ONLINE REMOTE VPN OPTIONS

GoToMyPC, LogMeIn, and TeamViewer are remote access and control solutions that perform extremely well. Each product is easy to set up and use, and all offer a similar set of features. The products are similar, but you should know about a few noteworthy differences.

Security

Although each technology handles security slightly differently, all are secure and can be safely used in any environment. LogMeIn and TeamViewer do offer a few more features than GoToMyPC. Both LogMeIn and TeamViewer provide 256-bit end-to-end encryption. GoToMyPC, on the other hand, provides only 128-bit. In addition, LogMeIn and TeamViewer allow you to restrict the times that your computer can be remotely accessed and specify the IP addresses from which it can be remotely accessed—functionality that is not offered by GoToMyPC. If you feel that you need to limit remote access times and/or restrict access to only certain IPs, then you will want to consider either LogMeIn or TeamViewer.

Wake-on-LAN Support

Wake-on-LAN is an extremely valuable feature. Most computers include power management capabilities, allowing it to turn off or go to a low power state when not used for a time. These computers are awakened when sent a specific string of bits in a “magic packet.”

When the computer receives the magic packet, it wakes up. If it was off, it will turn on. If it was in a lower power state, it will go to a full-power state.

File Sharing

LogMeIn provides file-sharing functionality that enables you to email a link to a file on your computer that the recipient can use to download the file (directly from your computer) at any time. To share files in this manner, you do not need to invite the person to share your desktop, and you do not need to be at your computer at the time the file is downloaded. This feature is especially useful if you frequently need to share files that are too big to email. This feature is not available with GoToMyPC.

Remote Printing

GoToMyPC, TeamViewer, and LogMeIn enable you to easily print a document on the host using the printer attached to the client.

Mac and Mobile Device Support

With all three products, you can use a Mac as the client. Only TeamViewer provides support for mobile devices in addition to some Internet of Things (IoT) equipment (see the website for the specific hardware and software specifications).

THE TOR APPLICATION

Tor is an application that uses “onion routing.” Generically, onion routing was designed as an architecture to limit a network’s vulnerability to eavesdropping and traffic analysis. It uses multiple proxy servers or relays to provide anonymous connections. Each proxy server knows only the details from the previous proxy server or the next proxy server.

The proxy servers provide anonymity for users by requesting access to resources and making it appear as if the proxy server is requesting the access, not the original user.

FYI

Data leakage is a common problem with peer-to-peer (P2P) networks such as BitTorrent. Users share data they did not intend to share. As an example, the Top Secret plans for the U.S. president’s helicopter were leaked through a P2P network and found on servers in Iran. Some people think that organizations forbid these types of applications to prevent piracy of copyrighted material. However, the primary reason is due to the inherent security risks that most people simply do not understand. Of course, there is nothing wrong with helping prevent the theft of intellectual property.

Tor was derived from the Onion Routing Project managed by the U.S. Naval Research Lab. However, Tor is not an acronym for The Onion Routing project. Instead, it is simply a brand name—similar to Kleenex for facial tissues. The *torproject.org* website still uses an onion as a logo; however, Tor is not all uppercase.

The goal of Tor is to allow users to browse the Internet anonymously. Instead of going directly to an Internet site, Tor uses the computers of other Tor users as relays or proxies. Any single Tor connection will go through multiple other computers.

Interestingly, even though the U.S. Naval Research Lab originally designed Tor, it is forbidden on most government systems; the primary reason is related to data leakage. Although the Tor network does provide a level of anonymity, the user never knows what other computers the request will go through. Data sent and received can be captured by any of these computers.

For example, in 2007, Dan Egerstad, a security professional in Sweden, collected usernames and passwords for 100 email accounts of users at different embassies. He just installed Tor on his system and then captured all the data that went through it. His computer was used as a proxy in the Tor network for thousands of users, and a simple protocol analyzer captured the data. More than the credentials, he also captured a significant number of sensitive email messages from embassies and Fortune 500 companies. Tor was used by Edward Snowden to capture confidential information that he later shared publicly, which was illegal due to the confidentiality of the information. A 2018 Forbes article exposed Tor as the mechanism used to steal from several blockchain services.

PLANNING A VPN IMPLEMENTATION

VPNs create a secure data link with a branch office, remote employee, business partner, or customer that will enable or require server access behind a firewall. VPNs can provide a secure and encrypted data stream between a firewall and a remote client or server.

This section provides the general installation and configuration steps for a permanent site-to-site VPN tunnel using OpenVPN, an open-source VPN for Windows, Linux, and other OSs.

OpenVPN supports preshared secret key encryption via Static Key mode or SSL/TLS security with client and server certificates. Nonencrypted tunnels using TCP or User Datagram Protocol (UDP) through proxies or NAT are also supported. For flexibility, OpenVPN works with the TUN/TAP virtual networking interface that most platforms support. TUN is short for *TUN*nel, and TAP stands for network *TAP*, both of which are virtual network kernel devices.

Before you attempt this VPN installation, keep the following in mind:

- The IPSec protocol on which VPNs are based will not tolerate network address translation (NAT) of its data packets. If your firewall does NAT, then you will have to disable it specifically for the packets that will traverse the VPN.
- The networks at both ends of the VPN tunnel must use different IP address ranges. For example, the organization's internal network may be using an IP address range of 192.168.0.1 to

192.168.0.254. The other network must use a different address range, such as 192.168.1.1 through 192.168.1.254. To avoid confusion, you may want to use completely different private address ranges for each network, such as 172.16.y.z. or 10.x.y.z.

- Permanent site-to-site VPNs require firewalls at both ends that use static IP addresses.

FIGURE 12-3 depicts a sample topology diagram of a VPN between two environments.




FIGURE 12-3 Sample VPN topology diagram.

Download and Verify Installation Files

Download the Windows installer .exe file from the OpenVPN Community Downloads site at [***https://openvpn.net/community-downloads/***](https://openvpn.net/community-downloads/). Because installation file packages are signed with a GPG key, you must perform signature verification before installing the software. Download the signature file from the Downloads page to the same folder as the installer file, and then verify the signature file.

Installation Tips

You must install the software on at least two computers or virtual machines—one for the server and one for the client. Installing the Windows version of OpenVPN is simple—just double-click the installer .exe file and follow the prompts. On the Choose Components installation screen shown in **FIGURE 12-4**, leave the default components selected and click Next. After a few more screens, click Finish.




FIGURE 12-4 Choose Components installation screen in OpenVPN.
Courtesy of OpenVPN.

Once you finish installing files, you need to modify the sample configuration files, copy key/certificate files to the appropriate location, and then test your new installation.

■ NOTE

These instructions use Windows as an example. You can get details on Linux installation from the OpenVPN website.

Configuring OpenVPN

OpenVPN typically requires a configuration file and key/certificate files, which you should gather and copy to OpenVPN's configuration directory at \Program Files\OpenVPN\config. However, a simple VPN with a single server and client can use a Static Key setup (secret key) rather than create a public key infrastructure (PKI) with key pairs and certificates. One caveat with the Static Key setup is that the secret key must be exchanged using a secure channel outside of the VPN. Consult OpenVPN's installation guide for additional information on Static Key and PKI requirements.

OpenVPN provides sample configuration files in the \Program Files\OpenVPN\sample-config folder. Copy the client.opvn (**FIGURE 12-5**) or server.opvn (**FIGURE 12-6**) file to the OpenVPN\config folder, depending on whether you are configuring the client or server. To edit a file, double-click it to open it in Notepad.

```
#####
# Sample client-side OpenVPN 2.0 config file #
# for connecting to multi-client server.  #
#
# This configuration can be used by multiple #
# clients, however each client should have #
# its own cert and key files.  #
#
# On Windows, you might want to rename this #
# file so it has a .ovpn extension  #
#####

# Specify that we are a client and that we
# will be pulling certain config file directives
# from the server.
client

# Use the same setting as you are using on
```

FIGURE 12-5 A portion of the client.ovpn sample file in OpenVPN.
© OpenVPN Inc.

#####

Which local IP address should OpenVPN
listen on? (optional)
;local a.b.c.d

Which TCP/UDP port should OpenVPN listen on?
If you want to run multiple OpenVPN instances
on the same machine, use a different port
number for each one. You will need to
open up this port on your firewall.
port 1194

TCP or UDP server?
;proto tcp
proto udp

"dev tun" will create a routed IP tunnel,
"dev tap" will create an ethernet tunnel.
Use "dev tap0" if you are ethernet bridging
and have precreated a tap0 virtual interface
and bridged it with your ethernet interface.
If you want to control access policies

FIGURE 12-6 A portion of the server.ovpn sample file in OpenVPN.
© OpenVPN Inc.

Contents of each configuration file include comments that guide you as to which items need to be modified. Directives or attributes preceded by a semicolon are ignored; removing the semicolon allows the directive or attribute to be read and acted upon.

In the server configuration file, edit the ca, cert, and key parameters to point to your PKI files if you decided to use that method. Also edit the remotedirective parameter to point to the hostname/IP address and port number of the OpenVPN server.

Ensure that parameters in the client configuration file match those in the server configuration file, such as dev (tun or tap) and proto (udp or tcp). Also ensure that the client configuration file points to the domain name or public IP address of the server.

In addition, ensure that UDP port 1194 is open on the server. The virtual TUN interface cannot be blocked on either the client or server. This interface might be named Local Area Connection *n* unless you rename it in Control Panel.

Running OpenVPN

At this point, OpenVPN is ready for use. To run OpenVPN:

- Right-click on an OpenVPN configuration file (.ovpn), and select Start OpenVPN from the pop-up menu.
- Issue the openvpn myconfig.ovpn command from a Windows command prompt.

You can also start OpenVPN as a service using Administrative Tools/Services in Control Panel.

Testing and Troubleshooting

To verify that the VPN is running, issue the `ping` command from the server and the client. For example, issue `ping 10.8.0.2` from a server command prompt to attempt to reach the client, and `ping 10.8.0.1` from the client to reach the server.

If this check fails, remember that firewalls can interfere with any VPN, so you should pay attention to your firewall settings.

Remember to open UDP port 1194 on the firewall, or set up a port forward rule to forward UDP port 1194 from the firewall/gateway to the machine running the OpenVPN server. Also ensure that the TUN/TAP interface is not firewalled.

Many firewalls work well with OpenVPN. However, you might need to do the following:

- Create a name for the remote VPNs in the zones file.
- Describe which IPsec interface to use based on the names in the zone file.
- Describe how the networks named in the zones file interact in the policy file.
- Define the public IP address of the remote sites in the tunnels file.

For additional troubleshooting steps, consult the Troubleshooting section on the OpenVPN website.

OpenVPN Private Tunnel

Many SOHOs may not have the staff or skillset required to set up a VPN. However, surfing the Internet from home or over public Wi-Fi can be safer using a simple personal tunnel that encrypts the connection and makes your IP address anonymous. OpenVPN's Private Tunnel product is one such option. Unlike the open-source OpenVPN, Private Tunnel is offered on a reasonably priced monthly subscription.

Once you sign up for a PrivateTunnel subscription and install the software, you choose a server in a region close to you and then connect to the service (**FIGURE 12-7**).




FIGURE 12-7 PrivateTunnel interface.
Courtesy of OpenVPN.

At this point you are connected to the Internet through Private Tunnel, which encrypts, hides, and protects your web browsing and network traffic. You can configure PrivateTunnel settings to launch automatically when Windows starts, or turn it on and off as needed.

VPN IMPLEMENTATION BEST PRACTICES

The VPN is only as safe as the machine where used. Before deploying a VPN, review the implementation best practices, listed as do's and don'ts, in TABLE 12-1.

T
A
B
L
E

1
2
-
1
V
P
N

i
m
p
l
e
m
e
n
t
a
t
i
o
n
b
e
s
t
p
r
a
c

DO DON'T

Pas swor ds	Do change the original password to a strong password you will remember.	Don't write down your password unless it will be stored in a safe.
Soft ware	Do buy or upgrade antivirus detection software. Do ensure your virus definitions are set to update automatically. Do check frequently for updated OS patches and application patches.	Don't go without antivirus software. Don't ignore OS and application updates/patches. Don't use unsafe applications such as peer-to-peer file sharing tools or applications of unknown origin.
Fire wall s	Do enable built-in firewalls. Do use external standalone firewalls whenever possible.	Don't go without either a built-in or standalone firewall.
Hard ware	If connecting via a wireless interface, do disconnect or disable the wired network interface.	Don't enable or connect more than one network interface while using a VPN-connected computer. Don't allow people to use

If connecting via a wired interface, do disconnect the wireless. Do use the VPN for work purposes only.

the computer who might do so unsafely.

Serv
ices
and
prot
ocol
s

Do disable any unneeded services or protocols.

Don't run default services and protocols unless each is specifically needed.

Additional steps you can use for the VPN server include:

- **Use strong authentication**—Ensure that only authorized clients can connect. Because the VPN server will have a public IP address, it is accessible from an Internet user anywhere in the world. If someone can easily log on to the VPN server, that person can easily access your Internet network.
- **Use strong encryption**—The two primary encryption protocols used in VPNs today are IPSec and SSL/TLS. Either of these is strong enough to protect a VPN, but evaluate other protocols for the best match to your needs.
- **Protect the VPN server behind a firewall**—Whether you are using a host-to-gateway or gateway-to-gateway configuration, you should not put the VPN server directly on the Internet. Instead, place it behind a firewall such as in a DMZ configuration. This will provide a layer of protection from Internet attacks.

CHAPTER SUMMARY

This chapter discussed the different types, design, configuration, implementation, and testing of VPNs. It also discussed the main VPN technologies available on the market and best practices for implementation. VPNs can provide remote clients access to your internal network in a host-to-gateway configuration. VPNs provide access between two offices in the same organization using a gateway-to-gateway model.

Many different VPN applications are available. Microsoft provides VPN solutions built into the server OS. Cisco and other vendors sell VPN appliances for easy installation and configuration. You can use Windows, Linux, or other OSs and install free VPN solutions such as OpenVPN.

VPNs are becoming a part of everyday life on the Internet. Many people use them to gain access to resources in their offices, such as email servers and other intranet resources. This trend is certain to become more popular as many companies are finding it less expensive for their employees to work

from home, relieving the need to lease additional office space.

Site-to-site VPNs will continue to be deployed as companies both small and large find it increasingly necessary to share access to the main networks with remote offices.

KEY CONCEPTS AND TERMS

Customer premise equipment (CPE)

Demilitarized zone (DMZ)

DirectAccess

Gateway-to-gateway VPN

Host-to-gateway VPN

Internet Key Exchange v2 (IKEv2)

OpenVPN

Remote Desktop Connection (RDC)

Remote Desktop Protocol (RDP)

CHAPTER 12 ASSESSMENT

1. What is another term for a VPN?
 - A. Bastion host
 - B. Dedicated server
 - C. Data-encrypted tunnel
 - D. DMZ
2. Which of the following is Microsoft's free remote software for Windows Server and Windows 10?
 - A. Remote Data Services
 - B. Remote Desktop Services
 - C. Remote Desktop System
 - D. Remote Dial-up System
3. Which of the following is one of the primary methods for deploying remote access VPNs?
 - A. IKEv2
 - B. IPSec
 - C. TCP
 - D. SSH
4. Which of the following is true of Remote Desktop Services?
 - A. It can host multiple, simultaneous sessions.
 - B. It can support dynamic addressing.
 - C. It can synch multiple proxy servers.
 - D. It can replace suspicious packets with benign packets.
5. Which of the following protocols does *not* support VPN use?
 - A. IKEv2
 - B. L2TP
 - C. NAT-T
 - D. PPTP

6. Which of the following is *not* a remote VPN option discussed in this chapter?
 - A. AdobeConnect
 - B. GoToMyPC
 - C. LogMeIn
 - D. TeamViewer
7. Which of the following products features the ability to awaken when sent a “magic packet”?
 - A. Awaken
 - B. TimeToWork
 - C. SystemOn
 - D. Wake-on-LAN
8. Which organization originally managed the Onion Routing Project?
 - A. Central Intelligence Agency
 - B. National Institute for Standards and Technology
 - C. National Security Agency
 - D. U.S. Naval Research Lab
9. Which command do you use to verify that an OpenVPN VPN is running?
 - A. ping
 - B. open VPN
 - C. VPN on
 - D. VPN verify
10. Which of the following tools is a method for encapsulating IPSec ESP packets into UDP packets for passing through routers or firewalls employing NAT?
 - A. Openswan
 - B. LogMeIn
 - C. NAT-T

D. Protocol 50 ESP



© Digital_Art/Shutterstock

Implementing Network Security

CHAPTER 13


Firewall Security Management

CHAPTER 14

Best Practices for Network Security Management

CHAPTER 15

Emerging Technology and Regulatory Considerations



© Digital_Art/Shutterstock

Firewall Security Management

FIREWALL SECURITY MANAGEMENT is about focusing on the essential security goals of the environment, as well as making sure that the deployed firewalls assist in fulfilling those goals. Following general best practices with care and due diligence will help to ensure compliance. Realize that despite the important—even essential—nature of firewalls, they are not the totality of security; networks need other security safeguards to create a complete solution.

Even with an appropriate, properly configured firewall, ongoing security concerns will invariably arise. Plan for the inevitable attacks, exploits, and threats to security. Hackers might attempt to use tunneling or virtual private network (VPN) services as a mechanism to bypass firewall filters. Be aware of these issues and be prepared to deal with each. Have an incident response plan in place if a firewall is breached.

Firewall management includes testing, monitoring, and troubleshooting. Many excellent tools can help simplify or automate these processes. The best tools require solid information and a

working knowledge of your environment. Detailed implementation plans can be a significant benefit to maintaining reliable security.

CHAPTER 13 TOPICS

This chapter covers the following topics and concepts:

- Best practices for firewall management
- Some security measures in addition to a firewall
- How to mitigate firewall threats and exploits
- Concerns related to tunneling through or across a firewall
- How to test a firewall's security
- Some important tools for managing and monitoring a firewall
- How to troubleshoot firewalls
- The procedures for proper firewall implementation
- How to respond to incidents

CHAPTER 13 GOALS

When you complete this chapter, you will be able to:

- Describe firewall management best practices
- Demonstrate the use of tools for managing and monitoring a firewall
- Troubleshoot common firewall problems

- Create a firewall incident plan

BEST PRACTICES FOR FIREWALL MANAGEMENT

Firewall management best practices are recommendations, guidelines, or standard operating procedures for obtaining reliable firewall security on a real-world budget. Best practices are usually not specific recommendations for products or tools. Rather, these are recommendations for philosophies, stances, or concepts on how to proceed. The following items are suggested best practices, some of which also apply to network security in general. These might not apply to every environment, although you should consider each, and adapt or adopt when appropriate.

Have a written firewall policy. A written policy dictates which firewall features to enable or disable. It also describes how a firewall should handle inbound and outbound network traffic, based on Internet Protocol (IP) addresses, protocols, applications, and other factors. A written firewall policy complements an organization's overall security policy. To write a firewall policy, learn about and thoroughly examine every communication, transaction, and service within, across, and through the IT infrastructure. You can write a comprehensive and effective firewall policy only when you know and understand your operating environment.

Have a plan. This is the foundation of every firewall implementation. Any other method, process, or procedure for configuring, using, and managing a firewall is at risk of failure

without a plan. Only a written firewall policy can guide you toward a successful and reliable firewall deployment.

Understand your organization's infrastructure. Determine which parts of the network need protection to create an effective plan. You must also understand your organization's technology use, assets availability, and resource consumption, in addition to understanding where everything resides and how users work with the infrastructure.

Decide where a firewall is necessary. Make the decision once you understand your operating environment thoroughly. In most cases, every host should have a local software firewall. Firewalls are necessary at every border communication point and at every transition between subnets of different trust, risk, or purpose. Firewalls provide communications control; deploy them liberally.

Perform a risk assessment. It is important to create a useful and relevant firewall policy. Only through a proper risk assessment can you determine the threats and risks facing your environment and its communications. This, in turn, will guide the configuration of a firewall for maximum benefit.

Review the policy regularly. Once a firewall policy is in place, the policy should be reviewed at least annually. Update policy rules to accommodate network changes, such as the implementation of new applications and addition of servers. Verify that all services are still properly protected. Evaluate whether prevention, deterrence, and response have been adequate and effective. If you discover deficiencies, fix them immediately.

Establish a no-exceptions policy. Every service, every transaction, and every communication must pass through one or more firewalls, filtered according to the guidelines of the firewall policy.

Terminate any communication found to take place without firewall filtering immediately.

Focus on establishing a philosophy of default deny rather than default allow. By blocking everything as a starting point, only those features, services, protocols, ports, applications, and users you deem safe and appropriate can proceed by an exception. Using a default-allow stance forces a never-ending stream of explicit denials as new compromises or malicious events arise. *Deny by default/allow by exception* is always the preferred security stance.

Maintain physical security. Monitor all personnel access to firewalls. No one, except the firewall administrator, should have physical access to the firewall devices. Firewalls are concentrations of security and often the main embodiment of communications security. The risk of compromise from unauthorized access is significant.

Limit and filter Internet connectivity. The Internet is a significant risk for most organizations, as an unrestricted or unfiltered Internet connection is a highway for malicious code, social engineering attacks, and intrusion attempts. If an Internet service, protocol, domain, or IP address is not essential to a necessary business task, block it.

Use defense in depth. A multiple-layered defense should be used wherever possible. Do not rely upon a single or individual firewall. Attempt to interlock and layer firewalls along the pathways of communication and transactions. If you do, numerous filtering events will shield each asset.

Configure firewalls to allow Internet Protocol Security (IPSec), VPNs, and other forms of encrypted communications. Use these

where appropriate, but do not universally allow encrypted communications. Encryption can obfuscate malicious activities. Strongly authenticated endpoints are less likely to be sources of malicious actions.

Keep it simple: security (KISS). Focus on designing and configuring firewalls that use simple and direct rule sets. The greater the complexity involved, the greater the chance for misconfiguring, incorrectly ordering rules, or creating loopholes in protection.

Prioritize. Placing a firewall on a network is only the beginning. Once in place, there is maintenance that must be performed to ensure security. Keeping access rules and lists up to date is important to the success of network traffic. The *deny-by-default/allow-by-exception* philosophy should keep minor nuisances to a minimum. Be sure to assess whether the firewall is functioning properly on a regular basis, even if everything appears to be fine.

Harden both internal firewall hosts and border firewalls. Use firewalls hardened against compromise, regardless of the network location. Provide consistent and thorough security throughout the IT infrastructure. Protecting the stability and reliability of firewalls will, in turn, protect the security of the network as a whole.

Always test new code. Before you deploy it onto a firewall, test it—regardless of the source of the new code or whether it addresses a **mission-critical** issue. All untested code is unauthorized, and it should be blocked from all production firewalls. Testing is relatively inexpensive, and recovering from intrusion damage can be disastrously expensive. You will learn more about firewall code later in this chapter.

Always make a backup. Do not allow the only copy of a firewall configuration to be in memory on the firewall. Be sure all backups are easily retrieved in an emergency. If using the cloud for storage, consider download speeds and how long it will take to restore a system. It may be faster to have a traditional media backup on-site for redundancy.

Lock; then watch. If an asset is worth the time and effort to secure, then it is also worth monitoring. Configure each firewall per the security policy, and then watch for attempted breaches of that security. Perfect security does not exist. To improve security, be prepared to respond when breaches occur. Failure to watch a secured asset means only that when the compromise occurs, you won't notice it immediately. Firewalls can be the focus of an attack, just like any other host. Be aware and be prepared for such an attack; this will give you the edge to respond promptly.

Maintain a firewall incident response plan. Bad things will occur. Failures will happen. Breaches will take place. Firewalls can fail. Malicious traffic goes overlooked. Nevertheless, be prepared. Evaluate and examine the realistic threats that face your environment. Plan for the worst. Define procedures to respond to any and all situations.

Focus on balancing security and usability. Firewall filtering should not make all work tasks more difficult. Likewise, essential work functions need not compromise security. It is important to find a balance between these two extremes. Focus on reducing the risk to the infrastructure while enabling users to perform authorized tasks with minimal hassle.

Develop a firewall checklist. Review it for completeness and accuracy periodically, such as every quarter. Confirm every element on the checklist frequently, such as once a week or once a day. You can automate this to an extent, but it often requires human effort to test and confirm that every security mechanism is in place, active, armed, and effective.

Test all uses of the firewall for sufficiency. Perform verification scans of all deployed firewall settings to ensure functionality. Improper installation or misconfiguration can render even well-designed safeguards worthless. Test every new security setting or rule upon installation and at every reconfiguration.

Perform regular vulnerability assessments. Use automated tools with updated databases of security tests and exploit simulations. These tools should confirm patches and updates, verify security configurations, and probe for known vulnerabilities to exploits and weaknesses. Quickly resolve any issues the scans uncover.

Perform penetration testing. After you have improved and tuned firewall security, this is the ultimate test. Hire or develop an ethical hacking team to test the strengths and weaknesses of the firewalls. Ethical hackers use the same tools and attack techniques as criminals, but without the intent to cause actual damage. Professional security assessment teams can customize attacks, modify exploits, and react in real time to fully stress security defenses.

This collection of firewall management best practices should serve as a starting point for the development of an effective firewall deployment. Other valid and useful guidelines deserve consideration, so do not assume this list of recommendations, or any list from any

source, is exhaustive. There are always new lessons to learn and new challenges to face, as well as new wisdom to obtain.

SECURITY MEASURES IN ADDITION TO A FIREWALL

A firewall is an essential part of any security infrastructure. Every host needs a firewall. A traditional network needs border firewalls. Firewalls are concentrations of security, and they are easily recognizable embodiments of an organization's security policy. However, a firewall on its own is an incomplete security solution.

A complete security solution requires many more components (**FIGURE 13-1**) than just a firewall. Proper understanding of the organization's essential processes, threats, and goals will guide you toward the creation an effective security plan. Without a proper risk assessment, however, recommendations for safeguards and countermeasures are just guesses.




FIGURE 13-1 Firewalls: Only part of the security puzzle.

Several common elements characterize most properly designed security systems. These include:

- Multifactor authentication
- Anti-malware scanning
- Full hard drive encryption
- Communications encryption
- Detailed logging and auditing of events
- Intrusion detection systems and intrusion prevention systems (IDSs/IPSs)
- Network segmentation and traffic management
- Private IP addresses and network address translation (NAT) proxies
- VPNs for remote access
- Network access control (NAC)

Although these are often excellent security components, all are not necessarily applicable and required in every environment. Addressing your environment's specific risks is always more preferable than deploying a security product just because it seems like a good idea.

MITIGATING FIREWALL THREATS AND EXPLOITS

Firewalls are one of the most important components of a complete security solution. However, firewalls are not without issues, threats, and concerns. Exploits will cause a firewall to fail. Remember that a firewall is software that sometimes includes dedicated hardware and a host OS. Simply calling it a firewall does not eliminate this concern.

Ultimately, firewalls are software code written by human beings. In either a host firewall or an appliance firewall, the logic and controlling mechanisms are software. Software code is designed and written by people. And whenever people are involved, mistakes or oversights are possible. Any software product is prone to flaws and errors in programming.

Coding errors are not as common in firewall products as in operating systems (OSs) or other forms of software. But do not think that errors do not occur. Most firewall programmers and vendors take extra care to thoroughly test every line of code to exacting quality control and security standards. Standard industry practice requires considerable pilot testing of firewall code before release into the commercial environment.

In spite of these precautions, firewalls have been released to market with software coding errors later discovered and exploited. Hackers are constantly using scanning, testing, and probing tools to discover exploitable weaknesses. When a vulnerability appears, hackers take advantage of the flaw. Some firewall exploits have

caused firewalls to freeze or crash, while others have given the hacker the ability to read or adjust filtering rules. In most cases, vendors quickly release updates and patches to correct these problems. When selecting a firewall, check the version and patching history. Firewall products with lots of patches might not be as reliable in the future as a product with fewer patches.

How long does the vendor typically take to release a patch once a flaw or exploit for the product becomes public knowledge? The longer the vendor takes to release a fix, the longer you will be left insecure, waiting for the update. Although termed “zero-day exploits,” once it is found, a hacker can utilize the vulnerability until patched.

When the firewall vendor becomes aware of an exploit, it quickly develops and releases a patch. An effective security- and patch-management system will enable you to quickly test and update the firewalls that protect your production environment. A delay in the release of a patch or its application keeps the window of opportunity wide open for hackers to compromise your environment.

A good policy is to avoid the first generation or first release of a firewall product. Version 1.0 is likely to include more programming errors than later releases. If the first version of a product seems better than any other product available, thoroughly test the product or wait at least six months after its initial release before deployment. This gives other people—good and bad—a reasonable time to find flaws, test exploits, and trigger release updates by the vendor.

After you have installed a firewall, always install every available patch and update from the vendor. Always test updates first—no exceptions—but then get those installed as quickly as your testing



process will allow. Install only full and final releases of a firewall patch. Never deploy alpha, beta, pre-release, release candidate, or test-build patches on a production firewall.

Another potential concern of firewall security is the filtering used. Each form of firewall filtering or traffic management is vulnerable in some way. The easiest is basic packet filtering.

Basic packet filtering uses a simple and static rule set. Numerous allow- and deny-explicit exceptions compromise the rule set's final deny-all rule. Rules can filter by the IP address and port number of the source and/or destination, as well as the protocol in use. A firewall testing and probing technique known as firewalking can potentially discover some or all of the rules on a static packet-filtering firewall.

Overlapping can cause the full or partial overwriting of datagram components, creating new datagrams out of parts of previous datagrams (**FIGURE 13-2**). An overrun can create excessively large datagrams. These and other forms of fragmentation attacks cause denial of service (DoS) or attempt to confuse intrusion detection or firewall filtering.

Normal Fragmentation



Overlapping Header and Data Sections




FIGURE 13-2 Fragmentation and overlapping.

Fragmentation is a supported function of IP packets. Packets fragment when they encounter network segments that support a smaller maximum transmission unit (MTU). When packets reassemble, the fragments usually return the data to its original configuration. However, overlapping attacks can abuse the fragmentation offset value in the IP header, causing reassembly overlap. Reassembly overlap can create new payloads or headers. For example, if a payload of “EVIL PAYS BLUE LOAD” fragments into “EVIL PAYS BLUE” and “LOAD,” an overlap reassembly could result in the creation of “EVIL PAYLOAD.”

Protections against fragmentation attacks include using modern IDS detection and firewall-filtering features, as well as performing sender fragmentation. Sender fragmentation queries the network route to determine the smallest maximum transmission unit (MTU) or datagram size. Then the sender pre-fragments the data to ensure that no fragmentation needs to occur en route.

The only reliable method of stopping these attacks is to deploy a dynamic filtering system that performs virtual reassembly. Virtual reassembly will piece together both fragmented packets and the original payload. The reassembled payload is analyzed, and only if it is deemed nonmalicious are the original packets transmitted to the internal network. Virtual reassembly is not a completely foolproof technique, but it does greatly reduce the risk of these exploits.

The most common method of exploiting and/or bypassing a firewall is internal code planting. Firewalls are often border sentries protecting internal systems from communications that originate from external entities. Unfortunately, some security administrators use only

inbound firewall filtering, which leaves outbound traffic uncontrolled and unfiltered.

In this situation, if a hacker can plant code internally and trick a user into running code, or an employee brings in code, outbound connections to malicious external entities might be possible. For example, a laptop or USB drive can contract malware and then connect to the internal network, placing the malware behind the firewall. Many hacker tools rely on this technique, such as Loki, Back Orifice, NetBus, and Netcat. A server of sorts is hosted on an external host, and the automatic connection client utility runs on an internal host. The client utility establishes the initial connection as an outbound connection (allowed by the firewall), and then the external host is able to send data or commands back to the internal client through this connection. Often, this form of attack results in a hacker gaining modest to complete remote control over the compromised internal host.

The firewall has to collect, analyze, and respond to each and every packet received on its interfaces; a well-managed DoS attack can consume all available bandwidth of the connecting segment to the firewall, as well as consume all of the processing capabilities of the firewall. This, in turn, prevents any legitimate traffic from reaching the network. Even with a firewall protecting the internal network, a DoS flooding attack can still successfully disconnect or interfere with external communications.

A firewall's vulnerability to DoS flooding is the one limitation or weakness that you cannot fix, improve, or repair by either upgrading the firewall or applying a patch. Upgrading to a stateful inspection firewall addresses fragmentation, firewalking, and even internal

planting of code. Patching will address programming bugs and buffer overflows. Although highly specialized tools are available that can perform partial upstream DoS detection and filtering, in general, there are no fixes to prevent flooding attacks from reaching an Internet-facing firewall.

Knowing these threats and exploits, as well as realizing the probability of others, should remind security administrators to be proficient at the basics. Security management is mandatory to maintain any semblance of security in any environment. Your essential long-term strategies for maintaining security include keeping systems current with patches, using a hardened configuration, staying knowledgeable about new exploitations, and monitoring the environment for successful and attempted compromise.

CONCERNS RELATED TO TUNNELING THROUGH OR ACROSS A FIREWALL

Hacker tunneling that can be accomplished through or across your firewall deployments are a serious threat to your organization's network security. Tunneling is the creation of a communication channel similar to the creation of a VPN. In some cases, it uses actual VPN solutions and protocols. Tunneling can create either a covert channel or a known channel that is unfiltered.

Tunneling uses two techniques. One method is to install a server component on an internal system, and then an external client component initiates the connection. This technique requires the firewall to allow inbound initiation connections to internal systems. A second method is to install a server component on an external system, and then use an internal client component to initiate the connection. This technique requires the firewall to allow outbound connections.

Most firewall deployments strictly limit or restrict inbound initiation of communications to specific ports and services. A tunneling setup using an inbound connection must "hijack" an existing open port or reconfigure the firewall to open another port for use by the tunnel. If the perpetrator of this security violation has the ability to reconfigure the firewall, you have much more serious security concerns to address than tunneling.

Firewalls commonly allow most outbound communications with little or no restrictions. Although it is not the best security stance, it is

convenient and easy. Proper use of a deny-by-default stance should limit both inbound and outbound connections to those explicitly allowed. If the firewall allows outbound connections, then internally initiated tunnels are extremely easy to establish.

Tunnels can potentially use almost any protocol, including IP, ICMP, TCP, UDP, and most application protocols. Specialized tunneling server/client applications can replace traditional payload content with whatever content the hacker wishes. Effectively, this can convert almost any protocol at any layer of the OSI model into an encapsulation or tunneling protocol (**FIGURE 13-3**).




FIGURE 13-3 A tunnel across a firewall.

Tunneling is an obvious abuse and violation of protocol rules. However, if a firewall is not investigating the payload of a frame, packet, segment, and so on, then a hacker can set up a tunnel using a protocol not designed to perform encapsulation.

The problem gets worse when the tunneling system employs encryption. This makes payload investigations more difficult, because the firewall is unable to decipher the content or purpose of the

encoded payload. To combat this complexity, set the firewall to either block all encryption, allow only encryption that originates or terminates at the firewall, or allow encryption only on certain ports (potentially limiting communications to and from specific IP addresses, as well).

■ NOTE

Once a tunnel is open, data can move in either direction. The directionality of the connection has little to do with which end of the tunnel is the controller and which is the ultimate target or victim.

Encrypted tunnels across a firewall are not inherently or automatically malicious in nature. VPNs are encrypted tunnels that commonly create security for personal and business communications across the Internet and within private networks. However, most VPNs use well-known VPN protocols, such as IPSec and Secure Sockets Layer (SSL)/Transport Layer Security (TLS), on standardized ports, with predictable header constructions and characteristics.

Unauthorized or rogue tunneling mechanisms often use odd protocols for encapsulation and can operate on any port, use invalid packet constructions, and violate communication standards.

In addition to the standard VPN protocols used for this purpose, such as IPSec, SSL, TLS, OpenVPN, PPTP, L2F, L2TP, and so on, several other legitimate and subversive services, products, and protocols can configure unauthorized tunnels. These include:

- **Loki**—Loki uses ICMP as a tunneling protocol.

- **Netcat**—Netcat can create TCP and UDP network connections to or from any port.
- **Cryptcat**—Cryptcat is a version of Netcat that creates encrypted connections.
- **NetBus**—NetBus is a malicious remote control tool.
- **Back Orifice**—Back Orifice is a malicious remote control tool.
- **SubSeven**—Subseven (also denoted as Sub7) is a malicious remote control tool.
- **Remote Desktop Protocol (RDP) and Remote Assistance**—These are native features of the Windows OS. RDP is disabled by default, and Remote Assistance requires an invitation.
- **Tor**—This is a double-blind encapsulation system that enables anonymous, but not encrypted, Internet communications.
- **JanusVM**—This is Linux software powered by VMware that creates SSH encrypted tunnels used in combination with Tor. It was last updated in 2009, but it is still used in some exploits, as unpatched or legacy systems are still vulnerable.
- **PacketiX VPN**—This is an encrypted web proxy service. It was last updated in 2012, but it is still used in some exploits, as unpatched or legacy systems are still vulnerable.
- **HotSpot Shield**—This is an encrypted web proxy service.
- **HTTP Proxy**—This is a web proxy service.
- **GoToMyPC, RescueAssist, LogMeIn, Remote Framebuffer, and similar services**—These are third-party remote desktop control services.
- **TigerVNC and TightVNC**—These are virtual network computing applications.

This list of products, services, solutions, and protocols used to create illicit tunnels across a firewall is not exhaustive. Several of these are

legitimate services with valid uses, but when a hacker employs them without permission, especially across nonstandard ports, it represents a serious threat. Any ability to breach a firewall to allow unfiltered, full-duplex communications is cause for concern and response.

The best defense against these tunneling exploitations is to strictly enforce *deny by default* for both inbound and outbound communications. Clearly define in the **acceptable use policy (AUP)** which tools are unauthorized and which are considered security breaches. Use network and host IDS/IPS monitoring. Deploy whitelist controls to prevent the installation of unapproved software.

Another item often included in lists of tunneling issues and concerns is the danger of mobile code. Beyond the scope of this book, mobile code is dangerous because it executes on the client device. Any malware or harmful instructions included in the code are executed on your system. Limit mobile code, such as ActiveX, Java, Flash, Silverlight, and JavaScript, in browsers to minimize the possibility that a web-only form of tunneling might use an authorized web client.

TESTING FIREWALL SECURITY

Firewall tests are an important and integral part of the build and management process. Standard security management practice is to test security to confirm proper configuration, performance, and strength against attacks and exploits. Failure to test a firewall is a serious breach of this best practice and leaves the network unnecessarily vulnerable. This is because a firewall is one of the most important parts of your security infrastructure. A firewall is a concentration of security controls and a network's first line of defense against inbound attacks. Failure to test a firewall thoroughly almost ensures that you will have a breach or intrusion.

Every update, change, or alteration to any aspect of your firewall—or the network segments connecting to a firewall—should trigger another round of firewall testing. Test a firewall as if you were practicing for the Olympics. Repeat the testing again and again, striving to push the limits and improve with each drill. Strive to ensure that your firewall deployment is the best it can possibly be.

TECHNICAL TIP

Testing a firewall involves the use of several tools and techniques. These include automated vulnerability assessment tools, exploitation frameworks, and rogue hacker attack tools.

One tool you can use is a simulated firewall test. Such a test uses an attack simulator to transmit attack packets to the firewall. You can locate the attack simulator inside or outside the firewall to simulate an internal attack or an external attack. An attack simulator can verify that a specific weakness is present on a firewall, without actually causing damage or interrupting production. Most simulator tests are secure by design. Keysight's Threat Simulator System is an example of a simulator that can be used for firewall tests.

Creating a virtualized network environment using a virtualization tool, such as VMware, is a great place to perform virtual firewall tests. A virtual intranet is created, a virtual firewall bastion host is set up, and virtual external systems become active. The administrator then works through a variety of scenarios of attack, both from internal attackers and external ones. Logically, the virtualized environment functions like the real one, but you can test it using techniques that might otherwise damage or interrupt the production network.

Laboratory tests run in nonproduction subnets where you have configured a duplicate of the production environment. The laboratory setup mirrors each system, including the firewall. Test anything in the lab environment that might interfere with production or might cause data loss or system damage.

Laboratory tests and virtualized testing are similar, but both are useful mechanisms; use them both, rather than forgoing laboratory tests in favor of using virtual testing exclusively. The actual physical devices of firewalls and systems might reveal weaknesses not present in the virtualized versions.

Any of these testing configurations can benefit from the use of **fuzzing tools**. Fuzzing tools use a brute-force technique to craft packets and other forms of input directed toward the target. Fuzzing tools stress a system to determine whether it will react improperly, fail, or reveal unknown vulnerabilities. Fuzzing tools can help discover coding errors, buffer overflows, remote exploit flaws, injection weaknesses, and more. The downside to using fuzzing tools is that the tools can take a significant amount of time and bandwidth to discover anything interesting.

Using a variety of testing methods to thoroughly test a firewall is an important part of the security management used to maximize your network's security effectiveness. Plan, design, deploy, test, and then fine-tune your firewall test.

IMPORTANT TOOLS FOR MANAGING AND MONITORING A FIREWALL

Proper security management requires that you manage and monitor the firewalls in your infrastructure. As with every type of security, you cannot just install and forget about the firewalls on the network. Every security safeguard must be locked down, tuned, and monitored over time. Once the firewall has been installed and configured, the work of managing it begins.

Firewall management is about understanding that new threats are coming from hackers on a constant basis, and that the defenses and settings that work successfully today might not be the best defense against new attacks tomorrow. Firewall management is the essential task of maintaining a firewall in spite of these new threats.

Another element of firewall management includes verifying that the configuration settings of the firewall remain in place and do not change without authorization. Device failures, electrical fluctuations, and unexpected reboots can delete or corrupt firewall configurations. Regularly verify that the desired settings remain active on the firewall.

As with all network appliances, applying updates and patches to the firewall is an essential management practice to ensure the device is up-to-date. Always test updates before deployment. Back up firewall configurations before applying new and tested updates. Once you apply an update, confirm that the application of the update did

not reset the configuration to the default settings or make any unwanted or unintended changes.

An area not often considered in firewall management is testing the resistance of the firewall to attacks. This can involve the use of a variety of automated and manual attack tools and techniques. Effectively, firewall testing is the application of penetration testing or ethical hacking against the firewall as the primary target. Using the tools and methodology of criminal hackers in an ethical and boundary-limited manner is an excellent way to verify the ability of a firewall to withstand known attacks.

Subtle changes in the network (increased traffic, traffic from different originating systems) can help in predicting an attack, though it might be a very subtle increase or change over time and not something that would meet a firewall rule. Firewall management includes the use of monitoring tools to watch over the performance and reliability of a firewall over time. Use care when applying these or any tools to the network. Properly research each tool and fully understand the ramifications of deployment. Always make a backup prior to launching any tool. Some recommended tools to consider include:

- **Nmap**—This is a network mapper, port scanner, and OS fingerprinting tool. It can check the state of ports, identify targets, and probe services.
- **Netstat**—This is a simple command-line tool to list the current open, listening, and connected sockets on a system. A socket is the combination of the IP address and the port for any packet.
- **TCPView**—This is a graphical user interface (GUI) tool to list the current open, listening, and connection sockets on a system as well as the service/program related to each socket.

- **Fport**—This command-line tool lists the current open, listening, and connected sockets on a system, as well as the service/program related to each socket.
- **Snort**—This is an open-source, rule-based IDS/IPS that can detect firewall breaches.
- **Nessus**—This is an open-source vulnerability assessment engine that can scan for known vulnerabilities.
- **Wireshark**—This is a free packet capture, protocol analyzer, and sniffer that can analyze packets and frames (**FIGURE 13-4**) as each enters or leaves a firewall.
- **Netcat**—This tool creates network communication links using UDP or TCP ports that support the transmission of standard input and output. It commonly creates covert channels to control a target system remotely or to bypass a firewall. It can test a firewall's ability to detect and block covert channels.
- **Cryptcat**—This tool offers similar capabilities to Netcat, but uses encryption.
- **Kali Linux or Armitage**—These are Linux distributions that include dozens of security and hacking tools, including Nessus and Metasploit. These tools can perform attacks against or through a firewall for testing purposes. The predecessor of these tool sets was the very popular Backtrack PenTesting suite.
- **Syslog**—This is a centralized logging service that hosts a duplicate copy of log files. It provides a real-time backup of every log on every participating host.




FIGURE 13-4 Wireshark.
Courtesy of Wireshark.

These are only a few of the excellent tools to manage and monitor a firewall. But more important than the tools you use is ensuring that the careful management, monitoring, and tuning of your firewalls actually take place.

TROUBLESHOOTING FIREWALLS

Firewall troubleshooting, like any form of troubleshooting, is more about the process than the actual result. Most firewall problems are relatively easy to troubleshoot if you have created a detailed troubleshooting procedure with extensive documentation. The foundation of successful troubleshooting is preparation.

When trouble arises in or around the security of the network, especially in regard to the firewall, you must act promptly to resolve the issues at hand. Troubleshooting is both an art and a skill: to systematically diagnose and eliminate problems, whatever the source. It's a fairly lengthy and tedious process. Knowing a few procedures and commonsense guidelines can improve your troubleshooting skills and help keep downtime and security breaches to a minimum.

You can never have too much useful troubleshooting information. Data, documentation, information, and resources are the primary tools you can use to find and resolve problems. Useful troubleshooting information includes:

- Complete hardware and software inventory (relative to firewalls)
- Paper and electronic copies of configuration settings
- Firewall policy
- Change documentation
- Previous troubleshooting logs
- Activity, error, and alert logs

- Maintenance logs
- Physical and logical network topologies
- Any information about the current problem

Once you have this information in hand, you can begin the overall troubleshooting process. Do not wait until a problem occurs to collect this essential information. Instead, maintain this collection of documentation as a regular, essential element of system maintenance and management. The best time to document the environment and configuration is when you do not urgently need it. When problems do arise, you want to be able to focus on resolving the issue, not collecting resources and tools. Preassemble the tools and resources you will need for troubleshooting so that you will be ready for action.

Murphy's Law dictates that when things go wrong—and they will—they will go bad at the least convenient time. To combat this, be prepared before problems arise. That way, when the firewall fails, locks up, freezes, lets through unwanted traffic, or falls to a hacker, you can focus on dealing with the data loss and downtime ramifications without having to collect the essential materials to resolve the problem.

Being prepared is always a solid first step. Next, develop a troubleshooting plan, technique, or procedure. Some problems are predictable and have a known method of resolution, but many situations will require off-the-cuff brainstorming to craft a new solution. The following are several commonsense elements that you need to integrate into your troubleshooting planning:

- **Have patience**—Keeping your cool and taking your time will pay off; it will allow you to find a solution quickly without making

mistakes, overlooking essential details, or making the problem worse.

- **Know your firewall thoroughly**—The more you already know about the firewall hardware and software, the more you will know how it functions. You can immediately use that knowledge to seek out a solution.
- **Isolate the problem**—Whenever possible, isolate elements or components of the firewall system that are functioning correctly to narrow the range of potential problem sources.
- **Simplify**—Disable or disconnect software and hardware not essential to the function of the firewall. This will reduce the complexity of the situation and may assist in uncovering the cause.
- **Focus**—Work to find a solution to the current, most critical problem. Do not waste time fixing, repairing, upgrading, resetting, or configuring any other problem or aspect of the firewall system until you have resolved the primary problem. You can become distracted by minor details that “only take a second” to address; make a list of these smaller issues and come back to them later.
- **Review change documentation**—Is a recent change responsible for the unwanted activity? If so, try to undo the change to see if the problem stops.
- **Review previous troubleshooting logs**—Consider whether the current problem is the same as or similar to recent problems already in the log. Try repeating successful solutions.
- **Update the troubleshooting log**—Record every action attempted, whether successful or not. Record it into the troubleshooting log, and use it as a journal. Think of something and then write it down. Then try the solution and write that down. Then test for effectiveness and write that down. Next, repeat the failure fix, and write that down. Repeat until resolved, writing down the successful solution and making note of any other

thoughts, ideas, or observations. Once the issues are fixed, be sure to remove any attempts that failed and were undone prior to implementation of the next idea to preserve the change log with only relevant and applicable information.

- **Try the quick and easy fixes first**—Try the fast and easy actions before the hard and complicated options. You might be lucky; if not, undoing easily attempted but failed solutions will be simpler than undoing more complex options.
- **Avoid destructive or irreversible solutions until last**—Attempting to use an irreversible fix is a poor idea early in the troubleshooting process. Only after reversible and/or safe solutions have failed should you attempt more drastic measures.
- **Try the free options before the costly ones**—Always try to perform repairs and fixes in-house, using tools and resources that you already own or can obtain for free. Delay purchasing new resources or hiring technical support until you have exhausted other options. Contact technical support to learn if this is a known issue for which there is an available fix, but delay engaging technical support for any pay-as-you-go service until you have tried to troubleshoot the problem—unless it is a mission-critical situation.
- **Let the problem guide and direct you**—The more you understand how your firewall operates and what the problem is, the more the problem will direct you toward the affected area or the source of the issue.
- **Make fixes one at a time**—Try only one fix or repair option at a time. Attempting multiple fixes at once is more complex and might mask the successful resolution.
- **Test after each attempt**—After each fix is made, test the repair to see if it was successful. Document each success and failure.
- **Reverse or undo solution failures**—If a fix does not resolve the issue, undo it to return to the previous state. Leaving failed fixes

in place may cause other problems, intensify the main problem, or hinder resolution of subsequent attempts to resolve the issues.

- **Repeat the failure**—Sometimes causing the failure to repeat can assist in identifying the cause. However, do so only when the repetition will not cause further harm or loss.
- **Perform a post-mortem review**—The most valuable result of a problem, especially a resolved problem, is your ability to learn something from the event. Always review the entire troubleshooting response process. Look for ways to improve the response to future problems.

Use these commonsense ideas to create an official how-to procedure that your IT and security staff can follow when a firewall problem arises. Firewall problems will not wait; strive to understand and resolve each quickly. Otherwise, your network will be at risk, or productivity can suffer. Neither is beneficial to the long-term viability of your organization.

RESETTING FIREWALL APPLIANCES

Regarding firewall appliances, at some point you may have to restore factory defaults. Many firewall appliances can be quickly reset to factory defaults with the simple press of a button. Some devices have a reset button prominently located and labeled so that anyone can casually revert the device to default status. However, such an obvious button has disadvantages, such as accidental or innocent hand movements.

Fortunately, a growing number of devices that still offer button-based reset options have receded the button

behind a tiny pinhole. A straightened paper clip (or an expensive tiny screwdriver) depresses the reset button. Often, a single quick press is insufficient to reset the device; a series of presses or a prolonged 10- or 30-second depression may be necessary to trigger a reset. This feature makes accidental depression impossible, but it also makes intentional resetting a bit of a challenge. Be sure to consult the user manual for instructions that may not be obvious or intuitive.

DETECTING FIREWALL THREATS AND RESPONDING TO INCIDENTS

Threats to network security come in various forms. Where ransomware is obvious and easy to detect, advanced persistent threats (APTs) are stealthier and can remain undetected for months or even years. However, most firewall breaches occur due to an incorrectly configured firewall.

When a firewall breach is detected through an automated alert or manual log inspection, the incident must be handled immediately. Most incident response solutions include six primary steps or phases:

■ NOTE

Because most enterprise web traffic today is encrypted, attackers exploit encryption to hide threats from firewalls. Next-generation firewalls are designed to inspect and decrypt SSL, TLS, and other encrypted traffic.

- **Preparation**—Select and train security incident response team (SIRT) members and allocate resources.
- **Detection**—Confirm actual breaches.
- **Containment**—Restrain further escalation.
- **Eradication**—Resolve the compromise.

- **Recovery**—Return to normal operation.
- **Follow-up**—Review the process and improve future responses.

A SIRT member, who may be the firewall administrator, is responsible for detecting firewall security issues and reconfiguring firewalls to address vulnerabilities that were exploited. An organization should have a written firewall policy and procedures for restoring a firewall to a working state.

The process often involves disabling the breached firewall. In a redundant setup, all Internet traffic flows through a backup firewall, while the affected firewall is reconfigured. If a backup firewall is not available, Internet service should be disabled temporarily to protect other systems. After reconfiguration, the primary firewall may be restored and brought back online.

Documentation must be updated throughout the process to describe the breach, actions taken, and the new firewall configuration. This documentation also provides a useful tool for preventing future incidents and quick recovery.

CHAPTER SUMMARY

Firewall management involves addressing all security concerns related to the security provided by a firewall. This includes following recommended best practices and knowing the limits of firewall security. Be prepared to mitigate firewall threats and risks, deal with tunneling across a firewall, create a firewall testing procedure, use tools for management and monitoring, troubleshoot problems, and create a firewall incident response plan.

KEY CONCEPTS AND TERMS

Acceptable use policy (AUP)

Fuzzing tool

Graphical user interface (GUI)


Mission-critical

CHAPTER 13 ASSESSMENT

1. How can you know if a firewall is functioning properly?
 - A. Review the documentation
 - B. Review the configuration
 - C. Review the log files
 - D. Review the test results
2. Which of the following is *not* a task completed with tunneling?
 - A. Providing access for external users to the internal network
 - B. Bypassing filtering restrictions for packets
 - C. Ensuring encryption of traffic
 - D. Using open ports for standard communications
3. Which of the following VPN tools provides anonymous, unencrypted tunneling systems?
 - A. Cryptcat
 - B. JanusVM
 - C. PacketIX VPN
 - D. Tor
4. Which of the following is the best option for resolving firewall compromises?
 - A. Add additional rules
 - B. Reinstall the configuration
 - C. Apply outstanding patches
 - D. Use anti-malware scanning
5. Part of troubleshooting is identifying open ports. Which of the following tools aids in scanning if ports are open or closed?
 - A. Fport
 - B. Netstat
 - C. Nmap

- D. Wireshark
6. Which of the following tools is primarily used for network vulnerability assessments?
- A. Fport
 - B. Nessus
 - C. Netcat
 - D. Snort
7. Which of the following is *not* a firewall management best practice?
- A. Develop a checklist
 - B. Build a written policy
 - C. Secure physical access
 - D. Use vendor default configuration
8. Which of the following is a firewall management best practice?
- A. Install the first generation or first release of a firewall product.
 - B. Upon firewall installation, install available updates from the vendor.
 - C. Install pre-release patches on production firewalls as soon as they are available.
 - D. Because firewall vendors test code for quality assurance, an administrator does not need to duplicate testing efforts before installing code on a firewall.
9. Which of the following is the best, first tool to use when troubleshooting firewalls?
- A. Crimping tools
 - B. Fuzzing tools
 - C. Information
 - D. Source code

- 10.** When a firewall breach is detected, what is the first step that should be taken?
- A. Enable firewall alerts
 - B. Document the incident
 - C. Disable the firewall
 - D. Install a new firewall



© Digital_Art/Shutterstock

Best Practices for Network Security Management

COMPUTER NETWORK SECURITY is not a final solution or a task to be completed. Security is a continuous journey.

Safeguards and infrastructures that worked before might offer little or no protection against future attacks. You must constantly develop and deploy new defenses against new exploits. This vigilance is the essence of network security management.

Network security management strives to maintain established security, adjust the infrastructure to future threats, and respond to

breaches in a timely manner. A variety of techniques and tools is used including incident response, host security, backup and recovery, checklists, and security assessments. Network security management is a complex but essential component of long-term, reliable security.

CHAPTER 14 TOPICS

This chapter covers the following topics and concepts:

- Essentials of network security management
- Important network security management tools
- Techniques for limiting network access and implementing encryption
- Important elements of physical security
- Techniques for preventing and deterring incidents, such as containment, honeypots, antivirus, and backup and recovery
- What is included in incident response
- Differences among fail-secure, fail-open, and fail-close options
- The basics of security assessment
- Why change management is important
- What is included in compliance auditing
- The basics of network security troubleshooting

CHAPTER 14 GOALS

When you complete this chapter, you will be able to:

- List examples of network security best practices
- Describe the importance of physical security
- Compose a procedure for incident response
- Enumerate key components of an effective network security installation
- Describe the methods of network security assessment

ESSENTIALS OF NETWORK SECURITY MANAGEMENT

Security is fragile. Hackers need to discover only a single flaw in your defenses to mount an attack. Changes to the infrastructure, whether physical or logical, could open new holes in the network. Additionally, users and personnel may intentionally or accidentally breach security.

Network security management best practices are recommendations, guidelines, or standard operating procedures for obtaining reasonable security on a real-world budget. Best practices are usually not specific recommendations for products or tools; instead, these are recommendations for philosophies, stances, or concepts to use.

Focus on the core security services when designing security: confidentiality, integrity, and availability. Failure to properly and adequately address these essential security services will result in damage, data loss, and downtime. Confidentiality is the prevention of unauthorized access while supporting authorized access. Integrity is the protection against unauthorized modifications. Availability is the assurance that resources are accessible in a timely manner.

Consider designing security, especially physical security, around three central functions: deterrence, detection, and delay. **Deterrence** is the use of security to convince the potential attacker that the efforts to compromise a system are not worth it. The attack may be perceived as too hard or too complex, the attempt too easy to detect,

and the consequences too severe. **Detection** is to watch for the attempts at breaching security so as to respond promptly. **Delay** is to slow down the attack so that even successful breaches give the defenders time to respond in order to apprehend the attacker or prevent further intrusion.

Another common trio of parameters for security design and implementation is prevention, detection, and response. **Prevention** is the use of safeguards to thwart exploitation or compromise. It is usually more efficient, easier, and cost effective to prevent intrusions and breaches than to react to them. Immediate detection of attempted and successful security breaches is important. The longer the time span between a malicious action and an authoritative response, the greater the likelihood the perpetrator will get away without consequence. **Response** means being prepared to contain damage, restrict further compromise, and effect repairs to return the system to normal.

The following collection of network security management best practices should serve as a starting point for the development of your effective security endeavors:

- **Do not rely upon single or individual defenses**—Attempt to interlock and layer defenses. Implement defense in depth or a multiple-layered defense wherever possible (**FIGURE 14-1**). By following a defense-in-depth design concept, you will protect each asset with numerous safeguards. As one defense tool interlocks with another, there is overlap—like medieval armor plates—which improves the overall security. The strengths and benefits of one countermeasure can supplement or compensate for the weaknesses and limitations of another. Even if the layered defenses are weak, this strategy can still provide enough of a delay for malicious activity to be noticed and stopped, or it may

act as a deterrent so a hacker moves to exploit a less time-consuming target.

- **Know the limitations of security products**—Each security mechanism addresses a single or small set of issues within a specific context. The presence of security in one location does not cause a magical blanket of security to exist in other locations. Be sure specific and relevant security solutions are in place where necessary. For example, use of Internet Protocol Security (IPSec) for network encryption does not imply that data stored on clients is encrypted.
- **Assume programs are inherently insecure**—Security is not often a priority or a requirement in programming. When possible, use secure programming quality assurance.
- **Develop a policy for testing all code**—Always test new code before deployment onto a production system. No matter what the source of the code, you must test it. Even if an internal programmer wrote the code, you must test it. Even if the code addresses a mission-critical issue, you must test it. Test all new code without exceptions. Any and all untested code is unauthorized and should be blocked from all production systems.
- **Develop and implement an efficient patch management system**—You need to promptly analyze, evaluate, and install every new patch released from your vendors, whether for hardware or software, which is effective **patch management**. Never shortcut the test-all-code-before-deployment policy, and never implement code based on fear. Instead, have a focused and efficient process for testing and approving new updates so that each can be rolled out to production systems quickly. Keep every system current on patches; remember, however, that pushing out a patch too soon and without proper testing is just as bad as delaying patch approval because you lack the time to examine it.
- **Harden internal hosts**—Do not focus exclusively on one area or issue of security to the exclusion of all others. Provide consistent and thorough security throughout the IT infrastructure. All

networked devices are at risk—perhaps from different threats, but still at risk.

- **Plan for handling failures, errors, intrusions, and downtime—**Focus on what to do when bad things occur. The goal should be a fast and efficient recovery. Failing to plan is planning to fail.
- **Assume that you may become a victim of a denial of service (DoS) attack—**Every communications system is vulnerable to DoS. Do not forget that physical damage can be an effective DoS.
- **Do not assume that a service or protocol is secured by some other layer or service—**Verify that the data traversing a network segment is encrypted or otherwise secured.
- **Learn from your mistakes—**When a problem is uncovered, a design flaw is revealed, or a process is shown to be ineffective, deal with it. Improve the environment to resolve the issue. Review the process and learn from it.
- **KISS—Keep It Simple: Security—**Security is complex enough without purposefully imposing additional complexity. Focus on designing security that the average user can comply with easily and simply. The more arcane and cryptic procedures become, the greater chance that users will misunderstand, fail to comply, or purposefully subvert. Simplicity encourages compliance.
- **Ensure usability of the system once safeguards are in place—**Focus on balancing security and usability. Absolute network security can create an unusable system, and ease of use can create an unsecure network; it is important to find a balance. The focus should be on reducing risk to the infrastructure while enabling users to perform authorized work with a minimum of hassle.
- **Prioritize—**Security concerns can always overwhelm available time, effort, or budget. Focus on the big-impact and big-result issues first rather than attempting to fix every distraction as it is found. Security is constantly changing; the goal is to maintain reasonable—rather than perfect—security.

- **Know your network**—Always be fully aware of the state of the organizational security. Do not make assumptions. If you are not positive that an aspect of the organization is secure, find out. Assuming that security exists leads to a false sense of safety. A lack of knowledge about the security status can lead to complacency. By assuming that nothing is wrong, you feel no urgency to investigate and rectify problems. Do not fall into the “I thought it was protected” trap; if you do not know or are not sure, take the time to investigate and be sure.




FIGURE 14-1 An example of defense in depth around an asset.

Part of this process is to create and maintain benchmarks of network performance. Without such measures, it may be difficult to identify whether an apparent lag on the network is an attack or the network is reaching 95 percent utilization. Part of network management is maintaining optimal performance for all authorized users.

- **Critically search your network for flaws**—Look for the weakest link. Every system or structure has a weak point, and a security structure is no different. It might be power, the users, a coding issue, processing capacity, or something that you never considered. In any case, some weak link exists in the security chain. Think like a hacker and look for it. Find it. Improve it. Then, go looking for the next weakest link.
- **Be realistic**—Focus on known, real, probable threats rather than on unknown, imagined, or possible threats. While it is possible to have snow in Houston, Texas, is it likely you will encounter a blizzard there during your lifetime? Probably not. Do not plan for emergencies that are almost impossible to occur. You don't have enough time, energy, or budget to protect against everything. Perform a risk assessment and focus on current, significant issues.
- **Create and refine procedures**—Develop a standardized, procedural-based process for hardening new systems. You should subject every new piece of host equipment—including switches, routers, firewalls, servers, and other devices—to a rigorous and thorough hardening process before it is deployed onto the production network. In fact, having a dedicated, isolated, new system-testing network will protect the existing production network from new systems, as well as protect new systems from threats from the existing production network. As you improve your processes, be sure to document these improvements for easy repetition next time.
- **Know the threats**—Stay current on security and vulnerability research. If you are unaware of the threats, you cannot protect

against them. You need to seek out this information, because most general news sources, even those focusing on technology, do not address most of the security vulnerabilities, exploits, or compromises that arise daily. Find good resources-related security and vulnerability research (**FIGURE 14-2**) and consult it regularly. Knowledge is often one of the best weapons in the fight for security.

- **Implement the principle of least privilege**—All users, including administrators, require only the privileges, access, and permissions necessary to accomplish their assigned work. This is referred to as the principle of least privilege. Any abilities beyond this minimum increase the potential for compromise and abuse.
- **Employ separation of duties**—Isolate and compartmentalize administrative privileges through the implementation of **separation of duties**. This is also known as split knowledge. Administrators are given administrative power over a limited area of the IT infrastructure. No single administrator has full or total power over the entire environment. This limits the scope of potential abuse and harm by disgruntled administrators, as well as by hackers who compromise administrative accounts. For example, the IT personnel who are responsible for maintaining the firewalls on the network do not have access to the network logs captured and maintained by nonfirewall devices. This prevents anyone of the firewall team from easily compromising the firewall rules and then covering their tracks by deleting or manipulating the network logs.

General +

Vulnerabilities +

Vulnerability Metrics +

Products +

Configurations (CCE)

Contact NVD

Other Sites +

Search +



JSON 1.1 Vulnerability

Feed Released!



CVSS Version 3.1

Official Support!



CVSS/CWE from CVE

List now Supported!

The NVD is the U.S. government repository of standards based vulnerability management data represented using the Security Content Automation Protocol (SCAP). This data enables automation of vulnerability management, security measurement, and compliance. The NVD includes databases of security checklist references, security-related software flaws, misconfigurations, product names, and impact metrics.

Last 20 Scored Vulnerability IDs & Summaries

CVSS Severity

CVE-2019-4451 — IBM Security Identity Manager 6.0.0 is vulnerable to cross-site scripting. This vulnerability allows users to embed arbitrary JavaScript code in the Web UI thus altering the intended functionality potentially leading to credentials

V3.1: **5.4 MEDIUM**
V2: **3.5 LOW**

FIGURE 14-2 National Institute of Standards and Technology (NIST): National Vulnerability Database.
Courtesy of NIST.

This list of recommendations is not exhaustive. You will learn about many other best practices throughout this chapter.

NETWORK SECURITY MANAGEMENT TOOLS

The best network security management tools are neither commercial nor open-source products or solutions. Instead, the best network security management tools are quite simple and obvious. The best tools are:

- A written security policy
- A security checklist
- A complete inventory of all hardware and software
- A physical cabling layout and device location map
- A logical organization, addressing, and subnetting map
- Complete configuration documentation for every device
- Change documentation and log
- Backup and restoration procedures
- An incident response, business continuity, and disaster recovery strategy
- Troubleshooting guidelines
- Hardware and software documentation
- Personal knowledge and skill
- Access to online resources

Security management is not about having the most expensive products or the most automated configuration. Instead, good security

management is rooted in a solid understanding of the infrastructure and having the tools to improve, respond, and repair as necessary. Security management should focus on protecting assets, supporting authorized activities, and responding to threats as each is discovered. It may be more efficient or cost-effective to use off-the-shelf security management products, but using those devices is no substitute for addressing the organization's core security concerns.

Security Plan and Security Policy

The foundation of any security endeavor is a written plan. To have a plan, you must thoroughly understand your organization's infrastructure, its mission and goals, and the processes necessary to produce its products and services. This means understanding the technology that your co-workers need and use, which assets are involved, which resources are consumed, where everything resides, and how users access the infrastructure. Consider larger and more complex environments in light of the typical IT infrastructure.

A written security policy establishes a documentation trail that everyone in the organization can follow and subscribe to. A written security policy has no substitute. Only a written security policy has the potential to succeed.

To write a comprehensive security policy, you must first thoroughly inventory and examine every component of the IT infrastructure. Only when you fully understand the environment can you create a comprehensive and effective security policy.

Every device connecting to the company network must be in full compliance with the security policy and current, with approved patches and updates. Block any device not in compliance with these restrictions from interacting with the network. This effectively is the function of a network access/admission control (NAC) service. But it should also be a written element of the security policy. Define “no exceptions” in writing, and then use technology to enforce this element of the plan.

TECHNICAL TIP

To create a useful and relevant security policy, perform a risk assessment. The process of understanding assets, vulnerabilities, threats, and likelihoods is at the heart of risk assessment.

Security cannot succeed without the solid endorsement and support of senior management. Executives should support the security policy, adhere to the plan, and stand behind the decisions made by the IT staff. When the security staff and the executives are focused on the same goal, solutions will follow. When these two groups are not working in concert, time, money, and effort go to waste, and network security fails.

Once you have designed and deployed a written security policy, regularly review the policy. Investigate whether the overall quality and reliability of the existing security is sufficient or in need of improvement. Verify that assets are still properly protected. Evaluate whether prevention, deterrence, and response have been adequate and effective. Wherever you discover a deficiency, act to improve and rectify the situation immediately.

Security Checklist

One method to maintain the efficacy of security is through regular verification and validation checks of every countermeasure, safeguard, security control, deterrent, prevention, and defense. This requires an inventory of all security measures. This inventory can then become a checklist.

Physical security and logical security each need a separate list focusing on each respective area of concern. A security guard or a security tech investigates every physical security measure to ensure that it is still in effect, active, and unmodified. Such tasks at the level of logical security are best left to security techs with proper administrative access.

A security checklist run-through should occur daily. Every single item on the checklist should be physically visited and inspected. If any issues or concerns arise, document each and immediately bring it to the attention of the security team. Employ remedial measures promptly and investigate to track down the root cause of the concern. Consider it like a pilot's preflight checklist. Because the same process is repeated with regularity, it could be easy to forget a step or misremember whether that step was validated today or yesterday. The checklist provides the means to document the date and time of each element that was checked.

A physical security checklist should include every security control deployed for facility control. These include:

- Checking every window lock
- Checking every door lock
- Checking every external wall

- Inspecting access points to raised floor areas
- Inspecting access points to drop ceilings
- Ensuring that cabinets or containers are locked
- Verifying that security cameras are pointed in the correct direction
- Verifying that all light bulbs are the correct type and are functioning
- Checking motion detectors
- Testing alarm systems
- Interviewing security guards and confirming compliance with procedures

Depending on the industry and specialization, many more components can make up the physical security of an organization. This list is just a sample of what the typical physical security checklist should contain. Look for any modification, destruction, or general failure of every element of physical security. A single bad window lock or a door that does not fully close automatically can result in easy access for a burglar intent on computer mischief.

A logical security checklist should include every security control deployed for computer and network control. These include:

- Checking authentication
- Checking authorization and access control
- Auditing systems
- Verifying firewalls and other filters
- Checking proxies and other communication management solutions
- Verifying encryption, including key management

- Updating antivirus software and scanners
- Backing up and storing archival information securely
- Verifying logging is functioning and that the results are secured

Many more elements exist in a well-designed security infrastructure for an exhaustive logical security checklist. Again, this list is just a sample of some common elements found in most secure organizations. Poor configurations, overlooked defaults, and out-of-date systems can leave gaps that a probing hacker will discover and exploit. Never assume you are secure. Check and verify everything regularly. No one else will do it for you.

A security checklist can be an effective tool in managing and maintaining network security, in spite of change, accidents, and human nature. Make it a point to start your list now and use it in real-world review immediately.

LIMITING NETWORK ACCESS AND IMPLEMENTING ENCRYPTION

Controlling access to your network, and the types of actions users may perform, greatly reduces the risk of a successful breach.

Encryption helps to protect sensitive data whether it is in storage or transit. Keep the following best practices in mind when securing your network:

- **Limit and filter Internet connectivity**—An unrestricted or unfiltered Internet connection is a pathway for malicious code, social engineering attacks, and intrusion attempts. If an Internet service, protocol, domain, or Internet Protocol (IP) address is not essential to a necessary business task, block it. Computer access to the organization's network and assets is for work tasks, not personal activities.
- **Implement multifactor authentication**—Single-factor authentication, especially password-based authentication, is no longer a truly secure or sufficient method of protecting the logon process. Adding at least one additional factor significantly increases the authentication security of the environment. Once multifactor authentication begins, allow no exceptions that fall back on single-factor solutions. Otherwise, intrusion efforts will focus on the single-factor pathway as it is the weakest link in the network's defenses.
- **Do not rely on authentication at session initiation alone**—Session hijacking is a serious threat on most commonly used protocols. Use solutions that support periodic midstream reauthentication, as well as communications encryption.

- **Eliminate or severely limit remote access**—When possible, avoid remote access. Limit access to local direct connections only. When remote access is necessary, require a virtual private network (VPN). Once you allow remote connectivity, you allow the possibility of remote hacking. Remote hacking attempts to breach security without the need to be physically at or within the target's facilities. Requiring VPNs for all remote connections reduces this threat by preventing open connections or communications vulnerable to eavesdropping.
- **Restrict access**—Focus on establishing a philosophy of default deny rather than default allow. By blocking everything as a starting point, only those features, services, protocols, ports, applications, and users that management deems safe and appropriate can be enabled by exception. A default-permit stance means you must create a never-ending stream of explicit denials as you detect new compromises or malicious events. *Deny by default/allow by exception* is always the preferred security stance.
- **Employ device and file encryption**—Any device that is portable or easily accessible by a nonemployee should use whole hard drive encryption or at least file encryption. Although it is not absolutely foolproof, encrypting an entire hard drive makes accessing the data on a portable device much more complex. If a system is stolen when powered off, the hard drive's encryption typically presents an insurmountable barrier. If the equipment is running when stolen, a hacker might bypass the hard drive encryption by hacking the encryption key out of active memory. However, this is an unlikely attack.
- **Encrypt network traffic**—Require encryption of all internal network communications. Use IPSec to secure all intranet communications. This is often an easy and cost-effective means to quickly reduce the risk of eavesdropping, man-in-the-middle, replay, and many other forms of network attacks. Well-encrypted data is much less likely to fall into the hands of hackers, thieves, and unauthorized personnel.

- **Implement whitelist application control**—A whitelist is a list of allowed exceptions against a background of default deny. By configuring a system with a whitelist, you block all executables unless each is on the list of allowed exceptions. A whitelist consists of the application name, its file name, and the file's hash value. This prevents spoofing by malicious or unapproved applications that simply rename the executable.

TECHNICAL TIP

Don't overlook the importance of authentication, authorization, and accounting. These services are essential to a secure infrastructure. You must control who is allowed onto the network (authentication); place limitations on what users can do (authorization); and hold all entities—even upper-level management—accountable for their actions, activities, and results (accounting).

Remember, the strongest security comes in layers, making it difficult for an attacker to break through each line of defense to gain access.

PHYSICAL SECURITY

There is no security without physical security. Logical protections can protect only against logical attacks. Physical protections protect against physical attacks. With just a few moments of physical contact, a hacker can overcome or bypass most logical security.

Physical security is the prevention of direct physical contact, access, or influence of an unauthorized person to a sensitive component of an IT infrastructure. Any well-designed network security solution will include physical security components. If you fail to address physical security adequately, you fail to understand the threats facing your organization.

Physical security includes a facility that resists forcible entry. Select, improve, or build facilities that prevent and deter unauthorized physical intrusion. If physical breaking and entering is possible and potentially profitable for a thief, it is more likely to take place.

Physical security should include the monitoring of all personnel activity. Solid physical access control includes keeping track of when an employee enters and exits every secured or sensitive area. Good physical access control blocks outsider access and enforces escorts for visitors, consultants, and other authorized nonemployees. This is an often-violated area of a security policy. Locks on the doors serve a purpose but may be bypassed if there is no monitoring or alarm. Consider the employees who prop open the back door to a facility when they leave to take a smoke break. If the doorway is out of their view, anyone could sneak inside. Even if they can watch from where

they are, how many would confront a stranger and attempt to deter their entry? Unfortunately, we live in a world with active shooters and those who would harm us because of where we work, our religious or political beliefs, or simply because they can. Bypassing physical security essentially renders all other security components ineffective.

All servers and important devices should be kept in locked and secured areas to prevent tampering and accidents. When labeling a device, it is helpful to identify an internal strategy for easy access that would not be recognizable to an outsider. Instead of affixing a label noting “web server” and the corresponding IP and Media Access Control (MAC) addresses, it would be more secure to simply denote the information in a code known to IT administration staff, but not to others.

Physical security should include card key entry, burglar alarms, motion detectors, and security cameras. Use the best available technology to both deter and detect physical security violations. Education in physical security is also essential. For example, our natural tendency is to hold a door for someone entering a facility, but that would defeat the key card required for entry and fail to provide an accurate accounting of who entered the building and at what time. This could be disastrous if there was an emergency or evacuation and an accurate head count was necessary.

TECHNIQUES FOR PREVENTING AND DETERRING INCIDENTS

Everyone in your organization is part of the physical security effort. All individuals are responsible for staying within the security boundaries. They are also responsible for reporting violations or suspicious activity. Security is about locking and watching. Not every employee can assist in the process of locking things down, but all can and should assist with watching for violations and suspicious behaviors.

Use all available technologies—in addition to human eyes—to look for security violations. Use auditing, intrusion detection systems/intrusion prevention systems (IDSs/IPSs), security cameras, and motion detectors where appropriate. Proper network security management includes proper design and auditing to watch for violations.

Violations can occur in the physical world, as well as within the logical world of computers, networking, and the Internet. Do not assume that just because a violation causes no direct physical harm, the issue is not worth knowing about or pursuing. A crime is still a crime, whether it is a physical crime or a logical one. Design and implement detection mechanisms to notice, log, and monitor suspicious or abnormal activities.

Imposing Compartmentalization

Compartmentalization is the element of infrastructure design that takes into account the likelihood of a security breach by malicious code or some other intruder. Compartmentalization distinctly separates a network into areas or zones of access. Often, these zones are around departments or entities that commonly interact over the network. Any entity that does not commonly interact is separated by default, using network security access control devices such as routers, switches, and firewalls.

The purpose of compartmentalization is to create small collectives of systems that support work tasks while minimizing risk. The typical risk is that if one system is compromised, the breach can quickly lead to other systems being compromised. This is especially true of malicious code, which distributes rapidly. It can also apply to remote hacker intrusion, as the hacker attempts to leap from system to system, searching out valuable resources.

Network compartmentalization is similar to the bulkheads on ships. In the event of a hull leak, the bulkhead doors of the affected compartment close and seal. This prevents the entire ship from being sunk by a single breach. The same is true for computer network compartmentalization. When you detect a breach, you can sever the links between network compartments to prevent any further malicious communication to or from the affected sections.

Using Honeypots, Honeynets, and Padded Cells

Another specific intruder trap you can add to a security infrastructure is a honeypot. A honeypot is a hacker or intruder trap. A honeypot is a system or network that appears to contain valuable information and be part of a legitimate organizational network; in fact, it is a specially designed trap (see **FIGURE 14-3**). The purpose of a honeypot is to trap hackers/intruders, develop a better understanding of new attacks, or simply serve as a decoy to deflect attacks from reaching the actual primary network.




FIGURE 14-3 A honeypot captures and allow analysis of attack traffic.

A honeypot can be a single system or a network of many systems constructed to look, act, and operate like a real network environment, containing resources that seem attractive but are actually fake. For a honeypot to be effective, it must look and act just like a real target. Otherwise, hackers will quickly uncover that it is fake and search harder to find the real targets. A honeypot often includes extensive focus auditing and recording capabilities to thoroughly capture all activities that take place within it.

Honeypots can be single systems or multiple networked systems. A network of honeypots is sometimes called a **honeynet**. Both honeypots and honeynets are decoy systems that are always operating in the hopes of attracting would-be hackers and intruders. Another form of honeypot is the **padded cell**. The padded cell is a system that waits for an IDS to detect attackers and then transfers the attackers to a special host where the attacker cannot do any damage to the production environment.

Configure and situate honeypots to not interfere with authorized users performing authorized activities. However, you need to position these along common pathways of access that an outsider or intruder might take, in much the same way a hunter will hide along a game trail in the forest. A honeypot should be attractive and interesting enough to persuade a hacker to attack it, but not too attractive as to trigger that voice in the back of the hacker's mind that says, "It's a trap."

FYI

Designing and implementing a honeypot is beyond the skills of most IT professionals. Crafting a honeypot takes a level of skill at deception and attention to detail to ensure that the data within the honeypot is benign and not real information or accounts. Honeypots require constant monitoring, as a compromise could potentially become a platform to launch more attacks. The U.S. government, think tanks, educational institutions, security research firms, and some large organizations operate honeypots.

Visit the OWASP Honeypot Project website at [**https://owasp.org/www-project-honeypot/**](https://owasp.org/www-project-honeypot/) for more information.

A honeypot has value in two primary circumstances. First, if the organization's primary goal is to learn about new hacking exploits, then a honeypot can serve as a laboratory for that purpose. Second, if the organization's primary goal is to prosecute criminals, then a honeypot may assist in the discovery, identification, and apprehension of suspects. However, if neither of these two purposes is a primary goal of the organization, then deploying a honeypot is generally not a worthwhile endeavor.

Antivirus, Monitoring, and Other Host Security Controls

Every host, whether a portable device, workstation, or server, needs protections from its users as well as from the network it communicates with. There are always risks, even within a private network, from authorized users. No system is immune to malware and malicious communications; basic filters and barriers to known issues prevent easily avoidable compromises and downtime. Precautions include equipping host devices with reasonable defenses against known potential avenues of compromise.

Every host needs an antivirus scanner. Although malicious code is not as serious a threat in some OSs as in others, hackers create malware for every platform. You can no longer safely claim immunity based on your OS selection alone.

An antivirus scanner needs to be current. The scanner engine—the core mechanism for malware detection—should be less than a year old. Typically, a deployed antivirus should have the current year in its name or as its build date. Any antivirus scanner more than a year old is too out of date with new and current scanning technologies, malicious embedding techniques, and cleanup and removal procedures.

An antivirus scanner needs to have its database of definitions updated at least once per day. This should be an automated process so that as soon as the vendor releases a new update, it is downloaded and applied. Using a new scanner with an old database is just as bad as using an old scanner. An antivirus product provides protection only against known threats, so do not allow your antivirus software to become outdated.

An antivirus scanner must be configured to perform constant, consistent, and automatic scans. Monitor the memory, processing stack, and any other active elements of a computer system in real time for symptoms of malicious code compromise. When malware reaches a host, it can infect that host and copy or distribute its code in moments. Without real-time monitoring for infections, serious damage can take place long before a periodic scan detects the breach.

An antivirus scanner should be configured to perform complete, systemwide, low-level scans across all memory and storage devices on a periodic basis, every day or at least once a week. By scanning the entire system on a regular basis, you can detect and remove malware that was able to gain access through some unmonitored or covert channel.

In addition to an antivirus scanner, hosts should benefit from an anti-malware scanner. An anti-malware scanner is also known as a spyware or adware scanner. These are companion scanning products to antivirus software in that these forms of malicious, nuisance, or suspicious code might not qualify as a virus or worm. These targets include spyware, adware, keystroke loggers, rootkits, hacker tools, backdoors, Trojans horses, malicious/abnormal cookies, suspicious mobile code, bots, zombies, and so on. Do not let the names fool you; all of these things represent the potential for real network damage.

Every host should be equipped with a software host firewall. A host firewall provides filtering services for data entering and leaving the local box. A host firewall can limit network connectivity for local applications, as well as allow or deny access to resources from external entities. Even if a hardware firewall is on the entry points to the network, each server still needs a software firewall to protect it. A

single miswritten or incorrectly ordered rule in the rule set could render the hardware firewall ineffective. A software firewall on the host adds protection against anything that gets through from the network or attempts to access the network from that particular host.

Most hosts can benefit from the added security of whole hard drive encryption. This ensures that all data on the drive is secured. Depending on the product and technology you employ, whole hard drive encryption can be unlocked through a boot password or the use of a hardware device, such as a USB drive or a **Trusted Platform Module (TPM)** chip. Whole hard drive encryption essentially prevents a hard drive from being read by another system if it is stolen.

Some forms of system and file damage occur from activities that are not malware-related. Using a hash integrity checking mechanism to watch for unauthorized file changes can assist in tracking down abnormal sources of compromise. Monitor every system file and device driver. Any change not associated with installation of a valid update is likely a symptom of malicious activity.

You should configure and review auditing on all hosts, not just servers. Malicious activities and communications can take place on any system. Configure every host to audit, monitor, and log local and network events. Then, using an IDS/IPS or other form of security auditing tool, scan the log files for symptoms of concern. What you do not know can harm you, so arm yourself with as much knowledge as possible.

Backup and Recovery

There is no excuse for not having a **backup** solution. Backups are the only insurance against data loss. A single data loss incident, such as a failed hard drive, could be enough to cause the failure of an organization. Do not lose data because of carelessness; implement a backup process immediately. Backups are not difficult. They just take a bit of planning. First, plan where you want to store the backups. The three main choices are online, off-site, or on-site.

Online or cloud backup storage is growing in popularity, availability, and cost-effectiveness. Online backups offer access to your data from any Internet connection, making recovery possible from anywhere. However, online backups put your data on someone else's hardware, making you dependent on that party's security, confidentiality assurance, and reliability. Plus, backup transfer speeds are dependent on the local Internet link speed, which is usually much slower than most local media backup options.

Off-site and on-site storage employ the use of backup media. Backup media can be tapes, optical disks, hard drives, USB drives, and solid-state drives/cards. Physical media require space for storage and are subject to physical threats, namely theft, damage, and destruction. Off-site storage is the better option when faced with major catastrophes. This protects your backup media from damage by problems occurring at your primary work location. On-site storage provides for quick recoveries around minor issues.

Midsized to large organizations often elect a multistage backup solution that provides both on-site and off-site storage benefits. One possible configuration is to use banks of rewritable high-speed optical storage devices to host a live, online, on-site backup. Then perform

periodic backups from the optical storage set to portable media, such as tapes, for secure storage off-site. This configuration provides for fast backup from the original systems to the optical media, then provides a convenient process for creating tape copies quickly. It also allows for fast restoration of data from the online, on-site optical backup, while providing major disaster protection with the off-site secure storage of tapes.

■ NOTE

When considering where to store backups, remember that these are confidential copies of the network data and resources, and they must be safeguarded as an asset. The device used for the backup must be secured and encrypted. Local backups aid in quick retrieval of data in the event of an accidental deletion or small event. Off-site storage should be considered in case an event compromises the entire facility. Consider the access to the backup and the cost for retrieval when making this important decision.

As much as possible, automate the backup process to ensure that it happens. If the act of backing up the network's data interferes with production, then install a secondary network to exclusively support data transfer for backups. This will require a second network interface in every system protected through the backup process.

The off-site storage location for backup media should be secure and reasonably protected from disasters, especially severe weather. If

your organization is large enough to have multiple locations, rather than the same building or even buildings within a few blocks of each other, then off-site storage can switch from third-party storage to internal storage at distant branch locations.

Remember to strictly enforce who is authorized to request and receive a backup from the vendor. The backup is a full copy of your organization's data; as such, it is a very valuable confidential asset. There is usually a cost involved with retrieval of data. Be sure to factor that into the cost of the backup solution engaged. Corporate legal counsel should review the service level agreement (SLA) with the vendor. Be sure that the selected vendor has enough resources to satisfy any obligation in the contract.

User Training, Security Awareness, and Ongoing Education

User training is an essential part of any security endeavor. Not every worker in an organization is automatically an IT or technology expert, nor should you expect them to be. User training and security awareness aim at providing all users with basic security knowledge, as well as job-specific security information.

To hold users accountable for their actions, first clearly define the network security policy boundaries and limitations. Having a set of rules and restrictions without informing personnel of its existence will not enforce policy and can lead to employee grievances and dissatisfaction. Training and awareness are necessary to educate users on their responsibilities and the consequences for violating organizational policy.

Without adequate security education for users, maintaining a secure environment is difficult, if not impossible. Technology cannot solve or remedy all security concerns, including the human factor. However, many organizations fail to support end-user training adequately. This results in users not understanding the importance, need, and benefits of security. Additionally, it will result in an increase in violations, most of which are benign and accidental, but incidents that you will nonetheless need to investigate and resolve.

A few common end-user security mistakes or problems caused by a lack of security training include:

- Opening email attachments from unknown sources or from a known source with an unexpected or unusual attachment
- Preventing updates and patches from installing, even when approved and recommended by the security staff

- Installing unapproved software on work computers, including games, screensavers, utilities, instant messaging clients, and browser plug-ins
- Installing software on work computers that you have not verified as safe and free from malicious code
- Failing to make backups of work data on work computers if stored locally
- Using a modem or wireless connection from a desktop or notebook computer while still connected to the company LAN
- Using a password storage/tracking utility that does not encrypt its database
- Walking away from computers while still logged in
- Connecting unknown and unapproved devices to a work computer
- Using portable media and storage devices from an external source on a work computer
- Installing remote control or remote access software on work computers without obtaining approval
- Using the same password on multiple systems
- Leaving portable devices in locations where they could be easily stolen, such as in a car's backseat

These are just some examples of common security problems caused by untrained users. You can avoid or at least minimize most of these issues with reasonable user security education. It is also common for IT staff to make security mistakes, even though they are much more knowledgeable about technology. Someone who is a “techie” or might be considered a computer geek does not necessarily possess security intelligence.

Some mistakes made by technology experts (who really should know better) include:

- Using the same password on multiple systems
- Failing to change the default password on a device
- Allowing new systems to go online before each is properly hardened and/or tested
- Failing to keep current with available patches and upgrades, especially those related to security
- Using remote system and device management mechanisms that are convenient but not secure
- Discussing passwords over the phone, including changing passwords based on over- the-phone requests
- Failing to properly check identity, authority, and permission before giving out information or access to resources
- Failing to implement a proper backup solution
- Not verifying and testing that backups are working properly
- Assuming something is secured or properly configured without specifically checking and verifying
- Allowing unnecessary and potentially insecure applications, services, and protocols to remain on a system
- Using security and network devices without changing default settings, such as firewalls, proxies, routers, and so on.
- Failing to understand all of the security configuration options on a new software or hardware product
- Putting new software or hardware into production before thoroughly testing and gaining approval
- Allowing anti-malware defenses to become out of date

- Allowing sensitive information to be communicated without an encrypted channel

Even a knowledgeable technical professional can learn how to be more secure through proper training and awareness education. Every single person throughout an organization has security responsibilities. These need to be defined, explained, and taught. Users will improve their behaviors when they understand what the risks are, what is at stake, and how their behavior will affect them if they fail to support security.

The goal of security-related training is user behavior modification. Users need to take steps to change their regular activities from those that place people and property in the workplace at risk to those that avoid risk. Risk is easy for most people to understand, but they need to be made aware of risk to trigger a change in their actions.

Good user training will cause an improvement in user compliance with the standards, policies, procedures, and guidelines of your organization. Often, the beginning of security training is awareness. Awareness is introductory, foundational, and ubiquitous security information that applies to all employees. Awareness aims at establishing a common baseline of security understanding for the entire organization.

The principal means of awareness training is in a classroom setting. It should, however, include a wide variety of communications, including online videos, interactive websites, posters, email reminders, regular memos or newsletters, wall banners, coffee mugs, sticky notes, manager review meetings, loudspeaker announcements, screensavers, mouse pads, and even voice-mail

messages. The point is to inform users about basic security essentials, and then reinforce those basics while they are at work. Remember to keep the message fresh by changing the visuals in posters and screensavers so employees do not disregard the message; when it has been viewed many times, it is ignored.

Awareness is important for all personnel, in all job positions, at every level of access, from the top to the bottom of an organization. Everyone should understand basic security issues. These often focus on general responsibilities, liability, seeking to avoid waste and fraud, reduction of unauthorized activities, and watching for abnormal or suspicious events.

It is important for employees to observe that security is important to the organization. This means that even top executives must be held accountable to the same basic principles enforced throughout the general employee population. If employees see evidence of compromise or compliance avoidance by senior management, then they get the impression that the rules are not important enough for anyone to follow. You do not want employees to be frustrated and confused by a “Do as I say, not as I do” management model.

USER EDUCATION

Every organization should have a written security policy, but that is not enough. In order to ensure that an individual knows the policy, the individual should be made to read the policy and sign a statement that he or she has read the policy and will abide by the it and any subsequent updates. These requirements apply not just

to salaried and hourly employees, but to anyone given electronic or physical access to the premises of the organization and to the assets of the organization in transit when the assets are transported outside of the physical facility. This could include contractors, suppliers, and others, and it requires constant monitoring. It could even require a test or security certification, background check, or other means of determining ongoing compliance. One means of encouraging users to comply with policies is to ensure that they are aware of the consequences of noncompliance. No one given permission to access assets of the organization is exempt from the security policy.

After awareness comes training. Training focuses on security issues and topics more closely related to specific job tasks. Training consists, therefore, of job-specific security information. Security training of this type assists users in accomplishing their individual work tasks while staying within the boundaries of the security infrastructure. Most organizations offer in-house awareness and training. This is common because such training directly reduces incidents, as well as the associated costs of handling and responding to internal accidental and ignorance-based breaches.

Beyond training is security education. This form of learning has a broader scope than a job description or the organization as a whole. The purpose of security education is to obtain extensive knowledge about security and related subjects, even if they do not directly apply

to current work responsibilities or tasks. Education is for the advancement of the individual, perhaps to improve his or her career outlook.

Education is usually obtained outside an organization. A company might perceive education as either a threat to employee retention or a benefit to keep employees happy. Most organizations want to improve the skills of their personnel. Either way, security education improves the knowledge and skill of the individual.

Security awareness, training, and education are beneficial for any security endeavor. Including and funding it in your organization's overall security solution will improve your odds of success. For organizations in compliance-based industries, it is an essential element to maintaining annual compliance.

INCIDENT RESPONSE

Incident response is the planned reaction to negative situations or events. Inevitably, security breaches, or at least attempts to breach security, are going to occur. When those events affect the organization or its abilities to perform its tasks in any way, incident response is triggered. The goals of incident response are to minimize downtime, minimize loss, and restore the network environment back to a secured normal state as quickly as possible.

Most incident response solutions include six primary steps or phases (see **FIGURE 14-4**):




FIGURE 14-4 The phases of incident response.

1. **Preparation**—Select and train security incident response team (SIRT) members and allocate resources (budget, duplicate hardware, and portions of employee work time/responsibilities).
2. **Detection and Analysis**—Confirm actual breaches.
3. **Containment**—Restrain further escalation.
4. **Eradication**—Resolve the compromise.
5. **Recovery**—Return to normal operation.
6. **Post-Incident Follow-up**—Review the process and improve future responses.

An incident response plan is an important element of network security management.

Preparation

Every incident response program requires a plan for responding to incidents. Defining procedures to respond to situations and defining SIRT roles is called **incident response planning**. The incident response plan outlines steps that must be taken to stop the incident without causing unnecessary outage.

The SIRT should be fully trained and have access to documentation needed to respond to an incident as it occurs. The quality of the preparation phase determines how well the SIRT can handle incidents.

TECHNICAL TIP

Do not overlook making a **business continuity plan** and a **disaster recovery plan**. Above and beyond security breaches are business-terminating events in the form of natural and manmade disasters. Such events as building collapse, flooding, earthquakes, hurricanes, fire, sabotage, blackouts, malware infection, and criminal activities can end your organization's operations. If alternate means of accomplishing mission-critical tasks are not available, your organization will cease to exist. Create a plan to recover from disasters and do business in the future.

Detection and Analysis

Security often focuses on prevention and deterrence. When an intrusion or security violation takes place in spite of precautions, the next stage or level of security must respond. The escalated stage of security is detection and analysis. Knowing if and when a security violation occurs is essential to a timely and adequate response to a breach.

Not all events are security violations. You may become aware of an incident from a security vendor's notification, a user reporting a problem, or an alarm, for example. Make sure you have the procedures in place to react to any type of incident, practice the procedures, and refine as necessary.

Containment

Once you detect a breach, you must immediately initiate a response. This stage of response is **containment**. The goal of containment is to prevent further spread of a known malicious event. Containment is a mindset, a plan of action, and an element of design. As a mindset, containment is the focus of first responders or members of the SIRT. Their goal is to prevent any expansion or spread of the compromise to other systems.

Containment can be as simple as unplugging a network cable from a compromised system. This effectively cuts off any communication to and from the suspect host. You can also impose containment by disabling user accounts, stopping services, adding entries to filtering devices, and so on. Whatever the vector or source of the malicious event, its source or target should be cut off from causing continued harm.

The act of containment should interrupt or interfere with the continued spread or operation of the unwanted event. This type of initial response is often effective against malicious code, ransomware, remote access, backdoor control, a compromised user account, social engineering, and many other forms of network or system compromise.

Eradication

After you have contained the incident and eliminated or blocked its source, the next step is to eradicate the problem thoroughly.

Eradication involves fixing the vulnerability that allowed the intruder to access the network or systems, such as cleaning malware from a computer. For some incidents, eradication is performed during recovery.

Recovery

Recovery involves returning the affected systems to normal status. You might need to rebuild systems from clean data backups or uninfected applications, harden systems again by applying patches and updates, and change passwords. In the case of a network breach, reevaluate the network perimeter security, such as firewall rulesets and router access control lists, and remediate vulnerabilities to prevent similar incidents.

Post-Incident Follow-Up

Learning from the incident lets you establish new controls and procedures to prevent or react to future incidents more effectively. Conduct a lessons-learned review of each incident to capture information the SIRT can use for future incidents.

Documenting every step of the incident response process provides a resource for updating your incident response plan, as well as security policies and procedures. For example, documenting multiple incidents related to malware may indicate a need for better controls.

ANTI-FORENSICS

Anti-forensics are a series of techniques designed to try to frustrate forensic investigators and their **digital forensic techniques**. These techniques are also employed to scrub systems of information in easy target areas that hackers may attack.

Forensic techniques are well known to the information security community and, by extension, to the hacker community. As a result, hackers have developed an arsenal of counter-techniques they use to foil investigations. If they can prevent the investigator from recovering usable evidence, they stand a better chance of avoiding discovery or prosecution. Some of these techniques are also used to hide rootkits and other malware.

Some examples of anti-forensic tools and techniques include:

- *Securely overwriting data*
- *Overwriting metadata*
- *Designing code that won't run when the system is in debugging mode*
- *Designing code that won't run in a virtual machine*
- *Preventing a system from entering safe mode or debugging mode*
- *Running all code from an external device like a USB drive*
- *Tampering with file system date and timestamps*
- *Running in read-only mode to prevent the system from updating file information.*

You should be familiar with the concepts associated with anti-forensics, but unless you plan to become a forensic investigator, you probably will not encounter any of these, except possibly securely overwriting data. As a cybersecurity professional, you may utilize some of the techniques to protect the confidentiality of corporate secrets. If the organization has computers that are no longer useful due to the age of the equipment, you may wipe the hard drives using a tool that protects the drive from data retrieval. By writing to the drive in random patterns, and then erasing the information and rewriting random values, the software

eliminates the possibility that any real data can be retrieved.

Another legitimate use for anti-forensics is cleaning the cache on computer systems. A hacker who is able to gain access to a system may be able to retrieve username and password information or other confidential information from the system RAM (random access memory). By running the cleaner program, the cache and RAM are wiped of contents, and benign strings of characters are written instead.

FAIL-SECURE, FAIL-OPEN, AND FAIL-CLOSE OPTIONS

Failure is not an option; it is a certainty. At some point in the life of your organization, failure will happen. Will you be prepared to handle it properly? How much harm will the failure cause? What you do in the event of a failure is essential to the success of the organization and any ramifications of the failure.

A failure can be due to hardware ceasing to operate properly because of an unknown flaw in materials or construction, the accidental discharge of static that damages circuits, or atypical weather that floods the building. A failure can be due to a malware infection that crashes a server or a zero-day exploit that compromises security from outside. A failure can be triggered by a loss of power due to a blackout by the energy provider, or an accidental unplugging of a necessary device.

Failure means problems, downtime, loss, and consequences. However, failures do not have to mean bankruptcy, going out of business, jail time, or catastrophic loss. The difference between calamity and mere inconvenience is planning. As with most aspects of security, proper and extensive planning can assist in designing an infrastructure that can prevent or effectively survive any issue or concern.

Planning for failure is planning how to respond to and recover from failures. The goal is to fail into a state of security or safety rather than into a state of insecurity or chaos. This is known as fail-

secure or fail-safe. A fail-secure state reverts to a condition where little or no harm, or at least no further harm, is likely to happen. Depending on the situation, a fail-secure state could be fail-open or fail-close.

To fail-open is to revert to a state of being open, available, or unlocked. In the case of the physical world, fail-open is important for the safety of personnel. Fail-open doors allow people to leave a building easily. However, such easily opened doors might represent a weakness that an intruder could exploit to access secured internal areas. In the case of IT, fail-open is to revert to a state of unfiltered communication or data access.

To fail-close is to revert to a state of being closed, unavailable, or locked. In the case of the physical world, to fail-close is to prevent a doorway or container from being opened during an emergency or compromise. In the case of IT, fail-close blocks access to communications and other digital resources.

Either state—fail-open or fail-close—can be the fail-secure or fail-safe option, depending on the circumstances. In any case, you should carefully consider and plan the type of default response you will use, especially during an emergency, intrusion, compromise, or system failure. Consider the devices and resources affected when the network is unavailable as you develop your response plan.

NETWORK SECURITY ASSESSMENTS AND TESTING

Perform verification scans of all deployed countermeasures to ensure correct functionality. Improper installation or misconfiguration can render a well-meaning safeguard worthless. Test every new security control when you install it and every time you reconfigure it. Test after any new device or code is added to that segment of the network.

TECHNICAL TIP

Numerous security groups and government/military agencies post security implementation guidelines, manuals, and checklists. You can use these documents, commonly known as **Security Technical Implementation Guides (STIGs)**, to review and assess your organization's status and state of security. A self-assessment attempts to take an external, unique, or independent viewpoint in evaluating security. This is why using external security guidelines, standards, and measurements can reveal oversights in an existing security infrastructure.

Use automated tools with updated databases of security tests and exploit simulations. These tools should confirm patches and updates, verify security configurations, challenge user accounts, and probe for

known vulnerabilities and exploit weaknesses. Quickly resolve any issues discovered by the scans.

Security Assessment

Security assessment is the judging, testing, and evaluation of a deployed security solution. The state of the world, in terms of security, is constantly changing. The hacking community is actively developing new techniques, methodologies, and exploits to compromise targets. The defenses and strategies that worked yesterday might not be as effective tomorrow.

Security assessment is the ongoing process of evaluating security so that you can improve it. Security is established through the act of performing a risk assessment. The results of the assessment are formulated into a policy that guides the implementation of the initial security infrastructure. From that point forward, security management takes over.

Security management is the process of watching for breaches, tuning existing security, and evaluating the need for improvements to security. It is this latter area of concern that is the focus of security assessment. Security assessment employs several techniques to appraise the effectiveness of deployed security.

Configuration Scans

A configuration scan probes a system to determine the current state of configuration and settings. This scan determines which available vendor patches are installed or missing. This scan also evaluates numerous system configuration settings against a database of recommendations. Depending on the tool you use to perform the configuration scan, you may find different levels of recommendations based on the function of the system or the general level of desired security.

One well-known example of a configuration scanning tool is OpenVAS. OpenVAS is sponsored by Greenbone Networks and is supported by a community of developers and users via GitHub (<https://www.openvas.org/>). OpenVAS is a great tool for small offices and end users to employ to quickly check their systems for well-known patches, assess password strength, and confirm security setting configuration compliance. However, this is just a first step; it should not be the only tool you employ. You will need other tools that focus on broader issues to obtain a complete and thorough configuration scan.

Vulnerability Scanning

Vulnerability scanning focuses on locating known exploitable weaknesses or vulnerabilities in deployed systems. You perform this type of scanning by using an automated tool with an updatable database of exploitations and test scripts. A vulnerability scan will reveal problems with configuration, installation, and product code. A scan of this type is only as reliable as the product and the currency of the testing database.

A popular and well-known vulnerability scanning product is Tenable Network Security's Nessus. Nessus is available in a free version for home use (scanning for up to 16 IPs) and a paid version for commercial use. Nessus is a powerful scanning engine with a massive database of exploitation plug-ins that offers excellent reporting capabilities (**FIGURE 14-5**).




FIGURE 14-5 A Nessus scan report.
© Courtesy of Nessus

NOTE

Nessus can be found online at

<https://www.tenable.com/products/nessus>.

Penetration Testing

Penetration testing (pen testing), also known as **ethical hacking**, is the application of hacking techniques, methodology, and tools in the hands of trusted ethical security experts. The purpose of penetration testing is to evaluate the resiliency of current security infrastructure and recommend improvements. Penetration testing is performed by a team of professionals who can customize attacks, modify exploits, and react in real time to fully stress security defenses.

TECHNICAL TIP

Hire or develop an ethical hacking team to test the strengths and weaknesses of the IT security, as well as the security of the facilities and the employees. In-house testers can pose a potential risk in that employees often develop blind spots to issues or may lack the knowledge to thoroughly test a device or network. If you use in-house testers, occasionally augment that team's efforts with an external third-party pen testing team. This is a common requirement for organizations under mandated legal compliance in several fields.

Penetration testing typically consists of five main phases or steps: reconnaissance, scanning, enumeration, attacking, and post-attack activities. The steps or components of criminal hacking and ethical hacking are the same. The differences are the goal and abiding by ethical restrictions and contractual boundaries.

Post-Mortem Assessment Review

Even with adequate security assessment using the main techniques (configuration scans, vulnerability scans, and penetration tests), you may still have a need for one additional form of security assessment. This additional form is the post-mortem assessment review.

A post-mortem assessment review is the self-evaluation performed by individuals and organizations after each security assessment task. The purpose of a post-mortem is to learn from mistakes. This will improve the process in future events and avoid the reoccurrence of the same mistakes. It is true that practice makes perfect. If you practice the same tasks repeatedly, each time improving upon the previous performance, you will eventually master the task.

post-mortem reviews provide incremental improvements on a consistent basis along with the rare revolutionary improvement. A review of this type can include numerous elements of appraisal focus, including:

- The process as a whole
- Each step performed
- The order of the steps
- Whether steps might be missing
- What other actions could be taken
- Whether participants were properly trained
- Whether additional resources were needed
- Whether resources were wasted
- Whether additional people were needed

- Whether too many people were involved
- Whether reporting was sufficient or insufficient
- What was missing from the assessment
- What could be improved next time

Many other queries could be posed, as long as the goal is to improve future applications of the assessment processes and tools. A post-mortem review is a beneficial process when used consistently.

CHANGE MANAGEMENT

Change management is a process that manages changes to systems or network configurations. Unmanaged and unauthorized changes can affect security operations or controls, open security holes, and result in outages. Even if a system is secured, it is possible for system administrators to make modifications; however, a change management process provides the format and necessary consideration to how a change may affect other systems.

You should implement a formal process to change network configuration settings. This includes adhering to a written policy that defines all roles, responsibilities, and procedures related to change management. The objective is to maximize the benefits for all people involved in the change and minimize the risk of failure.

A change control board should approve each change. The primary objectives of the change control board are to ensure that all changes are properly tested, authorized, scheduled, communicated, and documented. Membership of the board should include system administrators, asset owners (such as a database owner), and management so that the purpose of any change is known throughout the organization, and stakeholders have the opportunity to discuss how that change may affect their assets and how the network functions.

Documentation includes what was done, by whom (person/team, third-party vendor), why it was done, and any testing that occurred. The list should be listed in reverse chronological order (most recent

changes listed first; then older changes listed in order of completion) and easily accessible to administrators on the network. This is because the latest change or modification prior to an incident may have caused the incident, and administrators need to know what should be done (or undone) if the change affected normal operations. This document also aids in sharing the ramifications of undoing a change or of testing.

COMPLIANCE AUDITING

Compliance auditing is a type of assessment that judges how well an organization is accomplishing set goals or requirements. These goals and requirements can be internal or set by government, industry, and other regulatory agencies. Compliance auditing is an important part of maintaining a business, especially growing businesses, as it ensures that the organization is following all necessary security guidelines.

Compliance auditing may be a legal requirement for some industries, such as finance and medicine. Independent external auditors perform compliance audits to ensure that a target organization is fully abiding by the rules and regulations imposed by the government. This audit is a comprehensive investigation and review of the ongoing business processes. The audit requires a review of the security policy, access controls, risk management processes, and historical log files.

The focus of compliance audits varies based on industry, information type, and whether the organization is public or private. Auditors investigate an organization through documentation analysis, interviewing personnel, and combing through audit logs. Compliance auditing can examine recent security breaches, evaluate incident response, interview current and ex-employees, judge user access levels, interview executives over critical security concerns, and more.

Organizations are usually distinctly aware when compliance auditing is a mandated periodic occurrence. In those cases, companies

should prepare for audits by collecting the various types of information and creating the appropriate records as needed by the auditor. In fact, establishing a standard practice of producing and archiving the necessary information is prudent. The goals of these actions are not to manipulate the data, but to provide adequate access to the facts and historical activities.

Organizations that do not have mandated compliance audits should consider self-imposed audits. The process of thoroughly investigating the compliance level with the stated security policy can improve the long-term stability and security of every organization. The act of self-assessment and improvement is a common characteristic of most successful IT organizations.

■ NOTE

Ensuring that you are in full compliance with all federal and state laws and regulations is not only good security management; it will keep your company officials out of jail. Compliance requirements change frequently, and a compliance audit is the best way to assure that your organization is still in compliance with the regulations and requirements.

NETWORK SECURITY TROUBLESHOOTING

Security troubleshooting aims at recovering from problems related to the countermeasures. Problems will occur with the defense mechanisms. Downtime of a security control is as critical as downtime of a core business process or asset. Troubleshooting failures of security controls is an important part of network security management. As with most concerns, prevention is always preferable to repair or response. By working to ensure that failures do not occur, or at least do not occur as often, maintaining effective security will be easier and less costly in terms of both budget and manpower.

Network security troubleshooting is often about triage—deciding which issues or problems are of the most imminent concern. The more a security component affects a mission-critical process, the more important rapid response and repair become. When security is down, the previously protected assets are put at greater risk for compromise. Minimizing the length of time consumed by the response is important to minimizing long-term losses.

One of the most effective preventative techniques in network security troubleshooting is installing patches and updates. As with patches and updates to production systems, always test the new code thoroughly before deployment. Once tested and approved for application, apply updates when downtime would cause the least number of problems. Always be prepared with a redundant option if the updating process causes further security control problems.

Possible complications from the application of patches and updates include resetting to factory defaults, loss of some—but not all—configuration settings, and **bricking** (i.e, making the device nonfunctional) the control. If the security control resets back to factory defaults, then it will need to be fully reconfigured. If a recent configuration backup is available, restoration might be a swift repair. Otherwise, manual resetting will be necessary. To facilitate this process, always have complete documentation of all settings of all security controls. Screenshots and images often make the process faster, as the desired state in the image can be compared with the information on the monitor in the same font with the same appearance for quick comparison.

If you lose some but not all configuration settings, then the security control is unlikely to operate as you expect. Restore or reconfigure the settings promptly. The update may have added or modified settings that need testing and verification for function and compatibility. If a new feature interferes with a business task, you might find it necessary to disable the feature or roll back the update.

If the update process causes a complete failure of the security control, the real possibility exists that the product is useless. This is known as bricking—turning a useful device into a worthless brick. In some cases, a hard reset can revive a seemingly bricked device, while there may be more esoteric repair and recovery options for other devices. If the vendor does not provide an unbricking solution, search the Internet for user groups or discussion forums for a home-grown solution.

Configuration errors might be the cause of a security control malfunction. Configuration errors may be caused by human error,

oversight, or ignorance, as well as by updates, power fluctuations, and physical damage. When a security control is improperly configured, it does not provide the expected security. Troubleshooting this situation can be as simple as reverting to a previously saved configuration or manually reapplying the settings.

A more serious concern arises when configuration errors are not easily fixed. A recent patch or update may have rendered the product unstable, or a patch or update might be needed to stabilize it. If the configuration error reappears after every power cycle, then the device might be defective, it might need additional memory (or need defective memory replaced), or it might need to be attached to an uninterruptible power supply (UPS) to reduce the number of unplanned power fluctuations. Physical damage should be repaired expeditiously. If unrepairable, you may need to replace the device. Install preventative measures that will prevent the reoccurrence of the physical damage.

Power faults can take place for many reasons. If the building's power grid is not reliable, if sags and spikes occur on a regular basis, the focus of the long-term repair should be on improving the building's power distribution systems. Short-term responses can include surge protectors, UPSs, and generators.

If a device encounters a power fluctuation due to overheating, investigate whether the room where the device is operating has adequate HVAC service. The average temperature for a room dedicated to housing computer equipment should be at or below 70 degrees Fahrenheit with moderate relative humidity to avoid generating static electricity. In rooms where people work with computers, the temperature should remain below 80 degrees

Fahrenheit. Also, check the device to see if it has sufficient internal airflow to maintain appropriate internal operating temperatures. Do not forget to check filters and vents for accumulation of dirt or dust.

Power faults can include power supply or power grid variances. In this situation, the only responses are surge protectors, UPSs, and generators, as customers are unable to affect the quality of the power company's electricity by themselves. If switching power suppliers is an option, investigate this to determine the reliability of other providers, as well as the expense and hassle of switching. Sometimes actively and publicly "shopping" power suppliers can result in improved service from your present supplier.

In some situations, the problems with security are not with the security components, but with the overall infrastructure design. Perform a reassessment of the design on a periodic basis, such as once a year, to judge whether the current infrastructure continues to meet the security needs of your organization. Because security changes over time, the security design might need to evolve to meet the demands of new risks, threats, and concerns.

Static electricity is also a concern. **Electrostatic discharge (ESD)** or **static-electric discharge (SED)** can easily damage equipment, including security equipment. The amount of electricity discharged between someone's finger and a doorknob when you can see the spark jump is more than enough voltage to destroy most computer chips. Take precautions to prevent static damage.

Physical damage can be a concern requiring troubleshooting. Damage might be caused through intentional destruction, accidents, or Mother Nature. If the damage is superficial or cosmetic, the device can be returned to service. However, if the damage prevents the

security device from operating properly, then in most cases you need to replace it. Address the cause of the physical damage to reduce the likelihood of a reoccurrence.

Network security troubleshooting focuses on whether an intruder can bypass a security defense or restriction. A security control that can be bypassed is worthless. Once you discover that a security control can be bypassed, you need to investigate the method or mechanism of the bypass. Then, if possible, remove or block the method of bypass. If the flaw is a design or infrastructure concern, consider revising the design to remove the loophole.

From time to time, hackers find flaws in the programming of a security control. Once an exploit is written, the security control can be rendered useless through taking advantage of the flawed code. Vigilance in reviewing vulnerability databases and research, paying attention to vendor information, and watching the log files of the organization's network should alert you when an exploit succeeds.

Defending against an exploitation focusing on a security control is the same as when one focuses on any other aspect of a network. If possible, reconfigure the control to minimize the effectiveness of the exploitation. Consider removing the component until you can implement a new defense. Apply a patch or update from the vendor as soon as it becomes available.

A final concern in relation to network security troubleshooting is hardware failure. Although it is not a common occurrence, hardware can fail. Over the life of an organization, you are almost guaranteed downtime caused by hardware failure. When that hardware is a security control, the downtime is more severe in that it places the rest of the environment at greater risk.

Manage hardware failure by having replacements on hand, or be able to obtain replacements quickly. Reduce downtime through redundant infrastructure design. Monitoring ongoing performance metrics might enable the detection of a future hardware failure as performance degrades. However, abrupt hardware failures are difficult to predict. Thus, being prepared with alternatives and replacement parts is often the best troubleshooting solution.

CHAPTER SUMMARY

Network security management focuses on vigilance. Vigilance to remain current in terms of technology. Vigilance to remain knowledgeable about new threats and exploits and respond promptly to downtime and compromise. Vigilance to restore, repair, and update security on a regular, consistent basis.

A wide number of activities make up the realm of network security management, including fail-secure responses, maintaining physical security, detecting compromise, preparing for incident response, trapping intruders, host security components, backup procedures, user training, management tools, checklists, troubleshooting, compliance auditing, and security assessment.

KEY CONCEPTS AND TERMS

Anti-forensics

Antivirus scanner

Backup

Bricking

Business continuity plan

Change management

Compartmentalization

Compliance auditing

Containment

Delay

Detection

Deterrence

Digital forensic techniques

Disaster recovery plan

Electrostatic discharge (ESD)

Eradication

Ethical hacking

File encryption

Honeynet

Incident response planning

Padded cell

Patch management

Penetration testing

Prevention

Response

Security assessment

Security awareness

Security management

**Security Technical Implementation Guides
(STIGs)**

Separation of duties

Single-factor authentication

Static-electric discharge (SED)

Trusted Platform Module (TPM)

Vulnerability scanning


Whole hard drive encryption

CHAPTER 14 ASSESSMENT

1. Which of the following is *not* a network security management best practice?
 - A. Avoid remote access
 - B. Rely upon single or individual defenses
 - C. Separation of duties
 - D. Principle of least privilege
2. What is the purpose of physical security in an organization?
 - A. To prevent external attacks through the VPN
 - B. To prevent external attacks through the firewall
 - C. To prevent unauthorized access to facilities and equipment
 - D. To prevent the addition of new user accounts
3. What is the primary purpose of a post-mortem assessment review?
 - A. To identify necessary devices and budgets
 - B. To learn from mistakes
 - C. To reduce costs
 - D. To understand compliance
4. Which of the following techniques is *not* considered part of a network security assessment?
 - A. Incident response
 - B. Configuration scan
 - C. Penetration testing
 - D. Vulnerability assessment
5. Which of the following is *not* a common mistake that should be included in user training?
 - A. Bricking corporate computers
 - B. Leaving a computer while still logged in

- C. Installation of unapproved software on company computers
 - D. Opening unexpected attachments from email
6. Which of the following is the only insurance against data loss?
- A. Auditing
 - B. Backup
 - C. Configuration
 - D. Encryption
7. Which of the following is *not* a step or phase in an incident response plan?
- A. Acquisition
 - B. Containment
 - C. Eradication
 - D. Recovery
8. Which of the following is most important to the effectiveness of an antivirus scanner?
- A. Available bandwidth on the network
 - B. Commercial products versus open source
 - C. Platform of the scanning engine
 - D. Timeliness of the definitions database
9. Which of the following is *not* part of a complete and comprehensive security approach?
- A. Secure assets
 - B. Train users
 - C. Monitor for violation attempts
 - D. Implement single-factor authentication
10. Which of the following is *not* a network security management tool or technique?
- A. Change documentation
 - B. Equipment inventory

- C. Logical organization map
- D. Products that won awards



© Digital_Art/Shutterstock

Emerging Technology and Regulatory Considerations

IN YOUR CAREER as a security professional, you will make it your business to know how to secure your organization's IT assets, networks, and systems. You are already learning a great deal about the tools, technologies, processes, and procedures currently available. In your work you may apply firewalls to secure the perimeter of the network and virtual private networks (VPNs) to secure data in transit. A deep knowledge of the current security landscape is vital to your professional success. But you need to know

what's on the road ahead, too. In this chapter, you will learn about what's next.

Today's information security professionals face a daunting task. Not only do you need to understand how to secure a complex, diverse, and rapidly changing IT environment, but you also need to be aware of the challenges and threats to come. An in-depth understanding of security technologies, such as VPNs and firewalls, and of threats like malware and social engineering provides the foundation of your arsenal. To be truly successful, however, you also need the ability to identify and respond to trends in both the technologies and the threats under development. To keep up with attackers, you need to know what is coming and where to get reliable, current information.

Throughout your career as an information security professional, you will encounter situations in which you must apply security best practices to new technologies or new architectures. You may find that you are conducting investigations, trying to determine how your infrastructure was compromised. You may be monitoring new attacks and trying to develop new defenses. Whatever challenges you face in the future will require that you leverage the experience, understanding, and best practices you have learned to ensure that you are effective against new risks and threats.

CHAPTER 15 TOPICS

This chapter covers the following topics and concepts:

- What the future holds for network security, firewalls, and VPNs
- Resources for network security technologies, techniques, and threats
- Useful network security tools
- The impact of wireless and mobile connectivity on the network perimeter
- Potential uses of security technologies
- Network security–related regulations and compliance measures
- Specialized firewalls
- Emerging network security technologies

CHAPTER 15 GOALS

When you complete this chapter, you will be able to:

- Discuss the different types of integrated and specialized firewalls, including the advantages and disadvantages of each
- List additional sources of information related to network security
- Describe emerging IT and security trends and their impact on network security
- Identify challenges and advantages presented by the new technologies and emerging threats to network security
- Understand the types of regulations that affect network security

- Discuss the future of network security, firewalls, and VPNs

WHAT THE FUTURE HOLDS FOR NETWORK SECURITY, FIREWALLS, AND VPNS

This chapter focuses on the future of network security, firewalls, and VPNs.

How does a security expert discern what the future holds? A number of factors should influence your planning as you try to draw a road map for your information security strategy:

- **Historical progression of threats**—If you are familiar with how threats have progressed over time, you will have a better understanding of where those threats may be headed in the future. For example, 10 years ago, virtually all malware attacks were focused on operating systems (OSs). Today, malware attacks go after applications far more than OSs. Browsers are a particularly popular target for attacks today.
- **Your industry**—Each industry has a different focus when looking to the future. For example, a bank will have very different requirements from a shoe store. Understand your industry, network with your peers, and make sure you are looking at the proper framework for your planning.
- **Experts**—It is always a good idea to see what the people who predict the future of the industry think before doing your planning. Although experts can disagree over certain specific predictions, if you keep an eye on the major information sources (some of which are provided later in this chapter), you can develop a pretty good idea of where information security is headed.

- **Vendors**—Because they are frequently trying to sell you their next-generation solution, vendors usually have a biased view of the future. They can, however, provide valuable information for your planning. If the vendors you work with or follow are all targeting a particular threat or technology, they view it as an area to generate revenue. If an industry targets a specific area as a source of revenue, it's a good bet that a threat lurks there somewhere.
- **Role of cybersecurity in national security**—In March 2019, President Trump increased the cybersecurity budget by 5 percent, and in May of that year he spoke about his plan to “strengthen America’s cybersecurity workforce to secure our nation and promote prosperity.” Although the number of trained security professionals is still inadequate, public awareness, along with that of lawmakers, heads of organizations, and—importantly—employers has increased dramatically. Every major publication from *The New York Times* to *The Wall Street Journal* has stated that there are not enough—and will not be enough—trained security professionals to fill the future need within the workforce.

Here are some of the areas you should pay attention to as you plan your information security strategy.

Threats

One area that is evolving extremely quickly is the variety of threats to your infrastructure. Several years ago, the most prevalent threats were from unsophisticated attackers whose main goal was to accomplish something they could brag about in chat rooms with their buddies. Today, organized crime has zeroed in on the huge amounts of money to be made in computer hacking. You now see targeted attacks against specific industries and companies; viruses that target credit card numbers, bank account information, and Social Security numbers; and rootkits that, once installed on a computer, turn it into a host the attacker can control remotely to attack other systems and networks. Some attackers will threaten to crash networks with denial of service attacks or encrypt system data unless the system owner pays extortion money—the high-tech equivalent of a protection racket.

In the future, you will see more resilient networks that will mitigate the risk of traffic-based attacks, more secure OSs and applications to resist malware, intrusion prevention systems that will respond instantly to attacks, and cloud triage that can isolate the threat before it can damage your infrastructure.

Firewall Capabilities

Firewalls have been adding capabilities since appliances were first introduced. Early firewalls contained some limited filtering and NAT capabilities and not much else. You will learn about the wide range of today's firewall capabilities and specialties later in this chapter.

Encryption

Encryption is a constantly evolving standard. In 1977, the Data Encryption Standard (DES) was specified in the Federal Information Processing Standards (FIPS) Publications. It became a national standard. The standard has since changed from DES, a 56-bit algorithm, to 3DES, an effective 168-bit algorithm, to the Advanced Encryption Standard (AES), which supports a 256-bit algorithm. The problem with encryption is the constantly improving processing power of computers. As computers get faster and more capable, with bigger and bigger memory space, encryption algorithms become easier and easier to break through brute force and other techniques.

Encryption's popularity has grown as concerns about protecting data at rest, in transit, and while archived have come to the forefront of many industries. How many stories have you seen about the stolen laptop with 100,000 Social Security numbers on it or an application compromised by an attacker able to read data directly from a hard drive because it was stored in cleartext? Today, government agencies such as the Department of Defense require full-drive encryption for all laptops.

Recognizing the growing need for encryption, the industry is responding quickly. The AES standard was designated as the replacement for DES. It is designed to scale upward with longer keys. As you look into encryption solutions, keep in mind whether they support AES or an equivalent algorithm, and be sure that you encrypt your data everywhere it is vulnerable. The days of relying on your VPN as the only data protection are gone. You need to secure data with encryption everywhere it can be accessed.

Authentication

Another area where you can expect to see dramatic changes in future capabilities is in authentication, especially with respect to identity and access management. In the past, much of user security was based on a user ID and password. Years were spent trying to teach users what a strong password is, why they need a strong password, and why they need to change their passwords every 90 days. The effort has been largely ineffective. Users still choose poor passwords; even when they select strong passwords, often writing it down, the passwords can still be cracked with sufficient computer power.

The other challenge associated with authentication is actually a user-management issue. All too often in a complex environment, the creation, permissioning, management, and eventual retiring of user accounts does not work. Accounts retain permissions long after users move on to other roles, accounts remain on systems long after employees have left the organization, and, in many cases, no auditing of actions using privileged accounts takes place. Collectively, these critical tasks are known as identity and access management.

How will these challenges evolve in the future? One trend is moving away from passwords to tokens, smart cards, and biometric authentication. Multifactor authentication, where multiple proofs of one's identity are used, is the better way to protect systems. Two-factor authentication is one means of multifactor that requires two proofs of identity. The proofs come from three categories: something you know (user ID and password or pass phrase typically), something you are (a biometric measure like iris or retinal scan, hand geometry, or voice or signature recognition), or something you have (such as a token, automatic number generator, or authentication application).

This is occurring during the contrary movement to offer single sign-on solutions to systems that allow users access to multiple devices or services after entering one set of credentials instead of separate credentials for each device accessed.

On the identity and account management front, a number of solutions automate these activities, providing full account life-cycle management and the associated auditing capabilities many companies look for. The one challenge with these solutions is that, due to the sophistication of most corporate computing environments, these are complex to install and maintain. It is not just a matter of automating the creation of Microsoft Active Directory accounts in most cases. Companies have multiple application systems and authentication requirements. Making all of these components work together is not easy. Once you get them working together, however, you have removed a significant threat to your security landscape.

Metrics

One of the biggest complaints from chief information officers (CIOs) about information security is the lack of measures for the success of a program. Information security has moved from being an esoteric discipline practiced by a few misunderstood experts to a core business function vital to supporting the bottom line. While this evolution has been long overdue, today management expects you to quantify your contribution to the company and justify the expenditures made to support security. The application of **data analytics** principles to better understand what is happening on a network, including statistics on how it functions, throughput, bandwidth usage, and security anomalies, can provide essential perspective and aid in finding unauthorized actions.

The good news is that the industry has been moving in this direction for some time and has developed a number of performance metrics. The most popular is the **Information Technology Infrastructure Library (ITIL)**, which is a set of concepts you can use to formalize your security management practice and the associated reporting. In addition, numerous solutions allow you to automate not only these processes, but also the associated measurements.

Focus

Another area of significant evolution in information security is in the nature of what you are trying to protect. Initially, information security was all about keeping the hackers out of your network. As a result, companies invested significant amounts of money in firewalls and other network security technologies. Then focus shifted from the network to the host, and organizations attended to managing patches, hardening OSs, and installing host-based firewalls.

Once the industry secured the host, attackers shifted to threatening the applications running on those hosts. IT then started focusing on integrating security into the software development lifecycle, testing and evaluating code, hiring penetration testers to try to break code deliberately, and deploying firewalls and proxy servers specifically to secure applications.

The next shift in focus for information security is a growing focus on what is truly valuable: the data. Ultimately, all the measures developed to date, from the network firewalls, to the host hardening, to the penetration testing, have all been about securing the data. The industry is now moving toward a data-centric security model, which is a significant paradigm shift from previous models. A data-centric model forces companies to focus on classifying and applying values to the data they choose to store. Although it will undoubtedly be a painful process for many companies, this approach will ultimately yield a much more secure environment.

Securing the Cloud

Cloud computing is still a relatively new phenomenon in computing infrastructure that involves moving computing resources out to the Internet. This often means contracting with a third party to maintain the data and provide access to it so that an organization need not purchase additional infrastructure. Resources are then shared by multiple applications and, in many cases, shared by multiple corporations. Think of how the phone networks or the electrical grid operates. These clouds are typically built using virtualization that allows for exceptional efficiencies, but also opens up a number of new security challenges.

To leverage the true benefits of cloud computing, you have to trust the vendor providing your cloud. This requires a shift in focus from deploying security technologies to ensuring that your vendors are contractually obligated and physically able to keep your data secure. You also need to be able to evaluate vendors to determine trustworthiness. If you have the available resources, you should be auditing the vendor(s) to ensure that your data is consistently kept secure.

You will learn about the security challenges specific to virtualization later in the chapter.

Securing Mobile Devices

A rapidly growing sector of the end-user computing space is the use of mobile devices like smartphones, smart watches, netbooks, or tablet computers. These devices present some unique challenges. Think about how many people today receive email on their smartphone, or who use an iPad to review confidential documents while riding the train to work. How do you secure these devices?

The good news is that once the industry identifies issues like these, it concentrates resources and ultimately solves the problem. The computer security industry has very few “unsolvable” issues. Already, virus protection, mobile device management, and encryption applications are available for mobile devices. The challenge you will typically see in both current and future technology is that these types of devices are frequently overlooked or discounted when documenting security risks. Be sure to keep these devices on your list of risks. Personal devices possess an alarming amount of storage and processing capacity, which makes it easy for an employee to inadvertently store confidential information on a personal device.

Internet of Things

The term **Internet of Things (IoT)** describes how a wide variety of devices and sensors can connect and be accessed, typically over a wireless network. IoT devices or sensors are used in security systems, insulin delivery systems, wearable activity trackers, building automation, thermostats, light bulbs, electronic appliances, and much more. The use of IoT devices is rapidly increasing, which creates issues and challenges. Security professionals are concerned that IoT development and connectivity are outpacing security, privacy, and regulatory compliance. Many discussions today revolve around how to better protect these IP-connected devices and the networks to which they connect.

Mobile IP

The future will bring with it an increased clarity of the techniques, technologies, and protocols associated with making IP mobile, from the technology actually called **Mobile IP** to a potential host of other technologies, including **IP Multimedia Subsystem (IMS)**, and changes to the Domain Name Service and other fundamental technologies.

Bring Your Own Device (BYOD)

Bring Your Own Device (BYOD) will increase, as will the need to secure a growing number of mobile endpoints and different OSs and applications. The BYOD mobility aspect becomes even more important when it is considered in light of the discussion of mobility technologies and the mobility enablement of base technologies, and when you consider that employees, contractors, clients, and others need to connect securely, with potentially varying permissions, to the assets of the organization.

RESOURCE SITES FOR NETWORK SECURITY, FIREWALLS, AND VPNS

A variety of resource sites are available to you for more information on network security, firewalls, and VPNs. Although this list of online and offline resources is a good start, it is impossible to cover every quality source. Here are some sites and books covering the critical topics just to get you started. If you need additional information, you are only a search engine click or a bookstore trip away from the answer.

At press time, all websites listed here were valid; however, with the mergers, acquisitions, and other changes in the industry, some of these links may change over time.

Firewall Vendors

- Check Point Software—www.checkpoint.com
- Cisco—www.cisco.com
- Juniper Networks—www.juniper.net
- Fortinet—www.fortinet.com
- McAfee—www.mcafee.com
- Palo Alto Networks—www.paloaltonetworks.com
- Sophos—www.sophos.com
- SonicWALL—www.sonicwall.com
- WatchGuard—www.watchguard.com

Virtual Private Network Vendors

- Cisco—www.cisco.com
- Citrix—www.citrix.com
- F5 Networks—www.f5.com
- Juniper Networks—www.juniper.net
- Microsoft—www.microsoft.com
- OpenVPN—www.openvpn.net
- Openswan—www.openswan.org

Network Security Websites

- CERT (security research)—www.cert.org
- Data Breach Investigations Report (security research)—www.verizonenterprise.com/DBIR/
- Gibson Research Corporation (security testing and tools)—www.grc.com
- HackerWhacker (firewall testing)—www.hackerwhacker.com
- ISACA (standards/certifications)—www.isaca.org
- (ISC)² (standards/certifications)—www.isc2.org
- Krebs on Security (current issue and hack blog)—krebsonsecurity.com
- National Institute of Standards and Technology (NIST)/Computer Security Research Center (CSRC)—csrc.nist.gov
- MITRE (nonprofit cybersecurity think tank and research facility)—www.mitre.org/
- Open Web Application Security Project (OWASP) (Top 10 issues and several open- source projects)—www.owasp.org
- SANS (security research and training)—www.sans.org
- Symantec (security vendor)—www.symantec.com

Network Security Magazine Websites

- CSO Magazine—www.csoonline.com
- Data Breach Investigations Report (security research)—www.verizonenterprise.com/DBIR/
- SC Magazine—www.scmagazine.com/
- Search Security—searchsecurity.techtarget.com

TOOLS FOR NETWORK SECURITY, FIREWALLS, AND VPNS

As you review the information security technologies that you may encounter as your information security career progresses, you might wonder where you will find these tools.

In your studies, you have probably learned about a number of commercial and open-source solutions for network security. Numerous commercial and open-source solutions for network security tools, firewalls, and VPNs are available. Where will you find the tools of the future?

Commercial Off-the-Shelf (COTS) Software

Commercial software remains the choice of corporations everywhere, and this will not be changing any time in the future. A number of reasons account for the dominance of commercial solutions in the information security industry. These include:

- **The popularity of combining hardware and software into an appliance**—Many of the most popular firewall, VPN, vulnerability management, intrusion detection, and other security applications will be developed and distributed in an appliance format. Although some open-source solutions will allow you to use standard server hardware to create an appliance-like solution, a truly open-source appliance will probably be distributed as a commercial solution.
- **Companies rely on solutions they can support**—One of the ongoing challenges with open-source software is the lack of commercial support for the solutions. Highly technical people can work well within the open-source support community. Support from an open-source community is typically better than what comes from a commercial help desk, but corporations are generally reluctant to trust security-related solutions to that model.
- **The commercialization of open-source solutions**—Many of the most popular open-source solutions are being commercialized as their developers look to generate revenue from the products. Two examples are the Snort IDS solution and the Nessus vulnerability scanner. Both tools were widely used and extremely popular open-source security solutions, and both are now available commercially.

Open-Source Applications and Tools

It is common to see companies embrace commercial tools in their production environments, but you cannot discount the sheer innovation available in the open-source community. Most of the tools referenced in this book are open source, and you can expect to see continued development on most and new additions in the future.

Keep in mind a couple of other advantages of open-source tools as you build your security tool kit. First, working with open-source tools is a great way to build your information security skill set. Most security professionals cannot afford to purchase multiple commercial security applications to learn with, so leveraging open source is a cost-effective career builder.

The other value that open-source security tools bring to a security professional is that these are the same tools many of the people attacking your network will use. One of the most important skills you can develop is the ability to understand the hackers trying to get into your systems. Learning the tools and techniques they may use and developing parallel expertise will aid in your career. Commercial applications seldom offer the same learning opportunities. Attackers generally are not using commercial applications in their attacks, and they typically do not draw from the open-source community to help with learning.

THE VANISHING NETWORK PERIMETER

If you look back to the first uses of network security, security professionals essentially put a firewall between the corporate network and the Internet. That firewall marked the perimeter of the network. As networks became more complex, companies started finding they were connecting to the Internet in more locations, so more firewalls were installed. The perimeter remained the demarcation between the Internet and the corporate network. Today, the vanishing perimeter has been the subject of many presentations, white papers, and sales pitches, as vendors try to convince organizations to invest to protect their perimeter-less environments.

VPNs extended the internal network outside the nice, neat network perimeter defined by the firewall/Internet interface. Suddenly, your network could be in any hotel or coffee shop in the country, as your employees utilize the benefits of wireless and mobile computing. To make matters worse, high capacity drives and portable storage meant that not only was your network extended into places that you had never imagined, but your data storage was no longer contained within your office space. Critical information could be stored easily anywhere your employees wanted.

Unfortunately, that was not the end of the challenges you faced as an information security professional. In addition to your employees taking your network and data anywhere they wanted, your business discovered the benefits of interconnecting with customers, vendors,

suppliers, and partners. Each of these groups needed varying levels of access to your systems, networks, and data.

In addition to all the challenges with connections and data beyond your network, businesses also discovered the benefits of online commerce over the Internet. That seems harmless enough, until you realize that to do business on the Internet, you need to punch holes in the network perimeter to provide access from the Internet to your systems, networks, and data. It is a little like installing skylights in your roof: You want the benefit of the sunlight in the house, but you do not want leaks when it rains.

In a perfect world, upper-level management would have consulted with the network security team before they embarked on these projects. They would have given you an unlimited budget to deploy network zones separating each of the different types of data, with firewalls and IDS/IPS deployed at each level, and enough staff to monitor all the infrastructure to ensure that nothing gets past your security measures.

In reality, the first few connections were most likely put in without input from security, and these might not even be firewalled. It is so much easier than having to justify purchasing firewalls, set up firewall rules, monitor logs, and do all the other things an organization needs to do as it opens holes in its protective perimeter.

The Impact of Ubiquitous Wireless and Mobile Connectivity

The first thing to realize as you think about wireless is that you need to consider a variety of wireless technologies. Be sure not to focus too narrowly and miss something critical. Consider the types of wireless technologies that affect your environment:

- **Wireless deployed as a local area network (LAN) technology**—The deployment of wireless as part of the office LAN caused constant sleepless nights for security professionals who used to invest countless hours sweeping the network to search for unsecured wireless access points (WAPs). The real challenge with this technology, other than the fact that it permitted the bypassing of perimeter security controls, was that it was convenient. Setup was so easy for an employee—particularly setting up in the default configuration, with all the attendant security issues.
- **Public and home office wireless**—Public/home office wireless is a great opportunity for your employees. Take out your laptop, connect to the wireless, and you are on the Internet, ready to surf. You can connect to the corporate network, do your banking, shop online, post to social media, or just catch up on the news. What could be more convenient? As is often the case, convenience is the enemy of security. A number of security challenges are present with public/home office wireless. First, any unencrypted traffic sent over these wireless networks can be intercepted. Counting on the coffee shop (or an employee's cable company) to configure the wireless with the most advanced security capabilities enabled is not reasonable.
- **Mobile wireless**—The ubiquity of smartphones and other devices that can connect to cellular networks presents another challenge when discussing the impact of wireless connectivity on information security. Now employees do not need to search for a wireless hotspot before connecting to the network; they can access their mail from just about anywhere.

The good news is that the industry has made great strides in wireless security. We now have stronger authentication technologies like 802.1x, which requires connecting systems to authenticate using public key infrastructure (PKI) machine certificates, stronger encryption technologies, and significant improvements in user awareness. Wireless connectivity is a convenient solution for users that can be much more secure than in the past.

Making Wireless and Mobile Connectivity More Secure

Work with your users to make sure they understand the risks associated with wireless connectivity. From a corporate network perspective, you are probably running a VPN for business connections, so that data should be secure.

The best defense against wireless and mobile risks is a combination of policy and technology. First, your policy should restrict access to company systems and email to company-owned devices. That allows you some level of control over the types and configurations of devices with access to your data. Next, make sure the devices you are using leverage the appropriate security technologies, like encryption and antivirus scanning. Although it is virtually impossible to physically secure something as portable as a smartphone, you can ensure that the data on the device cannot be recovered if it is stolen or lost.

Finally, put the appropriate controls in place to reestablish your network perimeter. If you treat the risk of wireless and mobile technologies connecting to your network the same way you treat the risks posed by the Internet, you will mitigate the risks associated with the increasingly fluid perimeter you will be dealing with.

POTENTIAL USES OF SECURITY TECHNOLOGIES

One ever-expanding area of discussion when looking at security technologies like firewalls, VPNs, intrusion detection, vulnerability management, and others is the many uses for this technology.

Obviously, you use security technologies to secure your data, networks, and systems, which is certainly the primary use for security technologies. But you should be aware of some additional uses for this technology.

The core security concepts can be summed up with the acronym C-I-A. This stands for confidentiality, integrity, and availability.

- **Confidentiality**—Confidentiality deals with keeping information, networks, and systems secure from unauthorized access. This issue is particularly critical in today's environment in light of the high-profile leaking of people's personal information by several large companies. These breaches in confidentiality made the headlines largely because the thefts exposed people to the potential of identity theft.

Several technologies support confidentiality in an enterprise security environment.

These include:

- Strong encryption
 - Strong authentication
 - Stringent access controls
- **Integrity**—Integrity is the consistency, accuracy, and validity of data or information. One of the goals of successful information security is to ensure that the information is protected against any

unauthorized or accidental changes. The architecture should include processes and procedures to manage intentional changes, as well as the ability to detect changes.

- **Availability**—Availability is the third core security principle, defined as a characteristic of a resource as accessible to a user, application, or computer system when required. In other words, when users need information, it is available to them. Typically, threats to availability come in two types: accidental and deliberate. Accidental threats include natural disasters such as storms, floods, fire, earthquakes, and so forth with attendant power outages. This category also includes outages due to equipment failure; software issues; and other unplanned system, network, or user issues. The second category is related to outages that result from the exploitation of a system vulnerability. Some examples of this threat include a denial of service (DoS) attack or a network worm that affects vulnerable systems' availability. In some cases, one of the first actions you need to take following an outage is determining which category an outage fits into. Companies handle accidental outages very differently from deliberate ones.

As you continue to work in the information security field, another acronym you will encounter is **GRC**, which stands for **governance, risk, and compliance**. This is the best way to look at your security technologies from a business perspective. Look at how the security technologies you are deploying provide much more than just securing data:

- **Governance**—Governance includes the processes and procedures that ensure that employees are following your organization's security policy. For example, you have a policy that says that data in transit across a public network must be secured with at least 128-bit encryption. You would ensure that the only way to connect to the network is through a VPN that uses at least 128-bit encryption. This mechanism would be part of your governance efforts.

Governance is generally used to demonstrate to management, customers, and auditors that your information security program is operating as outlined in your policies, procedures, and practices. If you are a service provider, you can leverage your governance capabilities to show that you can be trusted to keep clients' information secure.

- **Risk**—One of the foundations of any information security program is a robust risk management practice. If you do not identify your risks, how do you know which security technologies to deploy and where? A risk is the likelihood or potential for a threat to take advantage of a vulnerability and cause harm or loss. Risk is a combination of an asset's value, exposure level, and rate of occurrence. A goal of security is to recognize, understand, and eliminate risk.

A risk management process starts with a risk assessment of your environment. Start with your critical data, applications, systems, and networks. Document the risks associated with each in a risk matrix. A typical risk matrix will contain the following:

- A description of the risk
- The likelihood that the risk will actually occur
- The impact of the risk
- A total risk score
- The relevant business owner for the risk
- One of the core security principles the risk affects: confidentiality, integrity, and/or availability
- The appropriate strategy or strategies to deal with the risk

The last bullet is the one that is critical with respect to security technologies; you leverage security technologies as part of the strategy to deal with the risks.

REGULATIONS AND COMPLIANCE

One of the fastest-growing sectors in the information security field is compliance. **Compliance** means ensuring that your organization obeys internal policies, as well as any applicable laws or other regulatory requirements. Regulatory requirements are the laws enacted by the federal government and individual states to establish what is acceptable in business. The United States lacks a single federal law or regulation that covers all aspects of cybersecurity and denotes what is acceptable behavior regarding data and customer information. You will find that your security program, in conjunction with the security technologies supporting it, are critical to ensuring compliance with a growing number of regulations.

The following are some of the regulations and laws that address network security and/or data protection:

- **Health Insurance Portability and Accountability Act (HIPAA)** —Congress passed the **Health Insurance Portability and Accountability Act (HIPAA)** in 1996. HIPAA is best known for its data protection rules that address the security and privacy of protected health information (PHI). PHI is any individually identifiable information about a person's health. Of the major rules under HIPAA, the Security Rule has direct implications on network security, requiring layers of administrative, physical, and technical safeguards.

The administrative safeguards include adopting a written set of privacy procedures, restricting employee access to electronic protected health information (EPHI), performing internal audits

for HIPAA compliance, and implementing procedures for addressing and responding to security breaches.

Physical safeguards include implementing physical access controls to facilities, along with maintaining related records and visitor records.

Technical safeguards include implementation of information systems that deter intrusion, ensuring the confidentiality and integrity of EPHI, requiring authentication of all system users, and monitoring system access, among others.

- **HITECH Act**—Enacted as part of the American Recovery and Reinvestment Act of 2009, the **HITECH Act** expanded the scope of privacy and security protections available under HIPAA and motivated the widespread adoption and standardization of health information technology. Originally, providers that sought to adopt electronic health record (EHR) systems could apply for monetary incentives to help pay for the transition to EHR platforms. Those incentives expired in 2015. Under the HITECH Act, the Department of Health and Human Services (HHS) may directly require business associates to comply with HIPAA. Compliance requires protecting protected health information (PHI), detecting breaches, and reporting violations of HIPAA to covered entities. Business associates are subject to HIPAA audits, and penalties can be levied for failing to comply with HIPAA Security and Privacy Rules.
- **Gramm-Leach-Bliley Act (GLBA)**—The Safeguards Rule within the **Gramm-Leach-Bliley Act (GLBA)** is concerned with information security policies. Financial institutions must develop and comply with a comprehensive information security policy that includes safeguards for the handling of sensitive customer information. The plan calls for regular monitoring and testing of the safeguards. Companies are also required to implement a continuous risk management program. The program must initially identify potential risks to the company's infrastructure and

information. After the company meets GLBA's initial risk identification requirement, risks must be reassessed should the firm's business or technology change. The company must also update its written policies and procedures, if needed.

- **Sarbanes-Oxley (SOX)**—The **Sarbanes-Oxley (SOX) Act** of 2002 was created to protect investors by requiring publicly traded companies to validate controls securing financial data. SOX accomplishes this by imposing harsh penalties and making corporate officers personally responsible for the disclosures. Two sections of the Act place significant constraints on IT security. Section 302, "Corporate Responsibility for Financial Reports," requires the chief executive officer (CEO) and chief financial officer (CFO) to personally certify the accuracy of financial reports. They must also ensure that internal controls are followed, assessed, and reported every quarter. Section 404, "Management Assessment of Internal Controls" requires the company to prove that controls are effective. This occurs through assessment and an audit by a public accounting firm. The results must be published in the company's annual report. SOX affects IT departments in the form of record retention policies and access to an organization's electronic records, such as email and accounting system data.
- **Family Educational Rights and Privacy Act (FERPA)**—**Family Educational Rights and Privacy Act (FERPA)** is a federal law that protects the privacy of student information. FERPA applies to the data maintained by educational institutions that accept federal funding. Those schools must implement security controls in IT systems to protect the privacy of electronic student records.
- **Payment Card Industry Data Security Standard (PCI DSS)**—The **Payment Card Industry Data Security Standard (PCI DSS)** ensures the confidentiality, integrity, and availability of cardholder data and transaction-processing functions. The latest PCI DSS standard addresses numerous security matters and PCI compliance requirements. Each requirement addresses specific

controls for securing IT infrastructures, such as installing firewalls, using encryption, implementing anti-malware, performing ongoing security risk assessments, security testing, and much more. Although the above regulations are governmental controls, PCI DSS is an industry standard and control. The power of the council that developed it and the threat of compliance is that if a company fails to comply, the ability to accept credit card transactions can be ended permanently. The council, comprised of the largest credit card-issuing companies, established the standard in an attempt to ward off governmental control in credit card usage.

- **General Data Protection Regulation (GDPR)**—The **General Data Protection Regulation (GDPR)** is a European Union (EU) regulation that protects citizens' privacy and information. Like HIPAA, noncompliance can result in financial penalties. GDPR applies to countries in the EU, as well as companies that handle data from EU organizations or EU citizens. One of the biggest changes GDPR introduced to information security is the explicit opt-out, where all citizens are automatically opted out of marketing and soliciting campaigns, and they automatically elect not to have data stored unless it is in a manner outlined for a requirement described in the regulation. Otherwise, the person must explicitly opt in, meaning they elect to receive solicitations and have their protected information stored. **Data sovereignty**, the concept that data is subject to the laws of a country in which it is stored, is becoming a challenge for businesses as their operations move to the cloud.

■ NOTE

Under HIPAA, covered entities may use PHI in certain ways. A covered entity may be a health plan, healthcare clearinghouse, or any healthcare provider that transmits PHI in an electronic form. HIPAA also applies to the

business associates of covered entities. A business associate is an organization that performs a healthcare activity for a covered entity, such as a lab processing a blood sample.

Although compliance with regulations affects different groups within an organization, IT departments are responsible for implementing compliance controls. Therefore, IT staff should stay up to date on the latest regulations. However, a sound governance, risk management, and compliance program within organizations can make the process much more efficient and effective.

SPECIALIZED FIREWALLS AVAILABLE

In this section, you will learn about some of the specialized firewalls that are available, such as:

- **Hybrid**—A hybrid firewall combines a number of different functions in a single appliance. You learned about firewalls with integrated VPN capabilities, but you can also expect to see firewalls combined with vulnerability management, IDS/IPS, antivirus/anti-spyware, content filtering, or a variety of other network and security-related technologies. When looking into hybrid firewalls, the key is to ensure that your specific security needs will be met without creating a performance bottleneck. The more functions you put on a single system, the more likely you will run into a performance bottleneck.
- **Data protection**—A data protection firewall is a hybrid firewall, but it deserves some additional explanation. The focus of a mature security program is to ensure that data is secure. **Data leakage prevention (DLP)** is a technology specifically designed to ensure that your data stays within your network. DLP technologies are systems that identify, monitor, and protect data in use, data in motion, and data at rest from inappropriate use, distribution, transmission, or other unauthorized actions. DLP technologies perform deep-content inspection within a scope defined by a central management console and are usually deployed at multiple locations within the environment to ensure full coverage. When combined with firewall technologies at the perimeter of the network, a data protection firewall becomes an invaluable tool in keeping your data out of email and social media accounts and within the internal network where it belongs.

- **Application**—An application firewall is specifically designed to control input, output, and/or access to an application. This type of firewall became very popular when the focus of attacks shifted from OSs to applications. An application firewall operates by monitoring the input, output, and system calls made in an application, and blocking any calls that do not conform to the firewall rules. An application firewall can block buffer overflow attacks, SQL injection attacks, and the majority of other application-focused attacks.

Application firewalls are either network-based or host-based applications. One drawback of application firewalls is that they are typically designed for specific types of applications (i.e., web applications or databases), whereas a standard firewall can secure many different types of applications—although not to the level an application firewall provides.

- **Database**—A database firewall is an application designed to control the input, output, and system calls made to a database. Although limited to database security, this firewall offers very granular security for databases. This is an essential tool in the prevention of SQL injection if your organization offers a website with an online catalog and ordering capabilities, and experiences a large number of online visitors to the website each day.
- **Ubiquitous**—This is likely the direction that firewalls will ultimately take in the future. The industry currently supports a variety of firewalls integrated into the network and the host that run on the perimeter and the core to protect demarcation between network zones. Now firewalls are moving into the virtualization area. Eventually, this intelligence and control will be pushed directly into the infrastructure and managed from a single console. Instead of controlling isolated ports on a firewall-by-firewall basis, you will be able to manage secure data as it flows into, out of, and through your environment—from the originating host to the perimeter of your network and at all connections in

between. This will allow you to reestablish the perimeter that disappeared with all the external connections to your network.

- **Cloud-based, cloud-deployed**—A cloud-based firewall is one where the firewall functions are performed in the cloud. The “cloud” is a term used to relay information about where and how data and services are stored and accessed. Not all cloud solutions are equivalent. For example, suppose your dinner plans include going to the cafeteria. Most will realize that the cafeteria offers a variety of food and beverage options, rather than a specific meal. This is how the cloud functions. Depending on which items are selected and ordered, the cloud experience can vary greatly from company to company in what is offered and in what is purchased. For our purpose, a cloud-based firewall is usually software deployed and offers a range of options that enable the company to protect data assets stored remotely in the cloud environment.

Specialized firewalls are becoming common in networks due to the commonality of these devices and the value of the enhancements.

With the move to the cloud for data storage and security services, many organizations are handing off this part of their security to their third-party providers. To ensure that your organization is receiving the level of security required and contracted, be sure to read the end-user license agreement (EULA) and service-level agreement (SLA) made with the cloud provider.

EMERGING NETWORK SECURITY TECHNOLOGIES

In addition to these optional enhancements to security technologies, you should also be aware of three more areas of emerging security technology.

Data Leakage Prevention

As the value of data continues to climb, DLP technology should see wider implementation. The technology is already heavily used in the financial industry. New government regulations will help drive the implementation in other sectors, most likely in health care. DLP is most concerned with the protection of data at rest (i.e., when data is stored) or transmitted between systems or networks. With more employees working remotely, it can be difficult to secure the entire work environment, as the endpoint used to connect to the corporate network could be a personal laptop in a coffee shop. The employee's personal laptop may be protected and secured; it could even be a company laptop. However, the medium used to connect—free Wi-Fi—is not secure. By using DLP, the data in the cloud is scanned for alteration and access. The systems accessing the data are assessed for compliance to corporate policy and may be blocked from access.

Biometrics

Biometrics identify a user based on anatomical characteristics such as a fingerprint, voice print, or iris patterns. These methods of identification have a number of advantages over passwords, tokens, or ID cards.

First, biometric authentication requires the person being authenticated to be physically present. Second, biometric security removes the need to remember complicated passwords. No more passwords taped under keyboards; you carry your password with you wherever you go. Finally, biometrics remove the need to carry a token or ID card with you. You do not have to worry about not being able to work when you leave your token on the kitchen table at home.

The areas where biometrics are currently being investigated include ATMs and computer networks. Biometric readers integrated into ATMs are becoming more popular because they tie a level of nonrepudiation—removing the ability to deny that you made the transaction—to the withdrawal of funds.

Virtualization Security

As virtualization and cloud computing continue to gain ground, a new generation of virtualization-aware security tools are under development. Antivirus, vulnerability management, data leakage prevention, and IDS/IPS technologies are all being developed to run against the underlying hypervisor layer. The **hypervisor** layer is the hardware or software on which virtual machines run. This gives the security applications direct access to the underlying data transport layer of the virtual environment, rather than forcing it to run against each virtual server, dramatically improving performance, visibility to security issues, and ease of use.

In multi-tenant environments (multiple companies sharing the same virtual environment), additional security tools are being developed to ensure that no access occurs between the virtual environments.

IP Version 6

IPv6 is the next-generation IP version and the successor to IPv4.

Although the main driving force for the redesign of Internet Protocol was the rapidly approaching exhaustion of IPv4 addresses, IPv6 has significant implications to future information security professionals.

IPv6 includes a native information security framework (IPSec) that provides for both data and control packets. So, what you currently do with a traditional VPN you will be able to do natively with any IPv6 device. At a high level, this means you can run your IPSec VPN without requiring a client, but the implications are significantly more profound. In a fully IPv6 environment, any connection can use an IPSec connection. Thus, any connection from a user to an application, a host to a host, or even a peer to a peer authenticates and encrypts as it passes across the network.

The thought of a network featuring nothing but secure connections seems like a security professional's dream configuration, but you should also consider some drawbacks before undertaking your IPv6 migration project. One of the challenges with encryption is that it not only secures data from attackers, but it also secures the data from authorized users. A number of security technologies like IDS/IPS, content filtering, network-based antivirus, data leakage prevention, and even firewall technologies rely on the ability to look at packets to determine how each should be handled. Once those packets are encapsulated in a secure IPSec connection, all of the security tools you have relied up on stop working.

With the limited deployment of IPv6, significant development on solutions to overcome these challenges has not occurred yet, but it will be a subject of great interest as the use of IPv6 expands.

VPNs, Firewalls, and Virtualization

Some Secure Sockets Layer (SSL) VPNs have the ability to provide a unique virtual VPN configuration for each individual user group.

Much like other types of virtualization, a virtualized SSL VPN allows you to separate the physical and logical use of the VPN. In the future, this capability could extend to IPSec VPNs, as well. This technology offers some unique abilities when configuring secure VPN contexts for different user groups, but you need to understand the additional complexity. A misconfigured virtual VPN context could expose parts of your network to groups that should not have access. For example, if you are using a virtualized VPN to provide customers access to a help desk ticketing system, and you inadvertently grant that context access to your intranet where all your pricing information resides, you expose the organization's proprietary data to unnecessary risk.

Virtual firewalls are a relatively new and growing area of technology. Currently, some firewalls on the market can partition into multiple virtual firewalls. Each virtual firewall appears to be a separate firewall with its own security policy, interfaces, and configuration. This allows you to use your firewall hardware more efficiently than you might otherwise, but once again, additional risks occur with the deployment of virtual firewalls.

First, firewalls exist that support this technology, but not all firewalls support all features in the virtualized environment. Before you deploy a virtualized firewall, be sure that it supports all the features you need to meet your business and security requirements. This is a very promising technology, but it is also very new. Sometimes early adopters can find themselves encountering issues the vendor did not discover during the quality assurance processes.

Next, you are relying on the logical segregation of firewalls rather than on the physical separation offered by multiple physical firewalls. While this technology remains new, do some testing before deploying virtual firewalls into a critical environment. Finding that there is a way to bypass the virtual security and move from one virtual environment to another would not be good if you are using the firewall to separate two customers connected to your network.

Finally, this model suffers from the same complexity challenges with respect to the virtual VPNs. Any time you have a solution that offers greater flexibility, you also open the possibility for greater complexity. Complex environments are almost always more difficult to secure, monitor, and manage than simple environments.

In spite of all the challenges associated with virtualizing security technologies like VPN and firewalls, a compelling business case exists for leveraging hardware more effectively in a virtualized environment. Be sure you understand the technology thoroughly before deploying it. Information security is seldom a forgiving field for learning as you go.

CHAPTER SUMMARY

In this chapter, you learned about a number of topics dealing with the future of security technologies and some of the significant emerging security threats. Future information security tools may come from the commercial or the open-source communities, and you learned about the pros and cons of both of those areas of development.

Governance, regulation, risk, and compliance are the pillars of good policy for guiding your use of security technologies. The obvious use for security technologies is to provide security; when all is said and done, however, information security needs to be a business enabler, created with an understanding of the business requirements and the bottom line.

Changes in networking technologies on the security perimeter can affect how network security operates. Firewall specialization includes hybrid, data protection, application, database, ubiquitous, and cloud-based firewalls.

KEY CONCEPTS AND TERMS

Compliance

Data analytics

Data leakage prevention (DLP)

Data sovereignty

Family Educational Rights and Privacy Act (FERPA)

General Data Protection Regulation (GDPR)

Governance, risk, and compliance (GRC)

Gramm-Leach-Bliley Act (GLBA)

Health Insurance Portability and Accountability Act (HIPAA)

HITECH Act

Hypervisor

Information Technology Infrastructure Library (ITIL)

Internet of Things (IoT)

IP Multimedia Subsystem (IMS)

Mobile IP

Payment Card Industry Data Security Standard (PCI DSS)

Sarbanes-Oxley (SOX) Act

CHAPTER 15 ASSESSMENT

1. Which of the following is *not* a benefit of virtualization's hypervisor?
 - A. Improved performance
 - B. Visibility to security issues
 - C. Ease of use
 - D. Deep-content inspection
2. Which of the following is *not* one of the best ways to consider security from a business perspective?
 - A. Compliance
 - B. Governance
 - C. Risk
 - D. Permissions
3. Which regulation was created to protect investors by requiring publicly traded companies to validate controls securing financial data?
 - A. Sarbanes-Oxley
 - B. Health Insurance Portability and Accountability Act (HIPAA)
 - C. HITECH Act
 - D. General Data Protection Regulation (GDRP)
4. Which of the following is *not* a type of specialized firewall?
 - A. Application
 - B. Data protection
 - C. Dynamic
 - D. Hybrid
5. Which of the following does *not* contribute to the erosion of the network perimeter?

- A. BYOD
 - B. Cloud computing
 - C. VPN
6. Which of the following specialized firewall types is designed to provide data leakage prevention?
- A. Application
 - B. Data protection
 - C. Database
 - D. Hybrid
7. Which of the following is *not* a biometric characteristic?
- A. Fingerprint
 - B. Height
 - C. Iris pattern
 - D. Voice pattern
8. Which of the following is *not* a core security principle?
- A. Availability
 - B. Confidentiality
 - C. Flexibility
 - D. Integrity
9. CIOs can be held accountable for security breaches in governmental compliance. When CIOs complain about security, which of the following is their top complaint?
- A. Lack of measures
 - B. The cost
 - C. Lack of flexibility
 - D. Downtime
10. Which of the following is an example of a biometric characteristic?
- A. Signature

- B. Token
- C. Password
- D. Smart card



© Digital_Art/Shutterstock

Answer Key

APPENDIX A

CHAPTER 1

Fundamentals of Network Security

1. D 2. D 3. C 4. B 5. B 6. C 7. C 8. C 9. A 10. A

CHAPTER 2

Network Security Threats

1. A 2. B 3. C 4. D 5. A 6. D 7. B 8. C 9. B 10. B

CHAPTER 3	Common Network Topologies and Infrastructures 1. B 2. A 3. A 4. C 5. C 6. A 7. D 8. C 9. B 10. A
CHAPTER 4	Network Design Considerations 1. B 2. D 3. A 4. C 5. D 6. C 7. B 8. D 9. D 10. A
CHAPTER 5	Firewall Fundamentals 1. C 2. C 3. A 4. C 5. A 6. C 7. A 8. C 9. A 10. B
CHAPTER 6	Firewall Implementation 1. C 2. B 3. B 4. C 5. D 6. C 7. A 8. C 9. A 10. C
CHAPTER 7	Firewall Deployment Considerations 1. D 2. A 3. C 4. A 5. C 6. C 7. D 8. A 9. A 10. C
CHAPTER 8	Configuring Firewalls 1. C 2. B 3. B 4. C 5. D 6. C 7. A 8. C 9. A 10. C
CHAPTER 9	VPN Fundamentals 1. D 2. B 3. A 4. C 5. C 6. A 7. D 8. D 9. C 10. D
CHAPTER 10	VPN Management 1. D 2. D 3. A 4. D 5. B 6. D 7. B 8. C 9. B 10. D
CHAPTER 11	VPN Technologies 1. B, C 2. B, C 3. C 4. D 5. D 6. B 7. B 8. A 9. A 10. B
CHAPTER 12	VPN Implementation 1. C 2. B 3. B 4. A 5. C 6. A 7. D 8. D 9. A 10. C
CHAPTER 13	Firewall Security Management 1. D 2. C 3. D 4. C 5. C 6. B 7. D 8. B 9. C 10. C

CHAPTER 14

Best Practices for Network Security Management

1. B 2. C 3. B 4. A 5. A 6. B 7. A 8. D 9. D 10. D

CHAPTER 15

Emerging Technology and Regulatory Considerations

1. D 2. D 3. A 4. C 5. C 6. B 7. B 8. C 9. A 10. A



© Digital_Art/Shutterstock

Standard Acronyms

APPEND
IX
B

3DES	triple data encryption standard
AAA	authentication, authorization, and accounting
ACD	automatic call distributor
ACK	acknowledgment

ACL	access control list
ADS	alternate data streams
AES	Advanced Encryption Standard
AH	Authentication Header
ALE	annualized loss expectancy
ANSI	American National Standards Institute
AP	access point
API	application programming interface
APT	advanced persistent threat
ARO	annualized rate of occurrence
ARP	Address Resolution Protocol
ASAv	Adaptive Security Virtual Appliance
ATM	asynchronous transfer mode
AUP	acceptable use policy
AV	asset value
B2B	business to business
B2C	business to consumer
BBB	Better Business Bureau
BCP	business continuity planning
BGP	BGP Border Gateway Protocol
BIOS	basic input/output system
BYOD	Bring Your Own Device
C2C	consumer to consumer
CA	certificate authority
CAP	Certification and Accreditation Professional
CAUCE	Coalition Against Unsolicited Commercial Email
CCC	CERT Coordination Center
CCNA	Cisco Certified Network Associate
CCPA	California Consumer Privacy Act
CERT	Computer Emergency Response Team
CFE	Certified Fraud Examiner
CGI	Common Gateway Interface

C-I-A	confidentiality, integrity, and availability
CIDR	classless inter-domain routing
CISA	Certified Information Systems Auditor
CISM	Certified Information Security Manager
CISSP	Certified Information System Security Professional
CLI	command line interface
CMIP	common management information protocol
COPPA	Children's Online Privacy Protection Act
CPE	customer premise equipment
CPU	central processing unit
CRC	cyclic redundancy check
CSI	Computer Security Institute
CTI	computer telephony integration
DBMS	database management system
DCOM	distributed component object model
DDoS	distributed denial of service
DD-WRT	DresDren-Wireless Router
DES	Data Encryption Standard
devfs	device file system
DHCP	dynamic host configuration protocol
DLLs	dynamic link libraries
DLP	data leakage prevention
DMZ	demilitarized zone
DNS	domain name system
DoS	denial of service
DPI	deep packet inspection
DRP	disaster recovery plan
DSL	digital subscriber line
DSS	digital signature standard
DSU	data service unit
ECC	elliptical curve cryptography
EDI	electronic data interchange

EF	exposure factor
EHR	electronic health record
EIDE	enhanced integrated development environment
EIGRP	Enhanced Interior Gateway Routing Protocol
EPHI	electronic protected health information
EPIC	Electronic Privacy Information Center
ESD	electrostatic discharge
ESP	encapsulating security payload
EU	European Union
EULA	end-user license agreement
FACTA	Fair and Accurate Credit Transactions Act
FAR	false acceptance rate
FBI	Federal Bureau of Investigation
FCC	Federal Communications Corporation
FDIC	Federal Deposit Insurance Corporation
FEP	front-end processor
FERPA	Family Educational Rights and Privacy Act
FIPS	Federal Information Processing Standards
FQDN	fully qualified domain names
FRCP	Federal Rules of Civil Procedure
FRR	false rejection rate
FTC	Federal Trade Commission
FTP	file transfer protocol
GDRP	general data protection regulation
GIAC	Global Information Assurance Certification
GLBA	Gramm-Leach-Bliley Act
GLBP	gateway load balancing protocol
GRC	governance, risk, and compliance
GUI	graphical user interface
HIDS	host-based intrusion detection system
HIPAA	Health Insurance Portability and Accountability Act
HIPS	host-based intrusion prevention system

HITECH	Health Information Technology for Economic and Clinical Health
HSRP	hot standby router protocol
HTML	hypertext markup language
HTTP	hypertext transfer protocol
HTTPS	HTTP over secure socket layer
HVAC	heating, ventilation, and cooling
IAB	Internet Activities Board
IAM	identity and access management
IANA	Internet Assigned Numbers Authority
ICMP	Internet control message protocol
ICS	industrial control system
IDEA	international data encryption algorithm
IDPS	intrusion detection and prevention system
IDS	intrusion detection system
IEEE	Institute of Electrical and Electronics Engineers
IEMI	intentional electromagnetic interference
IETF	Internet Engineering Task Force
IGMP	Internet group management protocol
IGRP	Interior Gateway Routing Protocol
IKE	Internet key exchange
IKEv2	Internet key exchange v2
IM	instant messaging
IMS	IP Multimedia Subsystem
InfoSec	information security
IoT	Internet of Things
IP	Internet Protocol
IPS	intrusion prevention system
IPSec	IP Security
IPv4	Internet protocol version 4
IPv6	Internet protocol version 6

IPX/SPX	Internetwork Packet Exchange/Sequenced Packet Exchange
IRC	Internet relay chat
IRS	Internal Revenue Service
(ISC)²	International Information System Security Certification Consortium
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
ISP	Internet service provider
ISS	Internet security systems
IT	information technology
ITIL	Information Technology Infrastructure Library
ITRC	Identity Theft Resource Center
IVR	interactive voice response
LAN	local area network
LCD	Liquid Crystal Display
LDAP	lightweight directory access protocol
L2F	layer 2 forwarding
L2TP	layer 2 tunneling protocol
MAC	Mandatory access code. Also Media Access Control.
MAN	metropolitan area network
MD5	Message Digest 5
MitM	man-in-the-middle
modem	MOdulator DEModulator
MPLS	multi-protocol label switching
MPPE	Microsoft point-to-point encryption
MTBF	mean time between failures
MTTF	mean time to failure
MTU	maximum transmission unit
NAC	network access control
NAS	network access server

NAT	network access translation
NII	National Information Infrastructure
NAT-PT	network address translation–protocol translation
NetBEUI	NetBios extended user interface
NFIC	National Fraud Information Center
NGFV	next-generation firewall virtual
NGFW	next-generation firewall
NGIPS	next-generation intrusion prevention system
NIC	network interface card
NIDS	network intrusion detection system
NIPS	network intrusion prevention system
NIST	National Institute of Standards and Technology
NMS	network management system
NNTP	network news transfer protocol
NPAS	network policy and access services
NSA	National Security Agency
NTFS	new technology file system
OBM	out-of-band management
OC	optical carrier
OpenVPN	open source virtual private network
Open WRT	Open Wireless RouTer
OS	operating system
OSI	open systems interconnection
OSPF	Open Shortest Path First
OWASP	Open Web Application Security Project
P2P	person to person
PAT	port address translation
PBX	private branch exchange
PC	personal computer
PCI	Payment Card Industry
PCI DSS	Payment Card Industry Data Security Standard
PGP	pretty good privacy

PHI	protected health information
PII	personally identifiable information
PKI	public key infrastructure
PNAC	port-based network access (admission) control
POP	Post Office Protocol
PPP	point-to-point protocol
PPTP	point-to-point tunneling protocol
PRNG	pseudo random number generator
QoS	quality of service
RADIUS	remote authentication dial-in user service
RAID	redundant array of independent disks
RAM	random access memory
RAS	remote access server
RBAC	role-based access system
RDC	remote desktop connection
RDP	remote desktop protocol
RDS	remote desktop services
RFC	request for comments
RFID	radio-frequency identification
RIP	Routing Information Protocol
RIR	regional Internet registry
ROI	return on investment
RRAS	routing and remote access service
RSA	Rivest, Shamir, and Adleman (algorithm)
SAN	storage area network
SANCP	security analyst network connection profiler
SANS	SysAdmin, Audit, Network, Security
SAP	service access point
SCSI	small computer system interface
SDD	Software Design Document
SED	static-electric discharge
SEM	security event management

SEO	search engine optimization
SET	secure electronic transaction
SGC	server-gated cryptography
SHA	secure hash algorithm
S-HTTP	secure HTTP
SIEM	security information and event management
SIM	security information management
SIRT	security incident response team
SLA	service level agreement
SLE	single loss expectancy
SMAC	social, mobile, analytics and cloud
SMFA	specific management functional area
SMTP	simple mail transfer protocol
SNA	systems network architecture
SNMP	simple network management protocol
SOHO	small office/home office
SOX	Sarbanes-Oxley Act of 2002 (also Sarbox)
SQL	Structured Query Language
SSA	Social Security Administration
SSCP	Systems Security Certified Practitioner
SSH	secure shell
SSID	service set identifier
SSL	secure sockets layer
SSL/TLS	secure sockets layer/transport layer security
SSO	single sign-on
STP	shielded twisted cable
STIG	Security Technical Implementation Guide
SYN	synchronize
SYN/ACK	synchronization acknowledged
TACACS	Terminal Access Controller Access-Control System
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol

TCSEC	trusted computer system evaluation criteria
TFTP	trivial file transfer protocol
TLS	Transport Layer Security
TNI	trusted network interpretation
TPM	trusted platform module
UDP	user datagram protocol
UNIX	Uniplexed Information and Computer Systems
UPS	uninterruptible power supply
URL	uniform resource locator
USB	universal serial bus
UTM	unified threat management
UTP	unshielded twisted cable
VLAN	virtual local area network
VOIP	Voice over Internet Protocol
VPN	virtual private network
VRRP	virtual router redundancy protocol
WAN	wide area network
WAN	wide area network virtual
VPN	private network
WAP	wireless accent point
WebGUI	web-based graphical user interface
WLAN	wireless local area network
WNIC	wireless network interface card
WORM	write-once read-many
W3C	World Wide Web Consortium
WWW	World Wide Web
XXS	Cross-site scripting



© Digital_Art/Shutterstock

Glossary of Key Terms

802.1x | Port or portal authentication. A mechanism commonly used by network devices, such as firewalls, routers, switches, and wireless access points, to perform authentication of users before allowing communication to continue across or through the device. The authentication can take place locally on the device or go to an authentications service, such as a credit card payment system, PKI, or directory service.

A

Access control | The process or mechanism of granting or denying use of a resource; typically applied to users or generic network traffic.

Acceptable use policy (AUP) | The structured document that outlines what is permitted to do and view when using company equipment and computers.

Access control list (ACL) | Mechanism defining traffic or an event to apply an authorization control of allow or deny against. Often used interchangeably with the terms *rule* and *filter* in relation to firewalls. An ACL focuses on controlling a specific user or client's access to a protocol or port.

Active hub | A device that consolidates wiring so that a single cable can attach to a server and includes the capability to regenerate the signal to decrease latency and increase the distance the signal can be sent.

Active threat | A form of threat that takes some type of initiative to seek out a target to compromise. These can be hackers, intruders, or automated worms. In any case, an active threat seeks out vulnerable targets. If you do not have reasonable security measures and the active threat discovers your system, you might be at risk for a compromise.

Advanced persistent threat (APT) | A network attack in which an unauthorized person gains access to a network and stays there undetected for a long period of time. The purpose of such an attack is to steal data, not to damage the network or organization. Sectors with high-value information, such as national defense, manufacturing, and the financial industry, are commonly the target of such attacks.

Adware | Unwanted software that displays advertisements. Often linked with spyware.

Agent | A malicious software program distributed by a hacker to take over control of a victim's computers. Also known as a bot or a zombie. Agents are commonly used to construct botnets.

Aggregation | Collecting disparate log information from various systems and locations into one location.

Alert | A notification from a firewall that a specific event or packet was detected. Alerts notify administrators of events that may need real-time human response or attention.

Algorithm | A set of rules and procedures, usually mathematical in nature. Algorithms can define how the encryption and decryption processes operate. Often very complex, many algorithms are publicly known; anyone can investigate and analyze the strengths and weaknesses of an algorithm.

Alternate data stream (ADS) | A feature added to the NTFS file system to support files from POSIX, OS/2, and Macintosh. ADS supports multiple resource forks for file objects. Hackers use ADS to hide files.

Annualized loss expectancy (ALE) | The calculation of the total loss potential across a year for a given asset and a specific threat. ALE calculations are part of risk assessment. $ALE = SLE \times ARO$.

Annualized rate of occurrence (ARO) | A probability prediction based on statistics and historical occurrences on the likelihood of how many times in the next year is a threat going to cause harm. ARO is used in the ALE calculation.

Anomaly-based detection | A form of intrusion detection system/intrusion prevention system (IDS/IPS) based on a defined normal, often defined using rules similar to firewall rules. All traffic or events that fail to match defined normal are considered anomalies and potentially malicious.

Anonymity | The ability for a network or system user to remain unknown. A number of tools and techniques provide anonymity when connected to a network, although the underlying network protocols make true anonymity very difficult.

Anti-forensics | A series of tools and techniques used to prevent forensic examination from identifying an attack or attacker.

Antivirus scanner | A software program that identifies the unique characteristics of malware (known as a virus signature) to detect and prevent infection of computer components.

Appliance | A hardware product that is dedicated to a single primary function. The operating system or firmware of the hardware device is

hardened, and its use is limited to directly and exclusively supporting the intended function. Firewalls, routers, and switches are typical appliances.

Appliance firewall | A hardened hardware firewall.

Application firewall | A type of firewall that filters on a specific application's content and session information.

Application gateway | See application firewall.

Application Layer (Layer 7) | The top or seventh layer of the OSI model. This layer is responsible for enabling communications with host software, including the operating system. The Application Layer is the interface between host software and the network protocol stack. The subprotocols of this layer support specific applications or types of data.

Application proxy | See application firewall.

Arbitrary code execution | An exploit that allows a hacker to run any command line function on a compromised system. Buffer overflow attacks and SQL injection attacks can often allow arbitrary code execution.

ARP spoofing | The falsification of ARP replies to trick the requestor into sending frames to a system other than its intended destination.

Asset | Anything you use in a business process to accomplish a business task.

Asset value (AV) | The cumulative value of an asset based on both tangible and intangible values. AV supports the SLE calculation.

Asymmetric cryptography | A means of encoding and decoding information using related but different keys for each process. A key used to encode cannot decode, and vice versa. Cryptography is based on algorithms that use either key pairs or some other special mathematical mechanism. Asymmetric cryptography that uses key pairs is commonly known as public key cryptography. Different keys serve different purposes. Different keys are used by different members of the communication session. Some systems use something different from keys altogether.

Attack surface | Portions of a software system that unauthenticated users can run.

Attacking | The act of targeting a computer or network and using available means to gain information or access to the system.

Auditing | Act of conducting an audit. Auditing can be the action of a system that is recording user activity and system events into an audit log. Auditing can also be the action of an auditor who checks for compliance with security policies and other regulations.

Auditor | Either an outside consultant or an internal member of the information technology staff. The auditor performs security audits, confirms that auditing is sufficient, and investigates audit trails produced by system auditing. In the case of regulatory compliance, auditors should be external and independent of the organization under audit.

Authentication | The process of confirming the identity of a user. Also known as logon.

Authentication, authorization, and accounting (AAA) | Programs used to control access to computer resources, enforce policies, audit usage, and provide billing information. Examples include RADIUS, TACACS, 802.1x, LDAP, and Active Directory.

Authentication Header (AH) | A protocol that provides integrity protection for packet headers and data, as well as user authentication.

Authenticity | The security service of the combination of authentication and access control (authorization) that provides either the identity of the sender of a message or controls who is the receiver of a message.

Authorization | Defining what users are allowed and not allowed to do; also known as access control.

Availability | When a system is usable for its intended purpose. The security service that supports access to resources in a timely manner. If availability becomes compromised, a denial of service is taking place.

Avalanche effect | A common feature of hash algorithms. This effect ensures that small changes in the input data produce large changes in the outputted hash value. A single binary digit change in a file should produce a clearly recognizable difference in the resultant hash value.

B

Backdoor | Unauthorized access to a system. A backdoor is any access method or pathway that circumvents access or authentication mechanisms.

Backup | The process of making copies of data onto other storage media. The purpose of a backup is to protect against data loss by having additional on-site or off-site copies of data that can be restored when necessary.

Banner | A message sent by a service in response to a valid or invalid query. A banner can confirm that communication is functioning properly or announce an error. Some banners disclose the product name and version number of the service.

Banner grabbing | The act of capturing or extracting banners from services. Hackers often perform banner grabbing after port scanning to learn which service is active on a port.

Bastion host | A firewall positioned at the initial entry point where a network interfaces with the Internet. It serves as the first line of defense for the network; also known as a sacrificial host.

Bastion host OS | A system designed, built, and deployed specifically to serve as a frontline defense for a network.

Behavioral-based detection | A form of IDS/IPS detection based on a recording of real-world traffic as a baseline for normal. All traffic or events that fail to match the normal baselines are considered abnormal and potentially malicious.

Bidirectional | A type of ring topology that requires two connections to each device to provide redundancy.

Blacklist | A type of filtering in which all activities or entities are permitted, except for those on the blacklist; also known as a block list.

Border sentry | A description often applied to firewalls positioned on network zone transitions or gateway locations.

Bot | A malicious software program distributed by hackers to take over control of victims' computers. Also known as *agent* or *zombie*. Bots are commonly used to construct botnets.

Botnet | A compromised system controlled by a hacker.

Botnet army | A network of zombie/bot/agent-compromised systems controlled by a hacker.

Bottleneck | Any restriction on the performance of a system. Can be caused by a slower component or a pathway with insufficient throughput. A bottleneck causes other components of system to work slower than their optimum rate.

Breach | Any compromise of security. Any violation of a restriction or rule whether caused by an authorized user or an unauthorized outsider.

Bricking | The process of making a device nonfunctional, either intentionally or unintentionally (such as when applying a patch that goes bad).

Bridge | A network device that forwards traffic between networks based on the MAC address of the Ethernet frame. A bridge forwards only packets whose destination address is on the opposing network.

Bring Your Own Device (BYOD) | A policy of allowing or even encouraging employees, contractors, and others to connect their own computers, smartphones, and other devices to their organization's networks. BYOD can save expenses and afford employees more autonomy, but it can also compromise security.

Router | A device that combines a bridge and a router. Also known as a bridging router.

Brute-force password attack | A form of password or encryption key cracking attack that tries all possible valid combinations from a defined set of possibilities (e.g., a set of characters or hex values). Brute-force attacks will eventually generate a valid solution given enough time, assuming the hacker uses the correct set of possibilities.

Buffer | A memory area that is designated to receive input for processing.

Buffer overflow | A condition in which a memory buffer exceeds its capacity and extends its contents into adjacent memory; often used as an attack against poor programming techniques or poor software quality control. Hackers can inject more data into a memory buffer than it can hold, which may result in the additional data overflowing

into the next area of memory. If the overflow extends to the next memory segment designated for code execution, a skilled attacker can insert arbitrary code that will execute with the same privileges as the current program. Improperly formatted overflow data may also result in a system crash.

Bump-in-the-stack | A term for a firewall that is implemented via software.

Bump-in-the-wire | A term for a firewall that is a separate hardware implementation.

Bus topology | A linear arrangement of computers and peripherals.

Business continuity plan | A plan to maintain the mission-critical functions of the organization in the event of a problem that threatens to take business processes offline. The goal of business continuity planning is to prevent the interruption of business tasks, even with a damaged environment and reduced resources.

Business task | Any activity necessary to meet an organization's long-term goals. Business tasks are assigned to employees and other authorized personnel via their job descriptions.

Bypass deployment | A configuration deploying a VPN without a firewall to filter traffic to the VPN.

C

Caching | Retention of Internet content by a proxy server. Various internal clients may access this content and provide it to subsequent requesters without the need to retrieve the same content from the Internet repeatedly.

California Consumer Privacy Act (CCPA) | Legislation that went into effect January 2020 requiring companies to protect consumer information stored on their systems.

Centralized logging system | A technique of storing or copying log events to a centralized logging server. This mechanism is used to create a redundant copy of all log files in a single warehousing location. A common example of this is syslog.

Certificate Authority (CA) | A trusted third-party entity that issues digital certificates to verify and validate identities of people, organizations, systems, and networks digitally.

Change management | An process for managing and tracking changes to systems or network configurations.

Channel | A communication pathway, circuit, or frequency dedicated or reserved for a specific transmission.

Chip creep | The slow movement of a chip out of its socket or solder points because of expansion and contraction caused by extreme temperature fluctuations.

Chokepoint | Similar to a bottleneck, but deliberately created within a network infrastructure. A chokepoint is a controlled pathway through which all traffic must cross. At this point, filtering to block unwanted communication or monitoring can occur.

Ciphertext | The seemingly random and unusable output from a cryptographic function applied to original data. Ciphertext is the result of encryption. Decryption converts ciphertext back into plaintext.

Circuit | A logical connection between a client and a resource server. May exist at Layer 3, 4, or 5 of the OSI model. Also known as a session or a state.

Circuit firewall | A filtering device that allows or denies the initial creation of a circuit, session, or state, but performs no subsequent filtering on the circuit once established.

Circuit proxy | See circuit firewall.

Client | A host on a network. A client is the computer system, which supports user interaction with the network. Users employ a client to access resources from the network. Users can also employ a client generically as any hardware or software product to access a resource. For example, standard e-mail software is a client.

Client/server network | A form of network where certain computers are designated as “servers” to host resources shared with the network. The remaining computers are designated as “clients” to enable users to access shared resources. Most client/server networks employ directory services and single sign-on. Also known as a domain.

Client-to-server VPN | A VPN created between a client and a server either within the same local network or across a WAN link or intermediary network to support secure client interaction with the services of a resource host. Also known as a host-to-host VPN.

Clipper chip | A chipset developed and promoted by the U.S. government as an encryption device to be adopted by telecommunications companies for voice transmission. It was announced in 1993 and was discontinued in 1996.

Closed source | A type of software product that is precompiled and whose source code is undisclosed.

Cluster | A logical division of data composed of one or more sectors on a hard drive. A cluster is the smallest addressable unit of drive storage, usually 512, 1,024, 2,048, or 4,096 bytes, depending on the logical volume size.

Cold calling | A tactic of pursuing and extracting information for the purpose of making a sale or performing a social engineering attack. A cold call presupposes little or no knowledge of the person answering the phone. It requires the caller to be able to pick up on vocal and word clues, be knowledgeable about human nature, and adapt quickly to changes in conversation.

Command shell | A software interface with a system that allows code execution. A command shell is often the focus of an attack. If a hacker gains access to a command shell, he or she can perform arbitrary code execution. Also known as a terminal window or a command prompt. For example, in Windows, the command shell prompt is usually “C:\>”.

Commercial firewall | A firewall product designed for larger networks. Usually, a commercial firewall is a hardware device.

Common Gateway Interface (CGI) script | The Common Gateway Interface (CGI) is a standard that defines how web server software can delegate the generation of webpages to a console application. Such applications are known as CGI scripts. They can be written in many programming languages, although scripting languages are often used.

Compliance | Ensuring adherence to applicable laws or regulatory requirements.

Compliance audit | A detailed and thorough review of the deployed security infrastructure compared with the organization’s security policy and any applicable laws and regulations.

Compression | Removal of redundant or superfluous data or space to reduce the size of a data set. Compression consumes less storage space and increases the speed of data transmission.

Concentrator | See smart hub.

Confidentiality | The security service of preventing access to resources by unauthorized users, while supporting access to authorized users.

Containment | Isolation of traffic or communication to prevent further escalation.

Content filtering | A form of filtering that focuses on traffic content. Application proxies perform most content filtering.

Contract worker | An outsider brought into an organization to work on a temporary basis. Contracted workers can be consultants, temporary workers, seasonal workers, contractors, or even day-laborers. Contracted workers potentially represent a greater risk than regular, full-time regular employees because they might lack loyalty,

not see the company as worthy of protection, might not be accountable after a project ends, and so on.

Cookie filter | A cookie is a small text file used by web browsers and servers to track web sessions. A cookie filter blocks the sending and receiving of cookies. Blocking cookies can reduce some threats of session tracking and identify theft, but can also disable many web-based services, such as online purchasing.

Corporate firewall | An appliance firewall placed on the border or edge of an organization's network.

Correlation | A relationship between log files using statistical measures.

Cost/benefit analysis | The final equation of risk analysis to assess the relative benefit of a countermeasure against the potential annual loss of a given asset exposed to a specific threat.

Countermeasure | An action undertaken to minimize a threat or thwart progression of a negative event.

Covert channel | An unknown, secret pathway of communication. Covert channels can be timing or storage-based.

Cross-site scripting (XSS) | The malicious insertion of scripting code onto a vulnerable website. The results of an XSS attack can include the corruption of the data on the website or identity theft of the site's visitors.

Cryptography | The art and science of hiding information from unauthorized third parties. Cryptography is divided into two main categories: encryption and decryption.

Customer premise equipment (CPE) | A customer premise equipment-based VPN. This VPN is also known as a VPN appliance.

D

Data analytics | Statistics and other tools used to apply understanding to data such as log files.

Data leakage prevention (DLP) | A distributed data protection technology that leverages deep analysis, context evaluation, and rules configured from a central console to ensure that confidential information remains secure while in use, in transit, and at rest.

Data Link Layer (Layer 2) | The second layer of the OSI model responsible for physical addressing (MAC addresses) and supporting the network topology, such as Ethernet.

Data sovereignty | Data is subject to the laws of the country where it is stored.

Database-based detection | A form of IDS/IPS detection based on a collection of samples, patterns, signatures, and so on stored in a database of known malicious traffic and events. All traffic or events that match an item in the database are considered abnormal and potentially malicious. Also known as signature, knowledge, and pattern-matching based detection.

DDoS mitigator | A device that inspects traffic entering a network.

Dead-man switch | A form of auto-initiation switch that triggers when the ongoing prevention mechanism fails. Common dead-man switches include firewalls and hand grenades. If the firewall stops functioning, the connection is severed. If a person dies while holding a live grenade, the safety latch opens and the grenade explodes.

Decryption | The process of converting cipher text back into plain text.

Dedicated connection | A network connection that is always on and available for immediate transmission of data. Most leased lines are dedicated connections.

Dedicated leased line | *See* dedicated connection *and* leased line.

Deduplication | Elimination of duplicate notifications pertaining to the same incident or device rule when aggregating logs from several devices.

De-encapsulation | The action of processing the contents of a header, removing that header, and sending the remaining payload up to the appropriate protocol in the next higher layer in the OSI model.

Default allow | A security stance that allows all communications except those prohibited by specific deny exceptions; also known as allow by default.

Default deny | A security stance that blocks all access to all resources until a valid authorized explicit exception is defined.

Defense in depth | A tactic of protection involving multiple layers or levels of security components. Based on the idea that multiple protections create a cumulative effect that will require an attacker to breach all layers, not just one.

Demilitarized zone (DMZ) | A type of perimeter network used to host resources designated as accessible by the public from the Internet.

Denial of service (DoS) attack | A form of attack that attempts to compromise availability. DoS attacks are usually of two types: flaw exploitation and flooding. DDoS (distributed denial of service) often involves the distribution of robots, zombies, or agents to thousands or millions of systems that are then used to launch a DoS attack against a primary target.

Deny by default/allow by exception | A security stance that prevents all communications, except those enabled by specific allow exceptions. Also known as default deny.

Desktop virtualization | Separates the personal computer environment from the physical machine by using a client/server model.

Detection | Identification of breach attempts prior to, during, or after access.

Deterrence | The use of security to make a target less appealing.

Deterrent | A form of security defense that focuses on discouraging a perpetrator with disincentives such as physical harm, social disgrace, or legal consequences. A deterrent can also be a defense that is complex or difficult to overcome, such as strong encryption, multifactor authentication, or stateful inspection filtering.

Dialer | A rogue program that automatically dials a modem to a predefined number. Sometimes this is to auto-download additional malware to the victim or to upload stolen data from the victim. In other cases, the dialer calls premium rate telephone numbers to rack up massive long-distance charges.

Dictionary password attack | A form of password or encryption key-cracking attack that uses a preconstructed list of potential passwords or encryption keys.

Digital certificate | An electronic proof of identity issued by a certificate authority (CA). A digital certificate is an entity's public key encoded by the CA's private key.

Digital envelope | A secure communication based on public key cryptography that encodes a message or data with the public key of the intended recipient.

Digital signature | A public key cryptography-based mechanism for proving the source (and possibly integrity) of a signed dataset or message. A digital signature uses the private key of a sender. Not the same as a "digitized signature," which is a digital image of handwriting.

Direct Access | A Microsoft solution that can be used as an alternative to a traditional VPN.

Directory service | A network service that maintains a searchable index or database of network hosts and shared resources. Often based on a Domain Name System (DNS). An essential service of large networks.

Disaster recovery plan | A plan to restore the mission-critical functions of the organization once they have been interrupted by an adverse event. The goal of disaster recovery planning is to return the business to functional operation within a limited time to prevent the failure of the organization due to the incident.

Disgruntled employee | A worker who feels wronged by his or her employer and who may take malicious, unethical, potentially illegal actions to exact revenge on the organization.

Distributed denial of service (DDoS) attack | An attack that uses multiple remotely controlled software agents disseminated across the

Internet. Because the denial of service attack comes from multiple machines simultaneously, it is “distributed.” DDoS attacks can include flooding, spam, eavesdropping, interception, MitM, session hijacking, spoofing, packet manipulation, distribution of malware, hosting phishing sites, stealing passwords, cracking encryption, and more.

Distributed LAN | A LAN whose components are in multiple places that are interconnected by WAN VPN links.

Diversity of defense | An approach to security similar to defense in depth, in that it supports multiple layers; it is unlike it, in that it uses a different security mechanism at each, or most, of the layers.

DNS poisoning | A form of exploitation in which the data on a DNS server is falsified so subsequent responses to DNS resolution queries are incorrect. DNS poisoning can wage man-in-the-middle attacks.

DNS spoofing | A form of exploitation in which unauthorized or rogue DNS server responds to DNS queries with false, spoofed resolutions. DNS poisoning can wage man-in-the-middle attacks.

Domain | A client/server network managed by a directory service.

Domain Name System (DNS) | A network service that resolves fully qualified domain names (FQDNs) into their corresponding IP address. DNS is an essential service of most networks and their directory services.

Domain registration | The information related to the owners and managers of a domain name accessed through domain registrar’s websites and whois lookups. A domain registration might include a physical address, people’s names, e-mail addresses, and phone numbers. This information is useful in waging social engineering attacks.

Downtime | Any planned or unplanned period when a network service or resource is not available. Downtime can be caused by attack, hardware failure, or scheduled maintenance. Most organizations strive to minimize downtime through security and system management.

Dual-homed firewall | A firewall that has two network interfaces. Each network interface is located in a unique network segment. This

allows for true isolation of the segments and forces the firewall to filter all traffic moving from one segment to another.

Dumb hub | *See* hub

Dumpster diving | A type of reconnaissance in which an attacker examines an organization's trash or other discarded items to learn internal or private information. The results of dumpster diving are often used to wage social engineering attacks.

Dynamic packet filtering | The process of automatically created temporary filters. In most cases, the filters allow inbound responses to previous outbound requests. Also called stateful inspection.

E

Eavesdropping | The act of listening in on digital or audio conversations. Network eavesdropping usually requires a sniffer, protocol analyzer, or packet-capturing utility. Eavesdropping may be able to access unencrypted communication, depending on where it occurs.

Edge router | A router positioned on the edge of a private network. Usually an edge router is the last device owned and controlled by an organization before an ISP or telco connection.

Education | The third and highest level of obtaining security knowledge that leads to career advancement. Security education is broad and not necessarily focused on specific job tasks or assignments. More rigorous than awareness or training.

Egress filtering | Filtering traffic as it attempts to leave a network, which can include monitoring for spoofed addresses, malformed packets, unauthorized ports and protocols, and blocked destinations.

Electronic Privacy Information Center (EPIC) | A public interest research group in Washington, D.C., established in 1994 to focus public attention on emerging civil liberties issues and to protect privacy, the First Amendment, and Constitutional values in the information age. It pursues a wide range of activities, including privacy research, public education, conferences, litigation, publications, and advocacy. It maintains two of the world's most popular privacy sites—epic.org and privacy.org—and publishes the online EPIC Alert every two weeks with information about emerging privacy and civil liberties issues.

Electrostatic discharge (ESD) | Static-electric discharge that can damage equipment.

Encapsulating Security Payload (ESP) | The second core IPSec security protocol; it can perform authentication to provide integrity protection, although not for the outermost IP header.

Encapsulation | The process of enclosing or encasing one protocol or packet inside another protocol or packet. Also known as “tunneling.” Encapsulation allows for communications to cross intermediary networks that might be incompatible with the original protocol.

Encapsulation is distinct from encryption, but many encapsulation protocols include encryption.

Encryption | The process of converting original data into a chaotic and unusable form to protect it from unauthorized third parties.

Decryption returns the data back to its original, usable form.

Encryption key | The mathematical formula solution that will decrypt a file.

Enumeration | The process of discovering sufficient details about a potential target to learn about network or system vulnerabilities.

Enumeration often starts with operating system identification, followed by application identification, then extraction of information from discovered services.

Eradication | Fixing vulnerabilities. Usually performed during recovery.

Ethical hacking | Also known as penetration testing, the application of hacking techniques with good intentions.

Exploit | An attack tool, method, or technique a hacker uses to take advantage of a known vulnerability or flaw in a target system.

Exposure factor (EF) | The potential amount of harm from a specific threat stated as a percentage. Used in the calculation of SLE.

Extranet | A type of perimeter network used to host resources designated as accessible to a limited group of external entities, such as business partners or suppliers, but not by the public. Often, access to an extranet requires the use of a virtual private network or VPN, especially when access originates from the Internet.

Extranet VPN | A VPN used to grant outside entities access into a perimeter network; used to host resources designated as accessible to a limited group of external entities, such as business partners or suppliers, but not the general public.

F

Fail-closed | A failure response resulting in no access as the device or port is closed.

Fail-open | A failure response resulting in open and unrestricted access or communication.

Fail-safe | A failure response resulting in a secured or safe level of access or communication.

Fail-secure | A failure response resulting in a secured or safe level of access or communication.

Fair queuing | A technique of load balancing that operates by sending the next transaction to the firewall with the least current workload.

Fallback attack | The process of using alternate attacks when the initial attack fails to provide access to a system or network.

False negative | An event that does not trigger an alarm but should have, due to the traffic or event actually being abnormal and/or malicious. This is the unwanted nondetection of a malicious event.

False positive | An event that triggers an alarm but should not have, due to the traffic or event actually being benign. This is the unwanted false alarm that wastes time and resources pursuing a nonmalicious event.

Family Educational Rights and Privacy Act (FERPA) | A federal law that protects the privacy of student information.

File encryption | A form of security protection that protects individual files by scrambling the contents in such a way as to render them unusable by unauthorized third parties.

File Transfer Protocol (FTP) | A protocol and a data exchange system commonly used over TCP/IP networks, including the Internet, but which is unencrypted and performs authentication and data transfer in plaintext.

Filter | A written expression of an item of concern (protocol, port, service, application, user, IP address) and one or more actions to take when the item of concern appears in traffic. A filter expresses the

intention to block or deny unwanted items of concern. Also known as a rule or ACL.

Filtering | The process of inspecting content against a set of rules or restrictions to enforce allow-and-deny operations on that content. Firewalls and other security components use filtering.

Firewalking | A hacking technique used against static packet filtering firewalls to discover the rules or filters controlling inbound traffic.

Firewall | A network security device or host software that filters communications, usually network traffic, based on a set of predefined rules. Unwanted content is denied, and authorized content is allowed. Also known as a sentry device.

Flaw exploitation attack | A form of DoS that uses a software specific exploit to cause the interruption of availability. Once you apply the appropriate patch, the system is no longer vulnerable to this particular exploit.

Flooding | An attack, usually resulting in a DoS, in which hackers direct massive amounts of traffic toward a target to fully consume available bandwidth or processing capabilities.

Footprinting | The act of researching and uncovering information about a potential attack target; also known as reconnaissance.

Fragmentation | This occurs when a dataset is too large for maximum supported size of a communication container, such as a segment, packet, or frame. The original dataset divides into multiple sections or fragments for transmission across the size-limited medium, then reassembles on the receiving end. Fragmentation can sometimes corrupt or damage data or allow outsiders to smuggle malicious content past network filters.

Fragmentation attack | An attack based on the abuse of fragmentation.

Frame | The collection of data at the Data Link Layer (Layer 2) of the OSI model, defined by the Ethernet IEEE 802.3 standard, that consists of a payload from the Network Layer (Layer 3) to which an Ethernet header and footer have been attached.

Fully qualified domain name (FQDN) | A complete Internet host name including a top-level domain name, a registered domain name,

possibly one or more subdomain names, and a host name. Examples include: **www.ittech.edu** and **maps.google.com**. A DNS is used to resolve FQDNs into IP addresses.

Fuzzing tool | Hacking and testing utilities that use a brute force technique to craft packets and other forms of input directed toward the target. Fuzzing tools place stress on a system, pushing it to react improperly, to fail, or to reveal unknown vulnerabilities.

G

Gateway | An entrance or exit point to a controlled space. A firewall is often positioned at a gateway of a network to block unwanted traffic.

Gateway-to-gateway VPN | A VPN model used to connect offices together, such as a main office and a remote office. It is also referred to as a site-to-site VPN.

General-purpose OS | An operating system such as Windows, Linux, Mac OS, or UNIX, which can support a wide variety of purposes and functions, but which, when used as a bastion host OS, must be hardened and locked down.

General Data Protection Regulation (GDPR) | A European Union regulation to protect citizen's privacy and information.

Governance, risk, and compliance (GRC) | A concept of applying security technologies in alignment with law and regulatory requirements.

Gramm-Leach-Bliley Act (GLBA) | Legal regulation concerning financial institutions to protect customer information.

Graphical user interface (GUI) | An interface using icons (pictures) instead of menus.

H

Hacker | A person who performs hacking. Modern use of this term now implies malicious or criminal intent by the hacker, although criminals are more correctly known as “crackers.” An “ethical hacker” obtains the permission of the owner of a system before hacking.

Hacking | The act of producing a result not intended by the designer of a system. Hackers may perform such acts out of curiosity or malice. Malicious hacking is known as “cracking,” but many people typically call all these actions “hacking,” regardless of intent.

Hacktivism | Politically or socially motivated hacking, seen by activists as a form of civil disobedience in the interest of free speech and human rights, but seen by its opponents as a form of cyberterrorism.

Hairpinning | A process by which malicious code can enter from a nonsecure network and make a hairpin—or sharp—turn. It can enter a secure network with little or no trouble because it is entering from a secure and verified endpoint. Hairpinning is a particular issue for organizations whose work-at-home employees want to connect to both an unsecure network and a VPN at the same time.

Hardening | The process of securing or locking down a host against threats and attacks. This can include removing unnecessary software, installing updates, and imposing secure configuration settings.

Hardware address | The physical address assigned to a network interface by the manufacturer. Also known as the MAC address.

Hardware firewall | An appliance firewall. A hardened computer product that hosts firewall software exclusively.

Hardware VPN | A dedicated device hosting VPN software. Also known as an appliance VPN. Hardware VPNs can connect hosts and/or networks.

Hash algorithm | A set of mathematical rules and procedures that produces a unique number from a dataset. *See* hash or hash value *and* hashing.

Hash or hash value | The unique number produced by a hash algorithm when applied to a dataset. A hash value verifies the integrity of data.

Hashing | The process of verifying data integrity. Hashing uses hash algorithms to produce unique numbers from datasets, known as hash values. If before-and-after hash values are the same, the data retain integrity.

Header | The additional data added to the front of a payload at each layer of the OSI model that includes layer-specific information.

Health Insurance Portability and Accountability Act (HIPAA) | Legal protection for healthcare records that went into effect in 1996.

Hierarchical File System (HFS) | A storage device file system developed by Apple Inc. for use on Macintosh computers. HFS supports multiple resource forks for file objects.

High availability | Percentage of time when a service or device is available; the higher the number, the better.

Hijacking | This attack occurs when a hacker uses a network sniffer to watch a communications session to learn its parameters. The hacker then disconnects one of the session's hosts, impersonates the offline system, and then begins injecting crafted packets into the communication stream. If successful, the hacker takes over the session of the offline host, while the other host is unaware of the switch.

HITECH | This 2009 Act expanded HIPAA with more specific requirements for computers and electronic records.

Honeynet | A collection of multiple honeypots in a network for the purposes of luring and trapping hackers.

Honeypot | A closely monitored system that usually contains a large number of files that appears to be valuable or sensitive, and serves as a trap for hackers. A honeypot distracts hackers from real targets, detects new exploitations, and learns the identities of hackers.

Host | A node that has a logical address assigned to it, usually an IP address. This typically implies that the node operates at and/or above the Network Layer. This would include clients, servers, firewalls, proxies, and even routers. The term excludes switches, bridges, and

other physical devices such as repeaters and hubs. In most cases, a host either shares or accesses resources and services from other hosts.

Host firewall | A software firewall installed on a client or server.

Host VPN | A VPN endpoint located on a host client or server. A host VPN relies on either a native feature of the operating system or a third-party application.

Host-to-gateway VPN | A VPN model where the remote client connects to the VPN server to gain access to the internal network.

Host-to-host VPN | A VPN created between two individual hosts across a local or intermediary network. Host-to-host VPNs are also known as client-to-server, remote-to-office, or remote-to-home VPNs.

Host-to-site VPN | A VPN created between a host and a network across a local or intermediary network. Also known as a remote access VPN.

HOSTS file | A static file on every IP-enabled host where FQDN-to-IP address resolutions can be hard-coded.

Hot standby router protocol (HSRP) | A proprietary standard that allows configuration of a pair of Cisco VPNs for fail-over performance.

Hub | A device that concentrates wired connections so only one cable connects to the next device.

Hybrid attack | A form of password or encryption key-cracking attack that combines dictionary attacks with brute-force attacks. A dictionary list provides seed values to a brute-force attack tool that makes modifications to the seed value. A very effective attack against users who mistakenly believe that changing a few characters or adding a few characters to a base password is actually improving the password's strength. For example, hybrid attacks may combine dictionary words with a digit or two to increase the likelihood of obtaining a successful result.

Hybrid VPN | A form of VPN establishing a secure VPN over trusted VPN connections.

I

ICMP redirect | An announcement message sent to hosts to adjust the routing table. ICMP type 5 messages are known as redirects. Hackers can use ICMP redirects to perform man-in-the-middle or session hijacking attacks.

Identity and access management (IAM) | The security discipline that enables the right individuals to access the right resources at the right times consistent with organizational policy.

Identity proofing | The act of authentication. Confirming the identity of a user or host.

IDS insertion | An attack that exploits the nature of a network-focused IDS to collect and analyze every packet to trick the IDS into thinking an attack took place when it actually has not. The common purpose of IDS injection attacks is to trick signature- or pattern-matching detection of malicious network events.

In-band management | A common method of data collection on a network.

Incident response planning | A predefined procedure to react to security breaches to limit damage, contain the spread of malicious content, stop compromise of information, and promptly restore the environment to a normal state.

Information Technology Infrastructure Library (ITIL) | A set of concepts and practices that provide detailed descriptions and comprehensive checklists, tasks and procedures for common IT practices. The Security Management section is based on the ISO 27002 standard.

Ingress filtering | Filtering traffic as it attempts to enter a network. This can include monitoring for spoofed addresses, malformed packets, unauthorized ports and protocols, and blocked destinations.

In-line proxy | A device usually installed between the firewall and the web server to filter requests.

Insertion attack | An exploit-based on the introduction of unauthorized content or devices to an otherwise secured

infrastructure. Three common insertion-based attacks include SQL injection, IDS insertion, and rogue devices.

Instant messaging (IM) | A form of near real-time text communication. Also known as chat, IRC, and SMS messaging.

Intangible cost and value | Costs or values not directly related to budgetary funds. They can include but are not limited to research and development, marketing edge, competition value, first to market, intellectual property, public opinion, quality of service, name recognition, repeat customers, loyalty, honesty, dependability, assurance, reliability, trademarks, patents, privacy, and so on.

Integrated Services Digital Network (ISDN) | A set of communications standards for simultaneous digital transmission of voice, video, data, and other network services over the traditional circuits of the public switched telephone network.

Integrity | The security service of preventing unauthorized changes to data.

Intentional electromagnetic interference (IEMI) | The result of an intentional discharge made to damage or destroy electronic equipment, ranging from cell phones to computers and servers.

Interception attack | Any attack that positions the attacker in line with a session between a client and server. Such attacks typically allow the hacker to eavesdrop and manipulate the contents of the session. Also known as a man-in-the-middle attack.

Intermediary network | Any network, network link, or channel located between the endpoints of a VPN; often the Internet.

Internal personnel | Any worker or person who is physically present within the building or who has authorization to remotely connect into the network. Internal personnel are the most common cause of security violations.

Internet Assigned Numbers Authority (IANA) | The entity responsible for global coordination of IP addressing, DNS root, and other Internet protocol resources.

Internet Control Message Protocol (ICMP) | A commonly used protocol found in the Network Layer (Layer 3). ICMP rides as the

payload of an IP packet. ICMP supports network health and testing. Commonly abused by hackers for flooding and probing attacks.

Internet Engineering Task Force (IETF) | The standards body for Internet-related engineering specifications.

Internet Key Exchange v2 (IKEv2) | The second and latest version of the IKE protocol.

Internet of Things (IoT) | Devices that can be connected and accessed via the Internet.

Internet Protocol (IP) | A protocol standard for communicating over the Internet.

Internet Protocol Security (IPSec) | A protocol designed to secure communication.

Internet Protocol version 4 (IPv4) | A protocol used for identifying unique addresses for computers and networks using dot delineated format. Because of the limit of potential addresses, this protocol is superseded by IPv6.

Internet Protocol version 6 (IPv6) | An addressing protocol not bound by the mathematical limit of IPv4.

Internet Relay Chat (IRC) | A real-time text communication system. Hackers commonly use IRC as a way to communicate anonymously and control botnets.

Internet Service Provider (ISP) | The vendor that provides Internet access to a home or business.

Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) | A legacy protocol developed by Novell for their NetWare networking product. Mostly replaced by TCP/IP.

Intrusion detection system (IDS) | A security mechanism to detect unauthorized user activities, attacks, and network compromise. An IDS can respond in a passive manner through alerts and logging or in an active manner by disconnecting an offending session.

Intrusion prevention system (IPS) | A security mechanism to detect and prevent attempts to breach security.

IP address | The temporary logical address assigned to hosts on a network. An IP address is managed and controlled at the Network

Layer (Layer 3) of the OSI model by IP (Internet Protocol). IPv4 addresses are 32-bit addresses presented in human-friendly dotted-decimal notation. IPv6 addresses are 128-bit address presented in a special hexadecimal grouping format.

IP Multimedia Subsystem (IMS) | An architectural framework for delivering IP multimedia services; IMS would carry packet communications in all known forms over wireless or landline, everything from traditional telephony to video on demand (VoD).

IPchains | The firewall capability built into most Linux distributions.

IPSec | IP protocol encryption services extracted from IPv6 to be used as an add-on component for IPv4. IPSec provides tunnel mode and transport mode encrypted Network Layer connections between hosts and/or networks.

ISO image | A drive image that can be used to install software or store files.

K

Kernel panic | A known issue in FreeBSD where the system freezes and stops processing instructions. This is similar to the Microsoft “blue screen.”

Key | The unique number used to guide an algorithm in the encryption and decryption process. A valid key must be within the keyspace of an algorithm.

Key exchange | The cryptographic function ensuring that both endpoints of a commutation have the same symmetric key. Key exchange occurs by simultaneous key generation or with a digital envelope.

Key pair | The set of associated keys including a public key and a private key used by public key cryptography. Only the public key can decrypt data encrypted by the private key, and vice versa.

Keyspace | The range of valid keys used by an algorithm. Keyspace is the bit length of the keys supported by the algorithm.

Keystroke logger | Malware that records all keyboard input and transmits the keystroke log to a hacker.

Knowledge-based detection | A form of IDS/IPS detection based on a collection of samples, patterns, signatures, and so on stored in a database of known malicious traffic and events. All traffic or events that match an item in the database is considered abnormal and potentially malicious. Also known as signature, database, and pattern-matching-based detection.

L

LAN-to-LAN VPN | A VPN between two networks over an intermediary network; also known as WAN VPN and site-to-site VPN.

Layer 2 Forwarding (L2F) Protocol | An early communications protocol that competed with Point-to-Point Tunneling Protocol.

Layer 2 Tunneling Protocol (L2TP) | An older protocol largely replaced by IPsec and SSL/TLS-based VPNs in production environments, but still in use in some older environments.

Leased line | A network communications line leased from an ISP or telco service. A leased line is usually a dedicated line between network locations or to the Internet.

Leetspeak | A somewhat secret form of communication or language that hackers use based on replacing letters with numbers, symbols, or other letters that somewhat resemble the original characters. For example, “elite” becomes “eleet” and then becomes “31337.”

Load balancer | A system or device (hardware or software) that takes the load coming into a set of servers and ensures that the load is balanced between or among the servers.

Load balancing | A network traffic management technique to spread the workload or traffic levels across multiple devices to maintain availability, uptime, and high performance at wire speed.

Local area network (LAN) | A network confined to a limited geographic distance. Generally, a LAN is comprised of segments that are fully owned and controlled by the host organization as opposed to using lines leased from telcos.

Log | A log is a recording or notation of activities. Many security services, applications, and network resources automatically create a log of all events. Also known as an event log or a log file.

Logging | The act of creating or recording events into a log. Similar to auditing and monitoring.

Logic bomb | Malware that acts like an electronic land mine. Once a hacker places a logic bomb in a system, it remains dormant until a triggering event takes place. The trigger can be a specific time and

date, the launching of a program, the typing of a specific keyword, or accessing a specific URL. Once the trigger occurs, the logic bomb springs its malicious event on the unsuspecting user.

Logical address | A temporarily assigned address given to a host. IP address is a common example of a logical address. Most logical addresses exist at the Network Layer (Layer 3) of the OSI model.

Logical topology | The way computers appear to behave versus how those computers are actually connected.

Loophole | An unexpected way around a defense or a policy.

M

MAC spoofing | The act of a hacker changing the MAC address of their network interface. Commonly used to bypass MAC filtering on a wireless access point by impersonating a valid client.

Malicious code | Any software that was written with malicious intent. Administrators use antivirus and anti-malware scanners to detect and prevent malicious code (also known as malware) from causing harm within a private network or computer.

Management interface | The command line or graphical interface used to control and configure a device. Often accessible through a console (CON) port on the device or through a logical interface across the network.

Man-in-the-middle (MitM) attack | This attack occurs when a hacker is positioned between a client and a server, and the client is fooled into connecting with the hacker computer instead of the real server. The attack performs a spoofing attack to trick the client. As a result, the connection between the client and server is proxied by the hacker. This allows the hacker to eavesdrop and manipulate the communications.

Maximum Transmission Unit (MTU) | The largest amount of data that a datagram can hold based on the limitations of the networking devices managing a given segment. As an MTU changes across a communication path, a datagram may be fragmented to comply with the MTU restriction.

Mean time between failures (MTBF) | A rating on some hardware devices expressing the average length of time between significant failures.

Mean time to failure (MTTF) | A rating on some hardware devices expressing the average length of time until the first significant failure is likely to happen.

Media Access Control (MAC) address | The physical address assigned to a network interface by the manufacturer. The MAC address is a 48-bit binary address presented in as hexadecimal pairs separated by colons. The first half of a MAC address is known as the

Organizationally Unique Identifier (OUI) or vendor ID; the last half is the unique serial number of the NIC.

Mesh topology | A network layout with increased redundancy and usually significantly increased set-up costs.

Metacharacter | A character that has a special meaning assigned to it and is recognized as part of a scripting or programming language. Metacharacters should be filtered, escaped, or blocked to prevent script injection attacks. Escaping metacharacters is a programmatic tactic to treat all characters as basic ASCII, rather than as something with special meaning or purpose.

Mission-critical | The state or condition of an asset or process vitally important to the long-term existence and stability of an organization. If a mission-critical element is interrupted or removed, it often results in the failure of the organization.

MITRE | The MITRE Corporation is a not-for-profit organization chartered to work in the public interest. It sponsors a vulnerability research, cataloging, and information organization:

[*http://cve.mitre.org/*](http://cve.mitre.org/).

Mobile code | A form of software transmitted to and executed on a client. Hackers can use mobile code for malicious purposes.

Mobile IP | A standard communications protocol designed to let mobile device users move from one network to another while maintaining a permanent IP address; this concept is also known as IP mobility. Mobile IP for IPv4 is described in RFC 5944. Mobile IPv6, designed to work with next generation of the Internet Protocol, is covered in RFC 6275.

Modeling | The process of simulating and testing a new concept, design, programming, technique, and so forth before deployment into a production environment. Modeling often occurs before piloting.

Modem | A device used to connect a computer or network to the Internet.

Monitoring | The act of watching for abnormal or unwanted circumstances. Commonly used interchangeably with logging and auditing.

Monkey-in-the-middle attack | Another term for man-in-the-middle attack.

Multifactor authentication | Authentication that requires multiple valid proofs of identity used in simultaneous combination.

N

National Information Infrastructure (NII) | The product of the High Performance Computing and Communication Act of 1991. It was a telecommunications policy buzzword, which was popularized during the Clinton administration under the leadership of Vice President Al Gore. It was a proposed advanced, seamless web of public and private communications networks, interactive services, interoperable hardware and software, computers, databases, and consumer electronics to put vast amounts of information at users' fingertips.

National Institute of Standards and Technology (NIST) | NIST is a nonregulatory federal agency within the U.S. Department of Commerce; its mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology. As part of its mission, the NIST performs vulnerability research, cataloging, and information distribution:

<http://nvd.nist.gov/>.

National Security Agency (NSA) | The National Security Agency/Central Security Service (NSA/CSS) is a cryptologic intelligence agency of the U.S. government, administered as part of the U.S. Department of Defense. It is responsible for the collection and analysis of foreign communications and foreign signals intelligence, which involves cryptanalysis. It is also responsible for protecting U.S. government communications and information systems from similar agencies elsewhere, which involves cryptography.

Native firewall | A firewall within an operating system or hardware device placed there by the vendor or manufacturer. Can also include firewalls not necessarily installed by default, but which you can add to a system through an update or patch installation.

Network access control (NAC) | A mechanism that limits access or admission to a network based on the security compliance of a host.

Network address translation (NAT) | A service that converts between internal addresses and external public addresses. This conversion is performed on packets as they enter or leave the network

to mask and modify the internal client's configuration. The primary purpose of NAT is to prevent internal IP and network configuration details from being discovered by external entities, such as hackers.

Network Layer (Layer 3) | The third layer of the OSI model. This layer is responsible for logical addressing (IP addresses) and routing traffic.

Network News Transfer Protocol (NNTP) | The protocol used by the USENET message service. USENET is a persistent message service that allows anyone to post and read messages from over 100,000 named, categorized, topical newsgroups.

Network security | The collection of security components assembled in a network to support secure internal and external communications. Network security depends upon host security. Network security operates to protect the network as a whole, rather than as individual systems.

New Technology File System (NTFS) | A file format developed by Microsoft commonly used on Windows systems. NTFS offers file security, large volume size, large file size, and alternate data streams (ADS).

Next-generation firewall (NGFW) | A device offering more than traditional firewall capabilities such as packet inspection.

Nmap | A network mapping tool that performs network scanning, port scanning, OS identification, and other types of network probing. Nmap is available at <http://www.insecure.org/>.

Node | Any device on the network that can act as the endpoint of a communication. This includes clients, servers, switches, routers, firewalls, and anything with a network interface that has a MAC address. A node is a component that can receive communication with, rather than one that communication only through or across. For example, network cables and patch panels are not nodes.

Nonauthenticating query service | Any communication exchange that does not verify the identity of the endpoints of a communication and accepts any properly formed response as valid. DNS and ARP are common examples. Hackers can easily spoof such a service.

Nondedicated connection | A network connection not always on and available for immediate transmission of data. A connection must be established through a negotiation process before the channel is open and ready for data transmission. Dial-up, ISDN, and DSL lines are nondedicated connections.

Nonrepudiation | A security service that ensures that a sender cannot deny sending a message. This service can be provided by public key cryptography, typically through a digital signature.

O

One-time pad | A form of cryptography in which each encryption key is used once before being discarded. Keys are pseudorandom and never repeat. Key length must match message length, so that each character is encrypted with a unique key character.

One-way function | A mathematical operation performed in one direction relatively easily; reversing the operation is impossible—or nearly so.

Open source | A type of software product that may or may not be precompiled and whose source code is freely disclosed and available for review and modification.

Open Systems Interconnection (OSI) Reference Model | A standard conceptual tool used to discuss protocols and their functions. The OSI model has seven layers. Each layer can communicate with its peer layer on the other end of a communication session. While the OSI model helps to discuss protocols, most protocols are not in full compliance with it.

Open VPN | A brand of VPN used in demonstration of the topics in this textbook.

Opportunistic hacker | A person who takes advantages of unique or abnormal situations to perform malicious actions, but who would not initiate such actions otherwise.

Optical carrier (OC) | A form of network carrier line, often leased or dedicated, which uses fiber optic cables for very high-speed connections. An OC-1 connection supports a throughput of 51.84 Mbps.

OS/2 | A multi-tasking operating system developed jointly by Microsoft and IBM. First released in 1987, it lost nearly its entire market share to Windows after the two companies ceased collaboration in 1990. IBM discontinued support in 2006.

Out-of-band communication | A method of communication through an alternative route, mechanism, or pathway than the current one employed (the current communication is known as “in band”).

Commonly used as a technique for secured data exchange or verification of an identity.

P

Packet | The collection of data at the Network Layer (Layer 3) of the OSI model. It consists of the payload from the Transport Layer (Layer 4) above and the Network Layer header. IP packets are a common example.

Padded cell | Specialized host used to place an attacker into a system where the intruder cannot do any harm.

Partition | A logical division of a hard drive that can be formatted with a file system.

Passive hub | *See hub*

Passive threat | Any harmful code or site that depends upon the user's actions to be accessed or activated. If users never visit an infected site or do not perform the risky activity, the threat never reaches them. A passive threat is similar to a virus in that it depends upon the activity of the user to activate, infect, and spread.

Patch management | The procedure of watching for the release of new updates from vendors, testing the patches, obtaining approval, then overseeing the deployment and implementation of updates across the production environment.

Payload | The nonheader component of a PDU/segment/packet/frame. The payload is the data received from the layer above that includes the above layer's header and its payload.

Payment Card Industry Data Security Standard (PCI DSS) | An industry standard for secure storage of credit card holder information and transaction-processing.

Penetration testing | *See ethical hacking*

Permission | An ability to interact with a resource that is granted or denied to a user through some method of authorization or access control, such as access control lists (ACLs).

Personal firewall | Typically, a software host firewall installed on a home computer or network client. Can also refer to SOHO hardware firewalls, such as those found on DSL and cable modems and wireless access points.

Phishing | An attack that seeks to obtain information from a victim by presenting false credentials or luring victims to an attack site.

Phishing can occur face to face, over the phone, via e-mail, on a website, or through IM.

Physical address | The hardware address assigned to a network interface by the manufacturer. Also known as the MAC address.

Physical Layer (Layer 1) | The bottom or first layer of the OSI model. This layer converts data into transmitted bits over the physical network medium.

Physical topology | The way computers and devices are connected in a network.

Piloting | Using a new service, device, configuration, software, and so on to a limited number of testing hosts before rolling out the new component to the entire production environment. Piloting often occurs after modeling. Also called beta testing.

Ping sweep | A network scan that sends ICMP type 8 echo requests to a range of IP addresses to obtain ICMP type 0 echo responses. A ping sweep can discover active systems and identify the IP addresses in use.

Playback attack | See replay attack.

Point-to-Point Protocol (PPP) | A protocol commonly used in establishing a direct connection between two networking nodes.

Point-to-Point Tunneling Protocol (PPTP) | An early proprietary protocol from Microsoft.

Pop-up blocker | A software tool that prevents or restricts websites from automatically opening additional tabs or windows without the user's consent. These additional windows are known as pop-ups or pop-unders. Pop-ups are commonly used as methods of advertising, as well as elements in social engineering and distribution of malicious code.

Port address translation (PAT) | An extension to network address translation (NAT) that permits multiple devices on a local area network (LAN) to be mapped to a single public IP address.

Port-based network access (admission) control (PNAC) | A form of network access control or admission control (NAC) used on

individual network access devices, such as firewalls, VPN gateways, and wireless routers, to offload authentication to a dedicated authentication server/service. Only after valid authentication are communications with or across the network device allowed.

Port forwarding | The function of routing traffic from an external source received on a specific predefined IP address and port combination (also known as a socket) to an internal resource server. Also known as reverse proxy and static NAT.

Port number | The addressing scheme used at the Transport Layer (Layer 4) of the OSI model. There are 65,535 ports, each of which can—in theory—support a single simultaneous communication.

Port scanning | A network scan that sends various constructions of TCP or UDP packets to determine the open or closed state of a port. Tools such as nmap are used to perform port scanning.

Post Office Protocol (POP) | An Application Layer protocol used by e-mail clients to receive messages from an e-mail server. The default TCP/IP port is 110, and it does not encrypt communications. The companion SMTP protocol sends messages to an e-mail server.

Presentation Layer (Layer 6) | The sixth layer of the OSI model translates the data received from host software into a format acceptable to the network. This layer also performs this task in reverse for data coming from the network to host software.

Prevention | The use of safeguards to thwart exploitation or compromise.

Principle of least privilege | The guideline that all users should be granted only the minimum level of access and permission required to perform their assigned job tasks and responsibilities.

Privacy | Keeping information about a network or system user from being disclosed to unauthorized entities. While typically focused on private information like a Social Security number, medical records, credit card number, or cellular phone number, privacy concerns extend to any data that represents personally identifiable information (also known as PII).

Private branch exchange (PBX) | A type of business telephone network. PBX systems allow for multiple phone extensions, voice

mailboxes, and conference calling. PBX systems require specialized equipment. PBX systems are largely being replaced by VOIP (Voice over IP) solutions.

Private IP address | The ranges of IP addresses defined in RFC 1918 for use in private networks that are not usable on the Internet.

Private key | The key of the public key cryptography key pair that is kept secret and used only by the intended entity. The private key decodes information encoded with its associated public key, encrypting information that can be decrypted only by its associated public key. This process validates the identity of the originator and creates a digital signature.

Privilege | An increased ability to interact with and modify the operating system and desktop environment granted or denied to a user through some method of authorization or access control, such as user rights on a Windows system.

Privilege escalation | The act of obtaining a higher level of privilege or access for a user account or a session. A tactic employed by hackers once they intrude into a network through the compromise of a normal user account.

Professional hacker | A criminal whose objective is to compromise IT infrastructures. Whether operating as individuals, offering mercenary hacking services, or functioning as members of a criminal ring, professional hackers focus time and energy on becoming effective cyber attackers. A professional hacker is someone who contracts out his or her hacking skills to others.

Proprietary OS | An operating system built exclusively to run on a bastion host device. Most appliance firewalls employ a proprietary operating system.

Protocol converter | *See gateway*

Protocol Data Unit (PDU) | The collection of data at the Session, Presentation, and Application layers (Layers 5–7) of the OSI model.

Proxy attack | *See man-in-the-middle attack.*

Proxy manipulation | An attack in which a hacker modifies the proxy settings on a client to redirect traffic to another system, such as

the hacker's own machine. The hacker may host a proxy server in addition to eavesdropping and manipulating the redirected traffic.

Proxy server | A network service that acts as a “middle man” between a client and server. A proxy can hide the identity of the client, filter content, perform NAT services, and cache content.

Pseudo random number generator (PRNG) | The mechanism of computer systems that produces partially random numbers using a complex algorithm and a seed value that is usually time-based. Computers are currently unable to produce true random numbers, and a PRNG approximates randomness.

Public IP address | Any address that is valid for use on the Internet. This excludes specially reserved addresses such as loopback (127.0.0.1–127.255.255.255), RFC 1918 addresses, and the Windows APIPA addresses (169.254.0.0–169.254.255.255). Organizations lease public addresses from an Internet Service Provider (ISP).

Public key | The key of the public key cryptography key pair that is shared with other entities with whom the holder of the private key wishes to correspond. The public key decodes messages encoded with its associated private key, originates messages that only the holder of the associate private key can decrypt, and creates digital envelopes.

Public key cryptography | A subset of asymmetric cryptography based on the use of key pair sets. Public key cryptography uses public and private keys to create digital envelopes and digital signatures.

Public key infrastructure (PKI) | A combination of several cryptographic components to create a real-world solution that provides secure communications, storage, and identification services. Commonly uses symmetric encryption, asymmetric/public key encryption, hashing, and digital certificates. In most cases, when PKI refers to authentication, digital certificates are used as credentials.

Public network | Any network that is accessible by entities from outside an organization. Most often, use of this term implies the Internet, but many other public networks exist.

Punch panel | A type of hub.

Punchdown block | A type of hub.

Pwned | A “leetspeak” word derived from a common IRC typo of “owned.” Used to mean hacking and taking over control of a computer or network.

R

Ransomware | A type of malware that encrypts a device drive until a ransom is paid.

Rainbow tables | Precompiled lists of hashed passwords used in attacks.

Reconnaissance | The act of learning as much as possible about a target before attempting attacks. Reconnaissance consists of collecting data about the target from multiple sources online and offline. Effective reconnaissance is done covertly, without tipping off the target about the research. Reconnaissance can also be called footprinting, discovery, research, and information-gathering.

Recreational hacker | Someone who enjoys exploring and learning about computer technology; he or she may put an organization's network at risk by bringing in unapproved software, experimenting on the network, or just trying an exploit to "see if it works."

Redundancy | The feature of network design that ensures the existence of multiple pathways of communication. The purpose is to prevent or avoid single points of failure.

Redundant array of independent disks (RAID) | A disk set management technology that gains speed and fault tolerance. RAID can provide some protection against hard drive failure, but it does not protect against software or data compromises, such as virus infection.

Regional Internet Registry (RIR) | The five regional organizations that oversee and monitor the allocation and registration of IP addresses (both IPv4 and IPv6). RIR consists of American Registry for Internet Numbers (ARIN), RIPE Network Coordination Center (RIPE NCC), Asia-Pacific Network Information Centre (APNIC), Latin American and Caribbean Internet Address Registry (LACNIC), and African Network Information Centre (AfriNIC).

Rekeying | The process of triggering the generation of a new symmetric encryption key and secure exchange of that key. Rekeying can take place based on time, idleness, volume, randomness, or election.

Remote access | A communications link that enables access to network resources using a wide area network (WAN) link to connect to a geographically distant network. In effect, remote access creates a local network link for a system not physically local to the network. Over a remote access connection, a client system can technically perform all the same tasks as a locally connected client, with the only difference being the speed or the bandwidth of the connection.

Remote access server (RAS) | A network server that accepts inbound connections from remote clients. Also known as a network access server (NAS).

Remote access VPN | Another name for host-to-site VPN.

Remote control | The ability to use a local computer system to remotely take control of another computer over a network connection. Often used for remote technical assistance.

Remote Desktop Connector (RDC) | A built-in application in Remote Desktop Protocol (RDP).

Remote Desktop Protocol (RDP) | A proprietary Microsoft protocol that provides a GUI.

Remote-to-home VPN | A VPN used to connect a remote or mobile host into a home computer or network. Also known as a host-to-host VPN.

Remote-to-office VPN | A VPN used to connect a remote or mobile host into office network workstation. Also known as a host-to-host VPN.

Repeater | A type of hub that rebuilds the signal received through a bit-by-bit process.

Replay attack | This attack occurs when a hacker uses a network sniffer to capture network traffic and then retransmits that traffic back on to the network at a later time. Replay attacks often focus on authentication traffic in the hope that retransmitting the same packets that allowed the real user to log into a system will grant the hacker the same access.

Request for comments (RFC) | A document that defines or describes computer and networking technologies. These documents are published by the Internet Engineering Task Force, the standards body

for Internet engineering specifications. RFCs exist for hardware, operating systems, protocols, security services, and much more.

Resources | Any data item or service available on a computer or network accessible by a user to perform a task.

Response | Action taken in the event of an attack or unexpected action on a system or network.

Return on investment (ROI) | A business evaluation technique to determine whether an investment will earn back equivalent or greater benefit within a specific time.

Reverse caching | A means of providing faster access to static content for external users accessing internal web servers.

Reverse proxy | The function of routing traffic from an external source received on a specific predefined IP address and port combination (also known as a socket) to an internal resource server. Also known as port forwarding and static network address translation (NAT).

RFC 1918 addresses | IP addresses that, by convention, are not routed outside a private or closed network.

Class A: 10.0.0.0–10.255.255.255;

Class B: 172.16.0.0–172.31.255.255;

Class C: 192.168.0.0–192.168.255.255

Ring topology | A topology that uses a closed cabling scheme.

Risk | The likelihood or potential for a threat to take advantage of a vulnerability and cause harm or loss. Risk is a combination of an asset's value, exposure level, and rate of occurrence of the threat. A goal of security is to recognize, understand, and eliminate risk.

Risk assessment | Risk assessment is the process of examining values, threat levels, likelihoods, and total cost of compromise versus the value of the resource and the cost of the protection. This involves the use of values and calculations, such as AV, EF, SLE, ARO, ALE, and the cost/benefit equation.

Risk management | Performing risk assessment, and then acting on the results to reduce or mitigate risk. Often risk assessment establishes a new security policy and then aids in revising it over time.

Rogue access point | An access point set up and configured by a hacker to fool users into connecting with it. The hacker may then use the connection to carry out an attack such as a man-in-the-middle attack.

Rogue device insertion | Physical attack where an imposter device is inserted into a network and connected to the network either wired or wirelessly.

Rogue DHCP | A false DHCP server providing IP addresses to networked devices.

Role | A collection of tasks and responsibilities defined by a security policy or job description for an individual essential productivity or security position.

Root cause | The actual cause of a problem, often buried behind several symptoms that appear to be causes.

Rootkit | A form of malware that hackers can upload and deploy on a target system. It often replaces multiple components of the host operating system with altered code.

Round-robin | A form of load balancing that hands out tasks in a repeating nonpriority sequence.

Route | The pathway communication travels within a wired network.

Router | A network device responsible for directing traffic toward its stated destination along the best-known current available path.

Router access control list | A tool used to limit network access and traffic while increasing network performance.

RRDool | A round-robin database tool intended to handle time-series data like network bandwidth, temperatures, CPU load, and so on. The data are stored in a round-robin database (circular buffer); thus, the system storage footprint remains constant over time.

Rule | A written expression of an item of concern (protocol, port, service, application, user, IP address) and one or more actions to take when the item of concern appears in traffic. Also known as a filter or ACL.

Rule set | The list of rules on a firewall (or router or switch) that determine which traffic is and is not allowed to cross the filtering device. Most rule sets employ a first-match-apply-action process.

S

Sacrificial host | A firewall positioned at the initial entry point where a network interfaces with the Internet serving as the first line of defense for the network. Also known as a bastion host.

Sarbanes-Oxley (SOX) Act | A federal requirement enacted in 2002 to protect investors by requiring publically traded companies to validate controls and financial data.

Scalability | The ability of a product or service to provide adequate performance across changes in size, load, scope, or volume.

Scaling | Changing the size of a network by making it larger or smaller.

Scanning | The act of probing a network using custom crafted packets. Scanning can determine the IP addresses in use and whether ports are open or closed. The tool nmap can be used to perform scanning.

Screening router | A router that can perform basic static packet-filtering services, in addition to routing functions. A screening router is the predecessor of modern firewalls.

Script kiddie | A new, inexperienced, or ignorant hacker who uses prebuilt attack tools and scripts instead of writing his or her own or customizing existing ones. Even though a derogatory term in the hacker community, “script kiddie” still describes a serious threat to network security.

Sector | A subdivision of computer storage medium that represents a fixed size of user-accessible data. Magnetic disks typically have 512-byte sectors; optical disks have 2,048-byte sectors. When a device is formatted, sectors are grouped into clusters.

Secure Shell (SSH) | A network protocol that allows data exchange using a secure channel between two networked devices. It is used primarily on GNU/Linux- and UNIX-based systems to access shell accounts. SSH was a replacement for Telnet and other insecure remote shells, which send information, notably passwords, in plaintext, rendering them susceptible to packet analysis. The

encryption used by SSH provides confidentiality and integrity of data over an insecure network, such as the Internet.

Secure Sockets Layer (SSL) | A security protocol that operates at the top of the Transport Layer (Layer 4) and resides as the payload of a TCP session. Netscape designed SSL in 1997 for secure web e-commerce, but it can encrypt any traffic above the Transport Layer. It uses public key certificates to identify the endpoints of session and uses symmetric encryption to protect transferred data. SSL v3.0 is the last version of SSL; TLS is replacing SSL.

Secured VPN | A VPN that uses encryption to protect the confidentiality of its transmissions.

Security assessment | The judging, testing, and evaluation of a deployed security solution.

Security awareness | Knowledge of industry-specific vulnerabilities, known threats, and ways to avoid lapses in security.

Security goal | Desired outcome for application of security practices and inclusion of tools and techniques.

Security Information and Event Management (SIEM) | A toolset for collecting relevant information from multiple systems and aggregating it for decision making.

Security objective | Sets of stated purposes or targets for network security activity. Standard objectives are confidentiality, integrity, and availability. Objectives are generally more oriented toward achieving or maintaining the goals, such as ensuring the confidentiality of resources.

Security policy | A written document prescribing security goals, missions, objectives, standards, procedures, and implementations for a given organization. Also identifies which assets need protection based on their value.

Security stance | An organization's filtering configuration; in essence, its answer to the question, "What should be allowed and what should be blocked?"

Security Technical Implementation Guides (STIG) | A security guideline, procedure, or recommendation manual.

Security through obscurity | A form of security based on hiding details of a system, or creating convolutions that are difficult to understand. Such strategies do not usually resist a persistent attack; they are used when true security is poorly understood or the perceived threat is insufficient to overcome the obscure methodology. For example, proprietary source encryption algorithms can be labeled security through obscurity, as no forum for peer review or for formal testing exists to examine whether the methodology is cryptographically sound.

Segment | The collection of data at the Transport Layer (Layer 4) of the OSI model. It consists of the payload from the Session Layer (Layer 5) above and the Transport Layer header. TCP segments are a common example. (Note: UDP segments are called datagrams, as they are connectionless, rather than connection-oriented).

Senior management | The individual or group of highest controlling and responsible authority within an organization. Ultimately, the success or failure of network security rests with senior management.

Separation of duties | An administrative rule whereby no single individual possesses sufficient rights to perform certain actions. Achieved by dividing administrative-level tasks and powers among compartmentalized administrators.

Server | A host on a network. A server is the computer system that hosts resources accessed by users from clients.

Service level agreement (SLA) | A contractual commitment by a service provider or support organization to its customers or users.

Session | A logical connection between a client and a resource server. May exist at Layer 3, 4, or 5 of the OSI model; also known as a circuit or a state.

Session hijacking | When a hacker is able to take over a connection after a client has authenticated with a server. To perform this attack, a hacker must eavesdrop on the session to learn details, such as the addresses of the session endpoints and the sequencing numbers. With this information, the hacker can desynchronize the client, take on the client's addresses, and then inject crafted packets into the data stream.

If the server accepts the initial false packets as valid, then the session has been hijacked.

Session Layer (Layer 5) | The fifth layer of the OSI model. This layer manages the communication channel, known as a session, between the endpoints of the network communication. A single transport layer connection between two systems can support multiple simultaneous sessions.

Shell code | The content of an exploit to be executed on or against a target system.

Signature | The unique characteristics of malware (known as a virus signature) to detect and prevent infection of computer components.

Signature-based detection | A form of IDS/IPS detection based on a collection of samples, patterns, signatures, and so on stored in a database of known malicious traffic and events. All traffic or events that match an item in the database are considered abnormal and potentially malicious. Also known as database, knowledge, and pattern-matching-based detection.

Simple Mail Transfer Protocol (SMTP) | An Application Layer protocol used by e-mail clients to send messages to an e-mail server and is also used to relay messages between e-mail servers. The default TCP/IP port is 25, and it does not encrypt communications. The companion POP protocol receives messages from an e-mail server.

Single-factor authentication | The use of only a single element of validation or verification to prove the identity of a subject.

Considered much weaker than multifactor authentication.

Single loss expectancy (SLE) | The calculation of the loss potential across of a single incident for a given asset and a specific threat. SLE calculations are part of risk assessment. $SLE = AV \times EF$.

Single point of failure | Any element of a system or network infrastructure, which is the primary or only pathway through which a process occurs. The compromise of such an element could result in system failure. Network design should avoid single points of failure by including redundancy and defense in depth.

Single sign-on (SSO) | A network security service that allows a user to authenticate to an entire domain through a single client logon process. All domain members will accept this single authentication. Local authorization is used to control access to individual resources. Such a single authentication can be more complex, since multiple logons for each individual server are not required.

Site-to-site VPN | A VPN used to connect networks. Also known as a LAN-to-LAN VPN and WAN VPN.

Slack space | The unused portion of the last cluster allocated to a stored file. It may contain remnants of prior files stored in that location. Hackers can hijack slack space to create hidden storage compartments.

Slideware | An industry term referring to any product that appears in a vendor's PowerPoint slide deck, but is not yet available in one of its products; also sometimes known as "vaporware."

Smart hub | See active hub

Sniffer | A software utility or hardware device that captures network communications for investigation and analysis. Also known as packet analyzer, network analyzer, and protocol analyzer.

Social engineering | The craft of manipulating people into performing tasks or releasing information that violates security. Social engineering relies on telling convincing lies to manipulate people or taking advantage of the victim's desire to be helpful.

Socket | The combination of an IP address and a port number as a complete address.

Software firewall | A host firewall installed on a client or server.

Software VPN | A VPN crafted by software rather than hardware. Software VPN may be a feature of the operating system or a third-party application.

SOHO (small office, home office) | Any small network, workgroup, or client/server, deployed by a small business, a home-based business, or just a family network in a home.

Solid-state drive (SDD) | A drive that eliminates many of the problems associated with magnetic drives by moving and restoring the data on the drive.

Source routing bridge | A type of bridge where routing occurs based on the use of a discovery frame that travels the network to determine possible paths to the destination.

Spam | Unwanted and often unsolicited messages. Spam is not technically malicious software, but spam can have a serious negative effect on IT infrastructures through sheer volume. Estimates vary, but spam may represent up to 95 percent of all e-mail (which implies that for every legitimate e-mail, there are up to 19 unrelated spam e-mails.)

Split tunnel | A VPN connection that allows simultaneous access to the secured VPN link and unsecured access to the Internet across the same connection.

Spoofing | The falsification of information. Often spoofing is the attempt to hide the true identity of a user or the true origin of a communication.

Spyware | An advancement of keystroke logging to monitor and record many other user activities. Spyware varies greatly, but it can collect a list of applications launched, URLs visited, e-mail sent and received, chats sent and received, and names of all files opened. It can also record network activity, gather periodic screen captures, and even recordings from a microphone or web cam. Can be linked with adware.

SQL injection | A form of website/application attack in which a hacker submits SQL expressions to cause authentication bypass, extraction of data, planting of information, or access to a command shell.

Star topology | A network layout with an independent cable to each device.

State | A logical connection between a client and a resource server. May exist at Layer 3, 4, or 5 of the OSI model. Also known as a session or a circuit.

Stateful inspection | The process of automatically tracking sessions or states to allow inbound responses to previous outbound requests. Also called dynamic packet filtering.

Static electricity discharge (SED) | A sudden and momentary electric current, usually of high voltage and low amperage, that flows between two objects. Commonly caused by low humidity environments. Humans, polyester, and plastics are prone to static buildup. SED can damage most computer components.

Static NAT | The static coding of a translation pathway across a NAT service. Also known as port forwarding and reverse proxy.

Static packet filtering | A method of filtering using a static or fixed set of rules to filter network traffic. The rules can focus on source or destination IP address, source or destination port number, IP header protocol field value, ICMP types, fragmentation flags, and IP options. Static packet filtering is therefore mainly focused on the Network Layer (Layer 3), but it can also include Transport Layer (Layer 4) elements. Static packet filtering focuses on header contents and does not examine the payload of packets or segments.

Subnetting | A logical networking technique that breaks IP addresses into smaller subsets, thus extending the available IP address pool or allowing for network segmentation.

Sunk cost | Time, money, and effort already spent on a project, event, or device. In economics, sunk costs are irrelevant to future decisions. Emotionally, however, people often use sunk costs as a rationalization to continue failing processes or procedures.

Switch | A device that provides network segmentation through hardware. Across a switch, temporary dedicated electronic communication pathways are created between the endpoints of a session (such as a client and server). This switched pathway prevents collisions. Additionally, switches allow the communication to use the full potential throughput capacity of the network connection, instead of 40 percent or more being wasted by collisions (as occurs with hubs).

Symmetric cryptography | Cryptography based on algorithms that use a single shared secret key. The same key encrypts and decrypts data, and the same key must be shared with all communication partners of the same session.

Synchronous dynamic random access memory (SDRAM) |

Dynamic random access memory (DRAM) that has a synchronous interface. Traditionally, dynamic random access memory (DRAM) has an asynchronous interface, which means that it responds as quickly as possible to changes in control inputs. SDRAM has a synchronous interface, meaning that it waits for a clock signal before responding to control inputs and is therefore synchronized with the computer's system bus.

Syslog | A standard for forwarding messages from a client device to a centralized server.

Systems Network Architecture (SNA) | A legacy networking protocol developed by IBM commonly used to support communications between mainframes. Mostly replaced by TCP/IP.

T

Tangible costs and values | Costs or values directly related to budgetary funds. They can include, but are not limited to: purchase, license, maintenance, management, administration, support, utilities, training, troubleshooting, hardware, software, updates/upgrades, and so forth.

Telco | Short for telecommunications company or corporation. Used to refer to any company that sells or leases WAN connection services, whether wired or wireless.

Telecommuting | The act of working from a home, remote, or mobile location while connecting into the employer's private network, often using a VPN.

Telnet | A protocol and a service used to remotely control or administer a host through a plaintext command line interface.

Terminal services | A modern form of legacy thin-client operation. A thin-client software utility connects to a central terminal server, which simulates remote control. A terminal service system can support multiple simultaneous terminal client connections. When terminal services are in use, the client workstation converts to thin-client status. All operations of storage and processing then take place on the terminal server.

Terminator | A hardware device used on a network with bus topology to stop signals from bouncing.

Thin-client computing | A legacy terminal concept used to control mainframes. Thin clients had no local processing or storage capability. Modern thin clients simulate these limitations and perform all operations on the terminal server, remote control server, or thin-client server.

Threat | Any potential harm to a resource or node on the network. Threats can be natural or artificial, caused by mother nature or man, or by the result of ignorance or malicious intent. Threats originate internally and externally.

Time synchronization | A benefit of SIEM that manages the disparity between system clocks on.

Token | In a ring network, the token controls traffic on the network cable by determining which systems can send information onto the line.

Topology | The arrangement and means of connection of computers and peripherals in a network.

Traceroute | A computer network tool used to show the route taken by packets across an IP network. An IPv6 variant, traceroute6, is also widely available.

Traffic congestion | The problem when too much data crosses a network segment. This results in reduced throughput, increased latency, and lost data.

Training | The second level of knowledge distribution offered by an organization to educate users about job task–focused security concerns. Training is more rigorous than awareness, but less rigorous than education.

Transmission Control Protocol (TCP) | The connection-oriented protocol operating at the Transport Layer (Layer 4) of the OSI model.

Transmission Control Protocol (TCP)/Internet Protocol (IP) | The suite commonly used for network and Internet traffic that uses the best of both protocols.

Transparent bridge | A bridge where LANs are connected and each device is unaware that the bridge is on the network.

Transparent proxy | A proxy most often installed immediately between the firewall and a web server.

Transport Layer (Layer 4) | The fourth layer of the OSI model. This layer formats and handles data transportation. This transportation is independent of and transparent to the application.

Transport Layer Security (TLS) | A security protocol that operates at the top of the Transport Layer (Layer 4) and resides as the payload of a TCP session. It uses public key certificates to identify the endpoints of session and uses symmetric encryption to protect transferred data. TLS 1.0 is the replacement for SSL 3.0.

Transport mode encryption | A form of encryption also known as point-to-point or host-to-host encryption. Transport mode encryption

protects only the payload of traffic and leaves the header in plain-text original form.

Trapdoor | A form of unauthorized access to a system. A trapdoor is any access method or pathway that circumvents access or authentication mechanisms. Also known as a backdoor.

Triple-homed firewall | A firewall that has three network interfaces. Each network interface is located in a unique network segment. This allows for true isolation of the segments and forces the firewall to filter all traffic traversing from one segment to another.

Trojan horse | A mechanism of distribution or delivery more than a specific type of malware. The Trojan horse embeds a malicious payload within a seemingly benign carrier or host program. When the host program is executed or otherwise accessed, the malware is delivered. The gimmick of a Trojan horse is the act of fooling someone (a type of social engineering attack) into accepting the Trojan program as safe.

Trust | Confidence in the expectation that others will act in your best interest, or that a resource is authentic. On computer networks, trust is the confidence that other users will act in accordance with the organization's security rules and not attempt to violate stability, privacy, or integrity of the network and its resources.

Trusted Platform Module (TPM) | A dedicated microchip found on some motherboards that host and protect the encryption key for whole hard drive encryption.

Trusted third party | A mechanism of authentication using a third entity known and trusted by two parties. The trusted third party allows the two communicating parties, who were originally strangers to each other, to establish an initial level of inferred trust.

Trusted VPN | A VPN whose components are wholly owned by the organization it serves.

Tunnel-mode encryption | A form of encryption also known as site-to-site, LAN-to-LAN, gateway-to-gateway, host-to-LAN, and remote access encryption. Tunnel-mode encryption performs a complete encapsulation of the original traffic into a new tunneling protocol.

The entire original header and payload are encrypted and a temporary link or tunnel header guides the data across the intermediary network.

Tunneling | The act of transmitting a protocol across an intermediary network by encapsulating it in another protocol. See encapsulation.

Two-factor authentication | A method of proving identity using two different authentication factors. Authentication factors are something you know, something you have, or something you are. Examples include a smart card (something you have) with a PIN (something you know); a biometric device (something you are) coupled with a password (something you know); or a proximity card (something you have) that activates a fingerprint reader (something you are).

U

Unidirectional | In a ring network the token only moves one way around the ringed cable.

Unified threat management (UTM) | The deployment of a firewall as an all-encompassing primary gateway security solution. The idea behind UTM is a single device can be designed to perform firewall filtering; IPS, antivirus scanning; anti-spam filtering; VPN endpoint hosting; content filtering; load-balancing; detailed logging; and potentially other security services, performance enhancements, or extended capabilities.

Uninterrupted power supply (UPS) | A power conditioning device that provides battery support to allow for logical closing of computer programs and powering down of the system in the event of a power loss.

Universal participation | The principle that for an organization's security policy to be effective, everyone must be forced to work within it and follow its rules.

Unpartitioned space | The area on a storage device not contained within a partition. Unpartitioned space is not directly accessible by the OS.

Upstream filtering | The management of traffic by a firewall or other filtering device located one or more hops away (upstream) from a private network.

URL injector | Malware that replaces URLs in HTTP GET requests for alternative addresses. These injected URLs cause a different webpage to appear in the browser than the one requested by the user's request. These replaced webpages could be advertisement sites, generate traffic to falsify search engine optimization (SEO), or lead to fake or spoofed sites.

User Datagram Protocol (UDP) | The connectionless protocol operating at the Transport Layer (Layer 4) of the OSI model.

V

Virtual firewall | A variety of firewall and firewall-like concepts using software to perform firewall activities.

Virtual local area network (VLAN) | A logical network created from devices on different LANs that function on the same network.

Virtual private network (VPN) | A mechanism to establish a secure remote access connection across an intermediary network, often the Internet. This allows inexpensive insecure links to replace expensive security links. VPNs allow for cheap long-distance connections established over the Internet. Both endpoints need only a local Internet link. The Internet itself serves as a “free” long-distance carrier. VPNs employ encapsulation and tunneling protocols, such as IPSec.

Virtual Router Redundancy Protocol (VRRP) | An industry-standard protocol using two VPN hardware units to provide redundancy for the connection.

Virus | Malware that needs a host object to infect. Most viruses infect files, such as executables; device drivers; DLLs; system files; and sometimes even document, audio, video, and image files. Some viruses infect the boot sector of a storage device, including hard drives, floppies, optical discs, and USB drives. Viruses are spread through the actions of users, and spread file-to-file (compare to worms).

VPN appliance | A hardware VPN device.

VPN concentrator | A VPN that decrypts information coming from a remote user over the Internet, and it encrypts information sent back over the Internet to the remote user.

VPN fingerprinting | A technique used by an attacker to identify the vendor and, in some cases, the software version of a VPN server.

Vulnerability | A weakness or flaw in a host, node, or any other infrastructure component that a hacker can discover and exploit. Security management aims to discover and eliminate such vulnerabilities.

Vulnerability management | The technology and business processes used to identify, track, and mitigate known weaknesses on hosts within a computing environment.

Vulnerability scanning | A form of investigation that aims at checking whether or not a target system is subject to attack based on a database of tests, scripts, and simulated exploits.

W

WAN VPN | A VPN between two networks over an intermediary network. Also known as LAN-to-LAN VPN and site-to-site VPN.

Wardialing | A method of discovering active modems by dialing a range of phone numbers.

Wardriving | A method of discovering wireless networks by moving around a geographic area with a detection device.

Weakest link | A security stance based on a repeating process of locating the least secure element of an infrastructure, securing it, and then identifying a new weakest link and securing it.

Whitelist | A type of filtering concept where the network denies all activities except for those on the white list. Also known as an “allow list” or “permissions list.”

Windows Defender Firewall | The software firewall available in Microsoft Windows 10.

Whois | A tool used to view domain registration information. Whois is a command-line function of Linux and Unix, but it is also a tool on most domain registrar websites.

Whole hard drive encryption | The process of encrypting an entire hard drive rather than just individual files. In most cases, whole hard drive encryption provides better security against unauthorized access than file encryption, because it encrypts temporary directories and slack space.

Wide area network (WAN) | A network not limited by any geographic boundaries. A WAN network can span a few city blocks, reach across the globe, and even extend into outer space. A distinguishing characteristic of a WAN is its use of leased or external connections and links. Often, telcos own these external connections.

Wire speed | The maximum communication or transmission capability of a network segment. Often used to describe a network device’s ability to perform tasks on traffic, while being able to maintain overall network transmission speeds without introducing delay, lag, or latency.

Wired topology | A physical connection of computers and devices using wires to transmit communication traffic.

Wireless topology | A connection of computers and devices that uses wireless transmissions (through the air) to transmit communication traffic.

Workgroup | A form of networking where each computer is a peer. Peers are equal to each other in terms of how much power or controlling authority any one system has over the other members of the same workgroup. All workgroup members are on equal footing because they can manage their own local resources and users, but not those of any other workgroup member.

Worm | Malware that does not need a host object; instead, a worm is a self-sustaining program in its own right. Worms are designed around specific system flaws. The worm scans other systems for this flaw and exploits the flaw to gain access to another victim. Once hosted on another system, the worm seeks to spread itself by repeating the process. Worms can act as carriers to deposit other forms of malicious code as they multiply and spread across networked hosts.

Wrapper | A tool used to create Trojan horses by embedding malware inside a host file or program.

Write-once read-many (WORM) | A form of storage device that can be written to once, but once written cannot be electronically altered. Examples include DVD-R, WORM tapes, and WORM hard drives.

Z

Zero-day exploit | New and previous unknown attacks for which there are no current specific defenses. “Zero day” refers to the newness of an exploit, which may be known in the hacker community for days or weeks. When such an attack occurs for the first time, defenders are given zero days of notice (hence the name.) Such attacks usually exploit previously unidentified system flaws.

Zeroization | The process of purging a storage device by writing zeros to all addressable locations on the device. A zeroized device contains no data remnants that other users could potentially recover.

Zombie | A malicious software program distributed by a hacker to take over control of a victim’s computer. Also known as a bot or an agent. Zombies are commonly used to construct botnets (or zombie armies).

Zombie army | A network of zombie/bot/agent-compromised systems controlled by a hacker. The network consists of the bots, agents, or zombies that intercommunicate over the Internet. Another term for botnet.

Zombie array | *See* zombie army

Zone of risk | Any segment, subnet, network, or collection of networks that represent a certain level of risk. The higher the risk, the higher the security need to protect against that risk. The less risk of a zone, the lower security need because fewer threats exist, or existing threats are less harmful. The flip side of risk zones is zones of trust.

Zone of trust | Any segment, subnet, network, or collection of networks that represent a certain level of trust. Highly trusted zones require less security, while low trusted zones require more security. The flip side of trust zones is zones of risk.



© Digital_Art/Shutterstock

References

- Beaver, Kevin. *Hacking for Dummies*. 3rd ed. Indianapolis: Wiley Publishing, Inc., 2010.
- Bhaiji, Yusuf. *Network Security Technologies and Solutions*. Indianapolis: Cisco Press, 2008.
- Bott, Ed. "The malware numbers game: How many viruses are out there?" The Ed Bott Report, April 15, 2012. <http://www.zdnet.com/blog/bott/the-malware-numbers-game-how-many-viruses-are-out-there/4783> (accessed May 2, 2013).
- Brenner, Laurie. "How to tell a firewall breach." It Still Works. <https://itstillworks.com/tell-firewall-breach-12210516.html> (accessed February 16, 2020).
- Brewster, Thomas. "How hackers abused Tor to rob Blockchain, steal Bitcoin, target private email and get away with it." Forbes, February 24, 2015. <https://www.forbes.com/sites/thomasbrewster/2015/02/24/blockchain-and-darknet-hacks-lead-to-epic-bitcoin-losses/#747eec2335e5> (accessed March 22, 2020).
- Cheswick, William R., Steven M. Bellovin, and Aviel D. Rubin. *Firewalls and Internet Security: Repelling the Wily Hacker*. 2nd ed. Boston: Addison-Wesley Professional, 2003.
- Cisco Systems. "IKEv2 Packet Exchange and Protocol Level Debugging," 2013. http://www.cisco.com/en/US/tech/tk583/tk372/technologies_tech_note09186a0080bf2932.shtml (accessed June 17, 2013).
- Cisco Systems. *Internetworking Technologies Handbook, Virtual Private Networks*. <http://www.cisco.com/en/US/docs/internetworking/technology/handbook/VPN.html>

- (accessed July 15, 2010).
- Cole, Eric. *Hiding in Plain Sight: Steganography and the Art of Covert Communication*. Indianapolis: Wiley Publishing, Inc., 2003.
- _____. *Network Security Bible*. Indianapolis: Wiley Publishing, Inc., 2009.
- Comer, Douglas E. *Internetworking with TCP/IP, Vol 1*. 5th ed. Upper Saddle River, NJ: Prentice Hall, 2005.
- Davies, Joseph, and Elliot Lewis. *Deploying Virtual Private Networks with Microsoft Windows Server 2003 (Technical Reference)*. Redmond, WA: Microsoft Press, 2003.
- Davis, Michael, Sean Bodmer, and Aaron LeMasters. *Hacking Exposed Malware and Rootkits*. New York: McGraw-Hill Osborne Media, 2009.
- Defense Information Systems Agency (DISA). "Information Assurance Support Environment (IASE), Security Checklists."
<http://iase.disa.mil/stigs/checklist/index.html> (accessed July 15, 2010).
- Defense Information Systems Agency (DISA). "Information Assurance Support Environment (IASE), Security Technical Implementation Guides (STIGS)."
<http://iase.disa.mil/stigs/stig/index.html> (accessed July 15, 2010).
- DeLaet, Gert, and Gert Schauwers. *Network Security Fundamentals*. Indianapolis: Cisco Press, 2004.
- Douligeris, Christos, and Dimitrios N. Serpanos. *Network Security: Current Status and Future Directions*. Hoboken, NJ: Wiley-IEEE Press, 2007.
- Dr. K. *Hackers' Handbook 3.0 (Expanded, Revised and Updated): Includes WiFi, Identity Theft, Information Warfare and Web 2.0*. New York: Carlton Books, 2009.
- Erickson, Jon. *Hacking: The Art of Exploitation*. 2nd ed. San Francisco: No Starch Press, 2008.
- Feilner, Markus. *OpenVPN: Building and Integrating Virtual Private Networks*. Birmingham, England: Packt Publishing, 2006.
- Ford, Jerry Lee Jr. *Absolute Beginner's Guide to Personal Firewalls*. Indianapolis: Que, 2007.
- Frenkel, Sheila, et al. "Guide to Ipsec VPNs: Recommendations of the National Institute of Standards and Technology." Special publication 800-77. Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology, 2005.
- Gerwitz, David. "How the President's secret helicopter plans wound up in Iran." CNN. March 2009. **<http://ac360.blogs.cnn.com/2009/03/10/how-the-presidents-secret-helicopter-plans-wound-up-iran/>** (accessed July 15, 2010).
- Harris, Shon. *CISSP All-in-One Exam Guide*. 5th ed. New York: McGraw-Hill Osborne Media, 2010.
- _____, Allen Harper, Chris Eagle, and Jonathan Ness. *Gray Hat Hacking, Second Edition: The Ethical Hacker's Handbook*. New York: McGraw-Hill Osborne Media, 2007.
- Huang, Qiang, and Jazib Frahim. *SSL Remote Access VPNs*. Indianapolis: Cisco Press, 2008.
- IETF RFC Repository. **<http://www.ietf.org/rfc.html>** (accessed July 15, 2010).
- Internet Engineering Task Force (IETF). "Internet Key Exchange (IKEv2) Protocol." December 2005. **<http://tools.ietf.org/html/rfc4306>** (accessed July 15, 2010).
- InfoSecurity. "Advanced malware detection—Signatures vs. behavior analysis." April 11, 2017. **<https://www.infosecurity-magazine.com/opinions/malware-detection-signatures/>** (accessed June 10, 2020).

- Internet Society. "State of IPv6 deployment 2018." June 6, 2018.
<https://www.internetsociety.org/resources/2018/state-of-ipv6-deployment-2018/>
(accessed July 27, 2019).
- ipv6-test.com**. Monthly statistics. <https://ipv6-test.com/stats/> (accessed August 5, 2019).
- Irwin, Mike, Charlie Scott, and Paul Wolfe. *Virtual Private Networks*. 2nd ed. Sebastopol, CA: O'Reilly Media, 1998.
- Katz, Jonathan, and Yehuda Lindell. *Introduction to Modern Cryptography: Principles and Protocols*. Chapman & Hall/CRC, 2007.
- Komar, Brian, Ronald Beekelaar, and Joern Wettern. *Firewalls for Dummies*. 2nd ed. Indianapolis: Wiley Publishing, Inc., 2003.
- Kozierok, Charles. *The TCP/IP Guide: A Comprehensive, Illustrated Internet Protocols Reference*. San Francisco: No Starch Press, 2005.
- Krawetz, Neal. *Introduction to Network Security*. Boston, MA: Charles River Media, 2006.
- Lewis, Mark. *Comparing, Designing, and Deploying VPNs*. Indianapolis: Cisco Press, 2006.
- Linux Home Networking. "Configuring Linux VPNs."
http://www.linuxhomenetworking.com/wiki/index.php/Quick_HOWTO_:_Ch35_:_Configuring_Linux_VPNs (accessed July 15, 2010).
- Liu, Dale, Stephanie Miller, Mark Lucas, Abhishek Singh, and Jennifer Davis. *Firewall Policies and VPN Configurations*. Rockland, MA: Syngress, 2006.
- Long, Johnny, Jack Wiles, Scott Pinzon, and Kevin D. Mitnick. *No Tech Hacking: A Guide to Social Engineering, Dumpster Diving, and Shoulder Surfing*. Rockland, MA: Syngress, 2008.
- Lyon, Gordon Fyodor. *Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning*. Sunnyvale, CA: Nmap Project, 2009.
- Mairs, John. *VPNs: A Beginner's Guide*. New York: McGraw-Hill Osborne Media, 2001.
- Maiwald, Eric. *Network Security: A Beginner's Guide*. 2nd ed. New York: McGraw-Hill Osborne Media, 2003.
- McClure, Stuart, Joel Scambray, and George Kurtz. *Hacking Exposed: Network Security Secrets and Solutions*. 6th ed. New York: McGraw-Hill Osborne Media, 2009.
- Merkow, Mark S. *Virtual Private Networks for Dummies*. Indianapolis: Wiley Publishing, Inc., 1999.
- Microsoft Docs. "About Remote Desktop Services." May 31, 2018.
<https://docs.microsoft.com/en-us/windows/win32/termserv/about-terminal-services>
(accessed March 15, 2020).
- Microsoft Technet. "DirectAccess Technical Overview."
<http://technet.microsoft.com/library/dd637827.aspx> (accessed July 15, 2010).
- Mitnick, Kevin D., William L. Simon, and Steve Wozniak. *The Art of Deception: Controlling the Human Element of Security*. Indianapolis: Wiley Publishing, Inc., 2003.
- _____, and William L. Simon. *The Art of Intrusion: The Real Stories Behind the Exploits of Hackers, Intruders and Deceivers*. Indianapolis: Wiley Publishing, Inc., 2005.
- MITRE Corporation. "MITRE's Common Vulnerability Database (CVE)."
<http://cve.mitre.org/> (accessed July 15, 2010).
- National Institute of Standards and Technology (NIST), National Vulnerability Database.
<http://nvd.nist.gov/> (accessed July 15, 2010).
- Nichols, Randall K., and Panos C. Lekkas. *Wireless Security: Models, Threats, and Solutions*. New York: McGraw-Hill Osborne Media, 2001.

- <http://www.jacco2.dds.nl/networking/certutil.html> (accessed June 21, 2010).
- OpenVPN. "2x HOW TO." <https://openvpn.net/community-resources/how-to/> (accessed March 16, 2020).
- Peltier, Thomas R. *Information Security Risk Analysis*. 3rd ed. Auerbach Publications, 2010.
- Philipp, Aaron, David Cowen, and Chris Davis. *Hacking Exposed Computer Forensics*. 2nd ed. New York: McGraw-Hill Osborne Media, 2009.
- PopTop + MSCHAPv2 + Samba + Radius + Microsoft Active Directory + Fedora Howto. http://www.members.optushome.com.au/~wskwok/poptop_ads_howto_1.htm (accessed July 15, 2010).
- PrivateTunnel.com**. "OpenVPN." <https://www.privatetunnel.com/> (accessed March 15, 2020).
- Real Time Enterprises, Incorporated. "Create a IPSEC + Certificates MMC." <http://support.real-time.com/open-source/ipsec/index.html> (accessed July 15, 2010).
- Rhodes-Ousley, Mark, Roberta Bragg, and Keith Strassberg. *Network Security: The Complete Reference*. New York: McGraw-Hill Osborne Media, 2003.
- Ruston, Chris, and Chris Peiris. *How to Cheat at Designing Security for a Windows Server 2003 Network*. Rockland, MA: Syngress, 2006.
- SANS. "The Top Cyber Security Risks." <http://www.sans.org/top-cyber-security-risks/?ref=top20> (accessed July 15, 2010).
- Scarfone, Karn, and Paul Hoffman. *Guidelines on Firewalls and Firewall Policy: Special Publication 800-41, Revision 1*. Washington: National Institute of Standards and Technology (NIST), 2009.
- Schneier, Bruce. *Schneier's Cryptography Classics Library: Applied Cryptography, Secrets and Lies, and Practical Cryptography*. Indianapolis: Wiley Publishing, Inc., 2007.
- _____. *Secrets and Lies: Digital Security in a Networked World*. Indianapolis: Wiley Publishing, Inc., 2004.
- Schudel, Gregg, and David J. Smith. *Router Security Strategies: Securing IP Network Traffic Planes*. Indianapolis: Cisco Press, 2008.
- SearchNetworking.com**. "Authentication, authorization, and accounting (AAA)" (n.d.) <http://searchsecurity.techtarget.com/definition/authentication-authorization-and-accounting> (accessed June 17, 2013).
- _____. "Network packet" (n.d.) <http://searchnetworking.techtarget.com/definition/packet> (accessed June 17, 2013).
- _____. "Port Address Translation (PAT)" (n.d.) <http://searchnetworking.techtarget.com/definition/Port-Address-Translation-PAT> (accessed June 17, 2013).
- Security Focus. "Embassy leaks highlight pitfalls of Tor." September, 2007. <http://www.securityfocus.com/news/11486> (accessed July 15, 2010).
- Shinder, Thomas W. *The Best Damn Firewall Book Period*. 2nd ed. Rockland, MA: Syngress, 2008.
- SmoothWall Community Forums. <http://community.smoothwall.org/forum/> (accessed July 15, 2010).
- SmoothWall Limited. <http://www.smoothwall.net/> (accessed July 15, 2010).
- SmoothWall Open Source Project. <http://www.smoothwall.org/> (accessed July 15, 2010).
- Snader, Jon C. *VPNs Illustrated: Tunnels, VPNs, and IPsec*. Boston: Addison-Wesley Professional, 2005.

- Stevens, W. Richard, and Gary R. Wright. *TCP/IP Illustrated, Volumes 1–3*. Boston: Addison-Wesley Professional, 1994.
- Stewart, James Michael. *CompTIA Security+ Review Guide: SY0-201*. Indianapolis: Sybex, 2008.
- _____, Ed Tittel, and Mike Chapple. *CISSP: Certified Information Systems Security Professional Study Guide*. 4th ed. Indianapolis: Sybex, 2008.
- Strassberg, Keith, Gary Rollie, and Richard Gondek. *Firewalls: The Complete Reference*. New York: McGraw-Hill Osborne Media, 2002.
- Sweeney, T. “Businesses lock in on VPN outsourcing options providers of virtual private network services put a new spin on the outsourcing spiel.” *InformationWeek*, April 3, 2000. <http://www.informationweek.com/780/vpn.htm> (accessed July 15, 2010).
- Volonino, Linda, and Ian Redpath. *e-Discovery for Dummies*. Indianapolis: Wiley Publishing, Inc, 2009.
- Vyncke, Eric, and Christopher Paggen. *LAN Switch Security: What Hackers Know About Your Switches*. Indianapolis: Cisco Press, 2007.
- Whitman, Michael E., Herbert J. Mattord, Richard Austin, and Greg Holden. *Guide to Firewalls and Network Security*. Florence, KY: Course Technology/Cengage Learning Inc., 2008.
- Whittacker, Zack. “NordVPN confirms it was hacked.” *Verizon Media*, October 21, 2019. <https://techcrunch.com/2019/10/21/nordvpn-confirms-it-was-hacked/> (accessed March 3, 2019).
- Yuan, Ruixi, and W. Timothy Strayer. *Virtual Private Networks: Technologies and Solutions*. Boston: Addison-Wesley Professional, 2001.
- Zwicky, Elizabeth D., Simon Cooper, and D. Brent Chapman. *Building Internet Firewalls*. 2nd ed. Sebastopol, CA: O’Reilly Media, 2000.



© Digital_Art/Shutterstock

Index

Note: Page numbers followed by *f* and *t* indicate figures and tables, respectively

A

AAA. *See* **authentication, authorization, and accounting**

acceptable use policy (AUP), **9, 363**

access control, **7, 117, 279**. *See also* **authorization**

access control list (ACL), **227**

access management, **408**

accidents, **45–46**

accounting, **117**

ACL. *See* **access control list**

active hub, **69**

active sniffing, **124**

active threat, **150–151**

addressing, **104–106**

ADS. *See* **alternate data streams**

Advanced Encryption Standard (AES), **315, 408**

advanced persistent threat (APT), **43, 48–49**

adware, **47**

agents, **60, 142**

aggregation, **216**

AH. *See* **Authentication Header**

alerts, **210–213**

ALE. *See* **annualized loss expectancy**

algorithms, **268**

allow by default, **224**

alternate data streams (ADS), **56**

American Registry for Internet Numbers (ARIN), **85**

annualized loss expectancy (ALE), **125**

annualized rate of occurrence (ARO), **125**

anomaly-based detection, **215**

anonymity, **301**

anti-forensics, **394**

anti-malware scanners, **357, 386**

anti-SPAM filters, **150**

anti-spyware scanner, **150**

antivirus scanners, **150, 386**

anycast address, [106](#)
AppleTalk, [103](#)
appliance firewalls, [89](#), [148](#), [161](#), [163](#)
application conflicts, [298](#)
application firewall, [166](#), [420](#)
application gateway, [166](#)
Application Layer (Layer 7), [312](#)
application-level firewall, [182](#)
application proxy, [166](#)
application proxy firewalls, [169](#)
APT. *See* **advanced persistent threat**
arbitrary code execution, [54](#)
ARO. *See* **annualized rate of occurrence**
ARP flooding, [124](#)
ARP spoofing, [56–57](#)
assets, [7](#), [33](#), [375](#), [375f](#)
asset value (AV), [125](#)
asymmetric cryptography, [269–274](#), [273f](#)
attacking, [43](#)
attack surface, [113](#)
audit capabilities, [168](#)
auditing, [7](#), [113](#)
auditors, [13](#)
AUP. *See* **acceptable use policy**
authentication, [7](#), [115–117](#), [168](#), [175](#), [186](#), [258](#), [278–279](#), [288](#), [295](#),
[300](#), [347](#), [381](#), [408–409](#)
authentication, authorization, and accounting (AAA), [115–117](#), [202–](#)
[204](#), [203f](#), [279](#)
Authentication Header (AH), [86](#), [313](#), [323](#)
authentication server, [138](#)
authenticity, [271](#)
authorization, [7](#), [116](#), [117](#), [279–280](#)
availability, [7](#), [301–302](#), [416](#)
avalanche effect, [269](#)
AV. *See* **asset value**

B

- backdoor, [11](#), [47](#)
- backdoor account attack, [299](#)
- Back Orifice, [360](#), [362](#)
- Backtrack, [365](#)
- backups, [356](#), [387–388](#)
- bandwidth for VPN, [320](#)
- banner, [42](#)
- banner grabbing, [42](#), [42f](#)
- basic packet filtering, [359](#)
- bastion host, [143](#), [144f](#), [177](#), [179](#)
- bastion host OS, [207–208](#)
- baud rate, [58](#)
- behavioral-based detection, [215](#)
- best practices, [284–289](#)
- bidirectional, ring topology, [78](#)
- biometrics, [117](#)
- BIOS/firmware flashing, [114](#)
- 128-bit key, [268](#)
- BitTorrent, [340](#)
- blacklist, [74](#)
- “Blue Screen of Death,” [298](#)
- border-crossing communications, [108](#), [109](#)
- border firewall, [139](#), [140f](#), [357](#)
- border sentry, [134](#), [146](#)
- botnet, [60](#), [142](#)
- bots, [60](#), [142](#)
- bottlenecks, [13](#)
- boundary network, [93–94](#)
- breach, [10](#), [34](#)
- bricking, [400](#)
- bridge, [73–75](#), [74f](#)
- Bring Your Own Device (BYOD), [122](#), [287](#), [411](#)
- brouter, [75](#)
- brute force attack, [116](#)

brute-force password attack, [116](#)
buffer, [53](#)
buffer overflows, [53–54](#)
build-it-yourself firewall, [170](#)
bump-in-the-stack, [143](#), [144f](#)
bump-in-the-wire, [143](#), [144f](#)
business continuity plan, [120](#), [392](#)
business tasks, [11](#), [218–219](#)
bus topology, [77–78](#), [77f](#)
BYOD. *See* **Bring Your Own Device**
bypass deployment, [292](#)
bypass VPN implementation, [292f](#)

C

- cable modem devices, [152](#)
- California Consumer Privacy Act (CCPA), [176](#)
- CA. *See* [Certificate Authority](#)
- centralized logging system, [209](#)
- Certificate Authority (CA), [278](#), [338](#)
- CGI scripts. *See* [Common Gateway Interface scripts](#)
- change management, [398–399](#)
- channels, [255](#)
- chief information officers (CIOs), [409](#)
- chip creep, [45](#)
- chokepoint, [136](#), [199](#)
- ciphertext, [262](#), [268](#)
- circuit, [134](#)
- circuit firewall, [167](#)
- circuit proxy, [167](#)
- Cisco, [302](#)
- clients, [16](#), [89](#), [121–122](#), [316](#), [320](#), [321](#)
- clients and users, [175](#)
- client security, [286–287](#)
- client/server network, [13](#), [89](#), [90f](#), [315](#), [324](#)
- client-side configuration, [331](#)
- client-to-server VPNs, [267](#)
- client virtualization, [323–324](#). *See also* [desktop virtualization](#)
- closed source software, [149](#)
- cloud backup storage, [387](#)
- cloud-based firewall, [420–421](#)
- cloud computing, [410](#)
- cloud implementation, [285](#)
- clusters, [58](#), [123](#)
- code testing, [356](#)
- coding errors, [358](#)
- cold calling, [62](#)
- command shell, [43](#)
- commercial firewall, [148](#)

- commercial hardware firewall, [149](#)
- commercial software firewall, [149](#), [160–161](#)
- commercial VPNs, [300–301](#)
- Common Gateway Interface (CGI) scripts, [176](#)
- communication encryption, [118](#)
- communications pathways, [106–110](#), [118](#)
- communications to block, [234–236](#)
- compartmentalization, [384](#)
- compliance, [201](#), [417–419](#)
- compliance auditing, [399–400](#)
- compression, [262](#)
- concentrators, [69](#)
- conditional trust, [5](#)
- confidentiality, [7](#), [86](#), [416](#)
- configuration errors, [401](#)
- configuration of VPN, [321](#)
- configuration scans, [396](#)
- configuring firewall pfSense, [186](#)
- Connection Protocol, [319](#)
- containment, [393](#)
- content filtering, [147](#), [167](#)
- contract workers, [37](#)
- cookie filters, [150](#), [151f](#)
- corporate firewall VPN termination, [264](#)
- correlation, [216](#)
- cost/benefit analysis, [125](#)
- cost-effective network security, [4](#)
- COTS software. *See* [commercial off-the-shelf software](#)
- covert channels, [57–59](#), [108](#)
- CPE. *See* [customer premise equipment](#)
- credential sharing, [298–299](#)
- cross-site scripting (XSS), [52](#)
- Cryptcat, [362](#)
- cryptocurrency, [34](#)
- cryptography, [268](#), [275–277](#), [275f–277f](#)
- customer premise equipment (CPE), [329](#)

cybersecurity role, **407**

D

- data analytics, [409](#)
- data at rest, [408](#), [420](#)
- database-based detection, [215](#)
- database firewall, [420](#)
- data-centric security model, [409](#)
- data-encrypted tunnel, [328](#)
- data encryption. *See* [encryption](#)
- Data Encryption Standard (DES), [313](#), [407](#)
- data integrity, [86](#)
- data leakage, [340](#)
- data leakage prevention (DLP), [419](#), [421](#)
- Data Link Layer (Layer 2), [312](#)
- data origin authentication, [86](#)
- data protection firewall, [419–420](#)
- data sovereignty, [419](#)
- DDoS attacks. *See* [distributed denial of service attacks](#)
- DDoS mitigator, [61](#)
- dead-man switch, [145](#)
- decryption, [268](#)
- dedicated application-specific proxy firewall, [169](#)
- dedicated connection, [257](#)
- dedicated leased lines, [259](#), [260f](#)
- deduplication, [216](#)
- default allow, [224](#)
- default deny, [224](#), [355](#)
- default password, [296](#)
- defense in depth, [14](#), [98–102](#), [99f](#), [195–196](#), [355](#), [375](#)
- delay, [374](#)
- demilitarized zone (DMZ)-based implementation, [294–296](#), [294f](#)
- demilitarized zones (DMZs), [10](#), [59](#), [93](#), [94f](#), [126](#), [145](#), [174](#), [179f](#), [229–230](#), [337](#)
- denial of service (DoS), [17](#), [59–60](#), [60f](#), [72](#), [174](#), [180](#), [240](#), [286](#), [299](#), [359](#), [416](#)
- deny by default/allow by exception, [135](#), [145](#), [224](#), [355](#), [363](#)

- desktop virtualization, 323–324
- DES. *See* **Data Encryption Standard**
- detailed implementation plans, 353
- detection, 374
- deterrence, 374
- deterrent, 32
- DHCP. *See* **Dynamic Host Configuration Protocol**
- dialer, 48
- dial-up modem connections, 262
- dictionary password attacks, 116
- dictionary password cracking, 117
- Diffie-Hellmann, 274
- digital certificates, 5, 278
- digital envelope, 271, 272*f*
- digital forensic techniques, 394
- digital signatures, 175, 272, 273*f*
- digital subscriber line (DSL) modems, 149, 152
- DirectAccess, 336–337
- directory services, 138
- disaster recovery plan, 120, 392
- disgruntled employees, 37
- distributed denial of service (DDoS) attacks, 60–61, 61*f*
- diversity of defense, 196–199
- divide and conquer, 99
- DLP. *See* **data leakage prevention**
- DMZs. *See* **demilitarized zones**
- DNS poisoning, 57
- DNS. *See* **Domain Name System**
- DNS spoofing, 57
- documentation, 305
- do-it-yourself firewall, 169
- Domain Name System (DNS), 57, 175
- domain registrations, 39
- domains, 8
- domains of IT infrastructure, 8–11, 8*f*, 44–46
- DoS attacks. *See* **denial of service attacks**

DoS. *See* **denial of service**

downtime, **7**

DSL modems. *See* **digital subscriber line modems**

dual-homed firewall, **163**

dual-stack mechanism, **86**

dumb hub, **69**

dumpster diving, **62**

duplicate servers, **122, 123**

dynamic addressing, **105**

dynamic filtering system, **360**

Dynamic Host Configuration Protocol (DHCP), **105, 154**

dynamic NAT, **323**

dynamic packet filtering, **165–166**

dynamic password token, **117**

E

easy-access management console port or interface, [115](#)
easy-access power switch, [115](#)
eavesdropping, [50–51](#), [107](#), [108](#), [255](#)
ECC. *See* **elliptical curve cryptography**
edge routers, [264](#)
education, [383](#), [391](#)
efficient network security, [4](#)
EF. *See* **exposure factor**
egress filtering, [164](#)
electricity consumption, [115](#)
electronic health record (EHR) systems, [418](#)
Electronic Privacy Information Center (EPIC), [176](#)
electrostatic discharge (ESD), [401](#)
elliptical curve cryptography (ECC), [274](#)
Encapsulating Security Payload (ESP), [86](#), [313](#), [323](#)
encapsulation protocols, [18](#). *See also* **tunneling protocols**
encrypted protocols, [106](#), [107](#)
encryption, [18](#), [118](#), [138](#), [206](#), [242–243](#), [244f](#), [258](#), [268–278](#), [347](#),
[355](#), [361](#), [375](#), [407–408](#)
encryption filtering, [147](#)
encryption key sets, [258](#)
encryption level, [321](#)
endpoint security, [121–125](#)
enhancements for firewall, [244–245](#)
enumeration, [38](#), [41–43](#), [42f](#)
EPIC. *See* **Electronic Privacy Information Center**
equipment selection for secure network design, [114–115](#)
eradication, [393](#)
ESD. *See* **electrostatic discharge**
Ethernet frame, [71](#), [71f](#)
ethical hackers, [356](#)
ethical hacking. *See* **penetration testing**
European Union (EU), [419](#)
experts, [406–407](#)

exploits, **8, 296, 358–361**
exposure factor (EF), **125**
external attacks, **137**
external entities threats, **36–38**
external-only communications, **110**
external threats, **16**
extranet, **93, 94f, 126, 145, 260**
extranet VPNs, **267, 338, 338f**

F

- factory defaults, [400](#)
- fail-close state, [395](#)
- fail-open state, [395](#)
- fail-safe/fail-secure response, [143](#), [145](#)
- fail-safe security stance, [200](#)
- fail-secure state, [395](#)
- failures, [394–395](#)
- fair queuing, [241](#)
- fallback attacks, [38](#)
- false negatives, [214](#)
- false positives, [214](#)
- Family Educational Rights and Privacy Act (FERPA), [418](#)
- Federal Communications Commission (FCC), [70](#)
- Federal Information Processing Standards (FIPS) Publications, [407](#)
- file encryption, [382](#)
- file sharing, [340](#)
- File Transfer Protocol (FTP), [47](#), [103](#)
- filtering, [18](#), [108–110](#), [134–139](#), [164–167](#)
- filters, [224](#)
- firewalking, [239–240](#), [359](#)
- firewall checklist, [356](#)
- firewall deployment, [186](#)
- firewall filtering, [163](#), [359](#)
- firewall implementation, [177–180](#), [178f](#), [179f](#)
- firewall limitation, [238–240](#)
- firewall logging, [147](#)
- firewall logs, [210–213](#)
- firewall management, [354–357](#), [364–366](#)
- firewall policy, [236–237](#)
- firewall rules, [204](#), [224–234](#)
- firewalls, [109](#), [124](#), [134–170](#), [180](#), [192–201](#), [318](#), [355](#), [356](#), [357f](#), [363](#), [407](#)
- firewall-to-firewall VPN, [264](#)
- firewall troubleshooting, [366–369](#)

- flaw exploitation attacks, [60](#)
- flexibility, [168](#), [182](#)
- flooding, [33](#)
- focus for information security, [409–410](#)
- footprinting, [39](#)
- forced universal participation, [200–201](#)
- formal change management, [304](#)
- Fport, [365](#)
- fragmentation, [53](#), [261](#)
- fragmentation attacks, [53](#), [359](#), [360](#)
- free software, [149](#)
- FTP. *See* **File Transfer Protocol**
- full mesh of leased lines, [259](#)
- full mesh topology, [80](#), [80f](#)
- future developments, [406–411](#)
- fuzzing tools, [364](#)

G

- gateway, [75–76](#), [75f](#), [136](#), [137](#)
- gateway-to-gateway VPN, [330](#)
- General Data Protection Regulation (GDPR), [419](#)
- general filter firewall, [145](#)
- general-purpose OSs, [207](#)
- goals of network security, [11](#)
- GoToMyPC, [333](#), [339](#), [363](#)
- governance, [416](#)
- governance, risk, and compliance (GRC), [416](#)
- Gramm-Leach-Bliley Act (GLBA), [418](#)
- granular access control, [318](#)
- graphical user interface (GUI), [246](#), [246f](#)
- growth scenario contingencies, [101–102](#), [101f](#)

H

- hackers, [4](#), [8](#), [32](#)
- hacker tools, [62](#)
- hacking, [32](#)
- hacktivism, [35](#)
- hairpinning, [286](#)
- hardening firewall, [355](#)
- hardening host, [24](#)
- hardening networks, [9](#)
- hardening servers, [377](#)
- hardening systems, [113–114](#)
- hardware failures, [44–45](#), [402](#)
- hardware firewalls, [17](#), [143](#), [148](#), [161](#)
- hardware/software platform, [299](#)
- hardware VPNs, [263](#), [296](#), [310–311](#)
- hash algorithm, [274](#)
- hashing, [116](#), [175](#), [272](#), [273f](#), [274–275](#)
- hash value, [274](#)
- headers, [134](#)
- Health Insurance Portability and Accountability Act (HIPAA), [417](#)
- heat, [115](#)
- HFS. *See* **hierarchical file system**
- hierarchical file system (HFS), [56](#)
- high availability, [242](#)
- hijack attack, [51](#)
- HITECH Act, [418](#)
- home routers, [153–156](#), [156f](#)
- honeynets, [384–386](#)
- honeypots, [34](#), [384–386](#), [385f](#)
- host, [24](#)
- host firewalls, [143](#), [160](#), [387](#). *See also* **software firewall**
- host software firewall, [140](#), [157–160](#)
- host-to-gateway VPN, [330](#)
- host-to-host VPN, [267](#)
- host-to-site VPN, [265](#)

host VPN software product, [258](#)
HotSpotShield, [363](#)
Hot Standby Router Protocol (HSRP), [302](#)
HSRP. *See* **Hot Standby Router Protocol**
HTTP Proxy, [363](#)
HTTPS. *See* **Hypertext Transfer Protocol Secure**
hub, [69–70](#), [69f](#)
hybrid attacks, [116](#)
hybrid firewall, [419](#)
hybrid system, [182](#)
hybrid VPN, [255](#), [255f](#), [256f](#), [258](#)
Hypertext Transfer Protocol (HTTP), [315](#)
Hypertext Transfer Protocol Secure (HTTPS), [174](#), [315](#), [317f](#)
hypervisor, [422](#)

I

IAM. *See* **identity and access management**
IANA. *See* **Internet Assigned Numbers Authority**
ICMP redirects, 57
ICMP. *See* **Internet Control Message Protocol**
identity, 408
identity and access management (IAM), 115
identity proofing, 271
IDS insertion, 52–53
IDSs. *See* **intrusion detection systems**
IEEE 802.3 Ethernet Type II frame, 71*f*
IEMI. *See* **intentional electromagnetic interference**
IETF. *See* **Internet Engineering Task Force**
IKEv2. *See* **Internet Key Exchange v2**
IM. *See* **instant messaging**
IMS. *See* **IP Multimedia Subsystem**
in-band management, 216
inbound communication, 226
inbound traffic, 108, 109
incident response, 369, 391–394
incident response plan, 356, 392
increased address space, IPv6, 85
individual firewall, 150–156
industry-standard protocol, 302
information gathering. *See* **reconnaissance**
information security professionals, 405
Information Technology Infrastructure Library (ITIL), 409
infrastructure, 354
ingress filtering, 164
in-line proxy, 136
in-person policy training, 291
insertion attacks, 52–53
insider attacks, 139–140
installer-induced security threats, 297
installing pfSense firewall, 183–186, 185*f*

- instant messaging (IM), [47](#)
- intangible costs and value, [125](#)
- integrity, [7](#), [175](#), [274](#), [416](#)
- intentional electromagnetic interference (IEMI), [45](#)
- interception attack, [56](#)
- intermediary network, [254](#)
- internal and external hackers, [37](#), [37f](#)
- internal code planting, [360](#)
- internal firewall, [355](#)
- internally connected VPN, [293](#), [293f](#)
- internal personnel, [36–38](#)
- Internet Assigned Numbers Authority (IANA), [83](#), [85](#), [104](#), [212–213](#)
- Internet-based Traceroute tools, [307f](#)
- Internet Café VPNs, [339](#)
- Internet Connection Sharing service, [157](#)
- Internet connectivity, [355](#)
- Internet Control Message Protocol (ICMP), [145](#), [230](#), [236](#)
- Internet Engineering Task Force (IETF), [312](#), [336](#)
- Internet-facing servers, [174](#), [178](#)
- Internet Key Exchange (IKE), [313](#), [323](#)
- Internet Key Exchange v2 (IKEv2), [334](#)
- Internet of Things (IoT), [410](#)
- Internet Protocol Security (IPSec), [86](#), [103](#), [312–314](#), [320](#), [323](#), [329](#), [334](#), [341](#), [355](#), [375](#), [422](#)
- Internet Protocol version, [309](#)
- Internet Protocol version 4 (IPv4), [24](#), [84–86](#), [85f](#), [103](#), [163](#), [322](#)
- Internet Protocol version 6 (IPv6), [84–86](#), [85f](#), [103](#), [104](#), [106](#), [163](#), [422](#)
- Internet relay chat (IRC) channel, [60](#)
- Internet service provider (ISP) devices, [48](#), [104](#), [152](#), [153](#)
- Internet service providers (ISPs), [76](#), [240](#)
- Internet threats, [16](#), [150](#)
- Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX), [103](#)
- intranet, [337](#)
- intranet VPNs, [337](#)

intrusion detection systems (IDSs), [10](#), [110–112](#), [111f](#), [112f](#), [139](#),
[214–215](#), [387](#)
intrusion prevention systems (IPSs), [10](#), [110–112](#), [111f](#), [214–215](#)
IoT hub, [70](#)
IP address, [142](#)
IP addresses, [104](#), [295](#), [341](#)
ipchains, [180](#)
IP forwarding, [163](#)
IP Multimedia Subsystem (IMS), [411](#)
IP Network Address Translator, [322–323](#)
IPSec. *See* **Internet Protocol Security**
IPSec VPNs, [295](#), [314](#), [317](#), [318](#), [334](#)
IPSs. *See* **intrusion prevention systems**
IPv4. *See* **Internet Protocol version 4**
IPv6. *See* **Internet Protocol version 6**
IPX/SPX. *See* **Internetwork Packet Exchange/Sequenced Packet
Exchange**
IRC channel. *See* **Internet relay chat channel**
ISO image, [184](#)
ISP devices. *See* **Internet service provider (ISP) devices**
ITIL. *See* **Information Technology Infrastructure Library**
IT infrastructure domains, [8–11](#), [8f](#), [119f](#)
IT infrastructure threats, [44–46](#)

J

JanusVM, [362](#)
job description, [7](#)

K

Keep It Simple: Security (KISS), [355](#)
kernel panics, [188](#)
key exchange, [271](#)
key pairs, [270](#)
keyspace, [269](#)
Keystroke logger, [47](#)
knowledge-based detection, [215](#)

L

- laboratory tests, [364](#)
- LAN Domain. *See* **Local Area Network Domain**
- LAN infrastructure security, [13–14](#)
- LAN. *See* **local area network; wireless local area network**
- LAN-to-LAN VPNs, [266](#)
- LAN-to-WAN Domain, [10](#)
- latency, [258](#)
- layered security approach, [100](#)
- Layer 2 Forwarding (L2F), [314](#)
- Layer 3 of the OSI model, [312](#), [317](#)
- Layer 7 of the OSI model, [312](#)
- Layer 2 Transport Protocol (L2TP), [312](#)
- Layer 2 Tunneling Protocol (L2TP), [314–315](#), [334](#)
- leased lines, [259](#)
- leetspeak, [44](#)
- Linux firewalls, [152](#), [160](#)
- load balancers, [136](#)
- load balancing, [241](#)
- load-related performance, [320](#)
- local area network (LAN), [9](#), [69f](#), [73](#), [90–91](#), [93f](#)
- Local Area Network (LAN) Domain, [9](#)
- local host, [118](#)
- location-aware anti-theft software, [119](#)
- location of equipment selection, [115](#)
- location of VPN, [321](#)
- log, [74–75](#)
- logging, [15](#), [147](#), [208–210](#)
- logging on. *See* **authentication**
- logical security checklist, [380–381](#)
- logical topology, [76](#), [81–84](#), [82f](#)
- logic bomb, [47](#)
- LogMeIn, [339](#), [363](#)
- Loki, [362](#)
- loophole, [232](#)

M

MAC address. *See* **Media Access Control address**

MAC spoofing, [57](#), [124](#)

Mac support, [340](#)

mainframe, [91](#)

malicious code, [7](#), [12](#), [14](#), [15](#), [46–49](#)

malicious code scanner, [139](#)

malicious hackers, [4](#), [8](#)

management, [13](#)

management interfaces, [227](#), [245–248](#)

man-in-the-middle (MitM) attacks, [56–57](#), [56f](#)

maximum transmission unit (MTU), [53](#), [239](#), [360](#)

mean time between failures (MTBF), [44](#)

mean time to failure (MTTF), [44](#)

Media Access Control (MAC), [56](#), [70](#), [105](#)

memstick, [184](#)

mesh topology, [80–81](#), [80f](#)

metacharacters, [52](#)

metrics, [409](#)

Microsoft Point-to-Point Encryption (MPPE), [334](#)

mission-critical issue, [356](#)

MitM attacks. *See* **man-in-the-middle attacks**

MITRE, [42](#)

mobile code, [48](#)

mobile connectivity, [415](#)

mobile devices, [410](#)

Mobile IP, [411](#)

mobile wireless, [415](#)

modeling, [98](#)

modem, [76](#), [256](#)

module-based VPN, [311](#)

monitoring, [4](#), [147](#), [208–210](#), [365](#)

monkey-in-the-middle attack, [56](#)

MPPE. *See* **Microsoft Point-to-Point Encryption**

MSTSC command, [333](#)

MTBF. *See* **mean time between failures**

MTTF. *See* **mean time to failure**

MTU. *See* **maximum transmission unit**

multicast address, **106**

multifactor authentication, **117, 122, 123, 263, 381**

multi-homed firewall, **136–137, 137f**

multiple LANs, **266f**

multiple-layered defense, **355, 375**. *See also* **defense in depth**

Murphy's Law, **288, 367**

N

NAC. *See* **network access control**

NAS. *See* **network access server**

National Information Infrastructure (NII), **176**

National Institute of Standards and Technology (NIST), **42–43, 378f**

national security, **407**

National Security Agency (NSA), **174**

National Vulnerability Database, **378f**

native firewall, **151, 152**

native operating system firewall, **151**

NAT-PT. *See* **Network Address Translation–Protocol Translation**

NAT. *See* **network address translation**

NAT-Traversal (NAT-T), **334**

NAT-T. *See* **NAT-Traversal**

natural disasters, **45**

Nessus, **365**

Nessus vulnerability scanning, **396–397, 397f**

NetBEUI. *See* **NetBios Extended User Interface**

NetBios Extended User Interface (NetBEUI), **103**

NetBus, **362**

Netcat, **362, 365**

Netscape, **315**

Netstat, **365**

network access control (NAC), **24, 120**

Network Access Protocol, **337**

network access server (NAS), **138, 205**

network address translation (NAT), **59, 146f, 166, 205–206, 309, 318, 322–323, 334**

Network Address Translation–Protocol Translation (NAT-PT), **163**

network analyzer, **73**

network and resource availability threats, **59–61**

network communication, **234–235**

network compartmentalization. *See* **compartmentalization**

network design, **10, 14, 98–102**

network devices, types of, **68–76**

network infrastructures, examples of, [87–94](#)
network interface controller (NIC), [70](#), [162](#)
network issues, internal and external, [14–16](#)
Network Layer (Layer 3), [164](#)
Network Layer of the OSI model, [312](#)
Network News Transfer Protocol (NNTP), [103](#), [178](#)
network performance, [330](#)
network perimeter, [413–414](#)
network security, [4](#), [11–12](#), [174](#), [373–403](#)
network security components, [17–24](#)
network security design, [98–102](#)
network’s firewalls, [204–205](#)
network topology, [68](#)
network traffic access control security policy, [135](#)
new technology file system (NTFS), [56](#)
next-generation firewall (NGFW), [162](#)
NII. *See* **National Information Infrastructure**
NIST. *See* **National Institute of Standards and Technology**
nmap, [41](#), [365](#)
NNTP. *See* **Network News Transfer Protocol**
node, [17](#)
node security, [121](#)
no-exceptions policy, [355](#)
nonauthenticating query service, [57](#)
non-dedicated connection, [257](#)
nonrepudiation, [8](#), [175](#), [188](#), [271](#)
NSA. *See* **National Security Agency**
NTFS. *See* **new technology file system**
N-tier deployment, [94](#), [94f](#), [196](#), [197f](#)

O

- OC line. *See* **optical carrier line**
 - off-site storage, [387](#)
 - off-the-shelf firewall, [169](#)
 - one-time pad encryption systems, [277](#)
 - one-way function mathematical operation, [271](#)
 - ongoing education, [388–391](#)
 - Onion Router application, The. *See* **Tor application**
 - onion routing, [340](#)
 - online remote VPN options, [339–340](#)
 - online storage, [387](#)
 - onsite storage, [387](#)
 - open-source applications and tools, [413](#)
 - open source software, [149](#), [160](#)
 - open-source VPNs, [300–301](#)
 - Open Systems Interconnection (OSI), [103](#)
 - Open Systems Interconnection (OSI) Reference Model, [312](#), [318](#)
 - OpenVAS, [396](#)
 - OpenVPN, [342–346](#)
 - OpenVPN Private Tunnel, [346](#)
 - operating system-based VPNs, [310–311](#), [328–330](#)
 - operating systems (OSs), [41](#), [298](#), [321–322](#), [406](#)
 - opportunistic hackers, [35](#)
 - optical carrier (OC) line, [257](#)
 - optical disc, [184](#)
 - OSI model. *See* **Open Systems Interconnection (OSI) Reference Model**
 - OSs. *See* **operating systems**
 - outbound communication, [226](#)
 - outbound traffic, [108](#), [109](#), [155](#)
 - out of band communication, [270](#)
 - out-of-band management (OBM), [216](#)
 - overlapping, [359](#), [360](#)
-

P

- package filtering firewall, [182](#)
- packet analyzer, [73](#)
- packet filtering, [359](#)
- packet filtering firewall, [164](#)
- packet header, [134](#)
- PacketiX VPN, [363](#)
- packet payload, [134](#)
- packet sniffer, [339](#)
- padded cells, [384–386](#)
- partial mesh topology, [80–81](#), [80f](#)
- partition, [58](#)
- passive hub, [69](#)
- passive threats, [150](#)
- password cracking, [116](#)
- password-protected homegroup, [158](#)
- patches, [289](#), [400](#)
- patch management, [263](#), [376](#)
- PAT. *See* [port address translation](#)
- payloads, [134](#)
- Payment Card Industry Data Security Standard (PCI DSS), [418–419](#)
- PBX. *See* [private branch exchange](#)
- peer systems, [155](#)
- peer-to-peer (P2P) networks, [340](#)
- penetration testing, [356](#), [397–398](#)
- performance, [168](#), [182](#), [320–321](#), [409](#)
- performance improvement, [240–242](#)
- perimeter network, [178](#), [179](#)
- permissions, [5](#)
- personal firewall, [148](#), [169](#)
- personal hardware firewall, [149](#)
- personal/individual VPN, [301](#)
- personally identifiable information (PII), [175](#)
- personal software firewall, [149](#)
- pfSense, [180–181](#)

- pfSense requirements, [181](#)
- phases of hacking, [38–44](#), [38f](#)
- phishing, [47](#), [317](#)
- physical access, [124](#)
- physical access management, [107–108](#)
- physical attacks, [141](#)
- physical damage, [401](#)
- physical security, [355](#), [382–383](#)
- physical threats, [44–45](#)
- physical topology, [76–81](#)
- PII. *See* **personally identifiable information**
- piloting, [103](#)
- PING, [305](#), [306f](#)
- ping command, [305](#)
- ping sweeps, [39](#)
- PKI. *See* **Public Key Infrastructure**
- plaintext protocols, [115](#)
- platform independence, [318](#)
- playback attacks, [51](#)
- plug-and-play configuration, [86](#)
- PNAC. *See* **port-based network access control**
- Point-to-Point Protocol (PPP), [314](#)
- Point-to-Point Tunneling Protocol (PPTP), [314](#), [334](#)
- POP. *See* **Post Office Protocol**
- pop-up blockers, [150](#)
- portability of equipment, [114](#)
- port address translation (PAT), [206](#)
- port-based network access (admission) control (PNAC), [202](#)
- port forwarding, [146](#), [146f](#), [155](#), [206–207](#)
- port numbers, [146](#)
- port scanning, [39](#)
- port validation, [306](#)
- post-attack activities, [43–44](#)
- post-incident follow-up, [393–394](#)
- post-mortem assessment review, [398](#)
- Post Office Protocol (POP), [103](#)

- power faults, [401](#)
- power switch, [115](#)
- PPP. *See* **Point-to-Point Protocol**
- PPTP. *See* **Point-to-Point Tunneling Protocol**
- preparation, incident response, [392–393](#)
- prevention, [374](#)
- principle of least privilege, [117](#), [194–195](#), [279](#)
- privacy, [7](#), [175–176](#), [301](#)
- private branch exchange (PBX), [254](#)
- private IP address, [105](#)
- private key, [271](#)
- private messages, [175](#)
- private network, [126](#), [127f](#)
- private networks, [142](#), [142f](#), [333](#)
- private VPN, [254](#)
- privilege control, [168](#)
- privilege escalation, [43](#)
- privileges, [5](#)
- PRNG. *See* **pseudo random number generator**
- proactive security management, [141](#)
- professional hackers, [35](#)
- proprietary OSs, [207](#)
- protocol analyzer, [73](#)
- protocol converter, [76](#)
- protocol encryption, [118](#)
- protocols, [102–104](#)
- proxies, [124](#)
- proxy attack, [56](#)
- proxy-based encryption, [206](#)
- proxy firewall, [147](#)
- proxy implementation, [182](#)
- proxy manipulation, [57](#)
- proxy servers, [176](#), [182](#), [340](#)
- pseudo random number generator (PRNG), [278](#)
- public/home office wireless, [414–415](#)
- public IP addresses, [104](#), [105](#), [142](#)

public key, **271**

public key cryptography, **7, 270**

public key infrastructure (PKI), **138, 334, 342, 415**

public networks, **254, 333**

punchdown block, **69**

punch panel, **69**

PuTTY application, **319*f***

pwned, **44**

R

- rack mountable equipment, [114](#)
- RADIUS-based authentication, [295](#)
- rainbow tables, [116](#)
- random access memory (RAM), [182](#)
- ransomware, [33](#)
- RAS. *See* [remote access server](#)
- RDC. *See* [Remote Desktop Connection](#)
- RDP. *See* [Remote Desktop Protocol](#)
- RD RemoteApp, [335](#)
- RD Web Access, [335–336](#)
- reconnaissance, [39](#)
- recovery, [387–388](#), [394](#)
- recreational hackers, [35](#)
- reduced management requirement, IPv6, [85](#)
- redundancy, [102](#), [120–121](#), [285](#), [302](#)
- redundant array of independent disks (RAID), [182](#)
- redundant array of inexpensive disks (RAID), [120](#)
- Regional Internet Registry (RIR), [104](#)
- rekeying processes, [277](#)
- remote access, [18](#), [91–92](#), [118](#), [255](#), [256](#), [259](#), [280](#), [333–335](#), [358](#), [382](#)
- Remote Access Domain, [10](#)
- remote access policy, [286](#), [290](#)
- remote access server (RAS), [92](#), [138](#), [205](#)
- remote access VPN, [265](#)
- Remote Assistance, [333](#), [362](#)
- remote connection, [114](#)
- remote control, [18](#), [91](#), [255](#), [256](#), [333](#)
- Remote Desktop Connection (RDC), [256](#), [332](#)
- Remote Desktop Protocol (RDP), [331–333](#), [362](#)
- Remote Desktop Services, [333](#), [335–336](#)
- remote hacking, [382](#)
- remote or mobile host, [118–120](#)
- remote printing, [340](#)
- remote-to-home VPN, [267](#)

- remote-to-office VPN, [267](#)
- remote VPN connection, [320](#)
- removable case, [114](#)
- removable media, [114](#), [140](#)
- repeater, [70](#)
- replay attacks, [51](#), [51f](#)
- Request for Comments (RFC), [312](#)
- RescueAssist, [363](#)
- research. *See* **reconnaissance**
- reset button, [115](#)
- resources, [4](#)
- resources sites, [411–412](#)
- response, [375](#)
- return on investment (ROI), [49](#)
- reverse caching, [206](#)
- reverse proxy, [146](#), [205–206](#)
- RFC 791, [84](#)
- RFC 1918, [105](#)
- RFC 2460, [85](#)
- RFC 4193, [105](#)
- ring topology, [78–79](#), [78f](#)
- RIR. *See* **Regional Internet Registry**
- risk, [8](#), [175](#), [416–417](#)
- risk assessment, [125–126](#), [354](#)
- risk management, [125–126](#)
- Rivest-Shamir-Adelman (RSA), [274](#)
- rogue access point, [57](#)
- rogue device insertion, [53](#)
- rogue DHCP, [57](#)
- ROI. *See* **return on investment**
- roles, [13](#)
- rootkits, [47](#)
- round-robin, [241](#)
- route, [76](#)
- router, [72–73](#), [72f](#), [73f](#), [154f](#)
- router access control lists, [107](#)

router anti-spoofing, **107**

routers, **107, 123**

RSA. *See* **Rivest-Shamir-Adelman**

rules, **134**

rule sets, **139**

S

- sacrificial host, [143](#)
- Sarbanes-Oxley (SOX) Act, [418](#)
- scalability, [101–102](#), [168](#), [259](#)
- scaling, [77](#)
- scanning, [39–41](#), [40f](#)
- scope/binding nature statement, [290](#)
- screened IDS/IPS solution, [112](#), [112f](#)
- screening routers, [136](#), [150](#)
- script kiddie, [35](#)
- search engine optimization (SEO), [48](#)
- sector, [58](#)
- secured VPN, [255](#), [258](#). *See also* **virtual private network (VPN)**
- secure network design. *See* **network design**
- secure remote access, [259](#)
- Secure Shell (SSH) protocol, [318–319](#)
- Secure Sockets Layer (SSL), [139](#), [328](#), [334](#), [422](#)
- Secure Sockets Layer (SSL)–based tunneling protocols, [334](#)
- Secure Sockets Layer (SSL)/Transport Layer Security (TLS), [103](#), [315–318](#)
- Secure Socket Tunneling Protocol (SSTP), [334](#)
- security, [86–94](#), [206](#), [216–218](#), [355](#)
- security assessment, [396](#)
- security association (SA), [313](#)
- security assurance, [168](#)
- security awareness, [388–391](#)
- security checklist, [380–381](#)
- security goals, [11](#), [98](#)
- security information and event management (SIEM), [215–216](#)
- security infrastructure, [199](#)
- security management, [141](#), [396](#)
- security mistakes, [389](#)
- security objectives, [6–7](#)
- security planning, [379–380](#)
- security policies, [12](#), [135](#), [141](#), [176](#), [178](#), [379–380](#)

- security stance, [150](#)
- security suite firewall, [158](#)
- Security Technical Implementation Guides (STIGs), [395](#)
- security technologies, [405](#), [415–417](#)
- security through obscurity, [100](#), [192–194](#)
- security troubleshooting, [400–402](#)
- SED. *See* **static electricity discharge**
- segment, [137](#), [137f](#)
- Select and train security incident response team (SIRT), [369](#)
- sender fragmentation, [360](#)
- senior management, [13](#), [102](#), [390](#)
- separation of duties, [194–195](#), [377](#)
- servers, [5](#), [122–123](#)
- service level agreements (SLAs), [91](#), [285](#), [421](#)
- service set identifier (SSID), [154](#)
- session, [134](#)
- session hijacking, [54–55](#), [54f](#), [382](#)
- shell code, [43](#)
- ShieldsUP! port scanning tool, [156](#), [156f](#)
- sieve firewall, [145](#)
- signature-based detection, [215](#)
- signatures, [46](#)
- Simple Mail Transfer Protocol (SMTP), [103](#), [235](#)
- simple network management protocol (SNMP), [125](#)
- simplicity, [195](#)
- simulated firewall test, [364](#)
- single-factor authentication, [381](#)
- single loss expectancy (SLE), [125](#)
- single point of failure, [99](#), [100](#)
- single sign-on (SSO), [89](#)
- site-to-site VPNs, [266](#)
- Skype, [304](#)
- slack space, [58–59](#), [59f](#)
- SLAs. *See* **service level agreements**
- SLE. *See* **single loss expectancy**
- slideware, [285](#)

small office/home office (SOHO), [21](#), [89](#), [150–156](#), [227](#), [301](#)
smart hub, [69–70](#)
SmoothWall software, [160](#)
SMTP. *See* [Simple Mail Transfer Protocol](#)
SNA. *See* [System Network Architecture](#)
sniffer, [73](#)
Snort, [365](#)
social engineering, [62–63](#), [139–140](#)
socket, [146](#)
software coding errors, [358](#)
software firewall, [148](#), [160](#)
software firewall products, [157–158](#)
software host firewalls, [387](#)
software VPNs, [267](#), [297](#), [298](#), [310–311](#), [321](#)
SOHO, [150–156](#). *See* [small office/home office](#)
solid-state drives (SDDs), [45](#)
solid VPN policy, components of, [290](#)
source routing bridge, [75](#)
spam, [33](#)
split knowledge. *See* [separation of duties](#)
split tunnel, [280](#), [286](#), [295](#)
spoofed addresses, [139](#), [164](#)
spoofing, [55–56](#), [55f](#), [164](#)
spyware, [47](#), [386](#)
spyware scanner. *See* [anti-malware scanners](#)
SQL injection, [52](#)
SSH protocol. *See* [Secure Shell protocol](#)
SSID. *See* [service set identifier](#)
SSL. *See* [Secure Sockets Layer](#)
SSL VPNs, [295](#), [314](#), [318](#), [323](#), [334](#)
SSO. *See* [single sign-on](#)
SSTP. *See* [Secure Socket Tunneling Protocol](#)
stability of VPNs, [298](#), [321–322](#)
star topology, [79](#), [79f](#)
state, [147](#)
stateful inspection, [147](#)

stateful inspection filtering, [165–166](#)
stateful inspection firewalls, [169](#), [360](#)
state management, [147](#)
static addressing, [105](#)
static electricity, [401](#)
static electricity discharge (SED), [45](#), [401](#)
Static Key, [342](#)
static NAT, [146](#), [323](#)
static packet filtering, [164–165](#)
steganography, [193–194](#)
STIGs. *See* **Security Technical Implementation Guides**
strong authentication, [117](#), [347](#)
strong encryption, [347](#)
subnet architecture, [183](#)
subnetting, [83](#), [104](#)
SubSeven, [362](#)
suite-member firewalls, [158](#)
sunk cost, [104](#)
switch, [70–72](#), [70f](#)
switch deployment, [71f](#)
switches, [124](#)
symmetric cryptography, [269–270](#), [270f](#)
symptoms, [304](#)
syslog, [209](#), [365](#)
System/Application Domain, [10–11](#)
Systems Network Architecture (SNA), [103](#)

T

TACACS+. See **Terminal Access Controller Access-Control System Plus**

tangible costs and value, 125

targets of hackers, 36

TCP/IP. See **Transmission Control Protocol/Internet Protocol**

TCP/IP suite, 87

TCP. See **Transmission Control Protocol**

TCPView, 365

telecommuting, 260

Telnet, 103

Terminal Access Controller Access-Control System Plus (TACACS+), 123

terminal services, 91, 333

terminators, 77

testing, 187–188, 344–345

testing firewall, 363–364

thin client computing, 91

third-party software firewalls, 157

third-party trust system, 5, 6f

threats, 4, 296, 406, 407

time synchronization, 216

TLS. See **Transport Layer Security**

token, 78, 79

tools, 412–413

topology, 68, 321

Tor application, 301, 304, 340–341, 362

TPM chip. See **Trusted Platform Module chip**

traceroute tool, 307f

traffic and trend analysis, 108

traffic congestion, 262

traffic loads, 321

training, 388–391

transaction security, 174

translation, 86

Transmission Control Protocol (TCP), [40](#), [40f](#), [236](#)
Transmission Control Protocol/Internet Protocol (TCP/IP), [102–103](#)
transparent bridge, [75](#)
transparent network security, [4](#)
transparent proxy, [136](#)
Transport Layer Protocol, [319](#)
Transport Layer Security (TLS), [139](#), [315–318](#)
transport mode encryption, [138](#), [268](#)
transport mode (host-to-host) of IPSec, [313](#)
trapdoor, [47](#)
triple-homed firewall, [163](#)
Trojan horse, [46](#)
troubleshooting, [187–188](#), [303–307](#), [344–345](#), [366–369](#)
trust, [5–6](#)
Trusted Platform Module (TPM) chip, [387](#)
trusted third party, [278](#)
trusted VPN, [254](#), [255f](#), [256f](#), [258](#)
trustworthy, [6–7](#)
tunneling, [86](#), [361–363](#)
tunneling protocols, [18](#), [258](#)
tunnel mode encryption, [138](#), [267](#)
tunnel mode of IPSec, [313–314](#)
two-factor authentication, [288](#)

U

- UAG. *See* **Unified Access Gateway**
 - ubiquitous firewall, **420**
 - UDP. *See* **User Datagram Protocol**
 - unauthorized software, **15**
 - unauthorized tunnels, **362**
 - unfiltered IDS/IPS installation, **112**, **112f**
 - unicast address, **106**
 - unidirectional, ring topology, **78**
 - Unified Access Gateway (UAG), **337**
 - unified threat management (UTM), **161**, **245**
 - uninterruptable power supply (UPS), **120**
 - universal participation, **201**
 - unpartitioned space, **58**
 - updates, **400**
 - UPS. *See* **uninterruptible power supply**
 - upstream filtering, **60**
 - URL injectors, **48**
 - usability, **101**, **356**, **376**
 - USB, **184**
 - User Authentication Protocol, **319**
 - User Datagram Protocol (UDP), **40**, **40f**, **236**, **341**
 - User Domain, **9**
 - user training, **302–303**, **388–391**
 - UTM. *See* **unified threat management**
-

V

Van Eck phreaking, [14](#)
vaporware. *See* [slideware](#)
vendors, [196](#), [197](#), [289](#), [296](#), [304](#), [407](#)
virtual firewall, [162–163](#)
virtual firewall test, [364](#)
virtualization, [162](#), [323–325](#), [410](#)
virtualization security, [421–422](#)
virtualized network environment, [364](#)
virtualized SSL VPN, [324–325](#)
virtual local area networks (VLANs), [83](#), [84f](#), [124](#), [185](#)
virtual private network (VPN), [14](#), [16](#), [18](#), [92–93](#), [92f](#), [93f](#), [103](#), [118](#),
[146](#), [180](#), [253–280](#), [283–307](#), [309–325](#), [328](#), [341–347](#), [353](#), [355](#)
virtual reassembly, [360](#)
Virtual Router Redundancy Protocol (VRRP), [302](#)
virus, [46](#)
VLAN. *See* [virtual local area networks](#)
VMware, [364](#)
voluntary compliance, [201](#)
VPN appliances, [258](#), [264](#), [329–331](#), [334](#)
VPN authorization, [279–280](#)
VPN concentrator, [313](#)
VPN installation, [341–342](#)
VPN link, [261](#)
VPN. *See* [virtual private network](#)
VRRP. *See* [Virtual Router Redundancy Protocol](#)
vulnerabilities, [8](#), [16](#), [287](#), [296](#), [297](#), [299](#)
vulnerability assessments, [356](#)
vulnerability management, [287–288](#)
vulnerability scanning, [396–397](#), [397f](#)

W

- wake-on-LAN, [339–340](#)
- WAN Domain, [10](#)
- WAN. *See* **wide area network**
- WAN VPN connections, [266](#)
- wardialing, [39](#)
- wardriving, [39](#)
- weakest link security stance, [196–200](#)
- Web-based policy training, [291](#)
- Web browsers, [150](#)
- Web server, [298](#)
- whitelist controls, [363](#)
- whitelists, [15](#)
- whois, [39](#)
- whole hard drive encryption, [382](#), [387](#)
- wide area network (WAN), [10](#), [90–91](#), [337](#)
- Windows Defender Firewall, [158](#), [159f](#), [160](#)
- Windows Firewall, [151](#), [152](#), [158](#)
- Windows Server 2019 Network Access, [329](#)
- wired networks, [14](#)
- wired topologies, [76](#)
- wireless access point (WAP), [53](#), [87](#), [414](#)
- wireless local area network (LAN), [329](#), [414](#)
- wireless networking, [4](#), [14](#), [301](#)
- wireless router, [72f](#)
- wireless topology, [81](#), [81f](#)
- wireless vs. wired, [50](#)
- Wireshark, [235](#), [339](#), [365](#), [366f](#)
- wire speed, [241](#)
- workgroups, [87–89](#), [88f](#)
- Workstation Domain, [9](#)
- worms, [46](#)
- WORM storage device. *See* **write-once read-many storage device**
- wrappers, [47](#)
- write-once read-many (WORM), [210](#)

written firewall policy, [354](#)
written security policy, [12](#), [379](#)

X

802.1x, [138](#)

XSS. *See* [cross-site scripting](#)

Z

zero-day exploits, [15](#)

zeroization, [15](#)

zombie army, [60](#)

zombies, [60](#), [142](#)

zone of risk, [126–127](#), [127f](#), [127t](#)

zones of trust, [126](#), [127t](#)