

LEARNING MADE EASY

Ixia Special Edition

Network Visibility

for
dummies®
A Wiley Brand



Compliments
of

íxía

Get total visibility
of network traffic

Deploy high-availability
monitoring

Enable fail-safe
inline security

Lawrence C. Miller, CISSP

About Ixia

Ixia provides testing, visibility, and security solutions that strengthen network applications and are used worldwide by enterprises, governments, service providers, and network equipment manufacturers (NEMs) to verify their product and service designs, optimize performance, and ensure security.

Ixia uses high-performance architecture to deliver solutions that can handle the volume and complexity of real-world IT and provide complete visibility across an organization's physical and virtual networks. Ixia solutions minimize security blind spots in any size network, and maximize security resilience by exposing vulnerabilities others miss. Ixia solutions also feature an innovative and easy-to-use interface that accelerates configuration changes and makes operations more efficient.

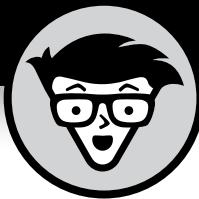
In the unpredictable and highly changeable world of information technology (IT), organizations need solutions that provides a complete understanding of user behavior, security vulnerabilities, network capacity, application performance, and information technology resiliency. Customers turn to Ixia to validate products and services designed for network connectivity, to test the integrity of their security infrastructure, and to monitor network performance and expand their infrastructure without interruption.

Ixia offers companies trusted environments in which to develop, deploy, and operate, to exceed their customer's expectations and achieve better business outcomes.

Network Visibility

for
dummies[®]

A Wiley Brand



Network Visibility

Ixia Special Edition

by Lawrence C. Miller, CISSP

for
dummies[®]

A Wiley Brand

Network Visibility For Dummies®, Ixia Special Edition

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

Copyright © 2017 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, For Dummies, the Dummies Man logo, 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. Ixia and the Ixia logo are trademarks or registered trademarks of Ixia. 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 978-1-119-37949-2 (pbk); ISBN 978-1-119-37959-1 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. Some of the people who helped bring this book to market include the following:

Project Editor: Martin V. Minner

Senior Acquisitions Editor: Amy Fandrei

Editorial Manager: Rev Mengle

Business Development Representative:

Karen Hattan

Production Editor: Antony Sami

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions.....	2
Icons Used in This Book.....	2
Beyond the Book.....	3
Where to Go from Here.....	3
 CHAPTER 1: Improving Your Security Posture	5
Recognizing the Challenges to Network Security.....	5
The Digital Warfare Attackers Are Waging	8
Recognizing Opportunities to Upgrade Network Security	10
Case Study: Compliance Mandates Drive Urgent Need for Inline Security	11
 CHAPTER 2: Focusing First on Visibility.....	13
Looking at the Problem: Too Many Blind Spots, Not Enough Visibility.....	13
Understanding Why You Need Network Visibility.....	14
Evaluating Current Approaches to Visibility.....	16
Achieving Total Visibility	17
Leveraging Network Taps to Deliver Superior Visibility.....	18
Getting Visibility into Virtual Traffic.....	19
 CHAPTER 3: Increasing Efficiency and Flexibility with a Security Fabric.....	21
Delivering the Right Data, to the Right Place	21
Defining Inline and Out-of-Band Tools	22
Protecting and Processing Your Data	24
Understanding the Role of a Security Fabric	25
Tool efficiency	25
Network availability.....	27
Security resilience.....	28
Intelligent data processing	30
Case Study: Retail Bank Gains More Value from Security Tools.....	31
Security fabric results in \$680,000 cost savings	32
 CHAPTER 4: Adding Application and Threat Intelligence to the Security Fabric.....	33
Streamlining Security With a Pre-Filtering Appliance	33
Incorporating Real-Time Expert Security Analysis.....	35

Performing Advanced Troubleshooting	36
Ensuring Regulatory Compliance	38
Case Study: Hosting Service Tackles Attack Traffic with Ixia ThreatARMOR	39
Intrigued by pre-filtering.....	39
Intrusions reduced by 80 percent	40
Out-of-band deployment upgraded to inline.....	40
CHAPTER 5: The Benefits of a Total Visibility Security Fabric	41
Maximizing the Value of Security and Monitoring Tools.....	41
Optimizing Incident Response to Reduce Mean Time to Repair	43
Enabling High-Availability and Fast Failover	43
CHAPTER 6: Building Confidence through Security Testing.....	45
Validating Your Designs and Configurations	45
Reducing Risk and Costs.....	46
Case Study: Leading Financial Institution Secures Its Network from Cradle to Grave	47
Validating the initial design	47
Implementing the design.....	48
Looking ahead.....	48
Staying Vigilant About Updates and Patches.....	49
Stress Testing Your Network Environment	49
CHAPTER 7: Ten Requirements for Security and Network Monitoring	51
Proven Lossless Monitoring.....	52
Visibility to Virtual and Cloud Resources.....	52
Fail-Safe Deployment for Inline Appliances	53
Maintenance Without Network Downtime	53
Non-Disruptive Scalability	54
High-Performance Data Processing.....	54
Context-Aware Visibility.....	55
Simple Management Interface	56
High Availability and Fast Failover.....	56
Integrated Application and Threat Intelligence	57
Remembering the Ten Key Requirements	57
GLOSSARY	59

Introduction

When designing their enterprise security architectures, many IT organizations focus on what their security appliances and monitoring tools can do, rather than how to provide those tools with a complete view of the data moving through the enterprise. But the fact is — you can't monitor what you can't see. Evaluating the traffic entering and leaving your network is your best defense against attacks and errors that can lead to network outages, loss of sensitive data, or customer frustration. Buying and deploying a growing array of sophisticated security tools is just the start. You must feed your tools all the relevant data they need, as quickly as possible, to keep your organization secure.

To establish strong network security and enable your security solutions to perform as expected, you need to implement strong visibility solutions that provide access to all your network traffic and provide exactly the right data — in real-time. Providing this data to your security tools securely and efficiently requires a resilient visibility architecture that is optimized for continuous, real-time network monitoring.

About This Book

This book describes how you can achieve truly resilient and effective IT security, and is conveniently organized into seven short chapters that:

- » Explore modern security challenges and ways to strengthen your organization's IT security posture (Chapter 1)
- » Explain why total visibility is fundamental to network security (Chapter 2)
- » Examine how to protect your organization with continuous network security monitoring (Chapter 3)
- » Describe the purpose of integrating a security and threat intelligence feed (Chapter 4)
- » Detail the benefits of a total visibility security fabric (Chapter 5)

- » Stress the importance of security testing (Chapter 6)
- » Summarize the important requirements for resilient network security (Chapter 7)

There's also a glossary at the end of the book — just in case you start feeling overwhelmed by any acronyms or technical terms used in this book!

Foolish Assumptions

It's been said that most assumptions have outlived their uselessness, but I'll assume a few things nonetheless!

Mainly, I assume that you know a little something about network security. As such, this book is written primarily for technical readers who evaluate, implement, or operate network security solutions.

If that describes you, you're in the right place! If it doesn't, keep reading anyway. It's a great book and you'll learn a few things about network security!

Icons Used in This Book

Throughout this book, I occasionally use icons to call attention to important information that is particularly worth noting. You won't find smiley faces winking at you or any other cute little emoticons, but you'll definitely want to take note! Here's what you can expect:



REMEMBER

This icon points out information or a concept that may well be worth committing to your nonvolatile memory, your gray matter, or your noggin' — along with anniversaries and birthdays!



TECHNICAL STUFF

You won't find a map of the human genome or the secret to cold fusion here (or maybe you will, hmm), but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon and is the stuff legends — well, nerds — are made of!



TIP

Thank you for reading, hope you enjoy the book, please take care of your writers. Seriously, this icon points out helpful suggestions and useful nuggets of information.



WARNING

This is the stuff your mother warned you about. . . well okay, probably not. But these helpful alerts do offer practical advice to help you avoid potentially costly mistakes.

Beyond the Book

There's only so much I can cover in 72 short pages, so if you find yourself at the end of this book, thinking "Gosh, this was an amazing book, where can I learn more?" just go to www.ixiacom.com.

Where to Go from Here

With my apologies to Lewis Carroll, Alice, and the Cheshire cat:

"Would you tell me, please, which way I ought to go from here?"

"That depends a good deal on where you want to get to," said the Cat — er, the Dummies Man.

"I don't much care where . . . ,” said Alice.

"Then it doesn't matter which way you go!"

That's certainly true of *Network Visibility For Dummies*, Ixia Special Edition, which, like Alice in Wonderland, is also destined to become a timeless classic!

If you don't know where you're going, any chapter will get you there — but Chapter 1 might be a good place to start! However, if you see a topic that piques your interest, feel free to jump ahead to that chapter.

Each chapter is individually wrapped (but not packaged for individual sale) and written to stand on its own, so feel free to start reading anywhere and skip around!

Read this book in any order that suits you (though I don't recommend upside down or backward). I promise you won't get lost falling down the rabbit hole!

4 Network Visibility For Dummies, Ixia Special Edition

IN THIS CHAPTER

- » Understanding IT transformations that create security challenges
- » Identifying events that can trigger a network security upgrade

Chapter **1**

Improving Your Security Posture

Organizations everywhere are generating an increasing volume of data and network traffic. Defending the enterprise against increasingly advanced cybersecurity threats, while simultaneously managing this deluge of data and traffic has, likewise, become increasingly challenging.

This chapter briefly reviews aspects of today's business and technological environment that make it harder to secure your organization from threats and vulnerabilities. This chapter also introduces some of the network lifecycle events that can trigger an upgrade.

Recognizing the Challenges to Network Security

Technological trends are enabling new opportunities for modern businesses to compete in today's global economy. At the same time, these trends provide obstacles to seeing and understanding the traffic flowing into and out of the enterprise:

- » **Expanding network perimeter:** The network perimeter, once clearly demarcated by a firewall between the enterprise



REMEMBER

network and the Internet, is disappearing as users become more mobile and the organization does more business in the cloud. Customers and employees are no longer tethered to offices, and they expect to have access to applications and data from anywhere, at any time, and on any device. Without a defined perimeter, it's harder to identify all the traffic flows that can impact the enterprise.

Gartner predicts that by 2018, 25 percent of corporate data traffic will bypass traditional security defenses and flow directly between mobile devices and the cloud.

- » **Increasing use of SSL encrypted traffic:** Research by Sandvine in early 2016 predicted that nearly three-fourths of all Internet traffic would be encrypted in 2016. Similarly, Google reported in the first quarter of 2016 that nearly 80 percent of its worldwide traffic (excluding YouTube) was encrypted. Unfortunately, attackers also use SSL encryption to hide threats and attack traffic. Thus, organizations can no longer arbitrarily pass SSL encrypted traffic to their internal networks without inspection. IT departments now commonly decrypt inbound and outbound SSL traffic, before it is delivered to its final destination, to identify risks and threats such as regulatory compliance violations, data leakage, malware, intrusion attempts, and attack communications.
- » **Growing volume and complexity in network traffic:** Network traffic is now largely comprised of structured and unstructured data, which includes significant amounts of voice and video traffic. The volume of network traffic can flood your existing security tools with more traffic than they were designed to handle and lead to oversubscription. The *Cisco Visual Networking Index: 2015-2020* notes that busy-hour Internet traffic grew more rapidly than average Internet traffic, at 51 percent in 2015, compared with 29 percent in average traffic. Busy-hour traffic is expected to increase by a factor of 4.6 between 2015 and 2020.
- » **Virtualization:** Virtualization has been one of the most disruptive IT innovations in the modern computing era. Virtualization technology allows IT to create compute, storage, memory, and networking resources using software running on standard, low-cost servers. In a virtual environment, IT resources, like compute servers, network switches, or storage systems are quickly and easily created on demand

on standard x86 platforms and just as easily decommissioned when the need is gone.

The impact of virtualization on the enterprise is an increase in “east-west” traffic: data that travels between these virtual resources on the same physical host or inter-blade traffic on the same server. The traditional methods of observing traffic as it passes between two physical locations along a network cable is not sufficient for monitoring virtual traffic. To solve this problem, you deploy to a software-based virtual tap to monitor inter-VM traffic and deliver it to your monitoring tools.

» **Cloud computing:** Virtualization has also enabled the cloud computing trend, including private, public, and hybrid clouds, infrastructure-as-a-service (IaaS), and platform-as-a-service (PaaS). Scalability and control over allocation of resources are major advantages of both virtualization and cloud. However, as organizations migrate workloads from their data centers to public clouds, it becomes more difficult to observe and monitor data flows, creating new blind spots.

And enterprises aren’t just using a single cloud environment or a single application in the cloud. Verizon’s 2016 *State of the Market: Enterprise Cloud 2016* study reports that more than 40 percent of businesses already use five or more cloud providers, and Okta’s 2016 *Business@Work* study found the average company uses between 10 and 16 public cloud apps.

Although the cloud offers unprecedented flexibility, it also forces companies to extend their traditional network perimeter — often into places where they don’t have visibility or control. These blind spots can rapidly become havens for security threats lurking in the shadows (which I cast a light on later in this chapter).

» **Cloud applications and software-as-service:** Trust and adoption of cloud-based software has grown along with technology advancements in processing power, storage solutions, memory, and networking. The *Cisco Visual Networking Index: 2015-2020* predicts that adoption of critical business applications as a service, a subset of software-as-a-service (SaaS), will accelerate through 2020. The first users of SaaS were consumers, followed by small and medium-sized



REMEMBER

businesses. Now, as SaaS solutions become more sophisticated and robust, larger enterprises are adopting these services as well. “Enterprises with big budgets, data centers and complex applications are now looking at cloud as a viable place to run core business applications,” said Dave Bartoletti of Forrester Research. Growing use of cloud applications will make them a target for cyber attacks.

Cisco’s 2015-2020 *Visual Networking Index* predicts that cloud apps will account for 90 percent of worldwide mobile data traffic by 2019.

» **Internet of Things (IoT):** Finally, the IoT trend has created a new market for billions of “smart, connected” devices and products — from security cameras to smart thermometers, appliances, and cars — that are connected to the Internet and collect, store, transmit, and share massive amounts of information. Many of these IoT devices will leverage new computing models — such as mobile edge computing (MEC) and “fog” computing — to extend the network perimeter still further. And unfortunately, in the race to be first to market, security in many of these devices and products is often overlooked or, at best, an afterthought.

IoT will be a massive data generator. Gartner predicts IoT will include 26 billion units installed by 2020, and IoT product and service suppliers will generate incremental revenue exceeding \$300 billion, mostly in services. For IoT to really thrive, users will need to embrace open standards that enable data access, security monitoring, and performance analytics.



REMEMBER

The Digital Warfare Attackers Are Waging

Attackers are increasingly using more sophisticated techniques to evade enterprise security. For example, Gartner predicts that more than half of cyberattacks in 2017 will use SSL encryption to hide attack communications. Some other advanced weapons in the modern attackers’ arsenal include:

» **Advanced persistent threat (APT):** A deliberate, focused attack that may go undetected for several years, perpetrated by an attacker (for example, a nation-state or organized

criminal organization) with extensive resources including money, people, time, and sophisticated tools (such as electronic surveillance equipment and satellite imagery). In addition to stealing data, an APT may be used to gradually destroy physical assets (such as nuclear centrifuges or integrated circuits).

- » **Botnet:** A network of infected endpoints (known as bots) working together and controlled by an attacker through command-and-control (C2) servers.
- » **Distributed denial-of-service (DDoS):** A coordinated attack, often from hundreds of thousands or millions of compromised endpoints, used to flood a target system or network with excessive traffic so that it cannot process legitimate traffic.
- » **Exploit:** Software or code that takes advantage of a vulnerability in an operating system or application and causes unintended behavior in the operating system or application, such as privilege escalation, remote control, or a denial-of-service.
- » **Hijacked IP address ranges:** IP addresses that are stolen from their legitimate owners, typically by corrupting the routing tables of Internet backbone routers. Once hijacked, they are used for malicious purposes, such as phishing and malware distribution.
- » **Malware:** Malicious software or code that typically damages or disables, takes control of, or steals information from a computer system. Malware broadly includes adware, anti-AV software, backdoors, bootkits, logic bombs, RATs, rootkits, spyware, Trojan horses, viruses, and worms. Advanced malware uses techniques such as metamorphism and polymorphism to avoid detection by signature-based tools.
- » **Phishing:** A social engineering technique in which an email that appears to be from a legitimate business, typically a financial institution or retail store, attempts to trick the recipient into clicking an embedded link in the email or opening an attachment containing malware or an exploit. The embedded link redirects the recipient's browser to a malicious website to enter sensitive personal information, such as account information. Alternatively, the malicious website may deliver malware or an exploit to the victim's endpoint in the background via the browser (known as a *drive-by-download*).

Recognizing Opportunities to Upgrade Network Security

Although having network visibility and control is critical to effective security, proposing a project to implement a visibility architecture may lack the “wow” factor of other IT projects. Don’t despair; several key events can provide a logical opportunity for IT to speak with stakeholders and executives about the need to upgrade security architecture and the benefits that can be obtained (see Chapter 5):

- » **Firewall migration and optimization:** The implementation of next-generation firewalls (NGFWs) or intrusion prevention system (IPS) appliances are great opportunities to propose a visibility architecture (discussed in Chapter 2) that can improve the performance of these important security systems. Maintenance windows are precious, so when it comes time to upgrade your firewall, an external bypass can help minimize the time spent configuring a resilient path for your new devices. Configuration of a typical NGFW or IPS appliance can take anywhere from two to four hours, but a bypass switch can be up and running in minutes.
- » **Geographic or data center expansion:** Expanding your network or data center to new geographies naturally leads to discussions about how to gain visibility across new network segments. Network taps have several advantages as a method for achieving network visibility for security and performance monitoring. Physical taps provide permanent inline access to all the traffic moving across a network link; virtual taps do the same for traffic moving between virtual resources. Taps operate at line-rate speeds and do not introduce delay or alter the content or structure of the data. They also fail open, so traffic continues to flow between network devices in the event a monitoring device is removed or power to the device is lost.
- » **Network upgrade:** Re-visiting your visibility architecture during a network upgrade (such as a cable, plant, or switch upgrade) is prudent. If you are upgrading capacity, a properly planned and deployed visibility architecture can ensure your existing security tools — even those designed for slower links — can still get the data they need when

they need it, and may help to reduce your switch port requirements.

- » **Major new service rollout:** Any time you are rolling out new services in your network, data center, or cloud environments, you should look at how the changes affect visibility and control. A large national communications carrier, for example, used the upgrade of its IPTV infrastructure to 100GE to revisit its visibility strategy. The company deployed new network taps on the faster links and a network packet broker to filter the volume of new traffic and deliver the appropriate data to support network monitoring and troubleshooting.
- » **Compliance mandates and audits:** Regulatory compliance and security audits are becoming increasingly common for organizations in practically every industry. Many of these mandates specifically require visibility of data in your network, and failure to comply can result in stiff penalties (See the case study “Changing compliance mandates drive urgent need for inline security”). Another common use case is complying with lawful intercepts, which occur when a company is required to provide detailed network data at the request of a legal authority, for the purpose of analysis or evidence.
- » **Recovery after a security incident:** In the immortal words of Sir Winston Churchill, “Never let a good crisis go to waste.” If your organization, or an organization in your industry, is unfortunate enough to suffer a security breach, make sure the “lessons learned” are actually learned — and implemented. Take the opportunity to improve visibility and control in your network.

Case Study: Compliance Mandates Drive Urgent Need for Inline Security

Financial institutions worldwide must comply with numerous compliance mandates, such as the Payment Card Industry’s Data Security Standards (PCI DSS), designed to mandate controls around cardholder data and reduce credit card fraud. In this case, a leading payment processing company in Europe, handling more

than 30 million transactions per day, urgently needed to demonstrate PCI DSS compliance, to avoid potential fines of up to \$10,000 per day.

The company planned to deploy Cisco firewalls and Sourcefire IPS to proactively inspect traffic entering its networks for cyber-attacks. However, direct inline deployment would have added numerous single points of failure on each network link and significantly increased the risk of network downtime.

Instead, the company decided to operate its new security solutions on an Ixia Security Fabric with redundant bypass switches and network packet brokers (NPBs) in a high availability configuration that enables continuous proactive security inspection.

Ixia iBypass switches monitor each appliance using high-speed heartbeat packets. If a device stops responding for any reason: power or port failures, traffic congestion, or planned maintenance, the bypass automatically re-routes traffic to maintain network availability. The new security architecture allows the company to more easily satisfy PCI-DSS requirements for continuous traffic monitoring and regular updating of security solutions. The security team was also pleased it could deploy new configurations across bypass switches and NPBs from a centralized management interface.

And the solution delivered a return on investment because the company was able to reduce planned capital purchases. Rather than having to install a firewall and IPS on each network link, the security fabric solution let the company aggregate traffic from multiple network links and deliver it for inspection by a single, high capacity firewall and IPS. The \$290,000 Ixia solution saved the company approximately \$3 million in additional firewall and IPS purchases — a full 10x savings.

IN THIS CHAPTER

- » Identifying blind spots in your network
- » Recognizing the benefits of total network access
- » Assessing current approaches to visibility
- » Using network and virtual taps to see enterprise traffic

Chapter 2

Focusing First on Visibility

Many IT organizations think more about the capabilities of their security appliances and monitoring tools than they do about how to provide those tools with a complete view of the data that moves across their enterprise. The fact is — you can't monitor what you can't see.

This chapter looks at how to achieve complete visibility of the traffic flowing into and out of your network, to safeguard your organization and support its mission.

Looking at the Problem: Too Many Blind Spots, Not Enough Visibility

Unfortunately, many IT organizations don't focus as much as they should on building an effective visibility architecture. They may believe they already have all the visibility they need or don't know how to evaluate a visibility solution.

Today's enterprise, however, is a complex entity, that is evolving beyond the data center. More and more, enterprises are moving

to a hybrid cloud model where they use both private and public cloud infrastructure to operate their business. The diverse hardware and software platforms used for mission-critical applications become “opaque containers,” into which IT may not have sufficient visibility.

The result is blind spots that limit visibility into enterprise applications and data flows. Achieving total visibility means having the ability to see into these containers so you know which of your resources are in use, who is using them, where bottlenecks might exist, whether your infrastructure is secure, and much more.

A 2016 Rightscale survey found that modern enterprises are using at least six clouds, on average. A single cloud network presents visibility challenges, but six or more cloud environments inevitably create significant blind spots.

Blind spots limit visibility into enterprise applications and data flows, and can lead to:

- » Application performance issues
- » Network downtime and service disruptions
- » Compliance issues
- » Data breaches

The goal of a total visibility architecture is to give you access to all the data that crosses your networks, so you can make informed decisions about how to best protect your business and its data, and ultimately deliver an excellent customer experience.



Visibility is the term used to describe a data distribution process that intelligently delivers raw, unprocessed data to your analytics and security tools. It is not a type of data analysis or monitoring application, or even deep packet inspection. Visibility, pure and simple, is data gathering and distribution.

Understanding Why You Need Network Visibility

All networks are inevitably exposed to increasingly complex and advanced security risks and threats. The key is to identify the risks and threats as quickly as possible and take effective action.

Visibility enables you to see data wherever it resides in your network and resolve issues, such as:

- » **Troubleshooting network performance:** In many cases, application performance is tied to network performance. When applications run slowly or stop working, you need real-time network diagnosis to pinpoint the root cause. A proper visibility solution allows you to immediately isolate the application, user, geo-location, and device experiencing problems — so you can identify and resolve the issue promptly.
- » **Protecting and securing the network:** Cyberattacks, including new malware instances and distributed denial-of-service (DDoS) attacks, continue to grow in number and sophistication. Every time a new attack occurs or a new vulnerability is discovered, a new firewall rule, intrusion prevention system (IPS), anti-malware signature update, or other network security countermeasure is reactively implemented to counter the new threat. Protection requires constant updates and changes to remain relevant. If you aren't proactively and continuously monitoring network traffic using a total visibility security fabric, you're leaving your organization vulnerable to cybersecurity threats.
- » **Monitoring application performance and reliability:** You need applications to perform not just today, but for the long term. Network-centric applications must be continuously and precisely monitored for reliability and performance. You've heard the phrase "Garbage in, garbage out." In the same vein, application monitoring tools are only as effective as the data they are fed.
- » **Optimizing performance of complex network infrastructure:** Today's networks have grown more complex with multiple data centers, data encryption, sophisticated security and analytic tools, and worldwide mobile users demanding instant access to applications and data. The monitoring tools you use will help you achieve excellent performance, but *only* if you are seeing all the data in a timely manner.
- » **Proactive monitoring for service-level agreements (SLAs):** The growing use of cloud environments for many of today's critical enterprise applications means you have an increasing number of sites and platforms to monitor, each with its own SLA in place. If you monitor those sites proactively, you can ensure operations are running smoothly and meeting SLA requirements.

Evaluating Current Approaches to Visibility

One of the most common ways to see what's on the network is to install a Switched Port Analyzer (SPAN) port on a network switch to send a copy of all the network packets for analysis and inspection.

The major drawbacks to SPAN ports include:

- » **Increased network traffic load** requiring additional memory and/or faster processors in the switch
- » **Difficult low-level troubleshooting** because some information (Layer 1 and some Layer 2) is stripped from the data streams before being sent to a mirrored port (on most switches)
- » **Limited visibility** because of dropped packets on lower priority mirrored ports when running in full-duplex mode at full line rate

Installing monitoring and security tools inline — directly in the network path — is another common method for viewing network packets and assessing security risks. Inline deployments can overcome the drawbacks of SPAN ports (described in the preceding bullet list) but have drawbacks of their own, including:

- » **Limited usage** because of the high cost to install inline tools along every segment in an enterprise network you need to monitor
- » **Network downtime** when inline tools suffer an outage for any reason and stop the flow of live network traffic
- » **Single points of failure** along production network segments unless expensive inline tools are deployed in a redundant, high-availability configuration



REMEMBER

Each of these approaches incurs significant performance, resiliency, complexity, and cost disadvantages, and still may not provide the full visibility that you need.

Achieving Total Visibility

Organizations today are moving away from connecting security tools directly inline or to a network tap or SPAN port. Instead, they are connecting their tools to a security fabric solution that gives them access to all the data they need to see — and nothing they *don't* need to see.

Full visibility of your network traffic is critical to the security fabric. An effective model for optimum network security and performance monitoring includes (see Figure 2-1):

- » Analytics, monitoring, and security tools
- » Total visibility into end-to-end network data
- » Global application and threat intelligence
- » Security fabric that ties it all together




FIGURE 2-1: A model to improve network visibility.

Without these components, IT is hampered in its ability to deliver reliable, fast, and secure networks.

The next section turns to the solutions that deliver total visibility of end-to-end network data: network and virtual taps.

Leveraging Network Taps to Deliver Superior Visibility

Organizations require 100 percent visibility into network traffic to ensure peak performance and security. But the larger a network grows, the harder it becomes to monitor. Traditional monitoring methods, such as port mirroring, can be costly and add layers of complexity. And some monitoring solutions reduce network availability. With networks growing faster than ever, now is the time for a high-performance, scalable monitoring solution.

A network tap is an inexpensive and permanent solution that can be used throughout the network to enable monitoring and analysis without interrupting data communication. Taps provide continuous, non-disruptive network access and have these characteristics:

- » Receive all traffic on a network link
- » Require little to no configuration and can be installed at any time
- » Are not IP addressable so they aren't vulnerable to remote attacker access
- » Do not introduce delay or alter the content of the data



TIP

After a tap is installed on a network segment, you can connect or disconnect various network and security monitoring tools to or from the tap any time, without disrupting network operation.

Unlike other deployment methods (such as SPAN ports and inline), network taps can be deployed on any connection and continuously create copies of all the traffic, regardless of traffic volume. With network taps, your monitoring infrastructure can cover network segments that extend beyond your network switches. Network taps use passive splitting or regeneration technology to transmit inline traffic to an attached management or security device without data stream interference as follows (see Figure 2-2):

1. The passive tap creates a permanent, inline access port to monitor full-duplex traffic.

2. The network signal is either split or regenerated so that the monitoring device has full access to the signal.
3. The monitoring device sees the same traffic as if it were also inline, including physical-layer errors.




FIGURE 2-2: A network tap deployment example.



TECHNICAL STUFF

Some signal degradation (insertion loss) occurs in fiber taps because a small amount of light is being diverted to create a replica of the data. When purchasing taps, pay attention to the amount of insertion loss in each vendor's solution. For example, an insertion loss that is 1dB less means the signal will be 25 percent stronger and you won't have to regenerate the signal as much, which can yield significant cost savings.

Getting Visibility into Virtual Traffic

The primary challenge to managing security and network performance in a virtualized data center or cloud is tied to achieving visibility and access to virtual traffic. On physical networks, traffic flowing between servers can be captured and analyzed in a fairly straightforward manner. In virtualized environments, the data may never traverse a physical switch or network, instead remaining in the same physical host, making monitoring difficult. Traffic passes from the virtual adapter to the virtual switch and back out again, without giving you a place to observe and monitor traffic flow.

When visibility is obscured, tools such as intrusion detection/prevention (IDS/IPS), data leak protection (DLP), and unified threat management (UTM) systems become far less effective. Network performance and de-bugging tools cannot see into

the virtual switch layer. Blind spots are favored “hangouts” for malicious intruders; performance issues lurk there unsuspected, as well.



TIP

In modern, largely virtualized data centers, as much as 80 percent of network traffic is east-west, as opposed to north-south (client-server) traffic that is typically inspected and controlled by a firewall.



TECHNICAL STUFF

Even in cases of east-west traffic between VMs on separate physical hosts (or traffic between a VM and a traditional physical server), visibility of traditional network monitoring and security solutions can be obscured. The loss of visibility occurs because virtual network interface cards (vNICs) that emulate multiple NICs associated with each VM on the physical host aggregate the traffic from all the VMs on that physical host and pass it through one or more physical NICs on that physical host.

A new approach has emerged to resolve this gap. *Virtual taps* are designed to monitor data passing between VMs and send packets of interest to any network or security monitoring tool — whether physical, virtual, or cloud-based. The goal of the virtual tap is to get packets out of the virtual environment in real-time with the least impact on the virtual switch.

Virtual taps deliver the following capabilities:

- » Complete visibility of virtual network traffic
- » No negative performance impact to the virtual environment
- » Granular enforcement of compliance requirements across converged, virtualized, and physical infrastructure
- » Integration with virtualization technologies without requiring architectural changes or a large footprint
- » Support for hypervisors including Linux Kernel-based Virtual Machine (KVM), Microsoft Hyper-V, OpenStack, and VMWare ESXi and NSX.



REMEMBER

Virtual taps work in conjunction with the other components of the visibility architecture to speed application delivery and effective troubleshooting and monitoring for network security, application performance, and service-level agreement (SLA) fulfillment — and to allow IT to meet compliance mandates.

IN THIS CHAPTER

- » Getting the right data to the right tools
- » Deploying active (inline) or passive (out-of-band) security tools
- » Identifying the key components of a security fabric
- » Knowing what to look for in a security fabric

Chapter 3

Increasing Efficiency and Flexibility with a Security Fabric

To effectively monitor network security and performance, IT needs to process all of the data provided by visibility solutions and zero in on the alerts that need attention. This chapter explains how a high-performance security fabric takes the data and applies high-performance processing and security intelligence, to optimize security and network monitoring.

Delivering the Right Data, to the Right Place

The goal of a security fabric is to provide all your security tools with full access to the specific types of traffic they are designed to monitor, no matter where that traffic is in your network, with complete resiliency. Resiliency, in this case, means that the security fabric delivers data to your tools continuously, even in the

event of a component failure. By utilizing a security fabric, you increase the effectiveness of analytics and security tools and optimize their data access. A security fabric intelligently routes and load balances the right data to the right tools, every time.

The larger a network grows, the harder it becomes to monitor. IT must deal with several key challenges when managing security and network monitoring tools, including:

- » **Delivering the right data to every tool:** This can be a challenge because different tools need different types of data. Some tools need packet data, for example, while others need NetFlow data.
- » **Managing the cost of tools:** Monitoring tools can be expensive and costs add up, especially when you need to monitor data on many network links. Increasing traffic volume can also mean the need to add processing capacity for certain tools.
- » **Keeping tools updated to match network capabilities:** Every time network technologies change or are upgraded, interoperability with security and monitoring tools needs to be reanalyzed and adjusted.
- » **Maintaining network security:** Continuous traffic inspection is necessary to prevent malware and other threats from entering the internal network. IT must consider how to achieve highly-resilient security.



TIP

Traditional monitoring methods, such as port mirroring, can be costly and add layers of complexity, and some monitoring solutions negatively impact network availability. With networks growing faster than ever, a high-performance, scalable monitoring solution is needed. The most cost-effective solution is to deploy a security fabric.

Defining Inline and Out-of-Band Tools

Enterprises use an array of multi-vendor security appliances, threat analysis tools, and network monitoring solutions to protect, secure, and optimize network traffic. Choosing between them, integrating them successfully, and monitoring them effectively is an operational imperative. With network traffic volumes

and threats increasing and evolving rapidly, the number of security tools that need reliable access to data streams is also growing. Network monitoring tools can be deployed either inline on the production network or out-of-band — to monitor copies of network traffic.

Inline tools enable real-time traffic inspection and active threat prevention. However, unless inline tools are deployed in a fail-safe manner, they introduce a potential point-of-failure in the production network. This occurs when an inline tool stops functioning for any reason, including power or port failures, software errors or configuration mistakes, or the need to take the tool offline temporarily for troubleshooting or maintenance. Additionally, tools deployed directly on a production network link may be underutilized when the traffic on that link is light or decreases.

Common inline security tools include:

- » Intrusion prevention systems (IPS)
- » Firewalls and next-generation firewalls (NGFWs)
- » Data loss prevention (DLP) systems
- » Unified threat management (UTM) systems
- » SSL decryption appliances
- » Web application firewalls (WAF)

Out-of-band tools perform passive traffic inspection, detection, and recording for routine analysis. This model is used extensively in detailed threat analysis, but does not enable any active prevention safeguards or countermeasures. Some organizations directly connect out-of-band security tools to a tap or SPAN port, but this greatly reduces scalability when the network segments that need monitoring exceed the available ports on a security tool.

Common out-of-band security tools include:

- » Intrusion detection systems (IDS)
- » Behavior analysis systems
- » Forensic tools
- » Data recording
- » Malware analysis tools

- » Log management systems
- » Packet capture (PCAP) tools



REMEMBER

A security fabric with total visibility can address the limitations of both inline and out-of-band tools.

Protecting and Processing Your Data

In addition to the visibility architecture (see Chapter 2) and the various security and network monitoring tools, the remaining components of a security fabric are:

- » **Bypass switches:** High-speed bypass switches deployed on the live network, in front of your inline tools, send traffic back and forth to inline security tools located off the network. Bypass switches stay in constant contact with each tool to confirm its readiness to receive traffic. With a bypass in place, you can direct traffic around any tool that is not responding for any reason, such as traffic congestion, tool failure, or being offline for maintenance. This simple switch protects the network from failures in the security infrastructure.
- » **Network packet brokers (NPBs):** NPBs have the ability to sort through raw data flows and send selected packets to specific tools, based on what they are designed to monitor and inspect. NPBs aggregate traffic from multiple monitoring points across your network and de-duplicate packets, so your tools do not perform redundant monitoring and waste processing capacity. NPBs intelligently filter the traffic to match the function of each tool. NPBs increase tool efficiency by limiting the actual volume of traffic each tool must process. With only relevant traffic to process, data congestion is reduced, false positives are minimized, and you can often handle the reduced volume with fewer monitoring devices.

Figure 3-1 shows the components of a security fabric in action. Data enters from outside the enterprise (the untrusted Internet) and flows through the bypass switch, which is constantly validating the availability of the tools, and then to the NPBs. The NPBs aggregate all of the data, filter according to IT-defined policies, and load balance the traffic across the tools. After inspection, traffic is sent back through the bypass switch and on to the trusted network.




FIGURE 3-1: Bypass switches and network packet brokers (NPBs).

Understanding the Role of a Security Fabric

A security fabric comprised of a visibility solution, along with bypass switches and NPBs, maximizes the effectiveness of your tools through the following:

- » Tool efficiency
- » Network availability
- » Security resilience
- » Intelligent data processing

Tool efficiency

Monitoring for performance, security, and compliance requires getting all the right data to the right monitoring tools. A security fabric can help your tools perform more efficiently by providing visibility to traffic from across your network, boosting tool performance, and reducing the risk of tool failure caused by traffic congestion. Important capabilities to look for include:

- » **Traffic aggregation:** Modern network design provides multiple paths through the network to increase network reliability, but this feature creates a challenge for monitoring.

Security tools require all data from a session to perform an accurate analysis. A security fabric aggregates traffic from multiple links to provide a complete view to your monitoring tools, which improves inspection and detection. Some security fabric solutions can provide the same traffic to multiple tools for concurrent analysis.

- » **Data filtering:** Using advanced security tools to sort through large volumes of traffic and find packets of interest is inefficient. Some tools are especially sensitive to being “oversubscribed.” They slow traffic noticeably or begin dropping packets randomly if the traffic flow exceeds their processing capability. A security fabric enables you to filter packets based on different criteria, and forward traffic to a tool only if it is relevant.
- » **Load balancing:** Ideally, your tools should normally operate at no more than 80 percent of their rated capacity, so intermittent data surges or microbursts do not overwhelm their ports and cause packet loss. This threshold can be difficult to maintain. The NPB component of a security fabric uses load balancing to sense and relieve overloaded monitoring tools by distributing traffic across multiple devices. The ability to establish multiple load balancing groups with unique filtering and traffic allocation rules is a powerful advanced feature. Look for solutions with the ability to keep session data together for more accurate analysis. Load balancing also enables tool upgrades and capacity scaling to be performed with no impact to either network availability or security monitoring.
- » **Intuitive drag-and-drop controls:** Because managing configurations, connections, and filter definitions can be complex, be sure to evaluate the management interface of the NPB. Well-designed solutions provide a consolidated view and a single managerial access point across the entire environment. A graphical drag-and-drop interface is the easiest to use and allows complete configuration with no special training or preparation. Configurations for traffic filtering, flow management, and load balancing can be easily replicated across multiple network paths, or even between data centers, to get security monitoring up and running quickly. Software and hardware updates can be performed quickly and efficiently across all common devices.



REMEMBER



TIP

Solutions that allow you to initiate tasks remotely from a web browser with drag-and-drop simplicity reduce the time you spend on security administration and the opportunity for configuration errors.

When evaluating security solutions, look for features such as zero-loss packet processing, user-defined packet filtering, and the ability to remove duplicate packets, trim away packet data that is not used, decrypt secure traffic, strip out unrecognizable packet headers, and mask packet data of a sensitive nature. Offloading these functions to an NPB can reduce the load on your monitoring tools by 50 percent or more, enabling more efficient processing and lengthening their service life.

Network availability

A well-designed security fabric strengthens security but does not allow security monitoring to slow or disrupt network response times. A strong security fabric maintains network availability, with features that include:

- » **Zero downtime maintenance:** Vendors sometimes embed bypass functionality inside a “fail-safe” security appliance. If the tool stops responding, the internal bypass routes traffic around the tool. An external bypass, however, can also protect network availability if you need to temporarily take the tool out of service — for example, to perform an upgrade or troubleshooting. Because the tool is completely separated from the live traffic flow, you can perform tool maintenance or any operational task, without having to wait for a scheduled maintenance window. The external bypass continues passing traffic along — even without a tool attached. An additional advantage is that the reliability of a simple external bypass switch is much greater than that of more complex tools with embedded bypasses. The general rule is that the more complex the tool, the shorter the mean time between failures (MTBF), and the greater the risk of failure to the entire system. Using an external bypass switch to route traffic around non-responsive tools provides maximum network protection and traffic inspection.
- » **Nanosecond heartbeat packets:** Bypass switches send very small heartbeat packets to your tools at a regular cadence to confirm their ability to respond. If a response is not received, the bypass can be configured to fail “open,” and

traffic flows on to the next device. The key characteristic of the bypass is the speed at which it can detect an issue and redirect traffic. You want this to happen as fast as possible to maintain network responsiveness. Although heartbeats are common to many bypass models, solutions where the heartbeat packets originate in the hardware, rather than in software, can be sent at very high frequency — one per nanosecond. Such frequent heartbeats enable the bypass to detect the failure instantaneously and react accordingly.

- » **Self-healing, negative heartbeats:** Another feature to look for in bypass switches is whether they continue to send heartbeat packets even after a tool stops responding. Bypasses with this feature know very quickly when the tool comes back online and can resume routing traffic through the tool. This self-healing feature enables fast and automatic recovery when tools come back online and limits the impact of tool outages.
- » **Traffic flow monitoring:** The ability to collect data on traffic moving through the bypass is another important characteristic in a security fabric. Some bypass switches keep track of traffic patterns in both the uplink and downlink directions, and use this information for reporting. The data, known as *bi-directional utilization and peak traffic indicators*, allows administrators to identify possible network or application anomalies before a network outage occurs. Another feature to look for is the ability to integrate this data into your existing management tools to streamline network management.

Security resilience

While the bypass switch is designed to fail-open and move traffic around any tool that stops responding, it is also important to think about how to maintain security inspection and threat prevention. A security breach that sneaks in during a tool outage can have a devastating impact. More organizations are requiring high availability in their security infrastructure, just as they do for their core network infrastructure. A security fabric supports several options for increasing security resilience:

- » **N+1 tool deployment:** Without a security fabric, each tool deployed — no matter how much capacity is being used — needs a redundant standby in place to take over in the event

of a failure. With a security fabric, tool capacity is shared across the network, meaning you can maintain full processing with $N+1$ number of tools, where N is the number of tools required to process all the data. The “plus one” tool provides the capacity needed to take over in the event a tool goes offline for service or replacement. Tool sharing typically saves the enterprise money when planning redundant tool capacity.

- » **High-availability security fabric:** To prevent unmonitored traffic from being passed into the trusted network if the bypass or NPB stop responding, you can deploy a security fabric with redundant bypass switches and NPBs. With complete redundancy of the bypass, NPB, and tools, your security fabric provides dual monitoring paths and can recover from the failure of any component to provide non-stop, fail-safe inspection of network traffic (see Figure 3-2).
- » **High availability with near-instant failover:** For organizations with the need for extremely fast failover, some NPBs can be deployed in active-active mode. In this configuration, the NPBs are connected with complete synchronicity and session awareness. Designs of this type provide significant improvement in throughput during normal operations, because both NPBs are actively processing traffic. If one NPB goes down, the active-active configuration means the other NPB can take over seamlessly with no lost packets, within several seconds.




FIGURE 3-2: An example of a high availability (HA) active-active security fabric.

Intelligent data processing

Your security tools are already working hard enough. Free your tools from non-core processing tasks and use the intelligent data processing capabilities of your security fabric NPBs to let your tools focus on what they do best:

- » **Packet deduplication:** Network analysis and security tools commonly receive duplicate packets because data from multiple taps is forwarded to the same tool. An NPB can eliminate all the duplicate copies without losing any original data. The best NPBs accomplish this at their full specified data rate. Some tools may require that you strip out certain data from the packet; the NPB performs this function as well, to offload these non-core tasks from your tools.
- » **Security-specific processing:** Your security posture can be strengthened even more through use of security-specific processing functions. One example is the ability of an NPB to strip off header data that may contain personally identifiable information (PII), such as credit card or Social Security numbers, protected health information (PHI), or other sensitive data, before passing it on to your analysis tools. Another key function is the ability to see and analyze all traffic — clear or encrypted — to ensure you have no blind spots. By decrypting SSL traffic in the NPB, you gain total visibility and offload the processing-intensive decryption work from your security tools.
- » **Context-aware processing:** The growing use of web applications and social media has widened the focus of information technology (IT) security beyond the network. Security attacks at the application layer are also a challenge, because malicious code can masquerade as a valid client request or normal application data.

Security fabric solutions can also provide context-aware data processing as an extended functionality. This feature uses deep packet inspection (DPI) to examine information in the headers of a message or the contents of the message itself, to either block traffic or forward it to a monitoring device. This technique allows organizations to develop rich data on the behavior and location of users and applications, to identify hidden applications running on the network, mitigate security threats from rogue applications and users, and improve network performance based on application data.

Context-aware data processing enables you to create granular security policies to:

- » Block or manage the use of specific applications for individuals or groups of individuals
- » Identify, allow, or block applications and websites to provide protection against threats and malware
- » Control employee Internet access to inappropriate or illicit websites
- » Manage bandwidth to prevent unintentional or even malicious activity by insiders



TIP

Another important function of context-aware data processing is the ability to produce detailed logs and reports, which can be examined offline or out-of-band for warning signs of impending or actual attacks.

Case Study: Retail Bank Gains More Value from Security Tools

A leading financial institution providing retail and corporate banking services needed a solution to secure customer data while also controlling costs.

The company wanted to improve security by sharing information from its Blue Coat SSL Visibility Appliance with other tools, such as a FireEye network security appliance. The best method to accomplish this goal was to daisy-chain the tools together and pass only the relevant information to the next tool, to reduce redundant processing. Although the concept was simple, the implementation was not — until the company learned about Ixia's Security Fabric.

Ixia helped the company deploy an external bypass switch into the network to provide extremely fast failover capabilities so that security tools could be added or removed at will, with no disruption to the network. Next, an Ixia NPB was inserted to aggregate traffic from across the network and selectively send data to the appropriate security tools. Serial chaining of tools enabled the company to take data output from one tool and pass only that data on to a second tool. This technique allowed in-depth inspection of suspect data.

The NPB also provided data aggregation, data filtering, and load balancing functions to make the company's security tools much more efficient.

Beyond efficiency and security, the company also wanted to deploy its security tools using a high availability solution. This was easily accomplished by deploying two Ixia iBypass switches and two Ixia NPBs in active-active mode as the foundational security fabric. The tools were deployed in an N+1 configuration for survivability. For the intrusion prevention system (IPS), five units were deployed in a load sharing configuration. If any one of the five units fails, the remaining four can handle the entire network load.

Security fabric results in \$680,000 cost savings

The company's second major requirement was to reduce costs. The original specification required eight IPS units at a cost of approximately \$200,000 each. Deploying a security fabric with two NPBs and two iBypass switches reduced the requirement to five IPS units, which created a \$600,000 savings. Other tools (like the Blue Coat appliances and web application firewalls) were also reduced for another \$200,000 in savings. The cost of the security fabric ended up being \$120,000. All together then, this solution resulted in a total savings to the company of \$680,000.

“The Ixia security fabric enabled us to do things we had never done before, like service chaining of security tools and high availability security inspection,” says the company’s network operations manager. “At the same time, we reduced our overall deployment costs by almost half.”

IN THIS CHAPTER

- » Reducing the number of security alerts to investigate
- » Leveraging third-party application and threat intelligence feeds
- » Being proactive about resolving performance and security issues
- » Maintaining regulatory compliance

Chapter 4

Adding Application and Threat Intelligence to the Security Fabric

This chapter looks at how integrating application and threat intelligence into your security fabric can enable a more rapid response to attacks and reduce the number of alerts your tools and security team must deal with.

Streamlining Security With a Pre-Filtering Appliance

Many companies have high-quality firewalls, intrusion prevention systems (IPS), security information and event management (SIEM), and an experienced security team in place. And yet, breaches still happen. In addition, the Ponemon Institute calculates that the average time to detect a breach is 170 days. Why does it take so long? One answer is because there's so much security "noise" that it's hard to zero in on the important evidence of a

successful or ongoing breach. It's not that your security tools are failing, it's just hard to follow up on all the alerts they issue. Eliminating hijacked IPs, untrusted countries, and malware and phishing sites cuts through that noise so you can focus on what's important.

Many enterprises are establishing a new line of defense by deploying a special-purpose security appliance — a *threat intelligence gateway* — to very quickly filter out traffic coming from, or bound for, sites already known to be engaged in cyberattacks. Even encrypted connections from these sites are automatically blocked (see Figure 4-1).




FIGURE 4-1: Addition of application and threat intelligence appliance to the security fabric.

By eliminating unwanted traffic from entering your network, these security intelligence solutions can:

- » **Increase tool efficiency:** By removing known bad actors from attempting remote connections to your data and resources, you can sharply reduce the number of alerts sent to your existing security solutions. Some next-generation firewalls (NGFWs) are limited in the number of rules they can process. It makes sense to reserve their capacity for identifying application-layer events, not to filter out known Internet threats.
- » **Eliminate false positives:** Instead of analyzing each attack, a threat intelligence gateway checks the IP source address against a constantly updated list and blocks any that appear

on the list. The key advantage is that it doesn't use signature-based analysis, so there are no false positives.

- » **Provide visibility to threat activity:** The gateway is managed from a web portal and some provide a dashboard to track blocking activity and allow operators to drill down and examine the details of each action taken.
- » **Increase staff productivity:** For traffic pre-filtering to deliver efficiencies, this new type of security appliance should be preloaded with filtering rules based on existing intelligence and capable of receiving real-time updates as information changes. Free of the need to maintain the rule sets and with fewer alerts to analyze, your security team will then be able to focus on looking for new and unknown attacks, and more efficiently execute their jobs as guardians of your networked resources.

Incorporating Real-Time Expert Security Analysis

Appliances designed to filter out known bad traffic rely on intelligence gathered from multiple sources and independently confirmed. One organization that is devoted to maintaining such a database is the Ixia Application and Threat Intelligence Research Center. The Center has been performing advanced security research for more than a decade, providing intelligence updates to customers in every industry around the globe. The intelligence produced by the Center is leveraged by many leading hardware and software providers to test their products before launch.

A lot of work goes on behind the scenes to collect, analyze, validate, and distribute the intelligence, with raw input data coming from a variety of sources:

- » **Third-party source feeds** are collected from different public and private streams, including data from the open source community and various security technology partners. The data from these feeds is not treated as malicious until the Center has individually validated each one.

- » **Internet scanning** is used to essentially surf the Internet looking for sites of interest. This is done in an automated fashion, running 24/7. Sometimes the system immediately finds proof of malicious activity, such as malware binaries. Those sites are automatically scanned, catalogued, and added to the intelligence database. Other sites may be found not to have malicious content, but to have known vulnerabilities such as unpatched web server software. There is a high likelihood such vulnerable sites will soon be compromised, so they are put into a special queue for frequent scanning.
- » **Honeypots** are a staple of intelligence collection. The Center has a global network of sites with services such as secure shell (SSH), file transfer protocol (FTP), remote desktop protocol (RDP), voice over IP (VoIP), and hypertext transfer protocol (HTTP). Any connection attempts to these services are logged and cross-referenced with data collected from other points in the network. Internet scanning sites may change locations frequently, so it's important to have a large and global network to detect their activity.
- » **Spam** is another resource used to collect inputs for analysis and validation. The Center uses a variety of commercial and internal feeds for collecting spam information that is then fed into a machine learning engine. The Center follows the hyperlinks embedded in spam emails and analyzes any discovered binaries for malicious content.
- » **Binary analysis** is often the final and key step in determining whether a site is engaged in malicious activity. The Center performs dozens of static and dynamic (or "sandbox") analyses of each target, looking for known signatures within the binary of known malware samples.

Performing Advanced Troubleshooting

Rapid identification of common characteristics is a powerful tool in resolving security and performance issues and improving mean time to repair (MTTR). Activating advanced application and threat intelligence processing in your security fabric, through the network packet broker (NPB), allows you to strengthen security, accelerate troubleshooting, and enhance performance.

As Zeus Kerravala, Principal Analyst at ZK Research, observes: “Problem identification is IT’s biggest challenge.” He explains that 85 percent of the MTTR is the time taken to identify that there is, in fact, an issue. Even worse, the MTTR clock starts ticking whether IT knows the issue exists or not.

A second component of problem resolution is identifying the location of the problem or problems. You can try to find the needle in the haystack — but how do you know which haystack you should even be looking at? Possible “haystacks” might include:

- » **Network equipment**, such as core and distribution routers and switches
- » **User or customer premises equipment (CPE)**, such as access routers and modems
- » **Business applications**, including server applications, software-as-a-service (SaaS), and mobile apps
- » **Virtual and cloud environments**, including public, private, and hybrid clouds



REMEMBER

Starting in the right place can save IT teams hours, or even many days, of time. Starting down the wrong path can have devastating effects for the business.

A security fabric that leverages application and threat intelligence can be used to capture critical information needed for the whole troubleshooting process. Geolocation capability can help quickly locate geographic outages and potentially narrow troubleshooting efforts to vendors that may be causing network disruptions. This level of information granularity also reduces application troubleshooting costs and allows you to optimize customer quality of experience. You need to know that something is happening, but you also need to know why it’s happening.

The details enable you to access empirical data to identify bandwidth usage, trending, and growth needs so that you can be proactive in managing resources and forecasting expansions. The use of metadata not only makes traditional troubleshooting efforts better; it allows IT to become proactive. Proactive troubleshooting is the Holy Grail for network administrators — to prevent a network problem from happening, or at least remediate the issue

before anyone notices it. The metadata can also be combined with trending and bandwidth consumption to anticipate problems before they happen and affect the network.

Ensuring Regulatory Compliance

Regulatory compliance has been top-of-mind for enterprises for many years now. All companies are looking for ways to strengthen compliance, reduce costs, minimize security risks, and avoid potential non-compliance penalties and fines. Application and threat intelligence can provide many benefits in this area as well. These include:

- » **Application monitoring** lets you know when employees may be using cloud-based services (like Box or Dropbox) or web-based email to transfer company files and bypass your security policies. These services can be accessed from anywhere on any device and may therefore bypass anti-malware scanning and other security safeguards, potentially exposing your entire organization to malware and exploits. Another risk is that once an employee is no longer employed by the company, he or she may still have access to those files because IT cannot restrict access to personal, cloud-based storage services.
- » **Data masking** protects critical data (such as credit card numbers, Social Security numbers, and other information) by masking it before it reaches monitoring tools. This ensures that the data is not exposed and various regulatory compliance requirements can be met.
- » **Data searching and validation** looks for keywords, phrases, or numbers in network traffic. Regex search strings can be created for phone numbers, credit card numbers, Social Security numbers, emails from certain IP addresses, or names. This feature allows critical data to be filled out and sent to a DLP tool for further analysis. The data can also be sent to other purpose-built tools, like a tool that uses credit card data and a checksum algorithm to validate if a 16-digit credit card number, for example, is valid.

Case Study: Hosting Service Tackles Attack Traffic with Ixia ThreatARMOR

Known for its low-cost, high-quality hosting services, this service provider had been suffering in recent years from a glut of cyber-attacks. Attack traffic sometimes caused servers to crash entirely, making it necessary to find new ways to protect user systems.

Essential security infrastructures were in place, such as firewalls and intrusion detection systems (IDS). However, unlike the networks of private enterprises, shared networks providing services to users are unable to adopt aggressive countermeasures for cyberattacks, because of the risk of blocking valid user communications. As a result, the network and security operations team had resorted to manual countermeasures to halt the massive attack traffic.

Intrigued by pre-filtering

The concept of automatically eliminating unnecessary traffic based on known threats was the same idea behind the manual countermeasures the service provider had been implementing.

The company contacted Ixia to observe a test run of Ixia's ThreatARMOR solution on actual user machines. The sensitivity of the deployment dictated that the solution initially would need to be tested outside the live production network. Because the company already had a network packet broker deployed, ThreatARMOR was deployed out-of-band on a network tap.

Many features impressed the service provider, but a key one was ThreatARMOR's unique Rap Sheets, which indicate why each instance of blocked traffic is flagged. Within one hour, Rap Sheets were being generated, classifying the remote IP address in varying categories such as hijacked, phishing, malware, and so on. When the information from the Rap Sheets was correlated with information in the Security Incident and Event Management (SIEM) system, it turned out a brute force SSH login had been going on for some time. The attack had evaded the next generation firewall (NGFW) because SSH Login attempts are normal and not inherently malicious. This case was an obvious exception, because an unauthorized breach of the company's network was taking place. As in many attacks, the attacker was not a single user, but a script

that had been written to try millions of passwords until the targeted system was compromised.

“I was thrilled with the level of sensitivity to the network operation. One reason is that ThreatARMOR employs a fail-open system, and even when trouble occurs, communication is never interrupted,” said the manager of the engineering department. “In contrast, most security devices adopt a fail-close system, which makes them inappropriate for use over a shared network,” he added.

Intrusions reduced by 80 percent

After full deployment of the ThreatARMOR appliance, the number of intrusions detected by the company’s IDS decreased from one million cases to 200,000 cases, and the company saw a reduction in the amount of spam email as well.

Out-of-band deployment upgraded to inline

After confidence was gained in the ThreatARMOR solution, the decision was made by the security team to place the solution inline in “blocking mode” in front of the NGFW. The value shown during the first few weeks of deployment was immediate and tangible and the security team easily justified the investment.

IN THIS CHAPTER

- » Getting the most out of your monitoring and security tool investments
- » Limiting risk by improving troubleshooting processes
- » Enabling security resilience with near-instant failover

Chapter 5

The Benefits of a Total Visibility Security Fabric

This chapter explains how a high-performance security fabric delivers complete, end-to-end visibility for both inline and out-of-band security tools, all in a single solution. This powerful solution combines 100 percent data access, resilience, and security intelligence to ensure that the right data gets to the right tools, every time, at high speeds.

Maximizing the Value of Security and Monitoring Tools

A security fabric can help you realize the full value of your existing investments and minimize new tool purchases in the following ways:

- » **Share tool capacity:** Many security appliances and monitoring tools are not fully utilized, which limits the value they provide. A security fabric, however, provides a way for the capacity of these devices to be shared across multiple network segments, allowing them to be more fully utilized.

If traffic volume on two network links can be aggregated and delivered to a single monitoring appliance with available capacity, IT does not need to purchase a separate appliance for each link. The ability to share tools can result in substantial savings in additional tool purchases.

- » **Filter out unnecessary traffic:** Getting the right data to the right tool is an important function of a security fabric. For example, you may have an intrusion prevention system (IPS) deployed inline on your production network to block malicious or suspicious traffic. However, as the volume of traffic entering your organization increases, your IPS may become overloaded. One way to improve efficiency is to use a network packet broker (NPB) to filter out certain low-risk traffic, such as voice, video, and music, that does not need to be inspected, and pass it directly to your internal network. Filtering can significantly reduce the load on your security appliances and help you delay new purchases, to stretch your security budget further.
- » **Data conditioning:** Network taps and bypass switches provide a complete and unfiltered flow of network packets. Most tools, however, are designed to inspect or analyze only certain types of data, meaning the packets must be stripped of extraneous information before being processed. Using sophisticated tools to perform functions such as deduplication, packet stripping, or data masking is a waste of the tool's processing capacity and can lead to oversubscription of tools. Some tools are especially sensitive to being oversubscribed and begin to noticeably slow or drop packets randomly if flow exceeds their capacity. Offloading these functions to an NPB makes the monitoring tools more efficient and lets them process more data. A typical guideline is that tools can become 60 percent more efficient when data conditioning functions are offloaded.
- » **Extend useful life:** Deploying a security fabric also lets you separate your security monitoring tools from the production network, to allow lower-speed monitoring tools to process data from higher-speed core network links. You can make independent decisions about when to upgrade your network and tool infrastructure and extend the life of your slower 1Gbps (or 10Gbps) tools, or even invest in newer, more powerful security tools, without waiting for the network to be upgraded. Having more flexibility over the timing of tool upgrades helps you make better use of your security budget.

Optimizing Incident Response to Reduce Mean Time to Repair

The longer IT takes to respond to incidents, the more risk is incurred and the more costly such incidents become. Rapid identification of common points of failure in the network can dramatically improve mean time to repair (MTTR) and minimize potential loss of valuable and/or sensitive data.

Although a security fabric doesn't solve the problems themselves, an NPB with advanced application intelligence can enable IT teams to identify the underlying issues more quickly. It does this by filtering packets received based on potential risk factors such as by application, or by geographic source or destination, and/or device. A subset of traffic can be sent from the NPB to a security tool for further analysis, or set to generate additional NetFlow metadata for further insight. This type of advanced intelligence gives the security fabric the ability to mitigate the risk of security incidents, accelerate troubleshooting to restore network availability, and provide customized performance management of critical applications.

Operations staff can be automatically apprised of issues that are hindering performance, thus greatly reducing troubleshooting time. And security personnel are tapped into information about where, when, and how attacks are occurring so they can more strategically plan a proactive defense.

NetFlow metadata can provide a lot of context about a security or performance incident and its origin — such as the connection's browser, device name, user location (latitude, longitude, city, state, and country), and Internet service provider (ISP) — to reduce MTTR.



TECHNICAL STUFF

Enabling High-Availability and Fast Failover

A security breach can have a devastating impact on an organization. Most organizations require high availability for their security infrastructure, as they do for their core network infrastructure.

With a modular security fabric, you can achieve very high uptime for security monitoring.

You can ensure non-stop inspection of network traffic with no disruption to network availability, using a security fabric configured with redundant bypass switches, NPBs, security appliances, and monitoring tools. This high-availability architecture can recover from almost any failure. Configure the primary and secondary NPBs with complete synchronicity in active-active mode to achieve the following advantages:

- » **Near-instant failover:** If speed of recovery in a failover situation is key, active-active mode provides the fastest possible recovery. Both NPBs are actively engaged in traffic processing with full synchronicity and session integrity. No time is required to activate a standby NPB and transfer processing.
- » **Share workload:** The cost of redundant architecture can be a substantial portion of the overall security budget. Many organizations find it easier to justify the cost of a second NPB if it will be used immediately in normal operations, sharing the workload with the primary device.
- » **Eliminate risk of incomplete failover:** With the secondary device in standby mode, there is some risk it will fail to boot up and will be unable to take over processing if the primary device fails. With both NPBs actively processing data, there is no chance of an incomplete failover.
- » **Maintenance without disruption:** With an active-active configuration, any maintenance that you need to do on an NPB can be performed without the need for a maintenance window. Traffic is temporarily routed through a single device until both devices are back online.



TIP

Periodically switching packet processing to a single NPB (failing over) to perform maintenance is a good way to prove that the existing configuration is still able to handle the entire volume of traffic.

IN THIS CHAPTER

- » Getting your security deployments right the first time
- » Maximizing security investments with a proof of concept
- » Keeping systems and networks updated
- » Knowing the limits of your network environment

Chapter 6

Building Confidence through Security Testing

Cyberattacks are inevitable. However, real-world security testing helps you find problems and discover unknowns in advance of costly security incidents. This chapter explains the importance of comprehensive security testing to verify three things: that your network security tools are properly configured before they are deployed, that they are correctly updated and patched after they are deployed, and that they are performing optimally while they are deployed.

Validating Your Designs and Configurations

Whether you're deploying new equipment or updating existing equipment, validating your design and configuration helps to ensure your security architecture and tools are installed properly,

configured correctly, and performing optimally, and that your production network:

- » Is properly hardened against security threats
- » Doesn't suffer any negative performance impacts
- » Isn't exposed to new vulnerabilities resulting from poor design, interoperability issues, or misconfiguration of new equipment or security tools

Verizon's 2016 *Data Breach Investigations Report* (DBIR) lists misconfiguration errors, such as "mistyping a firewall rule allowing access to a sensitive file server," for example, among the top five "miscellaneous errors" that lead to a data loss or breach.



TIP

In many organizations, security is often perceived as an encumbrance or bottleneck to productivity. Validating and testing your security designs and configurations in pre-production can help security organizations overcome this perception — and safely enable the business — by proactively ensuring that network/application performance and core business processes are not adversely affected.

Reducing Risk and Costs

During the design phase, security testing helps ensure that your security design meets your business and technical requirements, that your security infrastructure is configured correctly, and that the components are optimally deployed in the right capacity and locations throughout your network. Without knowing exactly how a security solution will perform in your network, your design options are typically limited to best practice recommendations, matching specifications to requirements, and making certain assumptions about performance and your environment. Getting your design correct the first time will potentially save costly re-deployments and additional equipment purchases in the future.

During a proof-of-concept evaluation, security testing can also help you validate the performance specifications and capabilities of any solution you are considering, and enable an "apples to apples" comparison of solutions from different vendors. Vendor data sheet specifications do not provide an accurate indication of

how devices will perform while processing actual traffic on your production network. To get meaningful performance numbers that lead not only to the purchase of the best equipment for your environment, but also to significant cost savings, enterprises must often conduct onsite head-to-head bakeoffs.

Perhaps most importantly, rigorous testing of your security tools reduces the risk that an error will expose a vulnerability that ends up leading to a successful cyberattack against your network. The costs to an organization resulting from a successful breach can include:

- » Losses caused by business disruptions, brand reputation damage, and loss of customers
- » Legal fees incurred through litigation, as well as regulatory fines and penalties
- » Remediation, including incident response and recovery, new security investments, and security tool re-deployments

Case Study: Leading Financial Institution Secures Its Network from Cradle to Grave

Many vulnerabilities and threat vectors are exploited at certain phases in the life of a network, such as installation, upgrades, and decommissioning. For this reason, a leading U.S.-based provider of international banking services implemented a lifecycle management approach to ensure a solid defense.

The network security team engaged Ixia to help on three fronts — network performance testing, network security validation, and inline security tool deployment.

Validating the initial design

Ixia's BreakingPoint solution was used to validate the new network architecture designs against various known security threats, such as distributed denial-of-service attacks (DDoS) and malware, in the company's pre-production lab. This involved testing the company's new next-generation firewalls, wireless access firewalls, intrusion prevention systems (IPS), load balancers, and

various proxy devices, to ensure they performed to specifications. “The BreakingPoint solution enabled us to test all of our security devices in the lab before we actually deployed them,” says the network operations manager. “I was able to see exactly how they would perform in the network and adjust my security architecture to maximize effectiveness.”

The company also deployed the Ixia PerfectStorm solution to consolidate its testing capabilities and use existing performance tools while leveraging new capabilities, all in a single solution. With PerfectStorm, the company can now test its multi-level, layered security design and components against the full stack of network traffic and run it out of a single port. “The combined BreakingPoint and PerfectStorm solution allows me to find about 120 flaws a year, that I can clearly isolate to specific equipment,” says the network operations manager.” This saves me weeks of troubleshooting time in the production network.

The Ponemon Institute calculates that a flaw in a production network costs \$7,600 to remediate, but only costs \$960 to fix, if it is detected during quality assurance (QA) testing. By identifying 120 flaws in the lab before they made it to the production environment, the company estimates it saved approximately \$800,000 in the first year with the PerfectStorm solution.

Implementing the design

Once the design was validated, it was deployed to the production network. Additional performance testing was then conducted with the Ixia IxLoad solution simulating real-world traffic using Internet, video, voice, storage, security, virtual private network (VPN), wireless, infrastructure, and encapsulation/security protocols. This capability enabled the company to verify its deployment and establish a known network performance level and rollback configuration.

Looking ahead

Change control is another critical phase of network lifecycle management, as any network configuration change can create performance and security risk. To mitigate the risk, the company is now using its lab environment to test the effects of any configuration change or software update and simulate what-if scenarios before deploying changes to the production network.

Staying Vigilant About Updates and Patches

Organizations are not always vigilant about patching and updating known vulnerabilities in their systems and networks. This is an unfortunate reality that opportunistic attackers routinely exploit. Attackers, unlike the rest of us, don't always need the latest and greatest shiny new thing. In its 2016 *Data Breach Investigations Report* (DBIR), Verizon concludes that for attackers "the oldies are still goodies."

For example, for its 2016 Annual Security Report, Cisco conducted a one-day "outside-in" scan of 115,000 Cisco routers and switches across the Internet, and found that 106,000 (92 percent) of these network devices were running versions of software that had, on average, 26 known vulnerabilities — their security patches and updates weren't current.

Vulnerabilities in aging equipment with outdated software are another easy attack vector for attackers. In the same report, Cisco found network devices in financial, healthcare, and retail industry verticals using software versions that were, on average, four to six years old. Many of these devices had reached their last day of support (LDoS), meaning they can't be updated or patched.

Stress Testing Your Network Environment

Testing the strength of your network and security architecture under realistic conditions is necessary to build confidence in the capabilities of your security solutions. The more comprehensive your test platform in terms of numbers of threats and application flows it can represent, the more realistic the results.

Beyond the sophistication of today's attacks, the massive scale of many attacks is often enough to overwhelm an enterprise or carrier network. According to Verizon's 2016 *Data Breach Investigations Report* (DBIR), the mean bytes per second and packets per second in denial-of-service (DoS) attacks in 2015 was 5.51 gigabits per second (Gbps) and 1.89 million packets per second

(Mpps), respectively. Among the DBIR’s recommended controls: “Understand the capabilities of your defenses.”

Modern distributed denial-of-service (DDoS) attacks often involve hundreds of thousands or millions of infected endpoints that flood a target network with traffic. A DDoS attack that successfully crashes a system can expose a software vulnerability that can be exploited, or may cause the system to “fail open”, thereby allowing all traffic — legitimate or otherwise — to pass. In some cases, it isn’t unheard of for a network team to disable certain security tools and safeguards in a panicked attempt to troubleshoot a “performance issue” before realizing the organization is the victim of a DDoS attack.



TIP

A high performance security test environment can emulate realistic two-way conversations and transactions exactly as real applications behave in operation. You can put your network to the test and launch thousands of live attacks using different pieces of malware, DDoS, and exploits.

A network resilient to attacks, misconfiguration, bottlenecks, changes from user behavior, and patching can be the difference between an inconvenient incident and going out of business for many organizations. Testing with simulated real workloads, at scale, provides advanced knowledge on how your organization and its technology will fare under attack, and identify its breaking points. With this knowledge in hand, you can adjust configurations, architectures, and policies to ensure defenses are working properly and will bounce back within a reasonable timeframe.

IN THIS CHAPTER

- » Establishing visibility as the cornerstone of security monitoring
- » Ensuring resilience and minimizing downtime
- » Growing seamlessly without disruptions
- » Simplifying management with a graphical interface
- » Leveraging application and threat intelligence feeds

Chapter **7**

Ten Requirements for Security and Network Monitoring

To protect the enterprise and its customers, IT must have complete visibility to the traffic flowing into and out of the network. With total visibility, your security and network monitoring tools have access to all the information they need to maintain a strong defense and enable an excellent customer experience.

This chapter summarizes ten features of a resilient visibility architecture that is optimized for continuous, real-time network monitoring.

Proven Lossless Monitoring

I can't state it often enough: You can't monitor what you can't see. According to a 2016 research survey by Enterprise Management Associates (EMA), nearly 30 percent of respondents said they were not completely satisfied their tools were receiving all the data they needed. This skepticism is for good reason. The Tolly Group ran a comparison between the two leading network packet brokers (NPBs) and found one NPB dropped packets at every tested data size and had no mechanism for reporting that fact. Some solutions, like switched port analyzer (SPAN) ports and command line interface (CLI)-based packet brokers, can clip data or just provide a summarized version. This is not sufficient. Make sure you validate existing and future NPB solutions at realistic speeds, to ensure you are getting the visibility you need. Ask whether the packet broker you are considering is a zero-loss solution. Missing data can result in delays and misdiagnosis in identifying attacks and resolving network problems.

Visibility to Virtual and Cloud Resources

Complete visibility to data flows across your enterprise network environment, and elimination of blind spots is fundamental to strong security. Unfortunately, a loss of visibility into the traffic flow between virtual machines or with cloud-based resources is common, and blind spots are often used by attackers for infiltrating the enterprise network. Lack of visibility leads to an elevated risk of threats and an inability to sufficiently monitor and troubleshoot critical events.

To achieve full network visibility, organizations must:

- » Deploy visibility infrastructure components designed specifically for virtual environments, rather than relying on retrofitted hardware solutions.
- » Find solutions that do not require architectural changes or negatively affect the performance of the virtual environment.
- » Support elasticity of virtual environments and maintain visibility to virtual machines as they are moved around for optimized performance.

Fail-Safe Deployment for Inline Appliances

When you deploy an active security monitoring appliance inline on the production network, you introduce a potential point of failure. If a device deployed inline stops operating for any reason, it can halt the flow of data through your network — negatively affecting customer experience, employee productivity, or damaging your organization's brand image.

You can decrease the risk of an outage due to single points of failure by installing a bypass switch in front of your security appliances to monitor their status and route traffic around any device that fails to respond. A single external bypass switch can protect multiple inline monitoring and security tools. Bi-directional heartbeats ensure uninterrupted network uptime during system, link, or power failures.

Important features to look for in a bypass switch include:

- » **Hardware-based:** A heartbeat that is based in the field-programmable gateway (FPGA) of the bypass, rather than in software, can achieve response times between one and three nanoseconds. This translates to nearly instant failover.
- » **Automatic recovery:** A negative heartbeat detects when a tool that has previously failed open or been taken offline is once again available, then automatically brings the tool back online, to minimize the outage and restore normal operation as quickly as possible.

Maintenance Without Network Downtime

Many network monitoring and security tools have internal bypass switches that keep traffic flowing in the event of a failure in the tool. However, the functionality of an internal bypass switch is limited to the tool in which it is installed. What happens when you need to take the tool offline, for example, for a repair or upgrade, or to re-deploy the tool elsewhere in the network?

With an external bypass switch, you can also route traffic around the tool if you need to temporarily take the tool out of service for troubleshooting or to perform upgrades. With no need to wait for an available maintenance window, the security staff can keep security tools updated with new releases as soon as they are available, reducing the chance that a bad actor will take advantage of an unpatched vulnerability.

Non-Disruptive Scalability

NPBs with automatic load balancing can relieve overloaded security and network monitoring tools by distributing traffic over multiple devices, to increase efficiency and reduce congestion. This feature is especially useful for reducing the effort involved in adding tools to achieve scalability. New devices can be easily added to an established load balancing group and the NPB automatically distributes traffic over all available devices, seamlessly increasing capacity, without disrupting the network or requiring a maintenance window.

To further streamline operations and provide more flexibility, look for NPBs with the ability to establish multiple load balancing groups with unique filtering and traffic allocation rules. This important capability enables IT to change the way traffic is filtered and decommission legacy or obsolete equipment without network downtime.

High-Performance Data Processing

A network and security monitoring infrastructure needs to operate at high performance levels, particularly as core network speeds move from 10 Gigabit Ethernet (GbE) to 40 and 100 GbE. Also important is the ability to handle weekly, daily, and hourly fluctuations in traffic load, so you are processing all of the relevant data.

Unfortunately, IT is often forced to make trade-offs in what it can process. According to a 2016 ZK Research study, 45 percent of respondents admitted to turning off features in their security devices in order to improve performance. One reason is that

security tools might be pushed to do non-core tasks (like data stripping or data masking), which can overload the tool’s CPU and reduce performance. As discussed in Chapter 5, the advanced data filtering of an NPB can reduce the processing load on your security and monitoring tools by ensuring that individual tools get only the data they need, when they need it.

Another performance issue involves NPBs that cannot support line rate processing when multiple features are activated, such as data deduplication with NetFlow or secure sockets layer (SSL) decryption. This is particularly a problem with CPU and software-based processing. For optimum flexibility and feature usage, look for hardware-based systems (known as a field-programmable gate array, or FPGA) that are purpose-built for packet processing at line rate. This high-performance capability eliminates the need for processing compromises and trade-offs. If you’re evaluating an NPB solution, be sure to validate what happens to performance when all the filters you need are activated at realistic traffic loads.

Context-Aware Visibility

Early detection of breaches using application data reduces the loss of sensitive data and reduces the cost of breaches. Context-aware visibility can be used to expose indicators of compromise (IoCs), provide geolocation of attack vectors, and combat SSL encrypted threats. Efficiencies can also be created by filtering data at the application layer to remove unnecessary data clutter at the tools in order to control costs.

To thwart security attacks, you need the ability to detect application signatures and monitor your network so you know when something unusual is taking place. Context-aware visibility lets you see rogue applications running on your network, along with visible footprints that attackers leave as they travel through your systems and networks. This capability is also valuable in capturing critical information needed for troubleshooting and faster problem resolution.

Simple Management Interface

As you configure connections and define filters in your visibility platform, management and control can quickly become complex, especially when you have many traffic sources and many tools.

With a command-line interface (CLI), filtering traffic isn't difficult if you have only one tool and one data source. But what if you have hundreds of data sources and many tools? CLI management can quickly become cumbersome and overwhelm the team that handles it. It is also prone to mistakes and errors that you may not uncover until you perform analysis. By then, it's too late.

Thus, an important feature in a total visibility security fabric is a centralized, graphical, drag-and-drop interface that simplifies the configuration and management of data sources and tools. You have complete, granular control of your data and the ability to create and update simple filters, overlapping filters, and dynamic filters, as your needs change. With a graphical interface, you can easily see the connections and data mapping, whether you have one or many data sources and one or many monitoring tools.

High Availability and Fast Failover

To ensure maximum resiliency and uptime for security and network monitoring, you should deploy a high-availability security fabric. With complete redundancy of the external bypass, the NPB, and the monitoring tools, your security fabric can recover from the failure of any device to provide non-stop, fail-safe network monitoring.

For the fastest possible failover, important in many financial service environments, for example, deploy packet brokers with the ability to work synchronously in active-active mode, with session integrity. In normal processing, the NPBs share the workload for greater efficiency. If one NPB goes down, the active-active configuration enables the secondary device to take over seamlessly with no lost packets. Sessions are kept complete. Adding new tools requires configuring only one NPB, because the two NPBs share processing logic.

Integrated Application and Threat Intelligence

Pre-filtering known bad IP addresses and traffic out of the data that flows to your security solutions enhances the performance of your security solutions and reduces the number of alerts the security team must follow up on.

Integrating up-to-date security intelligence from experts into your security fabric helps stop vulnerabilities, exploits, and threats (including zero-day threats) from affecting your business and overwhelming your security staff with security alerts.

Security intelligence can be used to trigger automated responses such as blocking malicious IP addresses, domains, and websites, or entire regions or countries, if desired.

By leveraging a worldwide network of security resources, you can stay a step ahead of new threats and threat actors, to provide proactive security protection for your organization.

Remembering the Ten Key Requirements

Don't forget these key features of a resilient visibility architecture that is optimized for continuous, real-time network monitoring:

- » Proven lossless monitoring
- » Visibility to virtual and cloud-based resources
- » Fail-safe deployment for inline appliances
- » Maintenance without network downtime
- » Non-disruptive scalability
- » High-performance data processing
- » Context-aware visibility
- » Simple management interface
- » High availability and fast failover
- » Integrated application and threat intelligence

Glossary

active monitoring: A type of network monitoring in which test traffic is injected onto a live network to verify network performance. See also **passive monitoring**.

active-active mode: A high-availability architecture in which both nodes are always active and traffic is load balanced between them. When a failure occurs in one node, the other node automatically takes over processing all traffic with little or no service delay. See also **active-passive mode**.

active-passive mode: A high-availability architecture in which the primary node is active and processing all traffic and the secondary node is passive, awaiting activation. A heartbeat packet passed between the nodes verifies that the primary (active) node is operational and activates the secondary to take over processing when required. The time it takes to recover is dependent on the speed of the heartbeat and the readiness of the secondary node. See also **active-active mode**.

advanced persistent threat (APT): A deliberate, focused attack that may go undetected for several years, perpetrated by an attacker (for example, a nation-state or organized criminal organization) with extensive resources including money, people, time, and sophisticated tools (such as electronic surveillance equipment and satellite imagery).

adware: Pop-up advertising programs that are commonly installed with freeware or shareware.

anti-AV software: Malicious software that disables legitimate antivirus software on an endpoint, thereby preventing detection and removal of other malware.

application and threat intelligence: Identifying data on the source of verified bad traffic that can be used to eliminate the traffic before it reaches an enterprise's firewall or other inline security appliances. The intelligence is generally collected and verified by a team of experts and offered to others as a service.

backdoor: Malicious code that enables an attacker to bypass normal authentication and access a compromised endpoint.

blind spots: Areas in a network architecture that network and security monitoring tools are unable to fully access.

bot: An individual endpoint that has been infected with malware and is capable of being activated remotely.

bypass switch: Hardware device sitting on the live network path that transmits traffic back and forth to an active inline security appliance such as an intrusion prevention system (IPS) or next-generation firewall (NGFW), for example. The bypass automatically routes traffic along in the event the device fails.

CLI: See *command line interface*.

command-and-control (C2) server: Servers and other infrastructure used by an attacker to direct the activities of infected endpoints (or bots) in a coordinated cyberattack. See also **bot**.

command line interface (CLI): A text-based user interface for managing a security or network monitoring tool.

data loss prevention (DLP): A software or hardware security solution used to ensure that sensitive information, such as Social Security numbers or financial information, is not transmitted outside an organization's network or copied to unauthorized media.

DDoS: See *distributed denial-of-service*.

deep packet inspection (DPI): A type of network packet filtering that examines the data portion of a packet (as well as the header in some cases) rather than simply inspecting packet headers.

distributed denial-of-service (DDoS): A coordinated attack, often from hundreds of thousands or millions of compromised endpoints (or bots), that are used to flood a target system or network with excessive traffic so that it cannot process legitimate traffic. See also **bot**.

DLP: See *data loss prevention*.

DPI: See *deep packet inspection*.

east-west traffic: Traffic moving within the data center, between servers or virtual machines, that does not traverse a network firewall or boundary. See also **north-south traffic**.

endpoint: A computing device such as a desktop or laptop computer, tablet, or smartphone. Endpoints may also include devices such as routers and servers, but more commonly are client devices.

fail-closed: A control failure that results in all access being blocked or denied. See also ***fail-open***.

fail-open (or fail-safe): A control failure that results in all access being permitted. See also ***fail-closed***.

field programmable gateway array (FPGA): An integrated circuit where the silicon is imprinted with design specifications of the purchaser to provide an advantage in processing power. In this context, it refers to the processing capability embedded in the circuits used in manufacturing Ixia network packet brokers.

FPGA: See ***field programmable gateway array***.

HA: See ***high-availability***.

Health Insurance Portability and Accountability Act (HIPAA): U.S. legislation passed in 1996 that, among other things, protects the confidentiality and privacy of protected health information (PHI).

heartbeat (packets): Very small network packets transmitted at regular microsecond intervals between the bypass switch and security appliances or other monitoring tools to confirm the tools are operating correctly.

high-availability (HA): A system or device that is designed to be extremely resilient with little or no downtime, typically achieved with redundant components that eliminate single points of failure.

HIPAA: See ***Health Insurance Portability and Accountability Act***.

hybrid cloud: A cloud computing deployment model that is comprised of public and private cloud infrastructures.

hypervisor: Software that operates between the hardware kernel and operating system and enables multiple “guest” operating systems (or VMs) to run concurrently on a single physical host. See also ***virtual machine***.

IDS: See ***intrusion detection system***.

indicator of compromise (IOC): An artifact, such as a malware signature, suspicious URL, or command-and-control (C2) traffic, that provides strong evidence of a cyberattack.

inline tools: Network and security monitoring tools that are deployed directly in a network segment and collect actual traffic on that segment. See also ***out-of-band tools***.

insertion loss: The degradation of signal strength (expressed in decibels, or dB) that occurs when a device is installed in a transmission medium (copper wire or optical fiber).

Internet of Things (IoT): The network of smart, connected devices such as home automation systems, smart cars, smart electricity meters, and personal technology.

intrusion detection system (IDS): A hardware or software application that detects suspected network or host intrusions.

intrusion prevention system (IPS): A hardware or software application that both detects and blocks suspected network or host intrusions.

IOC: See *indicator of compromise*.

IOT: See *Internet of Things*.

IPS: See *intrusion prevention system*.

malware: Malicious software or code that typically damages or disables, takes control of, or steals information from a computer system. Malware broadly includes adware, anti-AV software, backdoors, bootkits, logic bombs, RATs, rootkits, spyware, Trojan horses, viruses, and worms.

mean time between failures (MTBF): A measure of reliability in a hardware component or product, typically measured in thousands (or tens of thousands) of hours.

mean time to repair (MTTR): A measure of the average time required to resolve an issue.

network monitoring tools: A device or appliance used to monitor and/or troubleshoot network characteristics such as speed, latency, and packet loss.

MTBF: See *mean time between failures*.

MTTR: See *mean time to repair*.

N+1: A redundant configuration in which the total number of devices required to handle a specified workload (N) has one (+1) additional backup device to ensure resilience and high availability in the event of a failure in the required devices.

NetFlow: A network protocol developed by Cisco for collecting IP traffic information and monitoring network traffic.

network functions virtualization (NFV): A network architecture that uses virtualization technologies in carrier-grade or cloud provider equipment such as firewalls, IDS/IPS, load balancers, SBCs, and WAN accelerators.

network interface card (NIC): An adapter that permits a computer or other system to be connected to a network.

network packet broker (NPB): A device that provides network traffic visibility to multiple monitoring tools. Sometimes referred to as a network visibility controller (NVC).

network perimeter: A network boundary typically demarcated by a network (perimeter) firewall, such as between an enterprise network and the Internet.

network traffic: The data moving across a network at a given point in time. Network data in computer networks is mostly encapsulated in network packets, which provide the load in the network.

network tap: A hardware device that provides a way to access (or see) the data flowing across a computer network.

next-generation firewall (NGFW): A network security platform that fully integrates traditional firewall and network intrusion prevention capabilities with other advanced security functions that provide deep packet inspection (DPI) for complete visibility, accurate application, content, and user identification, as well as granular policy-based control. See also **deep packet inspection (DPI)** and **intrusion prevention system (IPS)**.

NFV: See **network functions virtualization**.

NGFW: See **next-generation firewall**.

north-south traffic: Network traffic that traverses a network firewall or boundary, such as traffic from a data center to the Internet or to a client endpoint. See also **east-west traffic**.

NPB: See **network packet broker**.

out-of-band tools: Network and security monitoring tools that receive a copy of network traffic from a SPAN port rather than directly inline. See also **inline tools** and **switched port analyzer (SPAN)**.

passive monitoring: A type of network monitoring in which copies of actual traffic, captured from a SPAN port or network tap, are used for analysis. See also **active monitoring**.

Payment Card Industry Data Security Standard (PCI DSS): A proprietary information security standard mandated for organizations that handle American Express, Discover, JCB, MasterCard, or Visa payment cards.

PCI DSS: See *Payment Card Industry Data Security Standard*.

Personally Identifiable Information (PII): Any personal data that can potentially be used to identify a specific individual such as full name, home address, date of birth, birthplace, Social Security number, passport number, driver's license number, and telephone number, among others, as well as email address and IP address (in some cases).

PHI: See *Protected Health Information (PHI)*.

phishing: A social engineering technique in which an email that appears to be from a legitimate business (typically a financial institution or retail store) attempts to trick the recipient into clicking an embedded link in the email or opening a attachment (containing malware or an exploit). The embedded link redirects the recipient's browser to a malicious website to enter sensitive personal information (such as account information). Alternatively, the malicious website may deliver malware or an exploit to the victim's endpoint in the background via the browser. See also **drive-by-download**.

PII: See *Personally Identifiable Information*.

private cloud: A cloud computing deployment model that consists of a cloud infrastructure that is used exclusively by a single organization.

Protected Health Information (PHI): Any information about health status, health care or health care payments that can be associated with a specific, identifiable individual, as defined by HIPAA. See also **Health Insurance and Portability Act**.

public cloud: A cloud computing deployment model that consists of a cloud infrastructure that is open to use by the general public.

resilience: A characteristic of a monitoring system, achieved most often through failsafe deployment and redundant components, that allows it to continue operating in the event of a failure in a system component.

SDN: See *software-defined networking*.

Secure Sockets Layer/Transport Layer Security (SSL/TLS): A transport layer protocol that provides session-based encryption and authentication for secure communication between clients and servers on the Internet.

security appliance: A type of standalone network device designed to protect the network from unwanted traffic. Common examples are firewalls, intrusion prevention, and antivirus solutions. Frequently, but not always, deployed inline to monitor live network traffic.

security fabric: An integrated solution with a centralized management interface that provides full access (visibility), aggregation, filtering, and processing of network traffic from across the entire organization, and delivers it to a wide variety of security and performance monitoring tools.

security information and event management (SIEM): Security technology that provides real-time collection and analysis of security alerts from various network hardware and applications.

service-level agreement (SLA): A contract between a service provider and its customers (internal or external) that formally defines the service that is being provided and specific service requirements such as performance, problem management, responsiveness and availability. The SLA also typically includes penalties for non-compliance, such as service credits or refunds.

SIEM: See *security information and event management*.

SLA: See *service-level agreement*.

software-defined networking (SDN): An open standards-based network architecture that decouples the network control and forwarding functions from the underlying physical networking infrastructure.

SPAN: See *Switched Port Analyzer*.

spyware: Malicious software that gathers information about a person or organization without their knowledge or consent.

Switched Port Analyzer (SPAN): A port on a network switch that is configured to mirror traffic from a source port to the analyzer port, in addition to the destination port.

tap: See *Test Access Port*.

Test Access Port (tap): An external hardware device inserted at a specific point in a network to monitor traffic along that network segment.

threat intelligence feed: Real-time information on global cyber threats that organizations can use to protect their networks from sources verified to have participated in previous cyberattacks or other suspicious activities.

unified threat management (UTM): A security appliance that integrates various security features such as firewall, anti-malware, and intrusion prevention capabilities into a single platform.

UTM: See *unified threat management*.

virtual local area network (VLAN): A network broadcast domain that is partitioned and isolated at OSI layer 2. See also *Open Systems Interconnection (OSI) model*.

virtual machine (VM): An operating system (OS) or application environment that emulates a physical computer.

virtual network interface card (vNIC): A software network interface card associated with a virtual machine to provide network connectivity. See also *virtual machine*.

virtual network tap: Software that provides full access (or visibility) to all the traffic crossing virtual machines and virtual network devices in a virtualized environment.

virtual traffic: See also *east-west traffic*.

virtualization: Computing technology that abstracts software from the underlying hardware kernel and other physical architecture.

virus: Malicious code that embeds itself within other software or programs and replicates itself.

visibility solution: Provides a way to see details about the traffic flowing across a network, in order to identify and take action on security threats or performance issues.

VM: See *virtual machine*.

vulnerability: A bug or flaw in software that creates a security risk which may be exploited by an attacker.

WAF: See *web application firewall*.

web application firewall (WAF): A firewall designed to protect web-based applications and web servers.

MOSTLY SECURED

Don't settle for mostly secure. Only Ixia offers total network visibility with a security fabric.



www.ixiacom.com

íxíá

Get total visibility in your security fabric

To keep your network defenses strong, you need visibility to all the traffic in your enterprise. You also need the ability to see inside that traffic — to find exactly the right data for each of your security and monitoring tools. And then you need to deliver that data with speed and efficiency, without fail. A security fabric pinpoints and delivers the right data, but you need to make sure it has access to all the traffic. This book describes how turning your focus to network visibility will strengthen security and help you improve network performance.

Inside...

- Improve your security posture
- Increase efficiency and flexibility
- Send the right data to the right tools
- Add application and threat intelligence
- Make the most of tool investments
- Accelerate troubleshooting
- Enable microsecond failover

ixia

Lawrence C. Miller has worked in information technology for more than 25 years. He has written more than 50 *For Dummies* books.

Go to **Dummies.com®**
for videos, step-by-step photos,
how-to articles, or to shop!

for
dummies®

A Wiley Brand

ISBN: 978-1-119-37949-2
Ixia collateral number:
915-8126-01-1071
Not for resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.