DAVID L. PROWSE

CompTIA APPROVED QUALITY CONTENT

Cert Guide

Learn, prepare, and practice for exam success

CompTIA®

Security+

SY0-501

CompTIA Security+

PEARSON IT CERTIFICATION

Save 10% on Exam Voucher

See Inside

FEATURES

Three Complete Practice Exams, More Than 30 Videos and 30 Interactive Exercises

FREE SAMPLE CHAPTER

SHARE WITH OTHERS

CompTIA® Security+ SY0-501 Cert Guide

Fourth Edition

David L. Prowse

PEARSON

800 East 96th Street
Indianapolis, Indiana 46240 USA

CompTIA® Security+ SY0-501 Cert Guide Fourth Edition

ISBN-13: 978-0-7897-5899-6
ISBN-10: 0-7897-5899-7

Library of Congress Control Number: 2017951236

Printed in the United States of America

1 17

Trademarks

Editor-in-Chief
Mark Taub

Product Line Manager
Brett Bartow

Acquisitions Editor
Michelle Newcomb

Development Editor
Eleanor Bru

Managing Editor
Sandra Schroeder

Senior Project Editor
Tonya Simpson

Copy Editor
Bill McManus

Indexer
Ken Johnson

Proofreader
Paula Lowell

Technical Editor
Chris Crayton

Publishing Coordinator
Vanessa Evans

Cover Designer
Chuti Prasertsith

Compositor
Studio Galou

Warning and Disclaimer

Special Sales

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

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

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

Contents at a Glance

Elements Available Online

Table of Contents

About the Author

David L. Prowse is an author, technologist, and technical trainer. He has penned a dozen books for Pearson Education, including the well-received *CompTIA A+ Exam Cram*. He also develops video content, including the *CompTIA A+ LiveLessons* video course. Over the past two decades he has taught CompTIA A+, Network+, and Security+ certification courses, both in the classroom and via the Internet. David has 20 years of experience in the IT field and loves to share that experience with his readers, watchers, and students.

He runs the website www.davidlprowse.com in support of his books and videos.

Acknowledgments

It takes a lot of amazing people to publish a book. Special thanks go to Eleanor Bru, Chris Crayton, Michelle Newcomb, and all the other people at Pearson (and beyond) who helped make this book a reality. I appreciate everything you do!

About the Technical Reviewer

Chris Crayton (MCSE) is an author, technical consultant, and trainer. In the past, he has worked as a computer technology and networking instructor, information security director, network administrator, network engineer, and PC specialist. Chris has authored several print and online books on PC repair, CompTIA A+, CompTIA Security+, and Microsoft Windows. He has also served as technical editor and content contributor on numerous technical titles for several leading publishing companies. Chris holds numerous industry certifications, has been recognized with many professional teaching awards, and has served as a state-level SkillsUSA competition judge.

We Want to Hear from You!

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

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

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

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

Email: feedback@pearsonitcertification.com

Mail: Pearson IT Certification
ATTN: Reader Feedback
800 East 96th Street
Indianapolis, IN 46240 USA

Reader Services

Register your copy of *CompTIA Security+ SY0-501 Cert Guide* at www.pearsonitcertification.com for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.pearsonitcertification.com/register and log in or create an account.* Enter the product ISBN 9780789758996 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

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

Introduction

Welcome to the *CompTIA Security+ SY0-501 Cert Guide*. The CompTIA Security+ Certification is widely accepted as the first security certification you should attempt to attain in your information technology (IT) career. The CompTIA Security+ Certification is designed to be a vendor-neutral exam that measures your knowledge of industry-standard technologies and methodologies. It acts as a great stepping stone to other vendor-specific certifications and careers. I developed this book to be something you can study from for the exam and keep on your bookshelf for later use as a security resource.

I'd like to note that it's unfeasible to cover all security concepts in depth in a single book. However, the Security+ exam objectives are looking for a basic level of computer, networking, and organizational security knowledge. Keep this in mind while reading through this text, and remember that the main goal of this text is to help you pass the Security+ exam, not to be the master of all security. Not just yet at least!

Good luck as you prepare to take the CompTIA Security+ exam. As you read through this book, you will be building an impenetrable castle of knowledge, culminating in hands-on familiarity and the know-how to pass the exam.

> **IMPORTANT NOTE** The first thing you should do before you start reading Chapter 1, "Introduction to Security," is check my website for errata and updated information, and mark those new items in the book. Go to www.davidlprowse.com and then the Security+ section. On my site you will also find videos, bonus test questions, and other additional content. And, of course, you can contact me directly at my website to ask me questions about the book.

Goals and Methods

The number one goal of this book is to help you pass the SY0-501 version of the CompTIA Security+ Certification Exam. To that effect, I have filled this book and practice exams with more than 600 questions/answers and explanations in total, including three 80-question practice exams. One of the exams is printed at the end of the book, and all exams are located in Pearson Test Prep practice test software in a custom test environment. These tests are geared to check your knowledge and ready you for the real exam.

The CompTIA Security+ Certification exam involves familiarity with computer security theory and hands-on know-how. To aid you in mastering and understanding the Security+ Certification objectives, this book uses the following methods:

- **Opening topics list:** This defines the topics to be covered in the chapter.
- **Topical coverage:** The heart of the chapter. Explains the topics from a theory-based standpoint, as well as from a hands-on perspective. This includes in-depth descriptions, tables, and figures that are geared to build your knowledge so that you can pass the exam. The chapters are broken down into two to three topics each.
- **Key Topics:** The Key Topic icons indicate important figures, tables, and lists of information that you should know for the exam. They are interspersed throughout the chapter and are listed in table format at the end of the chapter.
- **Key Terms:** Key terms without definitions are listed at the end of each chapter. See whether you can define them, and then check your work against the complete key term definitions in the glossary.
- **Real-World Scenarios:** Included in the supplemental online material are real-world scenarios for each chapter. These offer the reader insightful questions and problems to solve. The questions are often open-ended, and can have several different solutions. The online material gives one or more possible solutions and then points to video-based solutions and simulation exercises online to further reinforce the concepts. Refer to these real-world scenarios at the end of each chapter.
- **Review Questions:** These quizzes, and answers with explanations, are meant to gauge your knowledge of the subjects. If an answer to a question doesn't come readily to you, be sure to review that portion of the chapter. The review questions are also available online.
- **Practice Exams:** There is one practice exam printed at the end of the book, and additional exams included in the Pearson Test Prep practice test software. These test your knowledge and skills in a realistic testing environment. Take these after you have read through the entire book. Master one, then move on to the next. Take any available bonus exams last.

Another goal of this book is to offer support for you, the reader. Again, if you have questions or suggestions, please contact me through my website: www.davidlprowse.com. I try my best to answer your queries as soon as possible.

Who Should Read This Book?

This book is for anyone who wants to start or advance a career in computer security. Readers of this book can range from persons taking a Security+ course to individuals already in the field who want to keep their skills sharp, or perhaps retain their job due to a company policy mandating they take the Security+ exam. Some information assurance professionals who work for the Department of Defense or have privileged access to DoD systems are required to become Security+ certified as per DoD directive 8570.1.

This book is also designed for people who plan on taking additional security-related certifications after the CompTIA Security+ exam. The book is designed in such a way to offer an easy transition to future certification studies.

Although not a prerequisite, it is recommended that CompTIA Security+ candidates have at least two years of IT administration experience with an emphasis on security. The CompTIA Network+ certification is also recommended as a prerequisite. Before you begin your Security+ studies, it is expected that you understand computer topics such as how to install operating systems and applications, and networking topics such as how to configure IP, what a VLAN is, and so on. The focus of this book is to show how to secure these technologies and protect against possible exploits and attacks. Generally, for people looking to enter the IT field, the CompTIA Security+ certification is attained after the A+ and Network+ certifications.

CompTIA Security+ Exam Topics

If you haven't downloaded the Security+ certification exam objectives, do it now from CompTIA's website: https://certification.comptia.org/. Save the PDF file and print it out as well. It's a big document—review it carefully. Use the exam objectives list and acronyms list to aid in your studies while you use this book.

The following two tables are excerpts from the exam objectives document. Table I-1 lists the CompTIA Security+ domains and each domain's percentage of the exam.

Table I-1 CompTIA Security+ Exam Domains

Domain	Exam Topic	% of Exam
1.0	Threats, Attacks and Vulnerabilities	21%
2.0	Technologies and Tools	22%
3.0	Architecture and Design	15%
4.0	Identity and Access Management	16%
5.0	Risk Management	14%
6.0	Cryptography and PKI	12%

The Security+ domains are then further broken down into individual objectives. To achieve better flow and to present the topics in more of a building-block approach, I rearranged the concepts defined in the objectives. This approach is designed especially for people who are new to the computer security field.

Table I-2 lists the CompTIA Security+ exam objectives and their related chapters in this book. It does not list the bullets and sub-bullets for each objective.

> **NOTE** Chapter 19 gives strategies for taking the exam and therefore does not map to any specific objectives.

Table I-2 CompTIA Security+ Exam Objectives

Objective	Chapter(s)
1.1 Given a scenario, analyze indicators of compromise and determine the type of malware.	2, 13
1.2 Compare and contrast types of attacks.	7, 9, 14, 17
1.3 Explain threat actor types and attributes.	1, 17
1.4 Explain penetration testing concepts.	12
1.5 Explain vulnerability scanning concepts.	12
1.6 Explain the impact associated with types of vulnerabilities.	5, 12
2.1 Install and configure network components, both hardware- and software-based, to support organizational security.	6, 8, 10, 13, 15
2.2 Given a scenario, use appropriate software tools to assess the security posture of an organization.	13, 14, 18
2.3 Given a scenario, troubleshoot common security issues.	10, 11, 17
2.4 Given a scenario, analyze and interpret output from security technologies.	3, 4, 8
2.5 Given a scenario, deploy mobile devices securely.	3, 6, 9
2.6 Given a scenario, implement secure protocols.	6, 7, 13
3.1 Explain use cases and purpose for frameworks, best practices and secure configuration guides.	12, 18
3.2 Given a scenario, implement secure network architecture concepts.	6, 7, 9, 10, 13
3.3 Given a scenario, implement secure systems design.	3, 4
3.4 Explain the importance of secure staging deployment concepts.	5, 12
3.5 Explain the security implications of embedded systems.	3, 4, 18
3.6 Summarize secure application development and deployment concepts.	5
3.7 Summarize cloud and virtualization concepts.	4, 6
3.8 Explain how resiliency and automation strategies reduce risk.	12, 16
3.9 Explain the importance of physical security controls.	10

Objective	Chapter(s)
4.1 Compare and contrast identity and access management concepts.	10
4.2 Given a scenario, install and configure identity and access services.	10
4.3 Given a scenario, implement identity and access management controls.	10, 11
4.4 Given a scenario, differentiate common account management practices.	11
5.1 Explain the importance of policies, plans and procedures related to organizational security.	18
5.2 Summarize business impact analysis concepts.	16
5.3 Explain risk management processes and concepts.	12, 18
5.4 Given a scenario, follow incident response procedures.	18
5.5 Summarize basic concepts of forensics.	18
5.6 Explain disaster recovery and continuity of operation concepts.	16
5.7 Compare and contrast various types of controls.	1, 12
5.8 Given a scenario, carry out data security and privacy practices.	18
6.1 Compare and contrast basic concepts of cryptography.	14
6.2 Explain cryptography algorithms and their basic characteristics.	14
6.3 Given a scenario, install and configure wireless security settings.	9, 10
6.4 Given a scenario, implement public key infrastructure.	15

Companion Website

Register this book to get access to the Pearson Test Prep practice test software and other study materials plus additional bonus content. Check this site regularly for new and updated postings written by the author that provide further insight into the more troublesome topics on the exam. Be sure to check the box that you would like to hear from us to receive updates and exclusive discounts on future editions of this product or related products.

To access this companion website, follow these steps:

1. Go to www.pearsonitcertification.com/register and log in or create a new account.
2. On your Account page, tap or click the **Registered Products** tab, and then tap or click the **Register Another Product** link.
3. Enter this book's ISBN (9780789758996).

4. Answer the challenge question as proof of book ownership.
5. Tap or click the **Access Bonus Content** link for this book to go to the page where your downloadable content is available.

Please note that many of our companion content files can be very large, especially image and video files.

If you are unable to locate the files for this title by following the preceding steps, please visit http://www.pearsonitcertification.com/contact and select the "Site Problems/Comments" option. Our customer service representatives will assist you.

Pearson Test Prep Practice Test Software

As noted previously, this book comes complete with the Pearson Test Prep practice test software containing three full exams. These practice tests are available to you either online or as an offline Windows application. To access the practice exams that were developed with this book, please see the instructions in the card inserted in the sleeve in the back of the book. This card includes a unique access code that enables you to activate your exams in the Pearson Test Prep software.

NOTE The cardboard sleeve in the back of this book includes a piece of paper. The paper lists the activation code for the practice exams associated with this book. Do not lose the activation code. On the opposite side of the paper from the activation code is a unique, one-time-use coupon code for the purchase of the Premium Edition eBook and Practice Test.

Accessing the Pearson Test Prep Software Online

The online version of this software can be used on any device with a browser and connectivity to the Internet including desktop machines, tablets, and smartphones. To start using your practice exams online, simply follow these steps:

1. Go to www.PearsonTestPrep.com and select **Pearson IT Certification** as your product group.
2. Enter your email/password for your account. If you do not have an account on PearsonITCertification.com or CiscoPress.com, you will need to establish one by going to PearsonITCertification.com/join.
3. On the My Products tab, tap or click the **Activate New Product** button.

4. Enter this book's activation code and click **Activate**.
5. The product will now be listed on your My Products tab. Tap or click the **Exams** button to launch the exam settings screen and start your exam.

Accessing the Pearson Test Prep Software Offline

If you wish to study offline, you can download and install the Windows version of the Pearson Test Prep software. There is a download link for this software on the book's companion website, or you can just enter this link in your browser:

http://www.pearsonitcertification.com/content/downloads/pcpt/engine.zip

To access the book's companion website and the software, simply follow these steps:

1. Register your book by going to http://www.pearsonitcertification.com/register and entering the ISBN: **9780789758996**.
2. Respond to the challenge questions.
3. Go to your account page and select the **Registered Products** tab.
4. Click the **Access Bonus Content** link under the product listing.
5. Click the **Install Pearson Test Prep Desktop Version** link under the Practice Exams section of the page to download the software.
6. Once the software finishes downloading, unzip all the files on your computer.
7. Double-click the application file to start the installation, and follow the onscreen instructions to complete the registration.
8. Once the installation is complete, launch the application and click the **Activate Exam** button on the My Products tab.
9. Click the **Activate a Product** button in the Activate Product Wizard.
10. Enter the unique access code found on the card in the sleeve in the back of your book and click the **Activate** button.
11. Click **Next** and then the **Finish** button to download the exam data to your application.
12. You can now start using the practice exams by selecting the product and clicking the **Open Exam** button to open the exam settings screen.

Note that the offline and online versions will synch together, so saved exams and grade results recorded on one version will be available to you on the other as well.

Customizing Your Exams

Once you are in the exam settings screen, you can choose to take exams in one of three modes:

- Study Mode
- Practice Exam Mode
- Flash Card Mode

Study Mode allows you to fully customize your exams and review answers as you are taking the exam. This is typically the mode you would use first to assess your knowledge and identify information gaps. Practice Exam Mode locks certain customization options, as it is presenting a realistic exam experience. Use this mode when you are preparing to test your exam readiness. Flash Card Mode strips out the answers and presents you with only the question stem. This mode is great for late-stage preparation when you really want to challenge yourself to provide answers without the benefit of seeing multiple-choice options. This mode will not provide the detailed score reports that the other two modes will, so it should not be used if you are trying to identify knowledge gaps.

In addition to these three modes, you will be able to select the source of your questions. You can choose to take exams that cover all of the chapters or you can narrow your selection to just a single chapter or the chapters that make up specific parts in the book. All chapters are selected by default. If you want to narrow your focus to individual chapters, simply deselect all the chapters then select only those on which you wish to focus in the Objectives area.

You can also select the exam banks on which to focus. Each exam bank comes complete with a full exam of questions that cover topics in every chapter. The exam printed in the book is available to you as well as two additional exams of unique questions. You can have the test engine serve up exams from all banks or just from one individual bank by selecting the desired banks in the exam bank area.

There are several other customizations you can make to your exam from the exam settings screen, such as the time of the exam, the number of questions served up, whether to randomize questions and answers, whether to show the number of correct answers for multiple-answer questions, or whether to serve up only specific types of questions. You can also create custom test banks by selecting only questions that you have marked or questions on which you have added notes.

Updating Your Exams

If you are using the online version of the Pearson Test Prep software, you should always have access to the latest version of the software as well as the exam data. If you are using the Windows desktop version, every time you launch the software, it will check to see if there are any updates to your exam data and automatically download any changes that were made since the last time you used the software. This requires that you are connected to the Internet at the time you launch the software.

Sometimes, due to many factors, the exam data may not fully download when you activate your exam. If you find that figures or exhibits are missing, you may need to manually update your exams.

To update a particular exam you have already activated and downloaded, simply select the **Tools** tab and click the **Update Products** button. Again, this is only an issue with the desktop Windows application.

If you wish to check for updates to the Pearson Test Prep exam engine software, Windows desktop version, simply select the **Tools** tab and click the **Update Application** button. This will ensure you are running the latest version of the software engine.

Premium Edition eBook and Practice Tests

This book also includes an exclusive offer for 70 percent off the Premium Edition eBook and Practice Tests edition of this title. Please see the coupon code included with the cardboard sleeve for information on how to purchase the Premium Edition.

This chapter covers the following subjects:

- **Firewalls and Network Security:** In this section, you find out about one of the most important strategic pieces in your network security design—the firewall. Then we discuss other network security concepts such as packet filtering, access control lists, proxy servers, and honeypots.
- **NIDS Versus NIPS:** This section delves into the characteristics, advantages, disadvantages, and differences of network intrusion *detection* systems and network intrusion *prevention* systems.

This chapter is all about the network border, also known as the **network perimeter**. This should be a network security administrator's primary focus when it comes to securing the network because it contains the entrances that many attackers attempt to use.

Network Perimeter Security

Allow me to analogize for a few moments. I've said it before; as you read this book, you are building yourself an impenetrable castle of knowledge, culminating in hands-on familiarity and the know-how to pass the exam. But we can use the castle analogy for your network as well. Imagine a big stone castle with tall walls, an expanse of clear land around the castle, or perhaps a moat surrounding it (with alligators, of course), and one or more drawbridges. The tall walls are meant to keep the average person out, sort of like a firewall in a computer network—not perfect, but necessary. The open area around the castle makes it difficult for people to sneak up on your castle; they would quickly be *detected*, just like malicious packets detected by a network intrusion detection system. Or better yet, if you had a moat, people trying to cross it would have a difficult time, would be easy targets for your bowmen, and would probably be gobbled up by your pet alligators. This would represent a network intrusion *prevention* system, which not only detects threats, but also eliminates those threats to the network.

The drawbridge, or drawbridges, could be seen as network ports open to the network. As drawbridges are part of the castle wall, so network ports are part of the firewall. You, as the network security administrator, have the ability and the right to close these ports at any time. At the risk of taking this analogy even further, you might decide to set traps for people; like a pool of quicksand that has an open netted bag of pyrite suspended above it, or maybe a false entry to the castle that, after a long corridor, is walled off on the inside, ultimately trapping the unwary. In a network environment, these would be known as honeypots. Of course, every once in a while, legitimate traffic needs to enter and exit your network, too! To do this in a more secure fashion, you can set up proxy servers to act as go-betweens for the computers inside your network and the servers they talk to on the Internet: kind of like a sentry in the tower of the castle that would relay an outsider's messages to someone inside the castle.

The network perimeter is less tangible in an actual network environment (thus the previous use of superfluous metaphor). Networking devices are commonly located in a single server room or data center, or perhaps are located in a hybrid of in-house and cloud-based locations. Either way, they can be difficult to visualize. To better envision your network, one of the best tips I can give you is to map out your network on paper, or create network documentation using programs such as Microsoft Visio and by utilizing network mapping tools (more on these tools in Chapter 12, "Vulnerability and Risk Assessment").

So, before we end up playing *Dungeons & Dragons*, let's talk about one of the most important parts of your strategic defense—the firewall.

Foundation Topics

Firewalls and Network Security

Nowadays, firewalls are everywhere. Businesses large and small use them, and many households have simpler versions of these protective devices as well. You need to be aware of several types of firewalls, and you definitely want to spend some time configuring hardware and software firewalls. There are many free software-based firewalls and firmware-based emulators that you can download. A quick search on the Internet will give you several options.

The firewall is there to protect the entire network, but other tools are often implemented as well; for example, proxy servers that help protect users and computers by keeping them anonymous; honeypots meant to attract hackers, crackers, and other types of attackers into a false computer or network; and data loss prevention (DLP) devices to keep confidential data from leaving the network. But by far, the most important element in your network will be the firewall, so let's begin with that.

Firewalls

In Chapter 3, "Computer Systems Security Part II," we discussed personal firewalls—you remember, the kind installed to an individual computer. Now let's broaden the scope of your knowledge with network-based firewalls. Network-based firewalls are primarily used to section off and protect one network from another. They are a primary line of defense and are *extremely* important in network security. There are several types of firewalls; some run as software on server computers, some as standalone dedicated appliances, and some work as just one function of many on a single device. They are commonly represented as a sort of "brick wall" between a LAN and the Internet, as shown in Figure 8-1.

Key Topic

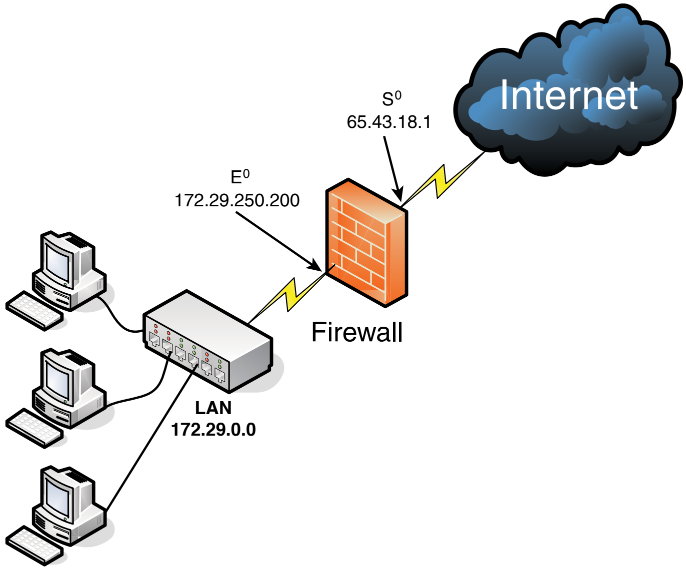

Figure 8-1 Diagram of a Basic Firewall Implementation

Just as a firewall in a physical building is there to slow the spread of a fire and contain it until the fire department arrives, a firewall in a computer network is there to keep fire at bay in the form of malicious attacks. Often, a firewall (or the device the firewall resides on) has NAT in operation as well. In Figure 8-1, note that the firewall has a local address of 172.29.250.200; this connects it to the LAN. It also has an Internet address of 65.43.18.1, enabling connectivity for the entire LAN to the Internet, while hiding the LAN IP addresses. By default, the IP address 65.43.18.1 is completely shielded. This means that all inbound ports are effectively closed and will not enable incoming traffic, unless a LAN computer initiates a session with another system on the Internet. However, a good security administrator always checks this to make sure; first, by accessing the firewall's firmware (or software application, as the case may be) and verifying that the firewall is on, and next by scanning the firewall with third-party applications such as Nmap (https://nmap.org) or with a web-based port scanning utility, as was shown in a Chapter 7 Real-world Scenario. If any ports are open, or unshielded, they should be dealt with immediately. Then the firewall should be rescanned for vulnerabilities. You can find more information on port scanning and vulnerability assessments in Chapter 12.

Important point: Firewalls should be used only as they were intended. The company firewall should not handle any other extraneous services—for example, acting as a web server or SMTP server. By using a firewall as it was intended, its vulnerability is reduced.

Generally, a firewall inspects traffic that passes through it and permits or denies that traffic based on rules set by an administrator. These rules are stored within **access control lists** (ACLs). In regards to firewalls, an ACL is a set of rules that applies to a list of network names, IP addresses, and port numbers. These rules can be configured to control inbound and outbound traffic. This is a bit different than ACLs with respect to operating systems, which we cover in Chapter 11, "Access Control Methods and Models," but the same basic principles apply: Basically, one entity is granted or denied permission to another entity. If you decide that a specific type of traffic should be granted access to your network, you would **explicitly allow** that traffic as a rule within an ACL. If on the other hand you decide that a specific type of traffic should *not* be granted access, you would **explicitly deny** that traffic within an ACL. And finally, if a type of network traffic is not defined in the firewall's rule set, it should be stopped by default. This is the concept of **implicit deny** and is usually a default rule found in a firewall's ACL. It is often added automatically to the end of a firewall's rule set (ACLs) and is also known as "block all."

Firewall rules should be specific. Here's an example of a firewall rule:

```
deny TCP any any port 53
```

This rule can be used to restrict DNS zone transfers (as they run on top of TCP and use port 53), but other DNS traffic will still function properly. The rule is specific; it gives the transport layer protocol to be filtered, and the exact port, and also states that it applies to *any* computer's IP address on the inbound and outbound side. Be careful with firewall rules and ACLs; they need to be written very cautiously so as not to filter required traffic.

> **NOTE** Traffic can also be passed to other computers and servers, or to specific ports. For a quick tutorial on setting up virtual servers and port forwarding on a typical SOHO router/firewall, see the following link: http://www.davidlprowse.com/articles/?p=916.

A lot of today's firewalls have two types of firewall technologies built into them: SPI and NAT. However, you also should be aware of a couple other types of firewall methodologies:

Key Topic

- **Packet filtering:** Inspects each packet passing through the firewall and accepts or rejects it based on rules. However, there are two types: stateless packet inspection and **stateful packet inspection** (also known as SPI or a stateful firewall). A stateless packet filter, also known as pure packet filtering, does not retain memory of packets that have passed through the firewall; due to this, a stateless packet filter can be vulnerable to IP spoofing attacks. But a firewall running stateful packet inspection is normally not vulnerable to this because it keeps track of the state of network connections by examining the header in each packet. It can distinguish between legitimate and illegitimate packets. This function operates at the network layer of the OSI model.

- **NAT filtering:** Also known as NAT endpoint filtering, filters traffic according to ports (TCP or UDP). This can be done in three ways: by way of basic endpoint connections, by matching incoming traffic to the corresponding outbound IP address connection, or by matching incoming traffic to the corresponding IP address and port.
- **Application-level gateway (ALG):** Applies security mechanisms to specific applications, such as FTP or BitTorrent. It supports address and port translation and checks whether the type of application traffic is allowed. For example, your company might allow FTP traffic through the firewall, but might decide to disable Telnet traffic (probably a wise choice). The ALG checks each type of packet coming in and discards Telnet packets. Although this adds a powerful layer of security, the price is that it is resource-intensive, which could lead to performance degradation.
- **Circuit-level gateway:** Works at the session layer of the OSI model, and applies security mechanisms when a TCP or UDP connection is established; it acts as a go-between for the transport and application layers in TCP/IP. After the connection has been made, packets can flow between the hosts without further checking. Circuit-level gateways hide information about the private network, but they do not filter individual packets.

A firewall can be set up in several different physical configurations. For example, in Chapter 6, “Network Design Elements,” we discussed implementing a DMZ. This could be done in a back-to-back configuration (two firewalls surrounding the DMZ), as shown in Figure 8-2, or as a 3-leg perimeter configuration.

Key Topic

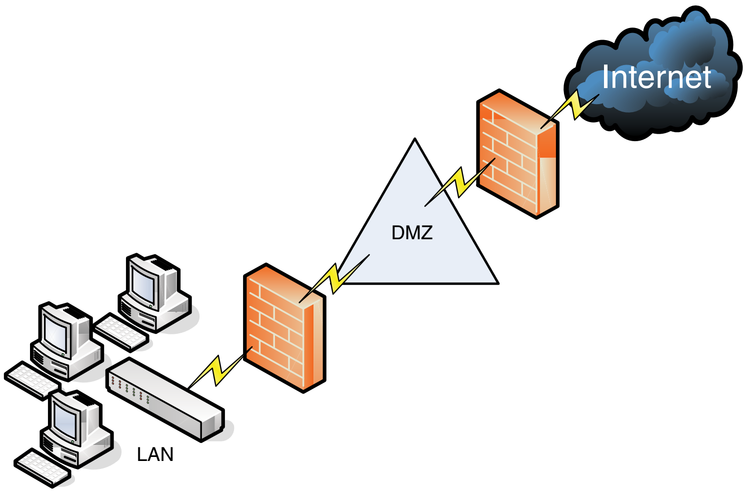

Figure 8-2 Back-to-Back Firewall/DMZ Configuration

Generally, there will be one firewall with the network and all devices and computers residing "behind" it. By the way, if a device is "behind" the firewall, it is also considered to be "after" the firewall, and if the device is "in front of" the firewall, it is also known as being "before" the firewall. Think of the firewall as the drawbridge of a castle. When you are trying to gain admittance to the castle, the drawbridge will probably be closed. You would be in front of the drawbridge, and the people inside the castle would be behind the drawbridge. This is a basic analogy but should help you to understand the whole "in front of" and "behind" business as it relates to data attempting to enter the network and devices that reside on your network.

Logging is also important when it comes to a firewall. Firewall logs should be the first thing you check when an intrusion has been detected. You should know how to access the logs and how to read them. For example, Figure 8-3 shows two screen captures: The first displays the Internet sessions on a basic SOHO router/firewall, and the second shows log events such as blocked packets. Look at the blocked Gnutella packet that is pointed out. I know it is a Gnutella packet because the inbound port on my firewall that the external computer is trying to connect to shows as port 6346; this associates with Gnutella. Gnutella is an older P2P file-sharing network. None of the computers on this particular network use or are in any way connected to the Gnutella service. These external computers are just random clients of the Gnutella P2P network trying to connect to anyone possible.

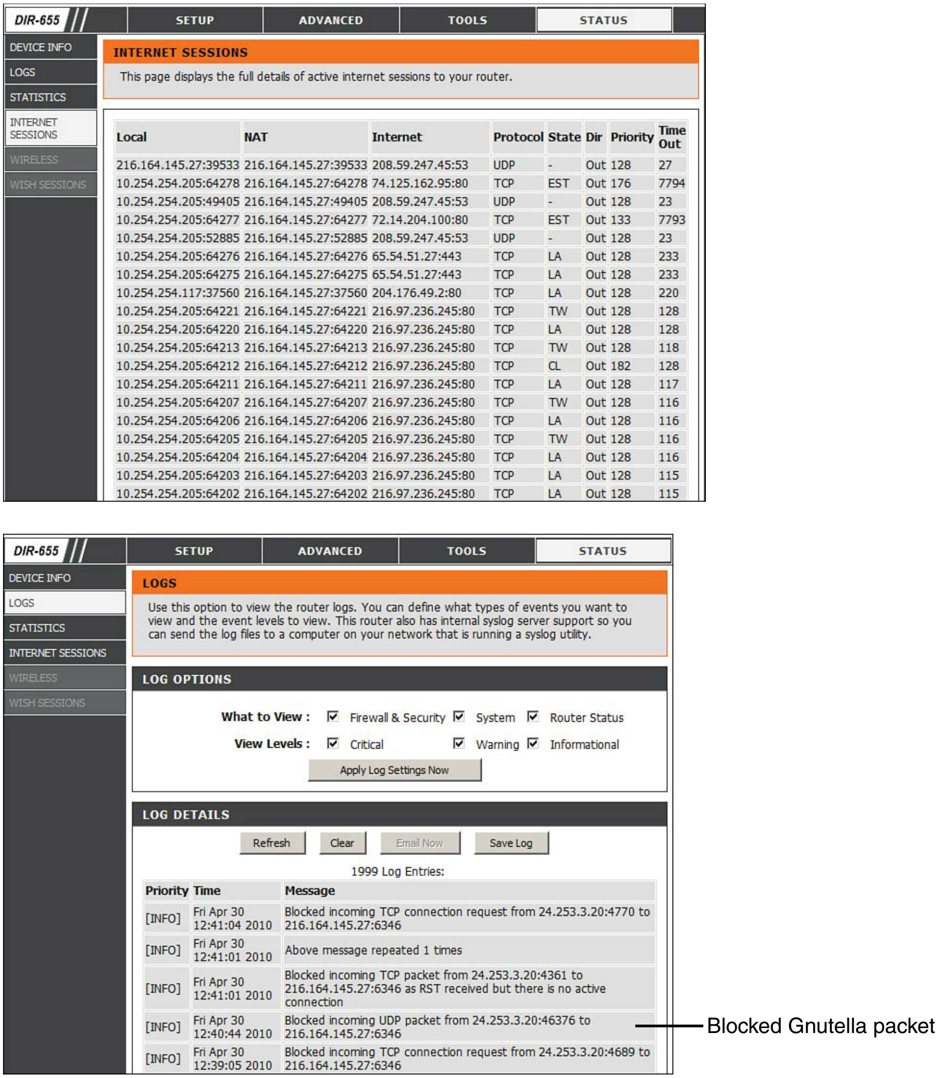

Figure 8-3 SOHO Router/Firewall Internet Sessions

It's good that these packets have been blocked, but maybe you don't want the IP address shown (24.253.3.20) to have any capability to connect to your network at all. To eliminate that IP, you could add it to an inbound filter or to an ACL.

So far, we have discussed host-based firewalls (in Chapter 3) and, just now, network-based firewalls. However, both of these firewalls can also fall into the category of **application firewall**. If either type runs protocols that operate on the application layer of the OSI model, then it can be classified as an application firewall. That means that it can control the traffic associated with specific applications. This is something a stateful network firewall cannot do, as this function operates at the application layer of the OSI model. Many host-based firewalls fall into this category, but when it comes to network-based firewalls, it varies. A basic SOHO router with built-in firewalling capabilities would usually not fall into the application firewall category. However, more advanced network appliances from companies such as Barracuda, Citrix, Fortinet, and Smoothwall do fall into this category. This means that they allow for more in-depth monitoring of the network by controlling the input, output, and access to applications and services all the way up through the application layer of the OSI model. These appliances might also be referred to as *network-based application layer firewalls*. Now that's a mouthful—just be ready for multiple terms used by companies and technicians.

Going a step further, some of the aforementioned network appliances have tools that are designed to specifically protect HTTP sessions from XSS attacks and SQL injection. These types of tools are known as **web application firewalls**. WAFs can help to protect the servers in your environment.

NOTE A firewall appliance needs more than one network adapter so that it can connect to more than one network; this is known as a *multihomed connection*. It might be dual-homed (two adapters), or perhaps it has more, maybe three network adapters, in case you want to implement a DMZ or another perimeter security technique.

Firewalls are often considered to be all-in-one devices, but actually they provide specific functionality as discussed in this section. Still, it is common to hear people refer to a firewall when they are really talking about another technology, or even another device. For example, many SOHO users have an all-in-one multifunction network device. This device has four ports for wired connections, plus a wireless antenna; it connects all the computers to the Internet, and finally has a firewall built-in. Because some users consider this to be simply a firewall, you should teach them about the benefits of disabling SSID broadcasting, and enabling MAC filtering. By disabling Service Set Identifier (SSID) broadcasting, the average user cannot connect wirelessly to the device. An attacker knows how to bypass this, but it is an important element of security that you should implement after all trusted computers have been connected wirelessly. MAC filtering denies access to any computer that does not have one of the MAC addresses you list, another powerful tool that we will cover more in Chapter 9, "Securing Network Media and Devices."

To make matters a bit more confusing, a firewall can also act as, or in combination with, a proxy server, which we discuss in the following section.

Proxy Servers

A **proxy server** acts as an intermediary for clients, usually located on a LAN, and the servers that they want to access, usually located on the Internet. By definition, *proxy* means go-between, or mediator, acting as such a mediator in between a private network and a public network. The proxy server evaluates requests from clients and, if they meet certain criteria, forwards them to the appropriate server. There are several types of proxies, including a couple you should know for the exam:

Key Topic

- **IP proxy:** Secures a network by keeping machines behind it anonymous; it does this through the use of NAT. For example, a basic four-port router can act as an IP proxy for the clients on the LAN it protects. An IP proxy can be the victim of many of the network attacks mentioned in Chapter 6, especially DoS attacks. Regardless of whether the IP proxy is an appliance or a computer, it should be updated regularly, and its log files should be monitored periodically and audited according to organization policies.
- **Caching proxy:** Attempts to serve client requests without actually contacting the remote server. Although there are FTP and SMTP proxies, among others, the most common caching proxy is the **HTTP proxy**, also known as a **web proxy**, which caches web pages from servers on the Internet for a set amount of time. Examples of caching proxies include WinGate (for Windows systems) and Squid (commonly used on Linux-based systems). An example of a caching proxy is illustrated in Figure 8-4. For example, let's say a co-worker of yours (Client A) accessed www.google.com, and that she was the first person to do so on the network. This client request will go through the HTTP proxy and be redirected to Google's web server. As the data for Google's home page comes in, the HTTP proxy will store or cache that information. When another person on your network (Client B) makes a subsequent request for www.google.com, the bulk of that information will come from the HTTP proxy instead of from Google's web server. This is done to save bandwidth on the company's Internet connection and to increase the speed at which client requests are carried out. Most HTTP proxies check websites to verify that nothing has changed since the last request. Because information changes quickly on the Internet, a time limit of 24 hours is common for storing cached information before it is deleted. Web browsers make use of a **proxy auto-configuration (PAC)** file, which defines how the browser can automatically choose a proxy server. The file itself and the embedded JavaScript function pose a security risk in that the file can be exploited and modified, ultimately redirecting the user to unwanted (and potentially malicious) websites. Consider disabling PAC files and auto-configuration in general within client web browsers.

Key Topic

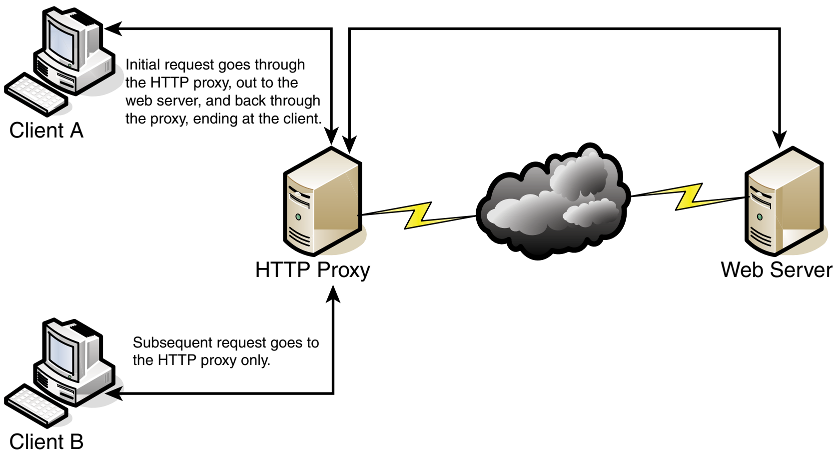

Figure 8-4 Illustration of an HTTP Proxy in Action

Other types of proxies are available to apply policies, block undesirable websites, audit employee usage, and scan for malware. One device or computer might do all these things or just one or two. It depends on the software used or appliance installed. Reverse proxies can also be implemented to protect a DMZ server's identity or to provide authentication and other secure tasks. This is done when users on the Internet are accessing server resources on your network. Generally, a proxy server has more than one network adapter so that it can connect to the various networks it is acting as a mediator for. Each of the network adapters in a proxy should be periodically monitored for improper traffic and for possible network attacks and other vulnerabilities. A proxy server might be the same device as a firewall, or it could be separate. Because of this, a multitude of network configurations are possible. Proxy servers, especially HTTP proxies, can be used maliciously to record traffic sent through them; because most of the traffic is sent in unencrypted form, this could be a security risk. A possible mitigation for this is to chain multiple proxies together in an attempt to confuse any onlookers and potential attackers.

Most often, a proxy server is implemented as a *forward proxy*. This means that clients looking for websites, or files via an FTP connection, pass their requests through to the proxy. However, there is also a *reverse proxy*, where *multiple* HTTP or FTP servers use a proxy server and send out content to one or more clients. These HTTP and FTP servers could be located in a server farm or similar grouping, and the reverse proxy might also undertake the role of load balancer in this situation. A reverse proxy can act as another layer of defense for an organization's FTP or HTTP servers. An *application proxy* might be used as a reverse proxy; for example, Microsoft's Web Application Proxy, which enables remote users to connect to the organization's internal network to access multiple servers. These are often multipurpose by design, allowing for HTTP, FTP, e-mail, and other types of data connections. However, it could be that you have a single application stored on several servers. Those servers can work together utilizing clustering technology. The clustering might be controlled by the servers themselves or, more commonly, a load balancer can be installed in front of the servers that distributes the network load among them. That load balancer in effect acts as a reverse proxy.

Regardless of the type of proxy used, it will often modify the requests of the “client computer,” whatever that client is, providing for a level of anonymity. But in some cases, you might need a proxy that does not modify requests. This is known as a *transparent proxy*. While it allows for increased efficiency, there is less protection for the client system.

Another example of a proxy in action is Internet content filtering. An **Internet content filter**, or simply a content filter, is usually applied as software at the application layer and can filter out various types of Internet activities such as websites accessed, e-mail, instant messaging, and more. It often functions as a content inspection device, and disallows access to inappropriate web material (estimated to be a big percentage of the Internet!) or websites that take up far too much of an organization’s Internet bandwidth. Internet content filters can be installed on individual clients, but by far the more efficient implementation is as an individual proxy that acts as a mediator between all the clients and the Internet. These proxy versions of content filters secure the network in two ways: one, by forbidding access to potentially malicious websites, and two, by blocking access to objectionable material that employees might feel is offensive. It can also act as a URL filter; even if employees inadvertently type an incorrect URL, they can rest assured that any objectionable material will not show up on their display.

Internet filtering appliances analyze just about all the data that comes through them, including Internet content, URLs, HTML tags, metadata, and security certificates such as the kind you would automatically receive when going to a secure site that starts with https. (However, revoked certificates and certificate revocation lists, or CRLs, will not be filtered because they are only published periodically. More on certificates and CRLs is provided in Chapter 15, “PKI and Encryption Protocols.”) Some of these appliances are even capable of malware inspection. Another similar appliance is the web security gateway. **Web security gateways** (such as Forcepoint, previously known as Websense) act as go-between devices that scan for viruses, filter content, and act as data loss prevention (DLP) devices. This type of content inspection/content filtering is accomplished by actively monitoring the users’ data streams in search of malicious code, bad behavior, or confidential data that should not be leaked outside the network.

As you can see, many, many options for security devices are available for your network, and many vendors offer them. Based on price, you can purchase all kinds of devices, from ones that do an individual task, to ones that are combinations of everything we spoke about so far, which are also known as *all-in-one security appliances* or unified threat management (UTM) devices (discussed in the upcoming "NIDS Versus NIPS" section).

NOTE Proxies, content filters, and web security gateways are examples of servers that probably face the Internet directly. These "Internet-facing servers" require security controls before they are installed. The two most important security controls are to keep the application up to date, and to review and apply vendor-provided hardening documentation. Remember to do these things before putting the proxy server (or other Internet-facing servers) in a live environment.

Honeypots and Honeynets

Honeypots and honeynets attract and trap potential attackers to counteract any attempts at unauthorized access of the network. This isolates the potential attacker in a monitored area and contains dummy resources that look to be of value to the perpetrator. While an attacker is trapped in one of these, their methods can be studied and analyzed, and the results of those analyses can be applied to the general security of the functional network.

A **honeypot** is generally a single computer but could also be a file, group of files, or an area of unused IP address space, whereas a **honeynet** is one or more computers, servers, or an area of a network; a honeynet is used when a single honeypot is not sufficient. Either way, the individual computer, or group of servers, will *usually* not house any important company information. Various analysis tools are implemented to study the attacker; these tools, along with a centralized group of honeypots (or a honeynet), are known collectively as a honeyfarm.

One example of a honeypot in action is the spam honeypot. Spam e-mail is one of the worst banes known to a network administrator; a spam honeypot can lure spammers in, enabling the network administrators to study the spammers' techniques and habits, thus allowing the network admins to better protect their actual e-mail servers, SMTP relays, SMTP proxies, and so on, over the long term. It might ultimately keep the spammers away from the real e-mail addresses, because the spammers are occupied elsewhere. Some of the information gained by studying spammers is often shared with other network admins or organizations' websites dedicated to reducing spam. A spam honeypot could be as simple as a single e-mail address or as complex as an entire e-mail domain with multiple SMTP servers.

Of course, as with any technology that studies attackers, honeypots also bear risks to the legitimate network. The honeypot or honeynet should be carefully firewalled off from the legitimate network to ensure that the attacker can't break through.

Often, honeypots and honeynets are used as part of a more complex solution known as a network intrusion detection system, discussed following a short review of data loss prevention.

Data Loss Prevention (DLP)

We mentioned DLP in Chapter 3. Let's discuss it briefly now as it relates to networks. **Data loss prevention (DLP)** systems are designed to protect data by way of content inspection. They are meant to stop the leakage of confidential data, often concentrating on communications. As such, they are also referred to as data leak prevention (DLP) devices, information leak prevention (ILP) devices, and extrusion prevention systems. Regardless, they are intended to be used to keep data from leaking past a computer system or network and into unwanted hands.

In network-based DLP, systems deal with data in motion and are usually located on the perimeter of the network. If data is classified in an organization's policy as confidential and not to be read by outsiders, the DLP system detects it and prevents it from leaving the network. Network-based DLP systems can be hardware-based or software-based. An example of a network-based DLP system would be one that detects and prevents the transfer of confidential e-mail information outside the network. Organizations such as Check Point offer DLP solutions, and there are some free open source applications as well. Going further, there are cloud-based DLP solutions available. But it all depends on where you store your data. If you store some or all of your data on the cloud, or if you have a large bring your own device (BYOD) or choose your own device (CYOD) population, then cloud-based DLP becomes an important part of your security strategy. Because the data—and the security of that data—is now external from the company, planning becomes even more vital. Some key elements of the security mindset include: 1) planning for the mitigation of security risks; 2) adequate understanding of the cloud-based provider, where and how data is stored, and their service-level agreement (SLA); 3) in-depth analysis of code and the types of data that will be stored in the cloud; and 4) strong authentication, auditing, and logging. If all this is planned for and implemented properly, it can build the organization's confidence in the cloud, which can lead to a smoother transition, and ultimately reduce risk. However, all this becomes a bigger conversation: We'll talk more about general mindsets when dealing with cloud-based companies in Chapter 16, "Redundancy and Disaster Recovery," and Chapter 18, "Policies and Procedures."

As for DLP, the monitoring of possible leaked information could become a privacy concern. Before implementing a system of this nature, it is important to review your organization's privacy policies. Leaks can still occur due to poor implementation of DLP systems, so it is essential to plan what type of DLP solution your organization needs, exactly how it will be installed, and how it will be monitored.

NIDS Versus NIPS

It's not a battle royale, but you should be able to differentiate between a network intrusion *detection* system (NIDS) and a network intrusion *prevention* system (NIPS) for the exam. Previously, in Chapter 4, "OS Hardening and Virtualization," we discussed host-based intrusion detection systems (or HIDSs). Although a great many attacks can hamper an individual computer, just as many network attacks could possibly take down a server, switch, router, or even an entire network. Network-based IDSs were developed to detect these malicious network attacks, and network-based IPSs were developed in an attempt to prevent them.

NIDS

A **network intrusion detection system (NIDS)** by definition is a type of IDS that attempts to detect malicious network activities, for example, port scans and DoS attacks, by constantly monitoring network traffic. It can also be instrumental in rogue machine detection, including rogue desktops, laptops, and mobile devices, as well as rogue access points, DHCP servers, and network sniffers. Examples of NIDS solutions include open-source products such as Snort (https://www.snort.org/), Bro (https://www.bro.org/), and many other commercial hardware and software-based products. A NIDS should be situated at the entrance or gateway to your network. It is not a firewall but should be used with a firewall. Because the NIDS inspects every packet that traverses your network, it needs to be fast; basically, the slower the NIDS, the slower the network. So, the solution itself, the computer/device it is installed on, and the network connections of that computer/device all need to be planned out accordingly to ensure that the NIDS does not cause network performance degradation.

Figure 8-5 illustrates how a NIDS might be implemented on a network. Often it is placed in front of a firewall. The NIDS detects attacks and anomalies and alerts the administrator if they occur, whereas the firewall does its best to prevent those attacks from entering the network. However, a NIDS could be placed behind the firewall, or you might have multiple NIDS points strategically placed around the network. If the NIDS is placed in front of the firewall, it generates a lot more administrator alerts, but these can usually be whittled down within the firmware or software of the device running the NIDS. Regardless of where the NIDS is located, a network administrator should monitor traffic from time to time; to do so, the computer, server, or appliance that has the NIDS installed should have a network adapter configured to work in **promiscuous mode**. This passes all traffic to the CPU, not just the frames addressed to it.

Key Topic

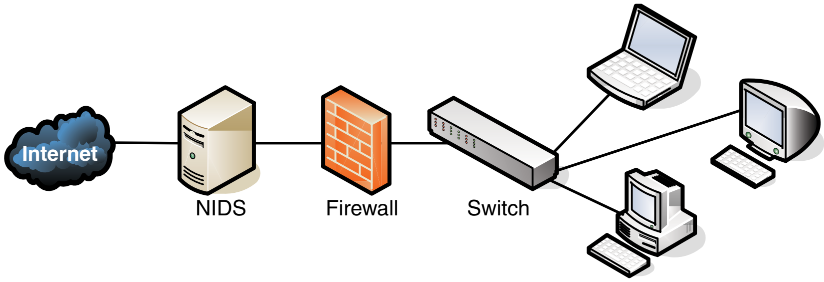

Figure 8-5 Illustration of NIDS Placement in a Network

The beauty of a NIDS is that you might get away with one or two NIDS points on the network, and do away with some or all the HIDS installed on individual computers, effectively lowering the bottom line while still doing a decent job of mitigating risk. A couple of disadvantages of a NIDS, aside from possible network performance issues, are that it might not be able to read encrypted packets of information and will not detect problems that occur on an individual computer. Therefore, to secure a network and its hosts, many organizations implement a mixture of NIDS and HIDS. If a NIDS is placed in front of the firewall, it is subject to attack; therefore, it should be monitored and updated regularly. Some NIDS solutions will auto-update. Finally, the biggest disadvantage of a NIDS is that it is passive, meaning it only *detects* attacks; to protect against, or *prevent*, these attacks, you need something active, you need a NIPS.

NIPS

A **network intrusion prevention system (NIPS)** is designed to inspect traffic and, based on its configuration or security policy, either remove, detain, or redirect malicious traffic that it becomes aware of. The NIPS (as well as the NIDS) is considered to be an *application-aware device*, meaning it can divine different types of packets, define what application they are based on, and ultimately permit or disallow that traffic on the network. More and more companies are offering NIPS solutions in addition to, or instead of, NIDS solutions. Examples of NIPS solutions include Check Point security appliances (https://www.checkpoint.com), and the aforementioned Snort, which is actually a NIDS/NIPS software package that should be installed on a dual-homed or multihomed server. Not only can a NIPS go above and beyond a NIDS by removing or redirecting malicious traffic, it can also redirect a recognized attacker to a single computer known as a padded cell, which contains no information of value and has no way out.

Like a NIDS, a NIPS should sit inline on the network, often in front of the firewall, although it could be placed elsewhere, depending on the network segment it protects and the network architecture. Whereas many NIPS solutions have two connections only and are known as perimeter solutions, other NIPS appliances have up to 16 ports enabling many points of detection on the network—these would be known as network "core" devices. Regardless of the solution you select, as packets pass through the device, they are inspected for possible attacks. These devices need to be accurate and updated often (hopefully automatically) to avoid the misidentification of legitimate traffic, or worse, the misidentification of attacks. If the NIPS blocks legitimate traffic, it would be known as a **false positive**, and effectively could deny service to legitimate customers, creating a self-inflicted denial-of-service of sorts.

If the IPS does not have a particular attack's signature in its database, and lets that attack through thinking it is legitimate traffic, it is known as a **false negative**, also bad for obvious reasons! Many IPS systems can monitor for attack signatures and anomalies. More information on signatures can be found in Chapter 4 and Chapter 13, "Monitoring and Auditing." Another type of error that can occur with NIDS and NIPS is a subversion error; this is when the NIDS/NIPS has been altered by an attacker to allow for false negatives, ultimately leading to attacks creeping into the network. This can be deadly because the NIDS/NIPS often is the first point of resistance in the network. To protect against this, some devices have the capability to hide or mask their IP address. They might also come with an internal firewall. It is also important to select an IPS solution that has a secure channel for the management console interface.

One advantage of newer NIPS solutions is that some of them can act as protocol analyzers by reading encrypted traffic and stopping encrypted attacks. In general, the beauty of a NIPS compared to a host-based IPS (HIPS) is that it can protect non-computer-based network devices such as switches, routers, and firewalls. However, the NIPS is considered a single point of failure because it sits inline on the network. Due to this, some organizations opt to install a bypass switch, which also enables the NIPS to be taken offline when maintenance needs to be done.

A vital NIPS consideration is whether to implement a fail-close or fail-open policy—in essence, deciding what will happen if the NIPS fails. Fail-close means that all data transfer is stopped, while fail-open means that data transfer (including potential attacks) are passed through. Let's consider an example. Say that the NIPS was protecting an individual server (or router), and had a certain level of control over that system. Now let's say that the NIPS failed. In a fail-close scenario, it would disconnect the system that it is protecting, stopping all data transfer. This is unacceptable to some organizations that require near 100 percent uptime. These organizations are willing to accept additional risk, and therefore are more receptive to a fail-open scenario. However, in this case, if the NIPS fails, it continues to pass all traffic to the "protected" system, which could include possible attacks. Sometimes, fail-open scenarios are necessary. In these cases, defense in depth is the best strategy. For instance, you might opt to have a firewall filter the bulk of traffic coming into the network, but have the IPS filter only specific traffic, reducing the chances of IPS failure. This layered approach can offer greater security with less chance of attacks passing through, but often comes with increased cost and administration.

Summary of NIDS Versus NIPS

Table 8-1 summarizes NIDS versus NIPS.

Key Topic

Table 8-1 Summary of NIDS Versus NIPS

Type of System	Summary	Disadvantage/Advantage	Example
NIDS	Detects malicious network activities	Pro: Only a limited number of NIDSs are necessary on a network. Con: Only detects malicious activities.	Snort Bro IDS
NIPS	Detects, removes, detains, and redirects traffic	Pro: Detects and mitigates malicious activity. Pro: Can act as a protocol analyzer. Con: Uses more resources. Con: Possibility of false positives and false negatives.	Check Point Systems solutions

The Protocol Analyzer's Role in NIDS and NIPS

You might be familiar already with protocol analyzers such as Wireshark (previously Ethereal) or Network Monitor. These are loaded on a computer and are controlled by the user in a GUI environment; they capture packets, enabling the user to analyze them and view their contents. However, some NIDS/NIPS solutions are considered to be full protocol analyzers with no user intervention required. The protocol analyzer is built into the NIDS/NIPS appliance. It decodes application layer protocols, such as HTTP, FTP, or SMTP, and forwards the results to the IDS or IPS analysis engine. Then the analysis engine studies the information for anomalous or behavioral exploits. This type of analysis can block many exploits based on a single signature. This is superior to basic signature pattern recognition (without protocol analysis), because with signature-based IDS/IPS solutions, many signatures have to be constantly downloaded and stored in the device's database, and they don't enable dynamic understanding of new attacks. However, as with any powerful analysis, like protocol analysis, a premium is placed on processing power, and the price of these types of IDS/IPS solutions will undoubtedly be higher.

NOTE There are also wireless versions of IDS: WIDS and WIPS. They monitor the radio spectrum for unauthorized access and rogue access points. However, these names might be incorporated into the concept of NIDS and NIPS by some organizations. Regardless, be sure to use an IDS (or IPS) for your wired and wireless connections!

Unified Threat Management

A relatively newer concept, **unified threat management (UTM)** is the culmination of everything we discussed in this chapter so far. As early as the year 2000, it was realized that the firewall was no longer enough to protect an organization's network. Other devices and technologies such as NIDS/NIPS systems, content filters, anti-malware gateways, data leak prevention, and virtual private networks were added to the network in order to better protect it. However, with all these extra devices and technologies come added cost and more administration. And so, UTM providers simplify the whole situation by offering all-in-one devices that combine the various levels of defense into one solution. The all-in-one device might also be referred to as a next-generation firewall (NGFW). Companies such as Cisco, Fortinet, and Sophos (to name a few) offer UTM and NGFW solutions; often this is a single device that sits last on the network before the Internet connection. They usually come with a straightforward web-based GUI, which is good news for the beleaguered security administrator who might be burning the midnight oil researching the latest attacks and prevention methods. There's a caveat to all this, and it is a common theme in network security: a single point of defense is a single point of failure. Get past the UTM, and your job as an attacker is done. Secondary and backup UTM devices, as well as server-based HIDSs, strike a balance and create a certain level of defense in depth, while still retaining a level of simplicity. Another consideration is that UTMs should be quick. If they are to take the place of several other devices, then their data processing and traffic flow requirements will be steep. The smart network administrator/security administrator will consider a device that exceeds their current needs and then some.

It was important to discuss each of the tools and technologies separately in this chapter so that you understand how to work with each. But keep in mind that many of these technologies are consolidated into a single solution, a trend that will likely continue as we move forward.

Chapter Summary

Well, it goes without saying that there are many potential attackers who would "storm the castle." The question presents itself: Have you performed your due diligence in securing your computer networking kingdom?

If you answered yes, then it most likely means you have implemented some kind of unified threat management solution; one that includes a firewall, content filter, anti-malware technology, IDS/IPS, and possibly other network security technologies. This collaborative effort makes for a strong network perimeter. The firewall is at the frontlines, whether it is part of a UTM or running as a separate device. Its importance can't be stressed enough, and you can't just implement a firewall; it has to be configured properly with your organization's policies in mind. ACLs, stateful packet inspection, and network address translation should be employed to solidify your firewall solution.

If you answered no, then prepare ye for more metaphorical expression. Remember that enemy forces are everywhere. They are lying in wait just outside your network, and they can even reside within your network—for example, the malicious insider, that dragon who has usurped the mountain and is perhaps in control of your precious treasure...your data. Analogies aside, this is all clear and present danger—it is *real*, and should be enough to convince you to take strong measures to protect your network.

Often, the act of securing the network can also provide increased efficiency and productivity. For example, a proxy server can act to filter content, and can provide anonymity, but also saves time and bandwidth for commonly accessed data. A honeypot can trap an attacker, thus securing the network, but the secondary result is that network bandwidth is not gobbled up by the powerful attacker. However, the same act can have the opposite effect. For example, a NIDS that is installed to detect anomalies in packets can slow down the network if it is not a powerful enough model. For increased efficiency (and lower all-around cost), consider an all-in-one device such as a UTM, which includes functionality such as firewalling, IDS/IPS, AV, VPN, and DLP. Just make sure it has the core processing and memory required to keep up with the amount of data that will flow through your network.

If you can find the right balance of security and performance while employing your network security solution, it will be analogous to your network donning the aegis, acting as a powerful shield against network attacks from within and without.

Chapter Review Activities

Use the features in this section to study and review the topics in this chapter.

Review Key Topics

Review the most important topics in the chapter, noted with the Key Topic icon in the outer margin of the page. Table 8-2 lists a reference of these key topics and the page number on which each is found.

Key Topic

Table 8-2 Key Topics for Chapter 8

Key Topic Element	Description	Page Number
Figure 8-1	Diagram of a basic firewall	257
Bulleted list	Types of firewalls	258
Figure 8-2	Back-to-back firewall/DMZ configuration	260
Bulleted list	Types of proxies	263
Figure 8-4	Illustration of an HTTP proxy in action	264
Figure 8-5	Illustration of NIDS placement in a network	269
Table 8-1	Summary of NIDS versus NIPS	271

Define Key Terms

Define the following key terms from this chapter, and check your answers in the glossary:

network perimeter, access control list, explicit allow, explicit deny, implicit deny, packet filtering, stateful packet inspection, application-level gateway, circuit-level gateway, application firewall, web application firewall, proxy server, IP proxy, HTTP proxy (web proxy), proxy auto-configuration (PAC), Internet content filter, web security gateway, honeypot, honeynet, data loss prevention (DLP), network intrusion detection system (NIDS), promiscuous mode, network intrusion prevention system (NIPS), false positive, false negative, unified threat management (UTM)

Complete the Real-World Scenarios

Complete the Real-World Scenarios found on the companion website (www.pearsonitcertification.com/title/9780789758996). You will find a PDF containing the scenario and questions, and also supporting videos and simulations.

Review Questions

Answer the following review questions. Check your answers with the correct answers that follow.

1. Which tool would you use if you want to view the contents of a packet?
   - **A.** TDR
   - **B.** Port scanner
   - **C.** Protocol analyzer
   - **D.** Loopback adapter

2. The honeypot concept is enticing to administrators because
   - **A.** It enables them to observe attacks.
   - **B.** It traps an attacker in a network.
   - **C.** It bounces attacks back at the attacker.
   - **D.** It traps a person physically between two locked doors.

3. James has detected an intrusion in his company network. What should he check first?
   - **A.** DNS logs
   - **B.** Firewall logs
   - **C.** The Event Viewer
   - **D.** Performance logs

4. Which of the following devices should you employ to protect your network? (Select the best answer.)
   - **A.** Protocol analyzer
   - **B.** Firewall
   - **C.** DMZ
   - **D.** Proxy server

5. Which device's log file will show access control lists and who was allowed access and who wasn't?
   - **A.** Firewall
   - **B.** Smartphone
   - **C.** Performance Monitor
   - **D.** IP proxy

6. Where are software firewalls usually located?
   - **A.** On routers
   - **B.** On servers
   - **C.** On clients
   - **D.** On every computer

7. Where is the optimal place to have a proxy server?
   - **A.** In between two private networks
   - **B.** In between a private network and a public network
   - **C.** In between two public networks
   - **D.** On all of the servers

8. A coworker has installed an SMTP server on the company firewall. What security principle does this violate?
   - **A.** Chain of custody
   - **B.** Use of a device as it was intended
   - **C.** Man trap
   - **D.** Use of multifunction network devices

9. You are working on a server and are busy implementing a network intrusion detection system on the network. You need to monitor the network traffic from the server. What mode should you configure the network adapter to work in?
   - **A.** Half-duplex mode
   - **B.** Full-duplex mode
   - **C.** Auto-configuration mode
   - **D.** Promiscuous mode

10. Which of the following displays a single public IP address to the Internet while hiding a group of internal private IP addresses?
    - **A.** HTTP proxy
    - **B.** Protocol analyzer
    - **C.** IP proxy
    - **D.** SMTP proxy
    - **E.** PAC

11. If your ISP blocks objectionable material, what device would you guess has been implemented?
    - A. Proxy server
    - B. Firewall
    - C. Internet content filter
    - D. NIDS

12. Of the following, which is a collection of servers that was set up to attract attackers?
    - A. DMZ
    - B. Honeypot
    - C. Honeynet
    - D. VLAN

13. Which of the following will detect malicious packets and discard them?
    - A. Proxy server
    - B. NIDS
    - C. NIPS
    - D. PAT

14. Which of the following will an Internet filtering appliance analyze? (Select the three best answers.)
    - A. Content
    - B. Certificates
    - C. Certificate revocation lists
    - D. URLs

15. Which of the following devices would detect but not react to suspicious behavior on the network? (Select the most accurate answer.)
    - A. NIPS
    - B. Firewall
    - C. NIDS
    - D. HIDS
    - E. UTM

16. One of the programmers in your organization complains that he can no longer transfer files to the FTP server. You check the network firewall and see that the proper FTP ports are open. What should you check next?

- **A.** ACLs
- **B.** NIDS
- **C.** AV definitions
- **D.** FTP permissions

17. Which of the following is likely to be the last rule contained within the ACLs of a firewall?

- **A.** Time of day restrictions
- **B.** Explicit allow
- **C.** IP allow any
- **D.** Implicit deny

18. Which of the following best describes an IPS?

- **A.** A system that identifies attacks
- **B.** A system that stops attacks in progress
- **C.** A system that is designed to attract and trap attackers
- **D.** A system that logs attacks for later analysis

19. What is a device doing when it actively monitors data streams for malicious code?

- **A.** Content inspection
- **B.** URL filtering
- **C.** Load balancing
- **D.** NAT

20. Allowing or denying traffic based on ports, protocols, addresses, or direction of data is an example of what?

- **A.** Port security
- **B.** Content inspection
- **C.** Firewall rules
- **D.** Honeynet

21. Which of the following should a security administrator implement to limit web-based traffic that is based on the country of origin? (Select the three best answers.)

- **A.** AV software
- **B.** Proxy server
- **C.** Spam filter
- **D.** Load balancer
- **E.** Firewall
- **F.** URL filter
- **G.** NIDS

22. You have implemented a technology that enables you to review logs from computers located on the Internet. The information gathered is used to find out about new malware attacks. What have you implemented?

- **A.** Honeynet
- **B.** Protocol analyzer
- **C.** Firewall
- **D.** Proxy

23. Which of the following is a layer 7 device used to prevent specific types of HTML tags from passing through to the client computer?

- **A.** Router
- **B.** Firewall
- **C.** Content filter
- **D.** NIDS

24. Your boss has asked you to implement a solution that will monitor users and limit their access to external websites. Which of the following is the best solution?

- **A.** NIDS
- **B.** Proxy server
- **C.** Block all traffic on port 80
- **D.** Honeypot

25. Which of the following firewall rules only denies DNS zone transfers?

   A. `deny IP any any`

   B. `deny TCP any any port 53`

   C. `deny UDP any any port 53`

   D. `deny all dns packets`

Answers and Explanations

1. **C.** A protocol analyzer has the capability to "drill" down through a packet and show the contents of that packet as they correspond to the OSI model. A TDR is a time-domain reflectometer, a tool used to locate faults in cabling. (I threw that one in for fun. It is a Network+ level concept, so you security people should know it!) A port scanner identifies open network ports on a computer or device; we'll discuss that more in Chapters 12 and 13. A loopback adapter is a device that can test a switch port or network adapter (depending on how it is used).

2. **A.** By creating a honeypot, the administrator can monitor attacks without sustaining damage to a server or other computer. Don't confuse this with a honeynet (answer B), which is meant to attract and trap malicious attackers in an entirely false network. Answer C is not something that an administrator would normally do, and answer D is defining a man trap.

3. **B.** If there was an intrusion, James should check the firewall logs first. DNS logs in the Event Viewer and the performance logs will most likely not show intrusions to the company network. The best place to look first is the firewall logs.

4. **B.** Install a firewall to protect the network. Protocol analyzers do not help to protect a network but are valuable as vulnerability assessment and monitoring tools. Although a DMZ and a proxy server could possibly help to protect a portion of the network to a certain extent, the best answer is firewall.

5. **A.** A firewall contains one or more access control lists (ACLs) defining who is enabled to access the network. The firewall can also show attempts at access and whether they succeeded or failed. A smartphone might list who called or e-mailed, but as of the writing of this book does not use ACLs. Performance Monitor analyzes the performance of a computer, and an IP proxy deals with network address translation, hiding many private IP addresses behind one public address. Although the function of an IP proxy is often built into a firewall, the best answer would be firewall.

6. **C.** Software-based firewalls, such as Windows Firewall, are normally running on the client computers. Although a software-based firewall could also be run on a server, it is not as common. Also, a SOHO router might have a built-in firewall, but not all routers have firewalls.

7. **B.** Proxy servers should normally be between the private network and the public network. This way they can act as a go-between for all the computers located on the private network. This applies especially to IP proxy servers but might also include HTTP proxy servers.

8. **B.** SMTP servers should not be installed on a company firewall. This is not the intention of a firewall device. The SMTP server should most likely be installed within a DMZ.

9. **D.** To monitor the implementation of NIDS on the network, you should configure the network adapter to work in promiscuous mode; this forces the network adapter to pass all the traffic it receives to the processor, not just the frames that were addressed to that particular network adapter. The other three answers have to do with duplexing—whether the network adapter can send and receive simultaneously.

10. **C.** An IP proxy displays a single public IP address to the Internet while hiding a group of internal private IP addresses. It sends data back and forth between the IP addresses by using network address translation (NAT). This functionality is usually built into SOHO routers and is one of the main functions of those routers. HTTP proxies store commonly accessed Internet information. Protocol analyzers enable the capture and viewing of network data. SMTP proxies act as a go-between for e-mail. PAC stands for proxy auto-config, a file built into web browsers that allows the browser to automatically connect to a proxy server.

11. **C.** An Internet content filter, usually implemented as content-control software, can block objectionable material before it ever gets to the user. This is common in schools, government agencies, and many companies.

12. **C.** A honeynet is a collection of servers set up to attract attackers. A honeypot is usually one computer or one server that has the same purpose. A DMZ is the demilitarized zone that is in between the LAN and the Internet. A VLAN is a virtual LAN.

13. **C.** A NIPS, or network intrusion prevention system, detects and discards malicious packets. A NIDS only detects them and alerts the administrator. A proxy server acts as a go-between for clients sending data to systems on the Internet. PAT is port-based address translation.

14. **A, B, and D.** Internet filtering appliances will analyze content, certificates, and URLs. However, certificate revocation lists will most likely not be analyzed. Remember that CRLs are published only periodically.

15. **C.** A NIDS, or network intrusion detection system, will detect suspicious behavior but most likely will not react to it. To prevent it and react to it, you would want a NIPS. Firewalls block certain types of traffic but by default do not check for suspicious behavior. HIDS is the host-based version of an IDS; it checks only the local computer, not the network. A UTM is an all-inclusive security product that will probably include an IDS or IPS—but you don't know which, so you can't assume that a UTM will function in the same manner as a NIDS.

16. **A.** Access control lists can stop specific network traffic (such as FTP transfers) even if the appropriate ports are open. A NIDS will detect traffic and report on it but not prevent it. Antivirus definitions have no bearing on this scenario. If the programmer was able to connect to the FTP server, the password should not be an issue. FTP permissions might be an issue, but since you are working in the firewall, you should check the ACL first; then later you can check on the FTP permissions, passwords, and so on.

17. **D.** Implicit deny (block all) is often the last rule in a firewall; it is added automatically by the firewall, not by the user. Any rules that allow traffic will be before the implicit deny/block all on the list. Time of day restrictions will probably be stored elsewhere but otherwise would be before the implicit deny as well.

18. **B.** An IPS (intrusion prevention system) is a system that prevents or stops attacks in progress. A system that only identifies attacks would be an IDS. A system designed to attract and trap attackers would be a honeypot. A system that logs attacks would also be an IDS or one of several other devices or servers.

19. **A.** A device that is actively monitoring data streams for malicious code is inspecting the content. URL filtering is the inspection of the URL only (for example, https://www.comptia.org). Load balancing is the act of dividing up workload between multiple computers; we'll discuss that more in Chapter 16, "Redundancy and Disaster Recovery." NAT is network address translation, which is often accomplished by a firewall or IP proxy.

20. **C.** Firewall rules (ACLs) are generated to allow or deny traffic. They can be based on ports, protocols, IP addresses, or which way the data is headed. Port security deals more with switches and the restriction of MAC addresses that

are allowed to access particular physical ports. Content inspection is the filtering of web content, checking for inappropriate or malicious material. A honeynet is a group of computers or other systems designed to attract and trap an attacker.

21. **B, E, and F.** The security administrator should implement a proxy server, a firewall, and/or a URL filter. These can all act as tools to reduce or limit the amount of traffic based on a specific country. AV software checks for, and quarantines, malware. Spam filters will reduce the amount of spam that an e-mail address or entire e-mail server receives. A load balancer spreads out the network load to various switches, routers, and servers. A NIDS is used to detect anomalies in network traffic.

22. **A.** A honeynet has been employed. This is a group of computers on the Internet, or on a DMZ (and sometimes on the LAN), that is used to trap attackers and analyze their attack methods, whether they are network attacks or malware attempts. A protocol analyzer captures packets on a specific computer in order to analyze them but doesn't capture logs per se. A firewall is used to block network attacks but not malware. A proxy is used to cache websites and act as a filter for clients.

23. **C.** A content filter is an application layer (layer 7) device that is used to prevent undesired HTML tags, URLs, certificates, and so on, from passing through to the client computers. A router is used to connect IP networks. A firewall blocks network attacks. A NIDS is used to detect anomalous traffic.

24. **B.** You should implement a proxy server. This can limit access to specific websites, and monitor who goes to which websites. Also, it can often filter various HTML and website content. A NIDS is used to report potentially unwanted data traffic that is found on the network. Blocking all traffic on port 80 is something you would accomplish at a firewall, but that would stop all users from accessing any websites that use inbound port 80 (the great majority of them!). A honeypot is a group of computers used to lure attackers in and trap them for later analysis.

25. **B.** The firewall rule listed that only denies DNS zone transfers is `deny TCP any any port 53`. As mentioned in Chapter 7, "Networking Protocols and Threats," DNS uses port 53, and DNS zone transfers specifically use TCP. This rule will apply to any computer's IP address initiating zone transfers on the inbound and outbound sides. If you configured the rule for UDP, other desired DNS functionality would be lost. Denying IP in general would have additional unwanted results. When creating a firewall rule (or ACL), you need to be very specific so that you do not filter out desired traffic.

Index

Numbers

A

C

D

E

F

G

H

I

J - K

M

N

O

P

Q - R

S

U

V

W

X - Y - Z