

Making Everything Easier!™

Sourcefire Edition

Next-Generation IPS

**FOR
DUMMIES®**
A Wiley Brand

Learn:

- What a next-generation IPS can do for your network security
- Why protecting against dynamic threats requires real-time network visibility and agility
- What you need to look for when purchasing a next-generation IPS

Brought to you by

SOURCEfire®

**David Stuart
Kevin Beaver, CISSP**



Sourcefire, now part of Cisco, is a world leader in intelligent cybersecurity solutions. Sourcefire provides a broad portfolio of integrated solutions that deliver unmatched visibility and continuous advanced threat protection across the entire attack continuum, allowing customers to act more quickly – before, during and after an attack. Sourcefire's innovation in open source security, as well as commercial next-generation network security platforms and advanced malware protection solutions has been trusted for more than 10 years. Sourcefire has earned a reputation for innovation, consistent security effectiveness and world-class research all focused on detecting, understanding and stopping threats.

Visit www.sourcefire.com.

Sample Sourcefire Awards & Recognitions:



Leader in Security Value Map for IPS



12 Hot Companies to Watch



Leader in Magic Quadrant for IPS



These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

***Next-Generation
IPS***

FOR

DUMMIES®

SOURCEFIRE EDITION

**by David Stuart and
Kevin Beaver, CISSP**

WILEY

Next-Generation IPS For Dummies®, Sourcefire Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2013 by John Wiley & Sons, Inc., Hoboken, New Jersey

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Sourcefire and the Sourcefire logo are trademarks or registered trademarks of Sourcefire, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-118-55497-5 (pbk); ISBN 978-1-118-55692-4 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Vertical Websites

Project Editor: Jennifer Bingham
Editorial Manager: Rev Mengle
Business Development Representative: Melody Layne
Custom Publishing Project Specialist: Michael Sullivan

Composition Services

Senior Project Coordinator: Kristie Rees
Layout and Graphics: Carl Byers
Proofreader: John Greenough

Business Development

Lisa Coleman, Director, New Market and Brand Development

Table of Contents

Introduction	1
How This Book Is Organized	1
Icons Used in This Book	2
Chapter 1: Understanding Next-Generation IPS	3
Defining Intrusion Prevention Systems	4
Elements of a Next-Generation IPS	5
Protection across the Network Stack	8
Common Network Threats	10
Spyware, Phishing, and Botnets	12
SYN Floods and Denial of Service (DoS) Attacks	14
Zero-Day Attacks	16
Advanced Persistent Threats (APTs)	17
Chapter 2: Know What You're Protecting	19
Real-Time Visibility Is Key	19
Understand Application Risks	22
Chapter 3: Next-Generation IPS Features	27
Common Functions of Any IPS	27
Protection Beyond a Typical IPS	28
Central Role-Based Management	34
Management Reporting	35
Maintaining Performance as You Grow	35
Lower TCO through IPS Automation	37
Chapter 4: Next-Generation IPSs, Virtualization, and Cloud Computing	39
Benefits and Risks of Virtualization	39
Securing Virtualization	42
Virtualizing Security	45
Securing the Cloud	46
Chapter 5: Deploying Next-Generation IPSs	51
Next-Generation IPS Deployment Scenarios	51
Using Integrated Policies to Your Advantage	54

Chapter 6: Selecting the Right Next-Generation IPS	57
Common Next-Generation IPS Selection Criteria	57
Industry-Specific Considerations	63
Hardware Considerations	64
Third-Party Testing	65
Chapter 7: Keys to Effective Next-Generation Security	67

Introduction



This book introduces you to next-generation IPS (NGIPS) solutions and shows how you can harness their power to improve the security in your organization's networks. We explain the benefits they can provide and why we think they are becoming a necessity and not a luxury. We also show you how to select the right next-generation IPS for your business.

How This Book Is Organized

This book is organized so that you don't have to read it cover-to-cover, front to back. You can skip around and read just the chapters that are of interest.

- ✓ In **Chapter 1, Understanding Next-Generation IPS**, we explain the differences between traditional intrusion prevention systems and today's next-generation intrusion prevention systems. We also discuss the complex threats they detect and protect against.
- ✓ **Chapter 2, Know What You're Protecting**, explains why protecting against dynamic threats requires real-time network visibility and agility. We cover awareness of network changes, application risks, and focusing on what's important.
- ✓ In **Chapter 3, Next-Generation IPS Features**, we explain in more detail the many features and functions found in next-generation IPSs, including application visibility/control, context and content awareness (including malware awareness), security automation, and user identification.
- ✓ **Chapter 4, Next-Generation IPSs, Virtualization, and Cloud Computing**, includes in-depth discussions of virtualization and cloud computing technologies, and the role that next-generation IPSs play to protect these new types of environments.

- ✓ In **Chapter 5, Deploying Next-Generation IPSs**, we cover deployment scenarios and considerations to help with your unique environment.
- ✓ **Chapter 6, Selecting the Right Next-Generation IPS**, is all about helping you get your next-generation IPS shopping list organized so that you can be sure to get the right NGIPS for your organization.
- ✓ In **Chapter 7, Keys to Effective Next-Generation Security**, we explain four essential elements for improving your network security and information risk management program.

Icons Used in This Book



This book uses the following icons to indicate special content.

You won't want to forget the information in these paragraphs.



These paragraphs provide practical advice that will help you craft a better strategy, whether you're setting up your software or planning to purchase.



Look out! When you see this icon, it's time to pay attention — you'll find important cautionary information you won't want to miss.

Chapter 1

Understanding Next-Generation IPS

.....

In This Chapter

- ▶ Understanding next-generation IPSs
 - ▶ Pulling together the fundamental elements of next-generation IPSs
 - ▶ Protecting all network layers
 - ▶ Seeing how next-generation IPSs fit into the big picture
-

Intrusion prevention systems (IPSs) have long been a critical part of an organization's overall network and systems protection strategy and a critical part of a *defense-in-depth* architecture. Without them, you're fighting the bad guys with one arm tied behind your back. But today, the world faces new security threats and business challenges that require additional visibility and agility. Next-generation IPSs are the ideal solution.

In this chapter, we look at the function and role of next-generation intrusion prevention systems and how they fit into an organization's network and information risk management program. We also discuss some of the threats that your network faces.

Defining Intrusion Prevention Systems

Before we talk about next-generation IPSs, we want to discuss traditional IPSs: devices or software programs used to detect signs of intrusions into networks or systems and take action. That action consisted of generating alarms and/or actively blocking intrusions.

Most IPSs take the form of purpose-built hardware appliances, software appliances that run on specialized network security devices, or less commonly, software programs that run within virtualized or cloud environments. It's what takes place behind the scenes that sets next-generation IPSs apart from their predecessors.

Both traditional and next-generation IPSs come in two basic flavors: network-based and host-based, each with its own FLA (four-letter acronym). They are:

- ✓ **HIPS (host-based intrusion prevention system).** An intrusion prevention system that is installed and run as a software program on a host and designed to block attacks against the host system.
- ✓ **NIPS (network-based intrusion prevention system).** You guessed it — typically a physical appliance, this is an IPS monitoring a network to block attacks.



With more than 400,000 registered users, open source Snort is a popular choice for intrusion detection and prevention, boasting a huge quality assurance (QA) team of both commercial and open source users. A Snort-based IPS features an open architecture, making it easy to inspect the quality of IPS rules and create custom rules for proprietary systems. More than 100 vendors have incorporated Snort into their network security devices.

The two modes of operation used by intrusion detection and prevention systems are *passive detection* and *inline prevention*. These modes are compared in Table 1-1.

Table 1-1 Comparison of Passive IDS and Inline IPS

<i>Passive Detection</i>	<i>Inline Prevention</i>
Connected to a tap or switch span port	Directly inline
Receives a copy of traffic	Traffic actually flows through system
Creates alerts	Creates alerts
Can't block attacks	Can block attacks
Detection errors can result in false alarms	Detection errors can result in service disruption
Device malfunctions will cause a cessation of alarms	Device malfunctions can result in service disruption

Elements of a Next-Generation IPS

Intrusion prevention systems have come a long way since the introduction of open source Snort in 1998. Although a “typical” IPS contains everything you need to bring the box online and start blocking attacks, a new breed of IPS technology has raised the bar in terms of what organizations should expect from their IPS investment.

Today’s information systems complexity, advanced malware, and relentless attackers are driving the evolution of security into a real-time threat detection game. To stay on top of security, you need to fully know what’s taking place on your network and hosts and you need to better understand the context of the events. Security awareness and controls built into next-generation IPSs provide this deeper insight into and management of what’s actually happening in and around your network including:

- ✓ Comprehensive visibility that continually monitors for changes over time
- ✓ Ability to understand and control application-layer activity
- ✓ Ability to detect, analyze, and block advanced threats such as malware

- ✓ Assessment of new threats to determine which ones really matter
- ✓ Automation of activities such as security policy tuning and response processes

These benefits of next-generation IPSs not only help manage risks but they also allow information security team members to focus their response efforts and work more efficiently when incidents do occur. It's IT productivity at its finest!

Improving security without impacting the business

Using single-pass, scalable architecture, next-generation IPSs allow you to layer on additional security functions as they're needed, with nominal performance impact. The agile engine of this technology has the ability to adapt to and incorporate new intelligence sources and respond quickly to new threat types. Next-generation IPSs simultaneously tie in access control, content filtering, file type monitoring, and more for each work stream you're trying to protect.



Next-generation IPS isn't a one-size-fits-all solution, but rather a set of fine-tuned technical controls that can be customized based on your specific business needs.

Next-generation IPSs can evaluate connections to hosts and meter suspect network activity via IP address blacklisting based on host reputation. In addition, file activity can be monitored and controlled in near real time by analyzing type, protocol, and direction of file transfers. Certain files (especially unknown files that could be malware) may need further analysis, which can be done via *sandboxing* (isolated file exercise and behavior analysis) or *lookups to the cloud* (checking at-large community intelligence for reputation). Such an approach allows for more fine-grained analysis and disposition rather than simple binary: good or bad.

Not your father's IDS

During the early 1990s, intrusion detection systems (IDSs) were notorious for generating hordes of false positives. Early IDSs lacked intelligence and an easy way to filter out the noise, so network engineers would have to spend hours upon hours to “tune out” the false positives. It was a laborious, manual task

that gave early IDSs a bad reputation. With today's security challenges, the last thing you need is such annoyances distracting you from more important work. Next-generation IPSs are far superior, to the point that false positives are now an anomaly and no longer a major headache.

You won't have to overhaul your network architecture when deploying a next-generation IPS. You can maintain your existing firewall and, over time, tweak your next-generation IPS controls to block the most risky web and application traffic that's often easily exploited. A next-generation IPS, working in conjunction with your current firewall, effectively negates the need for a next-generation firewall (NGFW). By so doing, organizations may be able to reduce costs, extend equipment life, and avoid difficult firewall rule conversions.

The important thing is to look at the bigger picture. Given that complexity is the enemy of security, simplicity needs to be a focal point. When you step back and see how a next-generation IPS can complement (and even reside within) your existing security infrastructure, it becomes clear that your job can get easier moving forward.

NGIPS management console

The central component of a next-generation IPS is a management console that provides centralized command and control over all IPS appliances. Typical features of an old school, traditional IPS management console include the following:

- ✓ Security event aggregation
- ✓ Centralized detection policy management

- ✓ Downloading, importing, and applying IPS rule updates
- ✓ User interfaces for viewing and filtering security events
- ✓ Reports, alerts, and dashboards
- ✓ Health monitoring, to monitor health and performance of IPS sensors and the IPS management console itself

Next-generation IPS management consoles do all of the things in the previous list and also provide additional functionality, including the following:

- ✓ Network forensics (for example, view full packet payload)
- ✓ Event correlation and impact assessment
- ✓ User identification and tracking
- ✓ Application monitoring and control
- ✓ Flow (for example, NetFlow, proprietary flow) storage and analysis
- ✓ Advanced detection policy management (for example, policy layering)
- ✓ APIs to enable streaming of events to external platforms, remediation to network infrastructure devices, and importing of external network and vulnerability intelligence
- ✓ Granular administrative access permissions
- ✓ “Manager of managers” hierarchy, enabling one IPS management console to manage multiple subordinate IPS management consoles



The next-generation IPS management console is typically offered on dedicated hardware appliances, but may sometimes be offered as Windows-based software (requiring server-class hardware) or as a VMware or other virtual machine.

Protection across the Network Stack

Today’s information system environments consist of network components, servers, desktops, personal devices, operating systems, applications, and databases — most of which are

accessible anytime and from anywhere. Be it mobile, virtual, or through the cloud, the environments you must protect are more dynamic and the threats you now face are more complex and widespread. This dynamic nature and complexity require greater visibility and agility.

Next-generation IPSs help protect your network across the board, rapidly adapting protections for each of the many components that make up modern, ever-changing IT environments. One of the ways it does this is by using contextual awareness to create a holistic view. This *full stack* visibility provides the information you need to know so you can make better, and more proactive, decisions for keeping things in check. It helps not only with general information security management, but also with internal audit and regulatory compliance.



The threats and your network environment are continually morphing — security has become a real-time game. You have to be prepared to respond.

Compliance versus security

In the quest to minimize security risks on your network, keep in mind the differences between compliance and security. *Compliance* is that boring, arguably Draconian, set of rules and regulations that government and industry bodies push on businesses to keep IT in check and data safe. But after ten years of following these regulations, reality is proving that compliance alone is insufficient.

Many organizations that are compliant with regulations (PCI DSS, HIPAA, or something else) are still getting hit with data breaches as if no security was ever in place. Despite what many companies think, compliance doesn't immediately translate into true *security*. With compliance, there's often a disconnect between what's on

paper and what's really happening in IT. In many situations, no formal risk assessment has been performed and no technical controls are in place — such as next-generation IPS.

For example, the Payment Card Industry Data Security Standard (PCI DSS) mandates the use of IPS technology on networks that process credit card transactions. Yet having such a control in place doesn't guarantee PCI compliance. Compliance never has equaled security and it never will. You actually have to *manage* your information risks. When you address information security at the right level with the proper techniques and tools, compliance will likely flow naturally as a result.

Imagine knowing at an instant whether or not a risk exists at any point along the network stack. Even with highly-sophisticated threats and advanced malware, next-generation IPSs can detect the slightest anomalies and immediately analyze the risk and impact to help piece the puzzle together. Having a means for contextual awareness in highly-dynamic IT environments is the best form of protection against all threats — and for the network as a whole.

The important thing to remember is that every situation is different. What matters to your competition or even a completely different type of business doesn't automatically translate into a high-priority security problem for you. Likewise, just because an auditor or business partner brings a security issue to your attention doesn't mean you're going to know the best plan of action. You need to be prepared to map out and respond to the threats impacting *your* environment.

Common Network Threats

Of course, older threats never seem to really disappear altogether. You must stay ever-vigilant against the plethora of threats that, while well understood, are still commonplace. These include the following attack types.

Worms

A *worm* is a program that is designed to self-propagate from one computer to the next. Typical worms are designed to discover nearby computers with specific features, particularly features with specific flaws that permit the worm to successfully attack the next computer and install itself there. Then the worm begins to scan for other nearby potential victims, and the cycle repeats itself until the worm can find no new victim computers to invade.

The primary characteristics of worms are:

- ✓ Self propagation; they travel automatically with no human intervention required.
- ✓ Exploits a vulnerability to install itself.
- ✓ Scans the network for additional potential victims.

Worms cause harm in three different ways:

- ✓ **Network traffic:** Worms have a tendency to flood networks with their probes for new victims, and for the traffic caused by their propagation.
- ✓ **System resources:** Worms consume resources on the victim system through their propagation operations. Worms can even consume resources on adequately protected systems if a worm's attack is persistent.
- ✓ **Harmful payload:** Individual worms may be programmed to do more than just scoot around on the Internet. In addition, they may be designed to hunt for specific data on infected systems, implant other malware, or intentionally harm data.

Trojan horses

A *Trojan horse* is another type of malware. Like a worm, a Trojan is designed to propagate itself from system to system. But unlike a worm, a Trojan requires human intervention to keep it moving.

A Trojan horse is so-named because it is disguised as something benign. For example, a Trojan may be embedded inside a computer program purported to be a game, screen saver, or other program. But once activated, a Trojan will do whatever harmful things that it was designed to do.

When activated, a Trojan may scan nearby networks for neighboring systems that are potential victims. Or, the Trojan may scan the user's system to look for valuable data, or install other malware that it is carrying.

Buffer overflows

A *buffer overflow* is a specific type of attack against a system, where the attack is designed to confuse the system into executing the attacker's instructions.

A buffer overflow attack works like this. An attacking program establishes a communications session with a specific component on the target system, and sends a specially crafted

message to the target system. The message deliberately sends too much data into the target system's input buffer. In a program that is vulnerable to a buffer overflow attack, the excess data will overwrite program instructions in the vulnerable program, and eventually the program will execute those instructions (thinking that it is executing its original instructions). Those new instructions usually contain code to open the target system and permit a partial or complete takeover of the target system.

Sound complicated? You bet it is!

A buffer overflow attack isn't easy to develop. It takes detailed knowledge of the target system's internal architecture (both software and hardware), as well as detailed knowledge of the program or service being attacked. That said, hackers who develop buffer overflow exploits often build a "kit" that makes it easy for others to exploit the same vulnerability.



Worms (discussed earlier in this chapter), Trojans, viruses, and other types of malware often use buffer overflows as a way of gaining a foothold in a new victim system.

Buffer overflows account for a significant portion of the attacks against systems on the Internet.

Spyware, Phishing, and Botnets

Here is another grab bag of attacks on systems and people.

Spyware

Spyware is a term ascribed to a wide range of techniques used to covertly obtain information from computers. Spyware most often takes on the form of computer code that is installed on a user's computer without their knowledge or consent; it gathers specific information and sends that information to a central source. Spyware may also alter the behavior of the victim's computer.

The activities performed by spyware include:

- ✓ Tracking sites visited with a browser
- ✓ Recording keystrokes and mouse clicks
- ✓ Changing browser settings (for instance, changing home page, default search engine, and so on)

Unlike other types of malware such as viruses and Trojans, spyware doesn't usually contain code for making copies of itself onto other computers.

Phishing

A pun on the word *fishing*, a *phishing attack* is an attack on computer users in an attempt to con them into performing an action that is intended to cause them harm. That harm may take the form of financial fraud or the installation of malware or spyware on their computer, for instance.

A typical phishing scam works like this:

- ✓ **The bait:** The scammer sends out large quantities of genuine-looking e-mail messages to intended victims in an effort to entice them to open an attachment or click a URL.
- ✓ **The hook:** Although most people ignore or don't receive (because of anti-spam) the message, a few believe it is legitimate, or they're just curious. They open the attachment or click on the link.
- ✓ **The harm:** The attachment installs malware or spyware on the victim's computer, which may steal information, install a key logger, or perform some other harmful action. If the user clicks a URL, the website may trick the user into believing they're logging into a legitimate website (such as an online banking site). If they type in their user ID and password, the scam artist will use these credentials to log in later and steal money from the victim. Also, the website may attempt to infect the user's computer with malware. The victim's computer may also be made a part of a botnet, which is discussed later in this section.



Phishing scams account for a significant portion of computer security incidents and malware infections by preying on a user's gullibility.

Botnets

A *botnet* is a collection of victim computers that have been commandeered into a *bot army*, a powerful computing resource awaiting instructions from its owner. Creators of botnets are typically financially motivated.

Here is how a botnet works. An individual or group will write a small software program — a bot — that will enable the computer it's running on to be remotely controlled. This bot will be packaged into a worm, malware program, or loaded on a malicious website, at which time a campaign of some sort (say, a phishing scam) will ensue to get the bot installed on as many computers as possible.

The owner of these bots, usually known as a *bot herder*, has a centralized “command and control” program that can be used to control all the computers that are running his bots. This control program can then be used to perform work on behalf of the bot herder, such as:

- ✓ **Spam:** A bot army can be used to send millions of spam messages — which themselves may contain malware intended to grow the bot army.
- ✓ **Denial of service attacks:** The bot army can be used to remotely attack a computer or network of the bot herder's choosing. Denial of service attacks are discussed later in this chapter.



Botnets range in size from hundreds to millions of computers. According to the BBC, as many as a quarter of all personal computers may be members of one or more botnets.

SYN Floods and Denial of Service (DoS) Attacks

Our next grab bag of attacks includes two common network-based attacks.

SYN floods

A *SYN flood* is an attack on a target system, specifically an attack in a key design attribute of the TCP/IP networking protocol.

In a SYN flood, the attacker is sending thousands of SYN packets to a target system. A *SYN packet* is ordinarily a message sent from another computer that wants to establish a network connection with the target. Upon receiving the SYN, the target system will reply with a SYN/ACK, at which point the conversation will begin.

An important fact to note is that the target computer will allocate resources (mainly, memory) in anticipation of the new connection. But in a SYN flood, the attacker sends thousands of SYNs and ignores all of the SYN/ACKs. The purpose of this is to flood the target system until it is incapable of communicating on any legitimate channels.

A SYN flood is a special type of a denial of service attack. These attacks are discussed in the next section.

Denial of service

A *denial of service (DoS) attack* is an attack on a target system where the objective of the attack is to partially or completely incapacitate the target system. The purpose of a DoS attack is to render the target system unusable for legitimate purposes.

The reason that an attacker would carry out a DoS attack could include revenge, jealousy, ideology, or economics.

Committing a DoS attack is akin to blocking the entrances to a business so that its customers are unable to patronize it.

There are two basic types of DoS attacks:

- ✓ **Flooding:** The most common form of DoS attack is one where the attacker sends such a high volume of messages to a target system that it either malfunctions or is otherwise unavailable for legitimate purposes.

- ✓ **Malfunction:** The other common form of DoS attack is one where a specially-crafted message is sent to the target system; the message causes the target system to malfunction or crash.

Another type of DoS attack is known as the Distributed Denial of Service (DDoS) attack. In a DDoS attack, the attacker is causing many different systems to flood a target system simultaneously. Such an attack can be nearly impossible to block if there are hundreds or thousands of different sources.



Botnets are often used to commit DDoS attacks.

Zero-Day Attacks

A *zero-day attack* is a brand new attack on a previously unknown vulnerability, or a new type of an attack on an existing vulnerability.

The term *zero day* comes from the number of days of warning between the time when the vulnerability is announced and when it is exploited. In other words, these are vulnerabilities for which no patches are available.

Encryption and other detection evasion

In the malware economy, the developers of malware consider their products successful if they're able to evade detection. Early attempts at this involved the release of several "variants" that were constructed differently from one another. However, this has proven ineffective in comparison to encryption.

Encryption is a popular way of hiding from signature-based detection

systems. This is particularly effective when each computer's copy of malware is encrypted with a different decryption key, making every copy of the malware unique. This can make detection by signature-based systems very difficult. Anomaly-based systems should have no trouble with encrypted malware, because the basic attack pattern is likely unchanged.

Zero-day attacks are significant because signature-based (exploit-based) IPS devices are generally defenseless against them. However, IPS systems that also use anomaly-based detection and leverage vulnerability-based rules (as opposed to exploit-based signatures) can protect effectively against zero-day attacks.

Advanced Persistent Threats (APTs)

One of the biggest threats impacting businesses is advanced malware. Often referred to as advanced persistent threats (APTs), this type of targeted malware can create serious problems for your business. All it takes to introduce advanced malware into your network is one unsuspecting user clicking a link in an e-mail or loading an infected Web page. Without the proper network and host-based controls and insight, your company has no way to curb such infections.



Even employees bringing their own devices to work (referred to as BYOD) opens up entirely new avenues for information threats and vulnerabilities — all of which translate into risks that you can't afford to not know about.

There is presently more hype and misinformation about APTs than practically everything else in this book combined. In truth, there is no silver bullet or single security device for defending against APTs. But a next-generation IPS is a strategic component of a defense-in-depth strategy that can help you get ahead in the game.

What is APT?

To understand what APT is and what it is not, we start with a short definition and then delve into the details.

An advanced persistent threat is information warfare, conducted by sophisticated adversaries who are determined to control information systems and gather intelligence on persons, organizations, and governments.

Does this definition scare you? Good! It should, because the actors who are responsible for these threats are financially motivated, patient professionals with research and development resources at their disposal. They're not looking for instant gratification, but instead are willing to go "low and slow" to patiently, systematically infiltrate the systems used by individuals and organizations.

So enough about the actors. What about the actual threats?

Advanced persistent threats are malicious, and they certainly fall into the class of malware. However, for highly sophisticated threats, you won't find signatures of this malware in anti-virus products or intrusion detection systems, because these threats are custom made for their specific targets.

Advanced persistent threats do consist of attacks that are detectable. However, these attacks may be subtle and take place over a very long period of time. Traditional defenses such as anti-virus, IPS, and firewalls may not see anything at all. The actors behind an advanced persistent threat don't want to set off any alarms.

Chapter 2

Know What You're Protecting

In This Chapter

- ▶ Understanding what's on your network — right now!
 - ▶ Obtaining insight into network hosts, applications, and users
 - ▶ Mapping application usage to business risks
 - ▶ Using file type and reputation to limit exposure
-

When it comes to managing information risks, you can't secure what you don't acknowledge. Visibility is absolutely required if you're going to keep your network in check. You must understand your networks, how they're configured, what they're being used for, who is using them, and how the environment is changing — and believe us, it's continuously changing.

In this chapter, we look at how visibility impacts your ability to make informed decisions about information security and respond to security incidents both efficiently and effectively.

Real-Time Visibility Is Key

One great thing about next-generation IPSs is that they can provide real-time visibility into your network to help you understand exactly the right information at exactly the right time — continuously. Specific areas of focus include the following:

- ✓ Hosts and devices
- ✓ Operating systems and browsers
- ✓ Applications, services, and protocols
- ✓ Users and identities

Constant surveillance allows you to understand what threats your changing network is vulnerable to and which events are significant at every moment — and which aren't. This kind of real-time visibility also allows you to evaluate changes in your environment and determine whether or not these changes pose a new risk — and what measures to take (block access, deploy additional signatures, and so on) given the context of what's taking place.

Without knowing everything you can about your network, you're simply reacting to every event seen or every new threat you hear about, without any context to help you filter and prioritize.

Hosts

In order to know what you're protecting you need to know which hosts are currently running on your network. From Windows to Linux to iOS-based systems, being able to discover and track the specific operating systems is important. Why? There's a good chance that certain hosts on your network are running on unsupported or older operating systems, which pose added vulnerability risks. You may even discover that systems are running platforms you never knew existed on your network. A new device type you have no knowledge of (an increasing likelihood with BYOD) may also pose new risks.

Gathering host profile information and having the ability to explore and drill down into specific hosts can provide you with invaluable information to ensure your network is made up of what you think it is. Real-time contextual awareness of hosts and devices in your network provides enhanced visibility to know what you're protecting (see Table 2-1).

**Table 2-1 Real-Time Contextual Awareness
of Hosts and Devices**

<i>Categories</i>	<i>Samples</i>
Threats	Attacks, anomalies
Users	AD, LDAP, POP3
Web applications	Facebook, Chat, eBay
Application protocols	HTTP, SMTP, SSH
File types	.doc, PDF, TAR, .exe
Client applications	Firefox, IE6, Chrome
Network servers	Apache 2.3.1, IIS4
Operating systems	Windows, Linux
Routers & switches	Cisco, Nortel
Wireless access points	Linksys, Netgear
Mobile devices	iPhone, Android
Network printers	HP, Xerox, Canon
VoIP Phones	Avaya, Polycom
Virtual machines	VMware, Xen, RHEV

Source: Sourcefire

Client-side applications

Knowing everything you can about operating systems and other high-level host information is only the beginning. With the proper visibility, you can gather information on specific applications running on each of your network hosts including the following:

- ✓ Web browser versions
- ✓ Plug-ins installed
- ✓ Third-party software such as Java and Adobe Reader
- ✓ Remote access software such as VPNs and Remote Desktop Protocol (RDP)
- ✓ Mobile OSs running on smartphones and tablets

When you know what's running on each client, you can create a more detailed picture of your network environment. Odds are that you'll uncover vulnerabilities you never knew you had.

Users

When discussing network hosts, you also usually want to know about *users* — you know, the folks who keep you in a job. Good visibility into the network will identify users and their activity from sources such as:

- ✓ User IDs captured from logins
- ✓ File and directory access
- ✓ Identity imports


Such information paints a bigger picture to help tie security events to individual users and speed up the remediation process when a security incident arises.

Understand Application Risks

Make sure software applications are on your radar. Regardless of how much visibility you have in other areas of your network, weak insight into your applications is a surefire way to overlook critical security risks.

Detect and classify applications

A logical first step to reduce your attack surface is to observe usage and then categorize applications, with emphasis on identifying high risk, low-business-relevance applications — those that you can and probably want to do away with. Proper visibility will provide information on common applications and their component parts (for example, Facebook Chat) at a granular level. This allows you to identify programs that may require further scrutiny and may be candidates for access controls. A summary view of detected applications is shown in Figure 2-1.



Source: Sourcefire

Figure 2-1: Content Explorer provides dynamically updated contextual views of the environment and easy drill down for details.

Determine risks and business relevance

Imagine having the ability to filter the important from the mundane. Proper visibility into network applications can spotlight unnecessary, unproductive, and risky applications while safely allowing legitimate applications that adhere to security policies.

Automatic categorization, based on prior experiences and application usage patterns, allows quick determination of the application type and inherent risk profile it represents for your organization. Having this information in hand facilitates discussion with user groups and business owners about potential application access blocking and controls.



Good visibility leads to better decision making. By focusing attention on activity and behavior, IT resources can be better utilized and threats can be significantly reduced or eliminated. Everyone wins.

Who's to say that's a risky app?

Many applications are running on any given computer on your network at any given time. The tricky part is determining which programs are relevant and which ones don't belong. You may find medium business-relevance applications running locally, via the intranet, or in cloud but they can still have a great impact if, say, they are known to harbor malware. Be careful when judging whether or not programs are risky. You have to know *all* the facts and be able to see the forest through the trees. The only realistic means for gathering all you need to know is by using a next-generation IPS that understands your environment at

this level and can automate the arduous data collection/analysis task.

Determination of security controls and policies must consider the user, group, and business function as well. It may be quite appropriate for a marketing department user to post content to the company's Facebook page; however, other groups may have no true business need. Likewise, some video streaming use within an organization may be appropriate for training/educational purposes, but you likely don't want employees on the job watching Netflix all day long.

Further, a next-generation IPS can correlate observed behaviors with detailed configuration and vulnerability profiles of targeted endpoints to prioritize response. For example, a Windows-based exploit that's targeting a Linux server might not mean much. The key is interpreting whether or not threats are relevant or irrelevant. This helps to eliminate event *noise* that can be a distraction and free up internal security resources to focus on more important things.

File and malware awareness

Next-generation IPSs provide enhanced content awareness, because they help you evaluate network file traffic so that you can better understand resource use and behavior. This information can also help you determine whether or not specific files contain malware. Appropriate policies can be created when you know the following:

- ✓ File type (office docs, archive files, multimedia, images, executables, system files, and so on)
- ✓ Direction of file transfer (inbound, outbound, any)
- ✓ File protocol (HTTP, SMTP, IMAP, POP3, FTP, SMB)
- ✓ File reputation (whether a file is malicious or exhibits suspect behavior)

Organizations with such visibility can reduce exposure to malware (by stopping propagation) and limit exposure to data leakage and institute appropriate use/compliance adherence.

Integrate within IPS policies

A next-generation IPS should be able to map out application usage, trends, and risks. But it must also be able to enact real-time enforcement based on your unique business needs. Your threat detection engine needs to be robust and flexible, because many applications will be unimpeded. Many organizations may not be ready to impose too many access controls — at least not until they better understand usage and the business relevance. So it's reassuring to know that what passes through will still receive thorough threat scrubbing. This affords a gradual approach where a company can walk (with application detection) before they run (with application control). Integration with threat detection policies (within the IPS) is an ideal way for users to strike the proper balance.

Chapter 3

Next-Generation IPS Features

In This Chapter

- ▶ Enhanced threat protection; going beyond your typical IPS
- ▶ Using application control to your advantage
- ▶ Protecting against advanced malware
- ▶ Automating key security analysis functions
- ▶ Maintaining performance as you grow
- ▶ Integrating third-party products into an IPS

In this chapter, we dive deeper into the key features of a next-generation IPS (NGIPS), with emphasis on capabilities related to enhanced threat and malware protection, application control, automation, and reduced total cost of ownership (TCO). We also discuss strategies for SSL (Secure Sockets Layer) inspection and integration with existing IT security products and infrastructure.

Common Functions of Any IPS

Virtually all of today's IPS devices share the following common functions:

- ✓ **Inline IPS and passive IDS modes.** Although passive IDS allows for threat detection, inline IPS threat *protection* is what most users today desire. However, when an IPS device is placed inline, be sure it supports fail-open ports that allow traffic to pass in the event of a network interface or security appliance failure. Some IPS providers offer fail-open ports on only a portion of their models.

- ✓ **Default detection policy.** Every IPS vendor should provide a detection policy comprised of the most common IPS rules to help get you started. However, your organization should never just rely on a default policy because it can't adapt dynamically to your changing network environment. Don't let IPS vendors fool you about this. *Tuning* is required to select the IPS rules that are most relevant for your organization (although this can be made easier via automation). With IPS, one size does *not* fit all.
- ✓ **Reports, alerts, and dashboards.** Most IPS providers offer a selection of reports, alerts, and dashboards — usually present in the management console. Reporting should be flexible, alerts should be offered through e-mail, syslog, and SNMP, and dashboards should be customizable based on the user's role in the organization. The managers who paid for IPS want to see their reports and dashboards, to know that the IPS is really working and providing business value.

Protection Beyond a Typical IPS

Traditional IPSs are *black boxes* that offer little visibility or context into the protection being offered. However, a next-generation IPS — especially one based on an open architecture — is different. The benefits offered by today's next-generation IPSs include network security capabilities beyond just detection and prevention:

- ✓ **Visibility.** Vendors with IPS offerings based on closed architectures require you to “trust” that they have the best protection for your needs — you have no visibility into how the detection engine works or whether their rules (or *signatures*) are designed to defend actual vulnerabilities or simply detect known threats. In contrast, a next-generation IPS features an open architecture with full visibility into the detection engine and rules. Next-generation IPSs are simply higher quality products that offer increased effectiveness and peace of mind. With the ability to continually monitor for changes over time, your network can achieve comprehensive coverage.
- ✓ **Custom rules.** Most typical IPS vendors will tell you that you can create custom rules, but few provide the means



to do it effectively. It's best to select an IPS vendor that makes it easy for you to create custom IPS rules by providing you with training opportunities and an easy-to-use wizard interface.

A benefit of being able to tweak your own rules and call your own shots is that you can adjust the settings based on your specific compliance requirements. Be it federal law, industry regulation, or business partner request, the ability to fine-tune your network security controls helps take some of the pain out of the compliance process.

- ✔ **Vulnerability-based protection.** Most IPS providers offer exploit-based signatures that detect a single variant of malware. A next-generation IPS puts in the extra effort to construct IPS rules to detect *any possible variant* of an exploit that targets an operating system or application vulnerability. Now you have the ability to assess new threats and determine which problem areas you need to focus on. This approach provides the best security and offers the greatest zero-day protection. It's better (more effective and efficient!) to determine and protect for all possible vulnerabilities of a faulty lock than it is to have to detect every possible key pattern that an intruder may try. To continue the analogy, exploit-based approaches attempt to detect only the known keys and may miss other vulnerabilities altogether.
- ✔ **Network behavior analysis.** Not all attacks come through the perimeter. Many are hand-carried on mobile computing devices right through the front door, thus bypassing a perimeter IPS. Network behavior analysis technologies baseline “normal” network traffic (using NetFlow or proprietary flow technology) and detect anomalies, including the spread of malware. Sophisticated next-generation IPSs can aid in determining trajectory and root cause of malware, which can help you pinpoint the source and collateral impacts in an organization to avoid reinfection.
- ✔ **Virtual IPS and management console.** A typical appliance-based IPS can't inspect traffic between one virtual machine (VM) and another on a VMware or other virtual server. Select a next-generation IPS that offers virtual IPS sensors and management consoles that are deployed with the hypervisor and can protect virtualization environments from within and defend cloud computing infrastructures that are increasingly virtualized.



The general trend in IT products is toward the capability to see inside the product to view and manage detailed configuration and operation. Make sure you select an IPS that gives you the capability to view and manage detection rules.

Granular application control

Most enterprises have documented acceptable use policies (AUPs) depicting operating systems and applications approved and/or restricted from use, but few organizations have the means to monitor and enforce these policies. A next-generation IPS helps IT to *reduce the surface area of attack*. How does this happen? First, by alerting IT to unauthorized uses of operating systems, applications, and devices. In addition, next-generation IPSs ensure application control through specific security measures such as allowing and blocking by specific category, including the following:

- ✓ User or group
- ✓ Application subfunctions (for instance, Facebook Chat)
- ✓ Type, tag, category, and risk rating



The ability to control types of applications by category or specific applications — or even functions within applications — is very beneficial to today's *app-centric* businesses. An example is allowing general social media applications while blocking portions of the applications — such as gaming or instant messaging. In addition, controlling user or group access allows further refinement of security policy. Such control allows you to fine-tune which applications are allowed and who can do what based on business needs.

With traditional IPSs, you can glean very little information regarding the ways specific applications are functioning on your network. Next-generation IPS technologies provide you with the ability to understand and control application-layer activity, such as application type, connection direction, user or group, host operating system, client-side environment, and so on to ensure full coverage across the network stack.

Many organizations are deploying next-generation firewalls for application control behind their traditional perimeter firewall. This is an incomplete solution because most

next-generation firewalls don't have access to information about vulnerabilities associated with the systems being targeted for attack. And worse, next-generation firewalls can complicate management given the inevitable disparate consoles and the requisite training they typically entail. We recommend the next-generation firewall approach when truly ready to refresh existing firewall infrastructure — and don't forget it's a big job to convert all the firewall rules that exist on your legacy firewalls!

An alternative architecture to next-generation firewalls places application detection and access control at the IPS itself, because IPSs excel at inline threat detection and prevention. This is often less disruptive, especially when application control is a function that can be enabled via software. Furthermore, application blocking can be controlled within a single console along with IPS threat detection that determines all protection rules. Certain next-generation IPSs allow varying threat inspection levels by type of application traffic, inspecting riskier traffic more thoroughly — to safely enable all forms of business processing.

Next-generation IPSs can detect incidents by observing individual network connections, for instance, and they can make alerting or blocking decisions based on what's considered normal for various types of activities. For example, a next-generation IPS may learn the sequence of typical events such as: a Web user logs into an application, and after logging in, issues commands to the application. The next-generation IPS may consider a user issuing commands without logging in to be an event that should be blocked, because this may be a sign of an intruder who is attempting to perform unauthorized transactions.

Protection against malware

Recent years have brought about a new wave of advanced malware often referred to as zero-day threats or advanced persistent threats (for more on these topics, see Chapter 1). Traditional IPSs have little to no knowledge or means to detect and prevent these sophisticated threats. Next-generation IPS can do both. From detection to analysis to blocking, a next-generation IPS is able to prevent or mitigate targeted attacks and complex network infections. Using

file-level controls, a next-generation IPS allows you to monitor content and determine what's bad and what's permissible. It also can watch suspect or unknown content (new content with no prior reputation) and then proceed to retroactively remediate or eliminate the infection.

Here's another analogy: Consider a padlock with a design weakness that makes it vulnerable to picking and pretend that lock is actually an application or host on your network. If the IPS were configured to detect only known lock-picking methods (attacks), then any new methods for picking the lock would go undetected. It would be better for an IPS to be familiar with the lock's vulnerability, so that it could detect any kind of an attack upon it.

Next-generation IPSs can detect incidents by comparing traffic patterns that the IPS considers normal with any new traffic patterns that emerge, and deciding whether these new patterns fall within acceptable patterns or not. A distinct advantage of this *anomaly-based detection* is the capability to detect incidents that may not be triggered by a traditional IPS rule or signature.

When attacks exploit new or previously unknown vulnerabilities, observation can lead to detection by pinpointing suspicious behavior: Does the suspect code call back to malicious hosts, does it modify certain system files (Windows registry), does it attempt privileged access escalation or something similar? The suspect code can be tagged, observed, tracked, and then later convicted once proven malicious. This form of advanced malware protection essentially creates a continuous protection capability even after the file has passed through your network.

SSL inspection

Many network security devices are blind to SSL-encrypted traffic, including most traditional network IPSs. This is because an SSL session is encrypted end-to-end, and the IPS in between typically sees only encrypted data. As the use of SSL grows within an organization — oftentimes comprising one-quarter to one-third of traffic — the potential of an SSL-encrypted attack rises.

To mitigate this risk, a next-generation IPS should either incorporate SSL inspection or be complemented by a dedicated SSL inspection appliance. The SSL inspection engine should decrypt SSL traffic, pass it to the IPS for inspection, and then re-encrypt the (clean) traffic before placing it back onto the wire — all with minimal added latency. When placed inline, the SSL inspection appliance should also feature fail-open ports (which allow traffic to pass), just as the IPS should.



Beware IPS/NGIPS providers that only offer on-board SSL decryption. Enabling SSL decryption on many IPSs can adversely affect the threat inspection throughput by up to 80 percent. Most organizations should offload SSL decryption to a standalone appliance, which not only decrypts traffic for the IPS, but for all network security devices placed behind it (such as ones that might perform other security functions such as forensics, data loss analysis, and so on). But regardless of whether SSL is decrypted by the IPS itself or a standalone appliance, ensure the SSL decryption capability also re-encrypts the original (clean) traffic before placing it back onto the wire to maintain confidentiality of the data and to maintain compliance with PCI and other regulatory or contractual requirements.

Third-party integration

A best-of-breed security device should integrate with other devices on your network to share intelligence, coordinate responses, and lower total cost of ownership. The following are common examples of how a next-generation IPS can integrate with popular third-party systems:

- ✔ **Security Information and Event Managers (SIEMs).** Stream security, compliance, and health events to your SIEM of choice (for example, ArcSight, Q1 Labs) for centralized security monitoring.
- ✔ **Vulnerability Management (VM) platforms.** Import vulnerability intelligence from popular VM platforms (for example, Qualys, Rapid7) for security event impact assessment and greater network visibility.

- ✓ **Network infrastructure devices.** Remediate to routers, switches, and NAC devices from leading network infrastructure providers (for example, Cisco, Juniper, Check Point) to quarantine hosts related to security and compliance events.
- ✓ **Network forensics.** Launch packet-level forensics queries directly from the IPS management console to leading network forensics devices (for example, NetWitness, Solera), saving both time and effort.

After you integrate your IPS into your SIEM and other platforms, you'll be humming right along at a level of security your organization has not experienced before.

Central Role-Based Management

At times, information security functions are overlooked or held back because people aren't doing their assigned jobs. This often happens when workers aren't sure what's expected of them — security functions must be clearly defined so everyone working in security understands their roles.

Role management can be made or broken by security products themselves. A next-generation IPS that provides the hierarchy of user roles allows network administrators and security managers to maintain their respective disciplines. You can set up your system so that certain activities are allowed or prohibited for certain people. For instance:

- ✓ The firewall administrator can create access control rules, but not security inspection policy.
- ✓ Security analysts can implement new threat inspection rules, but perhaps not modify which user groups gain access to which applications.
- ✓ Network managers may provision and configure network port settings without access to security settings.

Such features may seem unimportant now but they can have a tremendous productivity impact on your day-to-day security management over the long term.

Management Reporting

Managing information risks requires that you know your network. A next-generation IPS will have an executive dashboard that provides key information on how the system is working and where things currently stand with information risks including:

- ✓ Trends and high-level statistics
- ✓ Event detail and forensics
- ✓ Workflows
- ✓ Compliance



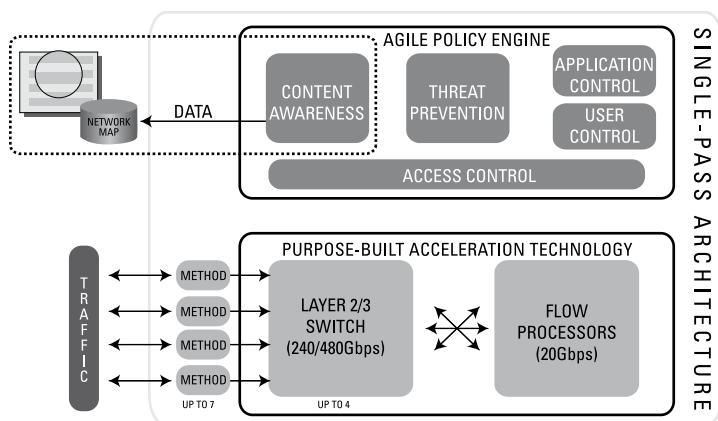
Both depth and breadth of information are needed to effectively detect and automatically respond to emerging attacks and security risks.

Addressing both general and focused information needs, driven by an individual's responsibilities or concerns is an absolute must. Executives will want to know one thing, mid-level managers another, and technical administrators an entirely different set of information altogether.

Maintaining Performance as You Grow

Look for a purpose-built next-generation IPS appliance that can scale throughput performance while delivering robust threat protection over a wide range of system load and connection speeds. If it isn't purpose-built from the ground up, it is less likely to perform well, especially under heavy loads. Poor performance is commonplace with many unified threat management (UTM) or other general-purpose architectures — they can't scale as more throughput and additional security functions are added.

By contrast, a single-pass design, as shown in Figure 3-1, allows for the efficient addition of multiple security functions (access control, threat detection, behavior analysis, host profiling, and so on) while maintaining high throughput performance.



Source: Sourcefire

Figure 3-1: A single-pass, hardware-accelerated design affords maximum scalability, threat effectiveness, performance, and security in a consolidated platform.



Agile security is something your organization needs in order to navigate today's information-risk realities. Agile security works by delivering a continuous process that responds to continuous change. Agile security is made up of four essential elements:

- ✓ **See.** Clarity and vision, reflecting the reality of your environment, as it currently exists.
- ✓ **Learn.** Applying intelligence and raw data to improve understanding and decision making.
- ✓ **Adapt.** Automatic evolution and modification of defense in response to change.
- ✓ **Act.** Decisive, flexible, and automated responses to events.

Agility requires insight into changing conditions and new attacks. This insight promotes rapid learning and better decision making. Agility also requires the ability to make quick and effortless changes to adapt to new threats. You also need to be able to act quickly and efficiently once you have new information. With the ability to properly learn and adapt, a next-generation IPS can help you get a handle on network security — once and for all.

Lower TCO through IPS Automation

Whether you work for a small, medium, or large organization, there never seems to be enough IT security resources to go around. IT security must know how to work more efficiently to defend today's dynamic network. A next-generation IPS makes it easier to do more with less:



- ✔ **Automated impact assessment.** IPS devices may generate hundreds of security events on a daily basis. When you take into account that a traditional enterprise may have a dozen IPS devices or more, sifting through thousands of security events each day is virtually impossible and can effectively render an IPS useless, because it will be ignored. A next-generation IPS, on the other hand, correlates threats against endpoint intelligence to reduce the quantity of *actionable* security events by 95 percent or more.

It's often difficult to find true cost of ownership and return on investment when it comes to security but a properly deployed and well-run next-generation IPS makes it plain as day and simple as pie.

- ✔ **Automated tuning.** Every network is different. Customize your IPS detection policy with rules that are relevant for your organization. If the detection policy is too small, the IPS will offer inadequate protection. And if it's too big, it can overburden the IPS, causing decreased network throughput and increased latency. A next-generation IPS can passively profile your network and automatically recommend rules to enable and disable at a user-defined interval (for instance, weekly or monthly).
- ✔ **User identity tracking.** What good is an IP address for an end-user device related to a security or compliance event if you don't know who is being attacked or who is violating a company IT policy? Instead of sifting through DHCP and Active Directory logs to manually cross-reference users with IP addresses, a next-generation IPS can place usernames and user identity at your fingertips. The time it takes to tie a user to a security event can be shrunk from one hour to under a second.

Reducing TCO through IPS automation

According to a SANS Institute white paper entitled “Calculating TCO on Intrusion Prevention Technology,” a multinational credit reporting organization with approximately 20,000 nodes and 7,500 employees saved more than \$230,000 per year in annual TCO reductions through

automated impact assessment, automated tuning, and user identification. By leveraging a next-generation IPS solution, organizations can recover their initial IPS investments in a matter of months by automating key IPS administrative tasks.

Chapter 4

Next-Generation IPSs, Virtualization, and Cloud Computing

In This Chapter

- ▶ Considering the benefits and risks of virtualization
 - ▶ Securing virtualization
 - ▶ Virtualizing security
 - ▶ Securing the cloud
-

Virtualization and cloud computing are revolutionizing information technology by facilitating a more efficient use of computing resources. These technologies aim to help organizations perform cheaper, better, and faster.

In this chapter, we discuss virtualization and cloud computing, and the relationship that each has with next-generation IPSs. We also discuss two approaches to virtualization and security. One is the process of securing virtualization, and the other is virtualizing security. Both are discussed in this chapter.

Benefits and Risks of Virtualization

Virtualization is the technology that permits an organization to run many separate instances of operating systems on a single server. This allows an organization to greatly enhance

the efficiency of its server hardware because grouping many separately running operating systems onto a single server means you have to buy fewer servers. Figure 4-1 illustrates virtualization.




Figure 4-1: Virtual servers.

Before virtualization, an organization whose environment required six servers had to purchase six separate hardware servers. With virtualization, the organization can purchase one server and install six virtual machines on that server.

The primary benefit of virtualization is that an organization can implement new virtual machines at will and with very little effort.

The primary risk of virtualization is that an organization can implement new virtual machines at will and with very little effort.

Yes, you read that right: The main benefit is also the main risk. What we're saying here is that, without proper safeguards, virtualization can introduce risks that can negate the benefits.

The benefits of virtualization include

- ✓ **Agility.** Virtualization allows an organization to respond more quickly to changing needs in its technical environment.

- ✔ **Rapid deployment.** With virtualization technology, you can build and deploy a new server in just a few minutes. No more running down to the local computer store for a server and loading an OS.
- ✔ **Improved system availability.** Virtualization enables an organization to implement servers that are more consistent with each other. Consistency breeds higher availability, because there are fewer differences between systems, which means systems engineers are less likely to make mistakes that cause unexpected downtime.
- ✔ **Energy savings.** Running many virtual servers on just a few physical servers means there are a lot fewer physical servers consuming energy.
- ✔ **Space savings.** The amount of space that servers consume is expensive, especially in commercial data centers that literally charge for rack space by the inch.

With these benefits, what's not to love? Sadly, there are also risks related to virtualization, and it's important to understand these risks, so that you won't make the same mistakes that others have made:

- ✔ **VM sprawl.** Because virtualization makes it so incredibly easy to deploy a new server, it can sometimes be tempting for an engineer to deploy a server and bypass the management processes that usually accompany the deployment of new servers. The result can be many unauthorized servers that are doing who-knows-what. VMs created outside of rigid deployment controls likely are unmanaged and may invite malware infection. For more on the topic, see the section "Controlling VM sprawl," later in the chapter.
- ✔ **Vulnerabilities.** One of the neat features of virtualization is the capability to *roll back* to an earlier *snapshot*, which is a fancy way of reverting to an earlier version of the virtual server. Doing so, however, can also result in the removal of critical security patches that can leave servers vulnerable to attack or malfunction.
- ✔ **Lack of separation of duties.** In the physical server world, there is more management and team coordination required to deploy a new server: Someone has to approve the hardware purchase, and network engineers

provide support by enabling the connection of a new server to the network. With virtualization, none of this coordination is necessary. A single individual can deploy a server without telling anyone.

✓ **Blind spots.** In the physical world, you can more easily observe the logical architecture and data flow in an environment, and control security with firewalls and IPSs where needed. With virtualization, however, servers that were once separated by firewalls or IPSs may end up on the same physical server, resulting in the loss of those network controls.

These risks may sound pretty scary — so is virtualization worth it? You bet it is. And IT management, aware of the cost savings realized with virtualization, will insist on it. So it's best to make your virtual systems secure.

Some next-generation IPS vendors even offer virtual versions of their systems, which can be incorporated into virtual environments, providing greater visibility and control of VM-to-VM traffic.



Ensure you have complete network security coverage so no weak links are introduced. Virtualization is an area where a lot of network traffic is generated but there's a disproportionately small amount of security resources dedicated to it.

Securing Virtualization

Like any information technology, virtualization needs to be secured. In other words, virtualization needs to be configured and managed in a way that results in the virtualization environment being free of vulnerabilities that could lead to compromised systems.

There are three main areas where virtualization needs security controls: with the people, processes, and the virtualization technology itself.

Virtualization: People

What we're getting at here is the fact that all personnel who design, implement, manage, or operate virtualized environments can do so only when they have the knowledge required to do it properly. Not only do personnel need to understand virtualization technology, but they also need to be familiar with the organization's policies and procedures regarding virtualization.



You can have all the right virtualization technology in place, but if personnel don't understand how to use it (or are unwilling to understand), your virtual environment will not be secure.

Virtualization: Processes

Like personnel, a virtualized environment will not be very secure unless the right business processes are in place. Some of the processes that we feel are important include:

- ✓ **Change management.** Changes to virtual machines, as well as changes to virtualization configuration, should be done under the control of a formal change management process. This will help ensure that the loop is closed when someone tweaks the environment. Just *how formal* this process should be depends on the organization's needs. However, you must set expectations. Under no circumstances should changes be made without *at least* informing all affected parties!
- ✓ **Technical standards.** Configuration settings for virtualization, as well as the virtual machines themselves, should be written down. Think of this as a process of establishing standards and then sticking to them not a one-time exercise. Sure, things need to change — in that case, you use your change management processes to manage your changes.
- ✓ **Audit.** Virtualization settings and virtual machines need to be examined from time to time, to ensure that they're being deployed and operated properly, no unauthorized activity is taking place, and vulnerabilities aren't lurking about waiting to be exploited.

Securing virtualization technology

Virtualized environments need to be properly designed and configured, so that they're free of vulnerabilities that may expose them to threats. Virtual environments should be designed and configured according to the following principles:

- ✔ **Least-privilege administration.** Each staff member who administers virtualization should have his or her own user ID, and each person should have only the privileges required to complete the necessary job functions.
- ✔ **Logging.** Administrative activities within the virtualized environment should be logged. This helps to identify who is performing what administrative functions. A documented history of administrative activities makes troubleshooting a lot easier.
- ✔ **Disable unneeded components.** Just as disabling unused ports and components on a server is good for security, this same principle applies to virtualization. The smaller the potential attack surface, the lower the risk.
- ✔ **Backup.** Certainly it should be obvious that all virtual machines in a virtualized environment should be backed up. But what may be less obvious is the need to back up virtualization configurations themselves if they're not contained in an OS being backed up.
- ✔ **Placement of IPS sensors.** Just as the placement of IPS sensors is critical in a traditional environment, it's also critical in a virtualized environment. This may necessitate both hardware IPSs as well as virtual IPSs that are installed within virtualized environments. This will help to protect VM-to-VM traffic even within individual hardware platforms.
- ✔ **Configuration standards.** Virtualization and virtual machines need to be configured according to a set of documented standards. Standards offer two main benefits. First, when properly circulated, reviewed, and approved, standards should represent a collective agreement on how systems should be configured. Second, standards (when enforced) help systems to be more consistent with each other.



When everyone is on the same page with virtualization, the greater the chances of it being successful and secure.

Controlling VM sprawl

VM sprawl is a result of deploying virtual machines without obtaining approval. Because engineers can unilaterally deploy VMs without obtaining approval, some enterprises are liable to experience uncontrolled growth of VMs and the chaos that results. Here we discuss some neat ways that IPSs can be used to control it.



Next-generation IPSs can help to control VM sprawl by detecting a VM by its virtual network card's MAC address. A next-generation IPS can be configured to generate an alert whenever it sees a new VM on the network. This procedure can help management keep an eye out for new VMs. Never send these alerts to the individuals who create VMs! Send them to other personnel, in order to prevent engineers from creating VMs on the sly.

Organizations that are zealous about controlling VMs can use their next-generation IPS to prevent new, unauthorized VMs from being able to communicate on the network. This is one important way that segregation of duties can be retained in a virtualized environment.

Virtualizing Security

Virtualization creates several new opportunities, including the capability to implement more than just operating systems in virtual environments. Besides OSs, you can also deploy network switches, firewalls, and next-generation IPSs as virtual machines, thus leveraging the cost-saving benefits that virtualization brings.

At first blush, it may appear that cost savings are the only motivator for virtualizing security devices. Sure, virtual versions of security devices may cost less than their physical counterparts, but sometimes using a virtualized security device is also the best way to approach a situation.

For example, imagine that an Internet-facing application is deployed in a virtualized environment. The application consists of a web server, an application server, and a database server. Compliance regulations require an IPS protecting the web server and a firewall protecting the database server. All these components can be incorporated into a single physical platform with the necessary detective and preventive controls in place to protect all these virtual components with as much confidence as though they were physically separate. Configured correctly, these components are every bit as secure as if they were *air-gapped* — completely separated.

Virtual IPS solutions can also be deployed to small remote offices (equipped with virtualization hosts). They are able to monitor both physical and virtual hosts for threats, without incurring the expense of physical IPS devices and the human costs to deploy them. Virtual IPS VMs can be *dragged and dropped* to protect virtually (no pun intended) any corner of the network with a few clicks of a mouse, saving both time and money.

Securing the Cloud

Cloud computing is all the rage these days. Whether they're providing cloud services or consuming them, enterprises are flocking to cloud environments faster than prospectors flocked to the Klondike in the 1896 Alaska gold rush.

Cloud computing is the term encompassing many technologies that enable an organization to enjoy a dynamically expanding and contracting computing environment. Organizations can build their own clouds, or buy services offered by external cloud computing providers.

In the context of intrusion prevention systems, *cloud computing* means the use of computers and networks as general-purpose, on-demand, and dynamically scalable computing environments that host applications and other computer-based services.

Organizations that wish to move their applications *into the cloud* generally want to outsource an application's infrastructure (computers and network devices), when they expect their need for computing resources to expand and contract

based on demand. Growing and managing a dynamic computing infrastructure is expensive and time consuming, and outsourcing frees the organization to focus on its core competencies.

One of our favorite sayings is, “You can’t outsource accountability.” This means that, even if you hire an outside organization to perform work, you’re still responsible for the outcome. In the context of cloud computing, an organization that outsources its infrastructure (and, possibly, applications and other services) to the cloud needs to make sure that its systems and data are protected from security threats.



Cloud computing doesn’t always mean “run by others.” An organization can have its own private cloud that’s hosted and managed internally.

The important thing to remember with cloud computing is that the same security technologies available to the traditional enterprise IT infrastructure — including next-generation IPS — are available to lock down your virtualized cloud environment. Whether you’re outsourcing to the cloud or have your own private cloud, make sure you’re using the right technologies for the risks you face and make sure that someone is accountable for system oversight and maintenance. The last thing you need is a false sense of security when proper management isn’t taking place.

The controls used to protect cloud-borne applications and data from threats are discussed in the remainder of this chapter. These controls are necessary, whether an organization is building and running its own cloud, or using the services from a cloud services provider.

Firewalls and next-generation firewalls

These access control devices are used to control the communications flowing to and from networks and specific endpoints by blocking unauthorized access as well as many types of intrusion attempts.

Intrusion prevention systems (IPSs) and next-generation IPSs

These systems watch for signs of malfunction, intrusion, and malware attacks. These devices detect and block the attacks that are permitted or find their way into organization networks. Of course, next-generation IPSs are the subject of this book!

Strict access controls

A well-designed access controls program is necessary to effectively secure a network, a system, or a cloud environment. Some of the characteristics of an effective access control system include:

- ✓ Formal access request process
- ✓ Least privilege access
- ✓ No shared accounts
- ✓ Access logging
- ✓ Strong password quality standards
- ✓ Periodic access reviews

Logging

Significant events at every layer of the cloud infrastructure need to be logged. Preferably, logging will be centralized for ease of management and the ability to correlate separate events and be able to see them as incidents.



Precise time synchronization is a key ingredient for accurate logging. Computers' time-of-day clocks are notoriously inaccurate; use NTP to synchronize all computer and network device clocks to well known standard time sources.

Change management

Change management is the formal process where all changes in an environment are formally requested, reviewed, scheduled, performed, and documented.

The heart of an effective change management process is a *period change review meeting*, where stakeholders discuss upcoming proposed changes. This helps ensure that changes will have the desired effect, be coordinated with the right parties, and help to reduce unscheduled downtime.

Configuration management

Developing good standards and using tools to ensure consistent configuration helps to make systems more resistant to intrusion and misuse. Configuration management tools can help to automate the settings on each virtual machine, enabling even instantaneous configuration changes across all systems in a virtualized environment.

Chapter 5

Deploying Next-Generation IPSs

In This Chapter

- ▶ Understanding the various deployment scenarios
 - ▶ Using integrated security policies for thorough network protection
-


This chapter discusses common next-generation IPS deployment scenarios that can fit into your enterprise. The most important thing is to understand what your business truly needs and then go about deploying your next-generation IPS in the proper ways. This chapter goes over some important deployment scenarios.

Next-Generation IPS Deployment Scenarios

A one-size-fits-all approach to network security is hardly ideal — purpose-built or best-of-breed solutions often better satisfy challenging requirements. There are many deployment scenarios where purpose-built next-generation IPSs make sense. Flexible and open solutions are needed to address new and advanced network security threats — especially when it comes to intrusion prevention systems. What works in one situation may fall flat in others.

Services and data protection

Application services and data are the core of what needs to be protected (see Figure 5-1). By deploying a next-generation IPS at the Internet edge or in a DMZ, you'll have the visibility into your network that is necessary to monitor for — and prevent — network attacks.




Source: Sourcefire

Figure 5-1: Augment existing firewalls with next-generation IPS's improved threat prevention and application visibility.

Client protection

You may wish to augment existing firewalls with improved threat prevention at the client level. Odds are you have a significant firewall investment that you're not ready to let go of. The next-generation IPS deployment scenario shown in Figure 5-2 helps strike a balance so you can use what you've got and still ensure the proper visibility and control for your network clients.




Source: Sourcefire

Figure 5-2: Implement application and user controls with next-generation IPSs to compliment firewall controls.

High-throughput passive IPS

As shown in Figure 5-3, you may have internally-hosted business applications where service availability and performance take precedent over threat protection. Nonetheless, your company needs to detect intrusions to meet compliance mandates. When that happens, you may wish to deploy a very high-throughput IPS in passive mode to detect and report on threats that may pose risks to or target core internal network infrastructure. Certain switches and routers are known to have vulnerabilities, but new threats targeting them may have evaded your perimeter protections. Knowing which critical applications and users are accessing core network infrastructure will shed light on potential sources of threats that can be addressed upstream via firewall access controls and intrusion prevention.




Source: Sourcefire

Figure 5-3: Core networks may require very high throughput threat inspection that only a next-generation IPS can deliver.

Private cloud connection

Figure 5-4 shows the scenario of a private B2B partner network. This configuration may not allow you to control the partner side of the network but, with strong threat inspection before entry into your environment, you can be assured that the traffic has been properly vetted.



Source: Sourcefire


Figure 5-4: Thorough threat inspection, which next-generation IPSs can provide, helps ensure secure connections to outside partner networks.



Make sure your next-generation IPS solution is architected for change — it's inevitable!

Inter-VM monitoring

The final common next-generation IPS deployment scenario is shown in Figure 5-5. By monitoring communications within your virtualization environment, you can find — and stop — applications and threats that would very likely go undetected otherwise. Virtual machine vulnerabilities are a choice target for attackers because they're often unprotected by traditional security infrastructure. A virtual next-generation IPSs can mitigate this shortcoming.



Source: Sourcefire

Figure 5-5: Inter-VM monitoring provides visibility and protection for virtualized environments that traditional protections don't address.

Using Integrated Policies to Your Advantage

Security policy enforcement can be quite difficult in less-integrated systems that feature multiple consoles and disparate technologies. With the proper next-generation IPS, you

can easily create singular, specific policies that incorporate all security elements that you now control — including networks, zones, applications, users, websites, file types, and host access. Such control includes traditional firewall rules (blocking HTTP over port 80), application rules (block BitTorrent peer-to-peer file sharing), block connections to malicious websites (that are known bot command-and-control hosts), except for certain user groups (development) when they occur within the trusted zone (an isolated development network segment). This level of control is also beneficial for permitted traffic threat inspection, perhaps applying stronger security settings such as a more robust signature set to inspect the traffic.

Chapter 6

Selecting the Right Next-Generation IPS

In This Chapter

- ▶ Developing next-generation IPS selection criteria
 - ▶ Understanding the unique requirements needed from enterprise and SMB organizations
 - ▶ Unraveling industry specific requirements
 - ▶ Exploring independent test labs
-

So here you are in the selection criteria section. You're probably thinking about getting a next-generation IPS for your organization now, or at least thinking about thinking about it. Or maybe you want to see what criteria other organizations use when they're ready to buy.

We suggest that you always develop objective criteria for any IT system, and then compare various products against your criteria.

In this chapter, we discuss selection criteria, starting with general requirements and moving to specialized requirements by company size and industry sector.

Common Next-Generation IPS Selection Criteria

When shopping for a next-generation IPS, you should consider many factors. Here are the primary characteristics that you need to be thinking about:

- ✓ **Detection effectiveness:** How well does a next-generation IPS detect unwanted traffic? The underlying technology may use a variety of means for detection including signature-based, anomaly-based, or both. Consider both *false positives* (a security finding that's flagged as a problem but really isn't) as well as *false negatives* (a security concern that's completely overlooked). Is the rule base visible so that you can examine rules or add more? At the end of the day, what percentage of relevant threats can this next-generation IPS catch? To find out more about how to determine effectiveness, see the section "Third-Party Testing" later in this chapter.
- ✓ **Awareness:** The solution must be capable of passively gathering information on the presence of network hosts and applications, analyzing behavior, correlating events, eliminating false positives, and ensuring policy compliance.
- ✓ **Scalability:** Rather than just consider what is needed today, also think about the future. What modes of change, growth, or future regulations or standards may require additional security devices, additional bandwidth, new technologies (such as virtualization and cloud computing), or different types of sensors (physical versus virtual)?
- ✓ **Automation:** Security is about sustainable and repeatable processes. The automation of controls, workflows, and analysis is essential. Although security demands increase every day, budgets and resource haven't kept pace — automation can alleviate much of the burden of event analysis and assessment.
- ✓ **Performance:** Make a purchasing decision with the long-term in mind. If any type of growth is anticipated at your organization, then you should select a security platform that will grow along with your needs without replacing hardware sooner than you're ready to.
- ✓ **Compliance:** Understand how your next-generation IPS investment may satisfy any relevant governmental and/or industry compliance regulations that affect your organization. In the case of PCI DSS, for example, some next-generation IPS solutions may satisfy more requirements than others.

- ✓ **Vision:** Consider whether you want to purchase a next-generation IPS from a market leader or from a company that just does what everyone else does. We suggest choosing a leader who is consistently respected for vision and execution, knowing that such providers tend to develop new kinds of detection and prevention long before the followers will even think of it.
- ✓ **Viability:** Buy a product from a company that will be in business for the long haul. There may be some advantages to buying some products from a startup or a garage outfit, but for something as strategic as IPS, we suggest buying from a company that you know will be around in five or ten years. Several years after purchase, you'll still want someone to answer the phone when you call.
- ✓ **Manageability:** There's nothing worse than a product that is difficult to figure out and operate. Most organizations want a fully-configurable next-generation IPS with the ability to tailor and tune the solution to their unique requirements — even if they don't plan on tinkering with the details too often.
- ✓ **Support:** Everyone gets stumped now and then, and every product is going to be prone to hardware or software trouble, no matter how good its quality program is. Choose a company that stands behind its product and is ready to offer whatever kind of help you need.
- ✓ **Cost:** Don't be afraid to understand and specify your spending limits. There's hardly anything worse than a bad purchase, especially if you realize that your investment can't provide the return you need.
- ✓ **Value:** Perhaps more important than costs, determine what is the cost per protected measure of inspected throughput (for example: \$/protected Mbps of throughput). This is often the best indicator of the value a solution can deliver and a good metric on which to compare alternatives.

In the rest of this section, we discuss requirements that are specific to small organizations, large (enterprise) organizations, and government.

Enterprise (large organization) buying requirements

Enterprises are typically those organizations with, say, 500 or more employees. They generally have multiple business locations, often in more than one country. Usually they have larger IT organizations with network engineers, system administrators, information security managers, IT operations managers, and other individuals and departments — in other words, a lot of people who get involved in things like IPS because it affects many people in the organization. Enterprises often prefer best-of-breed solutions that address all requirements proficiently, and are generally less willing to compromise.

In addition to the general requirements discussed earlier in this chapter, enterprises are generally also interested in some of these requirements:

- ✓ **Management:** Rather than just a single administrative user for their next-generation IPS, enterprises need an IPS that can support many users and different roles.
- ✓ **High throughput:** Processing speed is doubling every two years. Thus, you will continue to see network speeds grow. Next-generation IPS vendors, in particular, should have a broad range of products to support the smallest to the very largest network needs.
- ✓ **Forensics:** Enterprises need their IPSs to be able to provide forensics-quality information to support security events related to sophisticated threats or those that may find their way into the criminal justice system as evidence.
- ✓ **Fault tolerance:** Enterprises build high-availability, fault-tolerant infrastructures to support high-demand applications. These organizations need IPSs that can match the *five-nines* availability environments they support, meaning there is practically zero minutes of unscheduled downtime per calendar year.

- ✓ **Low TCO:** Although enterprises have larger operating budgets than smaller organizations, they also have greater demands for securing the network. Thus, enterprises must select a next-generation IPS that helps them to work smarter — not harder — by automating key functions, such as impact assessment, user identification, and IPS tuning.

Small-to-medium-business (SMB) buying requirements

Hats off to small and medium-sized businesses (those with fewer than 500 employees) that recognize their need for next-generation IPS and wade into the fray!

To the requirements listed earlier in this chapter, add one more: SMB customers are looking for ease of management. They don't have a deep staff to take training courses and spend weeks planning their IPS implementation. In the SMB world, the IT guy (or gal) who has 12 other jobs besides security just wants to open the box and have the system running in a couple of hours. This means a next-generation IPS that offers easy setup and easy-to-understand configuration without having to take a week-long class on managing the device. They just want to set it and forget it!

Two types of next-generation IPS users

Regarding how they approach security, there are generally two types of next-generation IPS users. First, there are “lean forward” users that truly care about security and are very cautious, but use their fear to inspire themselves to gain more knowledge. They see the big picture and meet security problems head on.

Then there are “lean back” users that are either in senior-level positions or simply don't have time to

spend monitoring and tuning the IPS. Some “lean back” users are driven by regulatory compliance and simply want to check the box to satisfy compliance. They tend to focus on the short-term.

It's not only important to gauge the organization's goals for next-generation IPS usage, but also understand the types of users that will interact with the platform.

Government buying requirements

Governments, especially leading national governments, are tough customers; primarily because they know what they want and they communicate this through a comprehensive set of requirements. In addition to the general requirements at the beginning of this chapter, plus the requirements of enterprises, governments often ask for these additional requirements:

- ✔ **Custom rules:** Some government organizations are required to “throw out” IPS rules provided by the manufacturer in favor of creating custom rules for proprietary systems. Selecting an IPS with an open architecture and easy-to-use rule-creation wizard is optimal for such organizations.
- ✔ **IPv6 support:** U.S. federal and other government regulations require all IT systems to be IPv6 compatible. In the case of an IPS, it must be capable of detecting and blocking IPv6 attacks and be managed on an IPv6 network.
- ✔ **Federal Information Security Management Act (FISMA) compliance:** This is a complete end-to-end security framework required of all federal information systems and supporting environments. FISMA requires federal agencies (and their service providers) to establish and carry out a security plan, maintain IT asset inventories, categorize information and information systems according to risk level, enact security controls, perform risk assessments, perform continuous monitoring, and undergo periodic certification and accreditation. A next-generation IPS is an essential tool for achieving FISMA compliance.
- ✔ **NIST compliance:** Government customers frequently cite various NIST (National Institute of Standards and Technology, the U.S. government’s IT standards setting organization) standards as part of their next-generation IPS selection criteria, especially NIST Special Publication 800-94, *Guide to Intrusion Detection and Prevention Systems*.
- ✔ **Evaluation Assurance Level (EAL):** Government customers may require that an IPS be tested and certified to a specific EAL standard: EAL testing is extremely expensive, so any vendor that claims EAL compliance is noteworthy for any government or non-government customer.

Industry-Specific Considerations

Organizations in some industries will impose additional requirements on security vendors, generally as a *pass through* — these organizations are asserting requirements on the suppliers that are imposed upon them by a third party.

Public utilities

Power, water, natural gas, and other public utilities rely on Supervisory Control and Data Acquisition (SCADA), Process Control Network (PCN), and Smart Grid technology for remote control and monitoring of utility equipment. These systems are almost always IP-based and frequently utilize the public Internet for transmission.

An IPS helps to secure SCADA, PCN, and Smart Grid systems by detecting and blocking intrusions that could include terrorist attacks. Leading IPS solutions may offer special SCADA, PCN, and/or Smart Grid rule sets and may also incorporate passive network intelligence collection for correlating threats without actively scanning the network.

Healthcare

Healthcare providers and other industry organizations subject to HIPAA requirements need to incorporate next-generation IPSs into their network infrastructure as part of their technical safeguards for patient record privacy and so on. These organizations' requirements will often resemble those required for most enterprises, as discussed earlier in this chapter.

Financial

Banks, credit unions, brokerages, and insurance companies are required to protect sensitive customer information from theft and abuse. These organizations will often impose enterprise level requirements, including enterprise scalability and management.

The Gramm-Leach Bliley Act and Basel II are the primary regulations requiring financial institutions to protect their systems and networks.

Retail

The retail industry has certainly had its fair share of abuse in recent years. With all the headlines of hackers and malicious employee abuse of payment information including credit card numbers, the Payment Card Industry Data Security Standard (PCI DSS) has been a primary focus for businesses both large and small.

Telecommunications

Common carriers, including telecommunications providers and Internet service providers, have the world's most extensive networks over which the world's Internet and private communications take place. Most of these organizations are under market or regulatory pressure to provide five-nines availability. Such organizations will require the most robust next-generation IPS platforms, including support for high-throughput environments, fault-tolerant hardware, and configurable/fail-open interfaces.

Hardware Considerations

Organizations shopping for next-generation IPSs need to understand what hardware features are important for them. Hardware-centric requirements will generally fall into these categories:

- ✓ **Purpose-built appliances:** Organizations may wish to specify whether they're looking for IPS software that they would install on their own servers, generic appliances, or a purpose-built appliance with IPS features built into the hardware. If you consider a next-generation IPS vendor with purpose-built appliances, ensure that this doesn't hinder the extensibility of the solution by verifying the availability of Virtual IPS offerings for VMware or other virtualization platforms.

Some IPS appliances rely on ASICs (application-specific integrated circuits) to accelerate certain network processing functions. Although ASICs make it easier for the vendor to achieve higher throughputs, it usually makes it more difficult for them to port their software to VMware and other virtual platforms. Even if you don't have a budgeted virtualization security project today, you will tomorrow. Be sure to select a next-generation IPS partner that offers both physical and virtual appliances, ideally with a common management console, so you don't eventually end up with two sets of administrative interfaces.

- ✓ **Hardened operating system:** Organizations' requirements may be as detailed as specifying the desired operating system that supports the next-generation IPS software. Most of today's next-generation IPS products incorporate a hardened Linux OS in their appliances.
- ✓ **Fault tolerance:** Organizations may specify various next-generation IPS fault tolerance features including redundant power supplies, disk drives, fans, and fail-open interfaces.
- ✓ **Fail open:** Organizations doing their homework will want a next-generation IPS appliance that fails open, meaning, in the event of a catastrophic hardware failure, network traffic will continue to flow through the IPS appliance uninterrupted. This feature requires special hardware not found in general purpose appliances.

Third-Party Testing

There are two independent test laboratories in particular that actively test security products: ICSA Labs (www.icsalabs.com) and NSS Labs (www.nsslabs.com). These companies evaluate leading IPS devices for accuracy, reliability, and performance. Organizations that are serious about the desired quality of their IPS systems should consider only products that have been independently evaluated by a reputable third-party testing organization.



Test reports on leading IPS products may be purchased from ICSA Labs and NSS Labs directly or can often be obtained at no charge from the IPS vendors themselves.

Chapter 7

Keys to Effective Next-Generation Security

In This Chapter

- ▶ Ensuring optimal defenses with a next-generation approach
- ▶ Examining four ways to improve how you manage information risks

Information security requires a comprehensive set of processes and tools. One gap and you're likely to overlook something meaningful or impactful — often both. A next-generation IPS plays an important role in the system. By leveraging real-time network, application, behavior, and user awareness to optimize effectiveness, a next-generation IPS empowers you to automate key security functions and create an agile and responsive threat protection system.

Here are four ways to ensure effective next-generation security in your business:

- ✓ **Give complex threats the visibility they need.** To thoroughly protect, security organizations need to fully understand their networks and act in real time to the frequent changes that occur. This understanding requires asset mapping, contextual awareness, cross-source correlation, and total network visibility. And you need the ability to continually analyze — and respond — to change as it occurs. Only when you continually observe and act on changes within your network can you eliminate blind spots that provide attackers the opportunity they seek.
- ✓ **Seek control without compromise.** Many of today's security controls eliminate unproductive and/or risky application access, and thereby reduce exploitable vulnerabilities. This allows organizations to reduce attack

surface area and minimize risks. However, this control needs to be selective and flexible to suit each unique network environment. Inflexible or overly-restrictive policies alone (example: Block all Facebook) will likely meet strong user resistance and/or lead to excessive false positives that become the bane of security and user organizations alike. A well-rounded solution provides the confidence and capability to tailor defenses for targeted threats while protecting completely without compromise to security posture.

- ✓ **Automate security for agility.** Threats evolve too rapidly for manually-tuned defenses to keep pace. IT consumerization, device mobilization, virtualization, and cloud-based computing create a fluid, boundless world to secure. Customers need the agility to stay protected despite the rapid changes and complexity; security automation is the key to keeping pace and discerning what really matters. New risks can be acted on quickly by tuning security defenses automatically — this can entail auto-applying additional signatures, auto-blocking unknown applications, files, sites, or users, auto-triggering authentication or remediation workflow, and so on. Automated event analysis and assessment can also reduce actionable events, concentrating security staff remediation efforts on the items of greatest importance. By automatically assessing changes and tuning defenses, organizations can adapt responsively to ensure they maintain their security posture and stay protected.
- ✓ **Maintain flexibility and openness.** Ensure that your chosen solution has the built-in flexibility and performance to grow and scale with your needs and address new security requirements as they emerge. Change is inevitable and your security solutions must be architected to adapt. Organizational factors such as who owns application security and who maintains control policy will influence the chosen path. A flexible security solution will accommodate varying managerial roles (and provide the correspondingly appropriate administrative privileges), so that administrators from each group can leverage the solution for their responsibilities. You should be able to easily tailor rules and implement custom defenses for specialized or proprietary applications. Many less flexible or closed solutions make this quite difficult if not impossible.

Harness the power of next-generation IPSs to improve security in your organization's networks!

To stay on top of security, you need to fully know what's taking place on your network and hosts and better understand threat context. Security awareness and controls built into next-generation IPSs provide you with deep insight into what's actually happening in and around your network.

- **Next-generation IPS basics** — *understand the differences between traditional intrusion prevention systems and today's next-generation intrusion prevention systems*
- **What you're protecting** — *why protecting against dynamic threats requires real-time network visibility and agility*
- **Examining the details** — *the many features and functions found in next-generation IPSs, including application visibility/control, context and content awareness, security automation, and user identification*
- **Working with virtualization and the cloud** — *understand how you can use next-generation IPSs with virtualization and cloud computing technologies*
- **Picking a next-generation IPS** — *get your next-generation IPS shopping list organized so that you can be sure to get the next-generation IPS that is right for your organization*



Open the book and find:

- Information on real-time visibility
- How to use application control to your advantage
- Tips on developing next-generation IPS selection criteria
- Keys to effective next-generation security
- Working with the cloud and virtualization

Go to **Dummies.com***

for videos, step-by-step examples, how-to articles, or to shop!

David Stuart is Director, Product Marketing at intelligent cybersecurity leader Sourcefire, Inc. and a 20-year high-tech veteran. **Kevin Beaver** is an Independent Information Security Consultant and author of *Hacking For Dummies*.

WILEY

978-1-118-55497-5
Not for resale