

Definitive GuideTM to *Complete Network Visibility*

How to Get High-Performing, Secure
Networks While Staying Within Budget



Jon Friedman

FOREWORD BY:
Bassam Khan

Product and Technical Marketing

Compliments of:

Gigamon[®]

About Gigamon

Gigamon® is the first and only company to deliver unified network visibility and analytics on all data-in-motion, from raw packets to apps, across physical, virtual and cloud infrastructure. We aggregate, transform and analyze network traffic to solve for critical performance and security needs, including rapid threat detection and response, freeing your organization to drive digital innovation. In short, we enable you to run fast, stay secure and innovate. Gigamon has been awarded over 75 technology patents and enjoys industry-leading customer satisfaction with more than 3,500 organizations, including 80 percent of the Fortune 100. Headquartered in Silicon Valley, Gigamon operates globally. For the full story on how Gigamon can help you, please visit www.gigamon.com.

Definitive GuideTM to *Complete Network Visibility*

How to Get High-Performing, Secure
Networks While Staying Within Budget

Jon Friedman

Foreword by Bassam Khan
Product and Technical Marketing



CYBEREDGE
P R E S S

Definitive Guide™ to Complete Network Visibility

Published by:

CyberEdge Group, LLC

1997 Annapolis Exchange Parkway

Suite 300

Annapolis, MD 21401

(800) 327-8711

www.cyber-edge.com

Copyright © 2020, CyberEdge Group, LLC. All rights reserved. Definitive Guide™ and the CyberEdge Press logo are trademarks of CyberEdge Group, LLC in the United States and other countries. All other trademarks and registered trademarks are the property of their respective owners.

Except as permitted under the United States Copyright Act of 1976, 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, without the prior written permission of the publisher. Requests to the publisher for permission should be addressed to Permissions Department, CyberEdge Group, 1997 Annapolis Exchange Parkway, Suite 300, Annapolis, MD, 21401 or transmitted via email to info@cyber-edge.com.

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. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. 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 CyberEdge Group research and marketing consulting services, or to create a custom Definitive Guide book for your organization, contact our sales department at 800-327-8711 or info@cyber-edge.com.

ISBN: 978-1-948939-10-2

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgements

CyberEdge Group thanks the following individuals for their respective contributions:

Editor: Susan Shuttleworth

Graphic Design: Debbi Stocco

Table of Contents

Foreword	v
Introduction	vii
Chapters at a Glance	vii
Helpful Icons	viii
Chapter 1: A Tool to Empower Tools	1
The Core Capabilities of NGNPBs	3
Use Cases for NGNPBs	7
Chapter 2: Use Case: Inline Bypass	11
Mix, Match, Bypass, and Toggle	11
Dynamic Load Balancing	12
Application Filtering	15
Inline Bypass: The Traffic Must Continue.....	15
From Inline to Out of Band	17
Chapter 3: Use Case: Offloading and Sharing Services	19
The Principle: Do It Once and Share.....	19
SSL/TLS Decryption	21
Be Ready for TLS 1.3	26
Generating and Distributing Metadata	27
De-duplication	31
And More.....	31
Chapter 4: Use Case: Visibility into Virtual and Cloud Environments	33
Behind a Veil: Visibility into Virtual and Cloud Environments.....	33
Challenges in Virtual Environments	34
Monitoring Virtual Environments.....	35
Challenges in Cloud Environments	38
Monitoring Applications on Cloud Platforms	39
Making Security a Technology Enabler Instead of a Blocker	41
Chapter 5: Use Case: Out-of-Band Security and Performance Monitoring Tools ...	43
The Proliferation of Out-of-Band Tools	43
Comprehensive Visibility for Analytics Tools	44
End-to-End Visibility for Performance Monitoring Tools.....	45
Unsampled Data	46
Packets and Metadata	47
De-duplication and VLAN Tagging	47
Overcoming the Limitations of SPAN Ports.....	47
Monitoring at Full Line Speeds	48
Chapter 6: Into the Future	49
How We Got Here: The Evolution of NGNPBs	49
Where We Are Going: Enabling New Initiatives.....	51

SIEMs: Going Beyond Logs	51
Harmonizing NetOps and SecOps	52
Data Persistence and Data Lakes.....	53
Empowering Threat Hunters.....	55
Chapter 7: Selecting the Right NGNPB.....	57
A Strategic Investment.....	57
Comprehensiveness of the Solution	57
Scalability and Capacity for Ultra-high-speed Networks.....	59
Integration with Cloud Platforms.....	59
Vendor Focus and Track Record	60
Closing Thoughts	60
Appendix: Key Features of NGNPBs	61

Foreword

Among the goals of digital transformation (DX) are streamlining business processes, cutting costs, improving productivity and introducing new business models that redefine industries or build new ones. The benefits are many — faster time to market, impactful customer experiences, shorter development cycles, organizational agility, to name a few — but the execution is unique to your organization’s circumstances. That’s not to say we cannot pinpoint a universal prerequisite for success: high-speed networks and efficient applications that run fast, respond to changing circumstances and stay secure.

To succeed in this DX world, moreover, NetOps, InfoSec, CloudOps and service provider operations teams depend on an array of network performance monitoring tools getting the information needed to do their jobs. Cybersecurity tools also play an ever-larger role today, given the network’s naked exposure to the wild-west public internet.

Unfortunately, these tools are only as good as the data that reaches them. And very often that data is too little — or too much. What’s needed is comprehensive network visibility across datacenters, remote locations, virtual and containerized, and public and private clouds. And you can gain such all-encompassing visibility only by using a next-generation network packet broker (NGNPB), such as the Gigamon Visibility and Analytics Fabric.

NGNPBs eliminate blind spots by capturing network traffic and metadata across physical, virtual and cloud environments. They filter network traffic and metadata, so tools receive only the types of data they are designed to process, which lets them handle much greater volumes of network traffic.

NGNPBs offer other remarkable features as well. They create decryption zones where multiple tools have access to decrypted SSL/TLS traffic. They offload packet deduplication and header stripping from network devices to give them more capacity. They provide inline bypass and load balancing for NetOps and cybersecurity tools so those tools don’t become bottlenecks or single points of failure.

In short, these hard-working tools offer sophisticated features that make entire ecosystems of NetOps and cybersecurity tools more efficient and more effective.

In a larger sense, NGNPBs empower NetOps, InfoSec, CloudOps and service provider operations teams to produce better outcomes with less effort at lower cost. Their tools produce more accurate results because they have access to complete network traffic from one source. Team members spend less time on repetitive manual tasks involved with collecting and filtering data, so they have more time for research, analysis and planning. And managers get more out of existing tools instead of blowing their budgets every time network traffic creeps upward.

This guide gives you a short but thorough overview of NGNPBs: what they are, their most important use cases, how they evolved from yesterday's network packet brokers, and how they benefit your teams and help them work together.

We at Gigamon hope you will find this information stimulating and useful. Please don't hesitate to contact us for more details, including examples of how NGNPBs are being used in organizations like yours.

Bassam Khan

Vice President of Product and Technical Marketing at Gigamon

Introduction

Unless you have taken a close look at network packet brokers recently, you are probably not aware of how rapidly they have evolved. In a few short years, network packet brokers designed to help network administrators monitor networks have morphed into next-generation network packet brokers (NGNPBs) that serve the needs of IT security groups, data analysts, network operations staffs, and others.

They provide visibility into network traffic across the entire enterprise...make security tools more effective...help network and application management tools ensure reliable performance...overcome concerns about security visibility in the cloud...save millions of dollars in hardware and software costs... help service providers slash monitoring costs to increase average profitability per user...and help enterprises take advantage of virtualization and cloud platforms such as VMware NSX, Cisco ACI, Amazon Web Services (AWS), Microsoft Azure, and OpenStack.

The goal of this guide is to give you a short, clear introduction to NGNPBs and all the capabilities you didn't know they had.

Chapters at a Glance

Chapter 1, “A Tool to Empower Tools,” discusses four key capabilities of next-generation network packet brokers.

Chapter 2, “Use Case: Inline Bypass,” describes how features like load balancing and inline bypass reduce costs and increase the availability of security and performance monitoring tools.

Chapter 3, “Use Case: Offloading and Sharing Services,” explores the benefits of offloading decryption, metadata generation, and other services from individual tools.

Chapter 4, “Use Case: Visibility into Virtual and Cloud Environments,” explains how NGNPBs provide comprehensive visibility into network traffic across physical and virtual environments and public and private cloud platforms.

Chapter 5, “Use Case: Out-of-Band Security and Performance Monitoring Tools,” highlights how NGNPBs enable out-of-band security and performance monitoring tools to be more effective.

Chapter 6, “Into the Future,” enumerates how NGNPBs enhance the effectiveness of SIEMs, data lakes, and threat hunting groups, and help NetOps and SecOps collaborate.

Chapter 7, “Selecting the Right NGNPB,” suggests criteria you can use to find the NGNPB that best fits your organization.

The Appendix summarizes many of the key features of NGNPBs.

Helpful Icons



TIP

Tips provide practical advice that you can apply in your own organization.



DON'T FORGET

When you see this icon, take note as the related content contains key information that you won't want to forget.



CAUTION

Proceed with caution because if you don't it may prove costly to you and your organization.



TECH TALK

Content associated with this icon is more technical in nature and is intended for IT practitioners.



ON THE WEB

Want to learn more? Follow the corresponding URL to discover additional content available on the Web.

Chapter 1

A Tool to Empower Tools

In this chapter

- Review factors that have been compromising the effectiveness of security, analytics, and performance monitoring tools
- Learn four key capabilities of next-generation network packet brokers
- Preview key use cases for NGNPBs

IT and network organizations rely on a wide range of tools to detect and block cyberattacks and to improve application performance and availability. These tools inspect network traffic to perform tasks like identifying malware, recognizing patterns that indicate attacks, and diagnosing network problems.

Unfortunately, the effectiveness of these tools has been compromised by the explosive growth of network traffic and the complexity of technologies such as virtualization, cloud computing, and mobility. In particular:

- ✓ Surging network traffic often overwhelms security tools, preventing them from doing their jobs, or even worse, turning them into bottlenecks that degrade network and application performance.
- ✓ Virtual environments and cloud computing platforms create blind spots where security, analytics, and performance monitoring tools cannot access or analyze traffic between applications, or between the elements of multi-tier applications.

- ✓ The growing use of encryption creates additional blind spots or forces tools to expend resources decrypting and re-encrypting network packets.
- ✓ Faster networks, combined with requirements to monitor traffic more thoroughly in more places, drive up the costs of acquiring, upgrading, and managing more tools.

As illustrated in Figure 1-1, high-speed networks and new computing models are forcing IT organizations either to accept gaps in security and network management or to dramatically increase their spending.



Figure 1-1: Faster networks and new technologies are driving up costs and causing gaps in security and network management.

Fortunately, there is a solution to these harsh trade-offs.

Next-generation network packet brokers (NGNPBs) are tools that enable other tools — security, analytics, and performance monitoring products — to function efficiently and effectively. They aggregate traffic across the network, optimize the flow of packets to individual tools, offload tasks from those tools, and use a variety of techniques to increase the availability of networks and applications. These capabilities allow enterprises to maximize network availability, improve security, and reduce costs, even as network traffic surges and computing environments become more complex.

The Core Capabilities of NGNPBs

NGNPBs have four core capabilities:

1. They collect and aggregate network traffic on behalf of security, analytics, and performance monitoring tools.
2. They intelligently filter the traffic, so each tool receives exactly what it needs and no more.
3. They offload services such as decryption, so tools can perform their primary tasks more efficiently.
4. They increase the resiliency and security of networks and security tools by balancing workloads and eliminating single points of failure.

Let's examine these capabilities.

Aggregation: collect all the traffic

NGNPBs operate under a “collect once, share everywhere” principle. They collect and aggregate network traffic across the enterprise, including traffic that is flowing:

- ✓ In and out of corporate datacenters (“north-south traffic”)
- ✓ Across systems, zones, and information silos within datacenters (“east-west traffic”)
- ✓ Between software instances running on virtual machines in virtualized environments
- ✓ Among application modules and services running in different tiers and different regions on public and private cloud platforms

The NGNPBs then share the traffic with all the security, analytics, and performance monitoring tools used in the enterprise. Each tool has access to all the traffic it needs to do its job. (Figure 1-2)

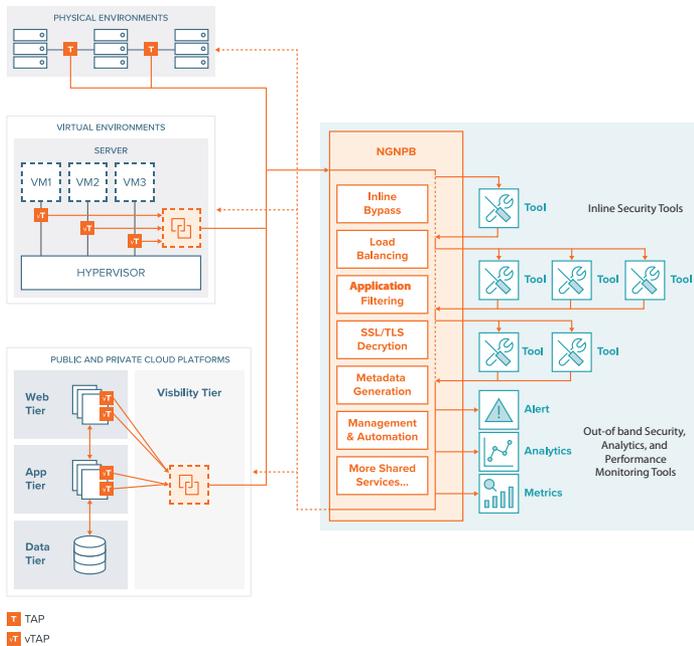


Figure 1-2: NGNPBs aggregate network traffic from different types of environments and share it with security, analytics, and performance monitoring tools.

Comprehensive visibility into network traffic improves security by ensuring that detection and protection tools don't miss indicators of compromise (IOCs), and security analysis tools have access to all network traffic to better detect patterns and trends. It also gives network and application performance monitoring tools data to track performance of transactions and network flows from end to end, so they can do a better job of troubleshooting problems and detecting trends.

Application filtering: send only what's needed

Only a few security products monitor emails. Not many tools are designed to inspect video traffic. Some traffic never needs to be inspected by security tools because it is covered by privacy regulations, or is trusted, or is considered low risk. Sending all traffic to every security tool wastes resources and can cause tools to become network bottlenecks.

NGNPBs can provide application filtering, the ability to differentiate among traffic types based on a wide range of criteria, including application type and the sources and destinations of the traffic. They use this information to route to each tool exactly the packets it needs – and no more.

This selectivity dramatically reduces the load on many tools, allowing them to perform better and preventing them from becoming bottlenecks that slow application performance. By filtering out low-risk and irrelevant communication, it also allows tools to inspect more high-risk traffic.

Offload services: optimize resources

Certain tasks are performed by multiple security tools and network devices. Examples include decrypting and re-encrypting SSL/TLS traffic, de-duplicating packets, and generating NetFlow (IPFIX) metadata. NGNPBs can perform these tasks once on behalf of all security, analytics, and performance monitoring tools in the enterprise. The performance hit for each task is only incurred once. The tools can perform their core tasks more efficiently, so fewer are needed, which reduces costs.

NGNPBs also improve security by ensuring that all tools, including those that do not have their own decryption capabilities, can inspect SSL/TLS traffic and detect hidden threats.

Increased resiliency: improve availability while lowering costs

NGNPBs provide a number of features that increase the performance and resiliency of networks and security tools, for instance, load balancing and inline bypass.

As a result:

- Multiple security appliances can share work and collaborate to handle spikes in network traffic.
- When inline tools fail, traffic can be routed around them automatically.
- Tools can be taken offline for maintenance and upgrades with no outages or network downtime.

- ✓ Security tools can be operated out of band under normal conditions to minimize network latency, but switched inline automatically during attacks to increase protection.

And more...

Later in this guide you will hear about features of NGNPBs such as load balancing, inline bypass, data masking, and centralized management. These provide additional benefits including:

- ✓ Reducing “forklift upgrades” of security appliances
- ✓ Slashing the number of interface network cards, test access points (TAPs) and Switched Port Analyzer (SPAN) ports the enterprise needs to purchase and deploy
- ✓ Enforcing privacy policies
- ✓ Making it easy to test security tools with production data.

Big Savings

A manager at one company using an NGNPB estimated that it allowed his organization to protect itself with as few as one-quarter of the security tools that would otherwise be needed.

The manager is quoted in a Total Economic Impact™ study conducted by Forrester Consulting on behalf of Gigamon. In the same study the Forrester analysts estimate

that by adopting an NGNPB (the GigaSECURE® Security Delivery Platform) a composite organization of 5,000 employees could save a net total of \$1.6 million over three years on security hardware, software, and staffing, with payback in seven months.

You can download the study at <https://insight.gigamon.com/forrester-tei-report.html>.

Use Cases for NGNPBs

Of course, IT organizations don't invest in technology for technology's sake. The next four chapters of this guide examine four use cases that have proved the most valuable for NGNPB users. Figure 1-3 provides a brief preview of those use cases.

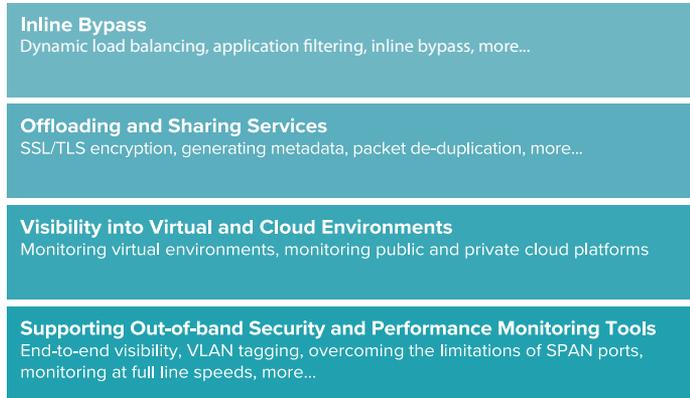


Figure 1-3: Four common key use cases for NGNPBs.

Inline bypass

Inline bypass refers to a collection of features that increase the availability and performance of networks and tools, including load balancing, bypassing failed security tools, sensing when tools go offline and when they come back online, and toggling tools between inline and out-of-band modes.

Inline bypass features are extremely valuable for inline tools, that is, tools that are deployed directly on the network and inspect packets as they come by in real time. Inline tools can block suspicious traffic immediately, but they can also become network bottlenecks and single points of failure. For these tools, inline bypass features provide advantages such as:

- ✓ Increasing availability and reliability by reducing the need for planned downtime and by providing automated failover when they go down

- ✓ Making better use of capacity to ensure that tooling keeps pace with network upgrades and organizations don't need to buy and manage as many units
- ✓ Dynamically rebalancing performance and security by allowing tools to work out of band during normal circumstances and automatically toggling them to inline mode when an attack is detected

In Chapter 2 we will examine these features and how they affect security, application performance, and network availability.

Services offloading

NGNPBs can offload a surprising number of services from individual security, analytics, and performance monitoring tools. In Chapter 3 we will highlight six of them:

1. SSL/TLS decryption
2. Metadata generation
3. Packet de-duplication
4. Header stripping
5. Packet slicing
6. Masking

Visibility into virtual environments and cloud platforms

Conventional security and analytics tools have little or no visibility into virtual environments and cloud platforms.

TAPs and SPAN ports cannot see traffic between virtual machines running on the same physical server. In addition, virtual environments are extremely dynamic, starting up and moving software instances too fast for administrators to react.

Public cloud platform providers don't provide easy access to network traffic flowing on their platforms, so enterprises can't put their own security and performance monitoring tools in the cloud. Many existing security tools and processes simply won't work with cloud-based applications.

In Chapter 4 we will examine some of the special requirements of providing visibility into virtual environments and cloud platforms and discuss how NGNPBs address them.

We will also look at how NGNPBs help CloudOps teams overcome concerns about security visibility on cloud platforms and ease migration from private to public clouds, and how service providers can gain visibility into subscriber activity at scale.

Support for out-of-band security and performance monitoring tools

Many security, analytics, and performance monitoring tools operate *out of band*, meaning that they inspect copies of network packets (not the original packets) and metadata about traffic flows. Although these tools cannot become network bottlenecks, they need to be scaled to be able to process all of the traffic they receive so as not to miss any vital information. To do their jobs effectively, they also need comprehensive visibility into network traffic from everywhere in the enterprise.

Performance monitoring tools are a class of out-of-band tools that have special requirements. Most NetOps teams use several products to track different types of performance, including network performance management (NPM), application performance management (APM), and customer experience monitoring tools. To provide accurate results they need end-to-end visibility into network traffic and access to comprehensive network metadata, as well as help dealing with features like VLAN tagging.

In Chapter 5 we will explore how NGNPBs provide end-to-end visibility and improve the flexibility and manageability of out-of-band tools. We will also review how they help NetOps teams use performance monitoring tools to do a better job tracking performance data, troubleshooting issues, and analyzing trends.

Chapter 2

Use Case: Inline Bypass

In this chapter

- Review how dynamic load balancing can reduce costs and make better use of capacity
 - See how logical bypass can work around tool failures
 - Learn the benefits of toggling tools between out-of-band and inline modes
-

Mix, Match, Bypass, and Toggle

Mix, Match, Bypass, and Toggle. The title of a dubious self-help book? A sleazy personal liability law firm? A forgotten 1960s dance craze?

No, these words summarize a collection of next-generation network packet broker features that go under the general heading of “inline bypass.” Those features make better use of the capacity of security tools and increase their availability and reliability.

Inline bypass features apply primarily to inline tools such as firewalls, next-generation firewalls (NGFWs), intrusion prevention systems (IPSS), web application firewalls (WAFs), and advanced threat protection (ATP) tools. These tools inspect the traffic and can automatically block sessions that they associate with IOCs or suspicious behaviors.

Dynamic Load Balancing

Load balancing techniques have long applied to network devices, databases, storage networks, and many other IT technologies. Next-generation network packet brokers apply these techniques to distribute traffic among inline tools.

Figure 2-1 illustrates a typical scenario: an NGNPB load balances network traffic across IPS appliances.

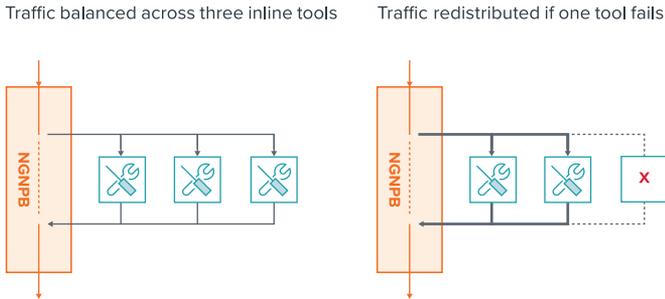


Figure 2-1: An NGNPB load balances traffic across IPS appliances.

Performance and cost

Load balancing smooths out network performance and reduces costs by making better use of available capacity.

If each IPS appliance shown in Figure 2-1 were inspecting a different network link in isolation, you might find situations where capacity utilization averaged 50 percent on each appliance but occasional surges in traffic overwhelmed one of the units. During the surge, this unit would become a bottleneck, slowing network performance and annoying application users.

With load balancing, the aggregate traffic can be shared equally among the appliances. Spare capacity on all three units is available to handle spikes in traffic so network and application performance is not affected. In addition, it might be safe to allow overall capacity to reach, say, 60 percent or 70 percent before upgrading any of the appliances.

Scaling by adding, not replacing

Let's say you are upgrading a network link from 10Gb to 40Gb. Without load balancing, you almost certainly would need to perform a “forklift upgrade,” throwing out the old IPS and replacing it with a new one with greater capacity.

With load balancing, you can scale by adding rather than replacing; that is, you can run a new system next to the old one and continue to obtain value from both.

Load balancing can even allow you to mix and match units of different sizes. For example, if unit A has twice the capacity of unit B, the NGNPB can send two-thirds of the traffic to the former and one-third to the latter.

Load balancing and more efficient use of existing capacity ensure that tooling keeps pace with network upgrades and you won't need to buy and manage as many appliances.

Active-active redundancy for high availability

In most situations without load balancing, high availability is achievable only with an N+1 configuration, where there is one inactive system on standby to take over if the primary system fails. With dynamic load balancing, all systems can be in use at all times. If one fails, the NGNPB will divide the traffic among the remaining systems (Figure 2-2). You don't need to pay for a stand-by device that sits idle most of the time.

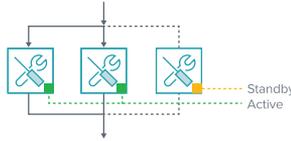
No downtime for planned maintenance

Nobody likes downtime for planned maintenance and upgrades: not the employees and customers it inconveniences, nor the administrators who have to perform the work during the late shift. Yet security tools, like all IT systems, must be maintained and upgraded.

NGNPBs provide an alternative. Administrators can direct all traffic to systems A and B, shut down system C and do their work, and bring system C back online, all without interrupting application availability.

Redundancy with an N+1 configuration

One system is on standby, inactive



High availability with an NGNPB

All tools are active

Traffic is redistributed if one tool fails

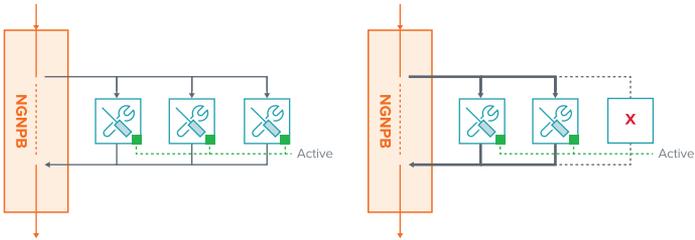


Figure 2-2: An N+1 configuration for redundancy vs. using load balancing for high availability.



When monitoring session traffic between a client and a server, many security, analytics, and performance monitoring tools need to see both sides of the conversation in order to identify suspicious behaviors, troubleshoot applications, and measure performance. This means that for a given session *all traffic in both directions* must pass through the same security tool. That can't be done with round-robin or other randomized load balancing techniques. A more effective strategy is to create hashes based on information in the header field of each packet, such as the source and destination IP addresses and TCP/UDP ports. Those hashes are used to assign packets to tools. Hashing ensures that the overall traffic flows are distributed evenly to each tool. Also, since every packet in a given client-server exchange will have the same source and destination addresses, they will have the same hash, so the NGNPB will assign all packets in that session to the same tool.

Application Filtering

Not all inline tools will inspect all network traffic. For instance, sending database replication traffic to a WAF is a waste: the WAF will use up processing cycles scanning traffic that it will ignore.

NGNPBs provide application filtering, allowing you to select traffic characteristics from Layer-2 to Layer-7 and use those criteria to determine what traffic is sent to each tool and what traffic will bypass it.

Inline Bypass: The Traffic Must Continue

We said at the beginning of this chapter that that the term inline bypass can refer to a collection of related features, but it also has a specific meaning: a feature that allows application traffic to continue flowing even when an inline security tool, or the NGNPB itself, goes down.

Logical bypass

Most non-technical executives and managers prioritize business processes over perfect cybersecurity. They are unhappy when applications slow down or become unavailable and are not receptive to explanations related to difficulties with the firewall, IPS, WAF, or ATP solution.

Fortunately, NGNPBs can keep application traffic flowing when an inline security tool fails because of a power outage or a software or hardware failure.

This capability is called logical bypass. The NGNPB monitors the health and performance of the tool with bidirectional heartbeat packets. If the tool goes offline or is overwhelmed by spikes in network traffic, the NGNPB arranges for the network traffic to bypass it, as shown in Figure 2-3. When the tool comes back online, the NGNPB automatically restarts traffic through it.

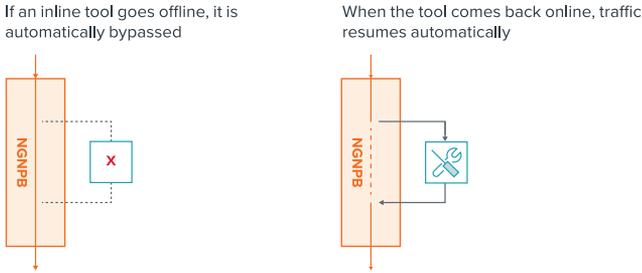


Figure 2-3: Logical bypass: when an inline security tool goes offline, the NGNPB senses the outage and keeps network traffic flowing around the tool.

Logical bypass can also be employed in a load balancing situation, as shown in Figure 2-4. In this scenario one tool in a pair goes down and the surviving unit does not have the capacity to pick up all the traffic. The surviving tool handles as much traffic as it can without slowing the network, and the NGNPB allows the rest to bypass the tool. This is also called “sampling” because only a portion of the traffic (a sample) is inspected.

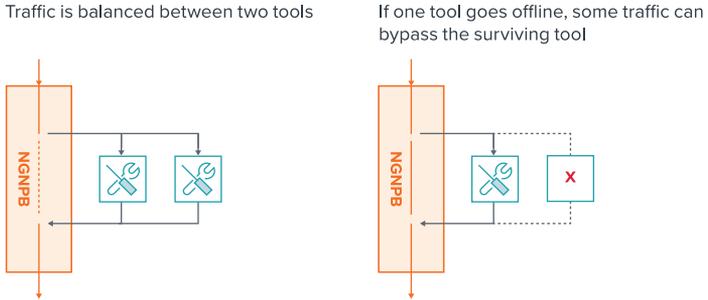


Figure 2-4: Logical bypass can also address a load balancing situation where one tool goes down by preventing the surviving tool(s) from being overwhelmed.

Physical bypass

Of course, we don’t want the NGNPB to become a single point of failure either. Physical bypass, sometimes called fail-to-wire, is another fail-safe mechanism. In the event of a power outage a physical relay is triggered and network traffic bypasses the NGNPB.



Physical bypass is a last resort. Look for NGNPBs that offer resilient designs and high availability configurations where one NGNPB unit can automatically fail over to another.

Daisy chaining

NGNPBs enable different types of inline tools to be daisy chained without creating single points of failure. With daisy chaining, upstream tools filter out large volumes of traffic, reducing the load on downstream tools. For example, a firewall might filter out traffic, helping a downstream IPS, which filters out more traffic to reduce the load on a distributed denial of service (DDoS) protection device.

But daisy chaining creates risks. In the example above, if the upstream firewall goes down, the traffic can't be inspected by *any* of the security tools in the chain. With an NGNPB, the firewall would be bypassed automatically and the downstream devices would continue to monitor traffic.

From Inline to Out of Band

NGNPBs provide a common platform to deliver traffic feeds to both inline and out-of-band tools. This flexibility greatly simplifies the management of security, analytics, and performance monitoring tools. It also enables a unique feature: the ability to switch inline tools back and forth between inline and out-of-band modes in milliseconds. Let's see when this capability might be useful.

Toggling when under attack

When fast response times and low latency are critical, you might deploy inline tools such as IPSs and APT systems in an out-of-band, "detect only" mode. This ensures that the tools will not become bottlenecks that slow performance.

However, when an ongoing attack is detected, an NGNPB can automatically toggle the tools to inline mode to block network traffic related to the attack (Figure 2-5). This arrangement maximizes application performance and availability, but gives priority to security when needed.

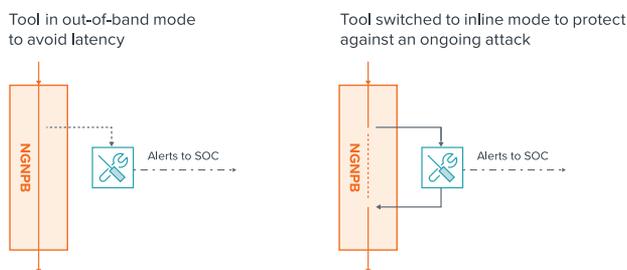


Figure 2-5: NGNPBs can switch tools from out-of-band to inline mode during an attack.

Testing and deploying tools

NGNPBs also make it easier to test and deploy new security and analytics tools. Tools being evaluated can be fed real production traffic in out-of-band mode, which increases the accuracy of the testing. Multiple tools can be tested side-by-side, with the same traffic, for true apples-to-apples comparisons. When a new inline tool is selected, it can be switched over to inline mode in seconds without causing a network outage.

The same approach can be used with security and analytics tools that require a learning period to observe network traffic. They can be deployed out of band, then switched to inline mode when the learning period is complete.



For more on inline bypass features, see the Gigamon feature brief, “[Why Security Tools Need Inline Bypass](#)” and the white paper “[Inline Bypass: Scaling Inline Threat Prevention Tools to Keep Pace with High-Speed Networks](#).”

Chapter 3

Use Case: Offloading and Sharing Services

In this chapter

- Examine the benefits of offloading tasks from tools
 - Learn how NGNPBs offload SSL/TLS decryption
 - Review other shared services offered by NGNPBs
-

The Principle: Do It Once and Share

In the typical enterprise several functions are duplicated or performed multiple times by security, analytics, and performance monitoring tools. These tasks include decrypting network traffic, generating metadata, and de-duplicating packets. In this chapter we will review the benefits of having next-generation network packet brokers perform these functions centrally, and then describe how those benefits play out for several shared services.

Offloading work

The most obvious benefit of performing a shared service is offloading work from the individual tools. It is clearly more efficient to perform an activity once and share the result with many tools, instead of repeating the same task two, three, or more times.

Into the Hardware

Optimization is another reason to perform tasks on a specialized platform rather than on devices built for something else. For example, tasks like decryption and re-encryption can be performed in

hardware with a specialized chip rather than in software running on a general-purpose chip. Optimization provides more “bang for the buck” and helps NGNPBs scale.

Ensuring consistency

Different tools may perform the same function using different algorithms, or export results in different formats. For example, network management tools deployed at different times might generate metadata using different versions of the NetFlow and IP Flow Information Export (IPFIX) standards. These variations can make it difficult for the NetOps team and SIEMs to correlate metadata from those tools. When an NGNPB takes over the function of generating metadata, the outputs are consistent and can be compared by analysts and tools.

Keeping technologies current

Most enterprises do not upgrade every tool every year. Not all vendors keep up with all the latest technologies, especially in areas that represent only a small part of their offerings. As a result, desirable new technologies such as support for strong cryptology methods may not be available to the enterprise quickly. But when a technology is updated in an NGNPB, it benefits all the tools that utilize the results.

Enforcing policies

Some services become even more useful when combined with other features of an NGNPB. For example, application filtering allows an NGNPB to apply policies selectively. Video streams from YouTube and Netflix don't need to be inspected by the same tools as most traffic. Special policies can be applied to internally developed applications. To comply with privacy regulations, the NGNPB can ensure that some types of application traffic are always encrypted and that personally

identifiable information (PII) is never included in specific traffic streams.

Simplifying scalability

As the volume of traffic increases, it is much easier and more economical to add capacity for a service in one place, the NGNPB, rather than upgrading multiple tools.

SSL/TLS Decryption

SSL/TLS encryption is becoming a de facto standard on the web for ecommerce, online banking, email, search, social media, voice-over-IP (VoIP), file storage, and many other public-facing applications. It is increasingly used for internal network traffic as well, as enterprises deploy more software to private and public clouds and make wider use of software-as-a-service (SaaS) applications.

Encryption Has Taken Over

Hypertext Transfer Protocol Secure (HTTPS) is now used for a majority of web traffic according to several metrics. According to Mozilla, traffic on HTTPS-encrypted websites grew from about 50 percent of all web traffic at the beginning of 2017 to over 70 percent by mid-2018. A Google Transparency Report

published in October 2017 showed that of the top 100 websites worldwide, 70 use HTTPS by default.

That's the good news. The bad news is that according to industry analyst firm Gartner, soon half of malware campaigns will use some type of encryption to hide malicious activities.

Implications for security tools

Unfortunately, encryption is also being widely adopted by cybercriminals and hackers to disguise their actions. Bad actors are using encryption to:

- ☑ Conceal malware downloaded from websites, sent through social media, and transferred as attachments to emails and instant messages
- ☑ Hide command and control (C2) traffic into and out of the corporate network

- ✓ Cloak the exfiltration of stolen data to remote websites
- ✓ Make it harder to detect and trace DDoS attacks

A wide range of security tools, including malware detection tools, NGFWs, web security gateways, IPSs, and data loss prevention (DLP) systems, can't do their jobs unless they are able to view traffic after it has been decrypted.

In addition, since internal communication is being encrypted, security tools now need decryption in order to inspect east-west as well as north-south traffic.

But decryption is a processor-intensive function that steals a large amount of resources from security tools. In a study of eight leading NGFWs, NSS Labs found that turning on SSL decryption degraded the performance of the firewalls by as much as 80 percent, and reduced transactions per second by as much as 92 percent. (NSS Labs: SSL Performance Problems.)

Implications for analytics and monitoring tools

Similar considerations apply to analytics and performance monitoring tools. In many situations encryption hides packet headers as well as payloads, making it difficult or impossible to read data about application type, source and destination addresses, ports, DNS lookups, certificates, and many other types of information used to measure performance, map data flows, find anomalous behaviors, and detect trends.

In fact, obtaining “pervasive visibility” into encrypted traffic is critical for:

- ✓ Security analytics and user and entity behavior analytics (UEBA) tools
- ✓ Cloud service monitoring tools
- ✓ Network and application monitoring tools

SSL, TLS, and HTTPS

Secure Sockets Layer (SSL) was the first widely used cryptographic protocol for establishing an encrypted link between web servers and browsers. It provided mechanisms for confidentiality (preventing traffic between the source and destination from being understood), authentication (proving that servers and clients are who

they represent themselves to be), and integrity (ensuring that the messages have not been tampered with).

The Internet Engineering Task Force has now deprecated SSL in favor of its successor, **Transport Layer Security (TLS)**. **HTTPS** is a slightly higher-level protocol used to piggyback HTTP on top of TLS.

Offloading SSL/TLS decryption to an NGNPB

The benefits of offloading SSL/TLS decryption to an NGNPB include increased visibility and reduced costs.



For inline tools, “decryption” is often used as shorthand for both *decryption* of network packets received from the source and *re-encryption* of the packets before they are sent to the destination.

Figure 3-1 is a diagram of a typical layout for decryption in part of an enterprise network. You may notice several non-optimal features of this arrangement:

- ☑ Not every tool has visibility into all the decrypted traffic it might want to inspect.
- ☑ The same traffic is decrypted and re-encrypted by several tools because the tools cannot share decrypted traffic with each other.

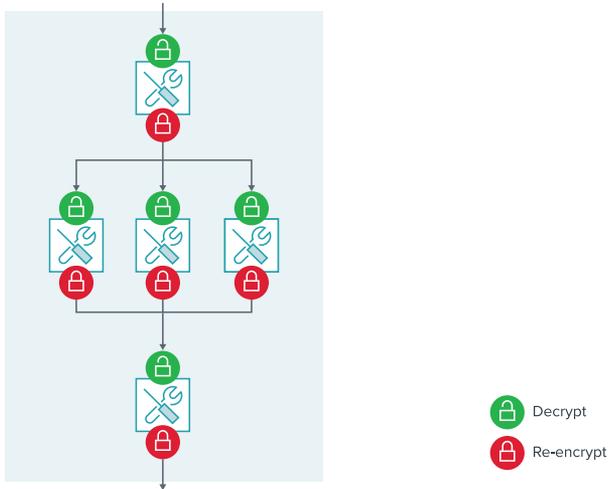


Figure 3-1: In a typical environment, tools do not have visibility into all the decrypted traffic, and decryption is performed multiple times on the same traffic.

Figure 3-2 shows how these issues could be addressed by an NGNPB. The NGNPB creates a “decryption zone” where SSL/TLS traffic from all TCP ports and applications is decrypted once and fed to multiple tools. In this scenario:

- ✓ Every tool has visibility into all the decrypted traffic it can usefully inspect, increasing the accuracy of the security and analytics tools.
- ✓ The same traffic is decrypted and re-encrypted only once, and that work is offloaded from the tools, speeding up performance, reducing latency introduced by decryption, and enabling fewer tools to handle the same traffic.

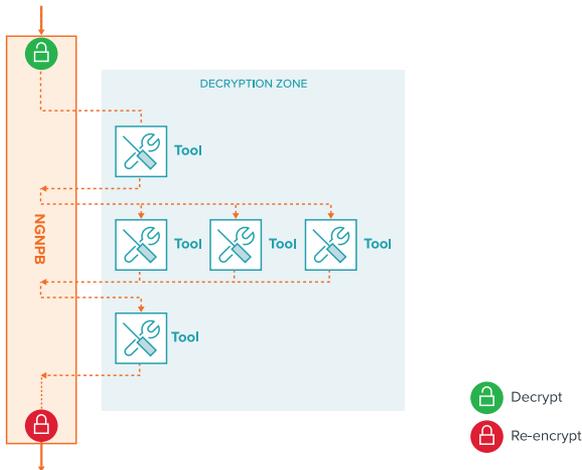


Figure 3-2: An NGNPB creates a decryption zone where tools have visibility into all decrypted traffic and decryption is performed once with no impact on the performance of the tools.

Supporting advanced cryptographic techniques

An NGNPB can support advanced cryptographic techniques, including not only RSA and static Diffie-Hellman (DH), but also elliptic curve (EC), Diffie-Hellman Ephemeral (DHE), and other types of perfect forward secrecy (PFS) public-key protocols. This support allows enterprises to utilize the latest encryption technologies without needing to upgrade every tool in the environment.

Checking for valid certificates

An NGNPB can check certificate revocation lists (CRLs) and use the Online Certificate Status Protocol (OCSP) to determine whether encryption certificates are invalid or have been revoked by the issuing certificate authority (CA). This helps authentication by making it more difficult for attackers to spoof legitimate websites.

Enforcing privacy regulations and other policies

The European Union’s General Data Protection Regulation (GDPR) and other privacy regulations mandate that PII must never be inspected or stored in unencrypted form outside of specific, tightly controlled applications.

NGNPBs can use application filtering intelligence and application metadata intelligence to selectively decrypt traffic using URL categories, IP addresses, whitelists, blacklists, and other criteria. You can use this filtering to ensure traffic that includes PII remains encrypted.

You can use filtering to enforce other policies, for example encrypting all traffic that leaves the network, but not encrypting traffic that remains inside the network.

Be Ready for TLS 1.3

Version 1.3 of the TLS protocol was published by the IETF in August 2018. It includes improvements that will speed up handshakes between endpoints, reduce the time needed for authentication, and prevent a hacker who acquires today’s secret key from using it to decrypt past and future traffic.

However, TLS 1.3 also eliminates the use of RSA and other non-PFS public key protocols and encrypts all certificate data used for handshakes. These changes will complicate decryption by security, analytics, and performance monitoring tools.

Enterprises can rely on the companies providing NGNPBs to incorporate state-of-the-art methods for coping with the challenges of TLS 1.3; they won’t have to worry whether their other tool vendors will adapt to the changes.

ON THE WEB



For more information on TLS 1.3 and its implications, see the Gigamon white paper, “[What Do You Mean TLS 1.3 Might Degrade My Security?](#)”

Generating and Distributing Metadata

Fans of World War II spy stories love to read about how Alan Turing and his colleagues at Bletchley Park in England broke the codes used by the German Enigma machines. But in fact, most signals intelligence during that war came from the humble field of traffic analysis: observing patterns in communication, even when the messages themselves were unreadable. For example, an increase in radio traffic might indicate an imminent attack. Messages sent on certain schedules, with certain call signs, in certain codes, might pinpoint the location of a motorized infantry division, or a fighter squadron, or a naval task force.

In that context, the code breakers were trying to find the “data” (the content of the messages), while the traffic analysts were finding valuable intelligence in the “metadata” (information about the messages).

Likewise, there are many situations where you can learn a lot just by knowing who is speaking, their tone of voice, and to whom the words are addressed. That applies even if you don’t understand a word of what is being said: think of the last time you were on vacation in a country where you didn’t know the language.

For network-related tools, metadata is primarily information in network packet headers about the sources, destinations, protocols, and characteristics of the packets and the traffic flows they are part of. Just as in WWII and on your trip to a foreign land, metadata can be a very valuable source of intelligence.

Metadata also has the advantage of taking much less storage space than data from raw packet capture. Organizations can afford to keep metadata going back years, which can be extremely useful for performing forensic and historical analysis.



If you are interested in traffic analysis for military intelligence, the National Security Administration has published a handy overview: [The History of Traffic Analysis: World War I – Vietnam](#).

Challenges of collecting and distributing metadata

Today most enterprises use routers, switches, and dedicated NetFlow generation appliances to collect metadata and distribute it in NetFlow or IPFIX formats. This approach has several shortcomings:

- ✓ Metadata generation can consume a lot of the devices' resources and slow their performance.
- ✓ Different network equipment vendors have their own versions of the NetFlow/IPFIX format.
- ✓ Many network devices can distribute metadata only to a limited number of destinations, sometimes only two.

To avoid performance issues on their network devices, many enterprises configure them only to sample, generating metadata from a fraction of the packets in the traffic stream. This saves money on network devices, but it can severely handicap tools by:

- ✓ Causing the tools to miss IOCs and other clues in the unsampled packets
- ✓ Preventing the tools from discovering anomalous patterns in the network traffic, for example, spikes in DNS queries and traffic streams using nonstandard port and protocol combinations

Metadata generation by NGNPBs

Just like routers and switches, NGNPBs can inspect network packet headers and generate metadata in NetFlow/IPFIX format. And as with other types of shared services, NGNPBs can not only offload that work from routers and switches, they often perform the work better.

Providing comprehensive visibility

Because NGNPBs manage network traffic flows from across the enterprise, they can generate metadata that provides a complete view of application and network activity.

Metadata generation can be performed selectively: NGNPBs can target specified applications and network flows for metadata generation to avoid consuming resources on irrelevant traffic.

Generating complete (unsampled) data

NGNPBs are powerful enough to inspect and generate metadata from 100 percent of the packets in the network flow. There is no need to settle for sampling. As a result, security and analytics tools are much less likely to miss valuable clues or fail to identify anomalies.

Supporting multiple export formats

An NGNPB can support a wide range of export formats, such as NetFlow version 5 and version 9, IPFIX, and Cisco Express Forwarding (CEF). This compatibility allows it to export data to many different types and vintages of security and analytics tools.

Finding botnets and compromised endpoints

NetFlow/IPFIX metadata such as source and destination IP and MAC addresses can help analysts identify anomalous network traffic and trace it back to external websites, especially when it is correlated with information from threat intelligence feeds and security tools. If these websites are confirmed as hosting botnets or being used for command and control traffic by attackers, metadata can indicate which endpoints on the corporate network might have been compromised by the bad actors. Metadata might also show the attackers moving data around inside the network.

This kind of analysis reduces the time to discovery of ongoing attacks and shows analysts what systems need to be remediated.

Identifying attacks using DNS metadata and HTTP response codes

Some NGNPBs include extensions to the standard IPFIX format that enable collection of metadata related to DNS lookups, HTTP request methods, and HTTP response codes. This metadata can help NetOps and security teams in many ways.

For example, DNS- and HTTP-related metadata might reveal:

- ✓ Rogue DNS servers on the network
- ✓ DNS tunneling for data exfiltration
- ✓ High-entropy DNS lookups that indicate command and control traffic from attackers
- ✓ DNS servers with low time to live (TTL) values, a common sign of attackers who move their servers to new domain names frequently to avoid being detected and blocked
- ✓ Sequences of HTTP requests indicative of SQL injection attacks and other attacks associated with OWASP top vulnerabilities
- ✓ Excessive numbers of 3XX (redirection) HTTP response codes in packet headers, indicating that attackers may be sending captured information to a compromised internal server being used as a staging area for exfiltration



TIP

Look for NGNPBs that can generate metadata for business and social media applications that you specify. That feature will help your analysts zero in on potential threats.



ON THE WEB

To explore how metadata can be used for security and network management and how NGNPBs help generate and distribute metadata, download the Gigamon feature brief [NetFlow and Metadata Generation](#) and the white papers [NetFlow Generation: The Security Value Proposition](#) and [Nine Metadata Use Cases](#).

De-duplication

Typically, network traffic contains large numbers of duplicate packets. SPAN ports, application delivery controllers (ADCs) and other network devices frequently make two or more copies of the same packet. Also, multiple TAP points can result in the same packets being collected and sent to the same security or performance monitoring tool several times.

Besides absorbing bandwidth, duplicated packets can throw off the results of performance monitoring tools by giving a misleading impression of the number of packets generated by applications.

An NGNPB can strip out redundant network packets. This activity:

- ✓ Reduces network traffic
- ✓ Offloads the work of performing de-duplication from the tools
- ✓ Prevents tools from wasting capacity by inspecting the same packet two or more times
- ✓ Frees up storage space in network and forensic recording devices that store packets and related metadata

And More...

Header stripping

NGNPBs can strip specified headers, tags, and encapsulations from packets. This is helpful when packets contain MPLS, VLAN, VXLAN, VN-TAG, and GRE headers and tags that are meaningless to security, analytics, and performance monitoring tools. Removing them increases the performance of the tools, and often produces more-accurate results because the tools do not ignore packets with unreadable header types.

Packet slicing

For security, analytics, and performance monitoring tools that only look at packet headers, packet slicing leaves the headers

intact and slices off the payloads. This process reduces the volume of traffic sent to the tools, cuts storage requirements, and ensures that sensitive information is not inspected or stored in violation of security and privacy regulations. Some NGNPBs offer a variety of ways to define how this slicing takes place in order to account for the characteristics of your applications and packets.

Masking

NGNPBs can also perform masking, which overwrites specific fields with a pattern of meaningless characters. Masking is often used to protect passwords, full or partial account numbers, Social Security numbers, and other PII.

Offloading these services to the NGNPB not only frees up capacity on the tools, it also ensures that the tasks are performed consistently across all network traffic in the enterprise.

Chapter 4

Use Case: Visibility into Virtual and Cloud Environments

In this chapter

- Examine the challenges of providing visibility in virtual and cloud environments
- Learn about virtual taps (vTAPs) and visibility nodes
- Understand how NGNPBs can integrate with management tools on cloud platforms and orchestrate adjustment to changes

Behind a Veil: Visibility into Virtual and Cloud Environments

One of the themes of this guide is the importance of comprehensive visibility. To be fully effective, security tools need to spot IOCs and malware everywhere on the network. Analytics tools must have a complete set of network data to build accurate baselines, detect anomalies, and identify suspicious behaviors. Performance monitoring tools must have end-to-end visibility into application and network flows to help NetOps teams develop meaningful metrics and troubleshoot problems.

However, comprehensive visibility has become much harder to achieve as enterprises migrate their computing workloads to virtual environments and cloud platforms. The methods developed for traditional datacenters no longer suffice.

The major challenges created by these changes include:

- ✓ An inability to use conventional TAPs and SPAN ports to collect network traffic and metadata
- ✓ Continuous changes in the number and location of application instances as applications scale up and down to meet demand
- ✓ A lack of interoperability between the virtual and cloud versions of traffic monitoring tools and the on-premises versions

These factors create major blind spots for security, analytics, and performance monitoring tools. They make it extremely difficult to correlate security and network data across different environments to detect the lateral movement of attackers. They also lead some organizations to deploy different network data acquisition and monitoring tool sets for each environment, resulting in duplication and more complexity.

One of the hallmarks of next-generation network packet brokers is their ability to overcome these challenges. They offer a single solution for collecting, processing, and distributing network traffic across physical, virtual, and cloud environments, a solution that can cope with the native elasticity of applications on virtual and cloud platforms.

Challenges in Virtual Environments

In a conventional datacenter environment where application modules run on separate physical servers, TAPs and SPAN ports can be used to capture east-west traffic between application modules. However, when application modules run in virtual machines on the same physical server, the TAPs and SPAN ports have no visibility into the traffic between them (Figure 4-1).

Adding to the complexity, when demand increases, the hypervisor may automatically start up new instances of the software on the same host or on a different host. This happens too fast and too often for human administrators to observe the changes and reconfigure tools to monitor the new data flows.

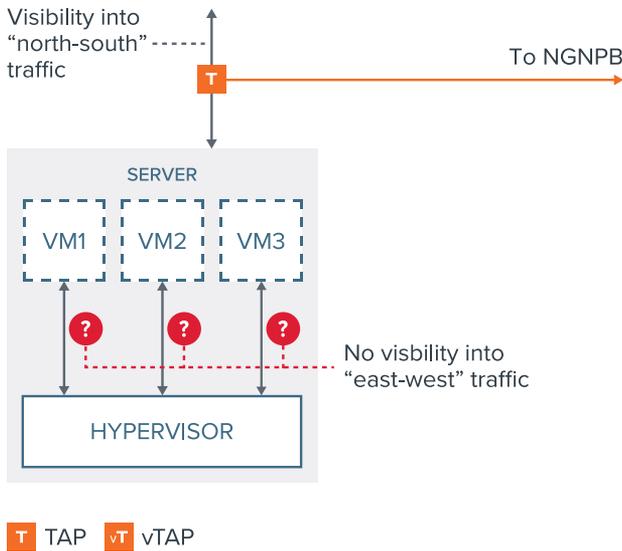


Figure 4-1: TAPs and SPAN ports have no visibility into east-west traffic between virtual machines in a virtual environment.

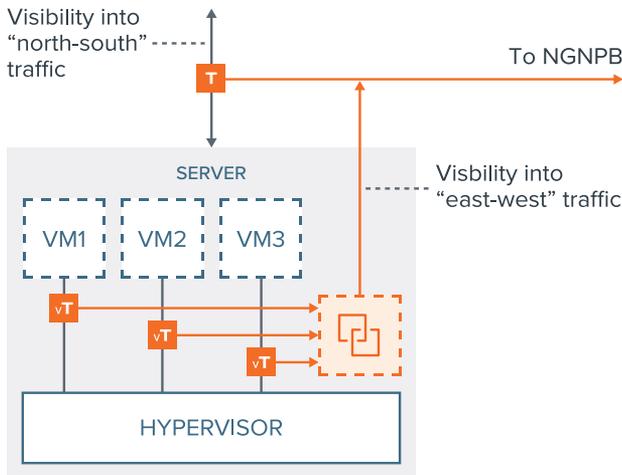
The tools available today to monitor activities in virtual environments often are not robust enough to meet the needs of multiple security, analytics, and performance monitoring tools. Also, adding new tool sets for virtual environments means yet more products to acquire, operate, and integrate with existing systems.

Monitoring Virtual Environments

Let’s look at how NGNPBs solve these problems.

vTAPs acquire east-west traffic

To obtain visibility into east-west traffic, NGNPBs often deploy lightweight software agents called vTAPs (virtual TAPs), to monitor network traffic into and out of individual virtual machines. As shown in Figure 4-2, the vTAPs send copies of network packets to a visibility node, which aggregates the traffic from the vTAPs, applies targeting policies created by an administrator, and forwards the packets and metadata to the central NGNPB platform.



T TAP **vT** vTAP

Figure 4-2: Visibility nodes use vTAPs to monitor traffic between VMs. Packets and metadata are aggregated and forwarded to the central NGNPB.

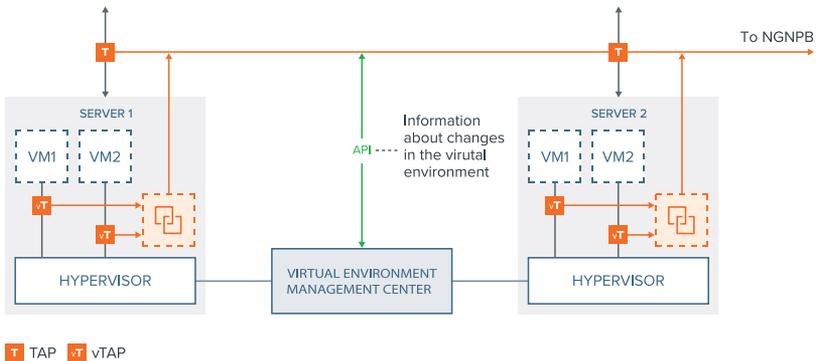
Platform vendors are starting to create their own virtual TAP services that can mirror traffic from each compute instance to a designated IP address – which can be the NGNPB.

NGNPBs integrate with the management center

NGNPBs can integrate with the management centers of virtual platforms, for example, vCenter in a VMware environment (Figure 4-3). This integration allows the NGNPBs to be notified when dynamic changes in the environment occur, so they can take appropriate actions. For example:

- ✓ When application instances are spun up on new virtual machines, an NGNPB can begin to monitor traffic to those VMs.
- ✓ When application instances are spun up on new hosts, an NGNPB can deploy visibility nodes on those hosts and start monitoring traffic there.

- ✓ If VMs are moved from one host to another, an NGNPB can disable monitoring on the old hosts and enable monitoring in the locations where the VMs have landed.



T TAP vT vTAP

Figure 4-3: The NGNPB is integrated with the management center of the virtual platform so it can react automatically to dynamic changes in the number and location of the VMs.



Administrators would find it extremely difficult to keep up these changes using manual methods. For example, if a VM were moved from one hypervisor to another, the VM administrator would have to disable the existing vSwitch port mirror sessions and create a new port mirror session on the destination hypervisor. An NGNPB automates these tasks, providing continuous visibility into the traffic to the VM, freeing up the administrator for other tasks.

One tool for physical and virtual environments

An NGNPB can be a single solution for acquiring, aggregating, and distributing network traffic and metadata for both physical and virtual environments. It provides comprehensive visibility across both environments, and performs the same shared services on all traffic, including decryption, metadata generation, and de-duplication.

Challenges in Cloud Environments

When enterprises migrate applications to AWS, Microsoft Azure, and other infrastructure-as-a-service (IaaS) platforms they gain agility and scalability and no longer have to install and manage hardware and system software. However, they have to cope with a lack of access to east-west network traffic, highly dynamic environments, and duplicate tool sets, just as they do when running applications in virtual environments.

In fact, the cloud situation is even more challenging. An enterprise has direct access to the virtual environments in its own datacenter. But public cloud platform providers put tight controls on access by clients, because they have to worry about privacy and security for many parties.

Figure 4-4 illustrates why it is difficult to monitor applications on cloud platforms. In this example, east-west traffic flows back and forth between instances of the application modules in a three-tier web application. The enterprise can't put TAPs or SPAN ports between the tiers.

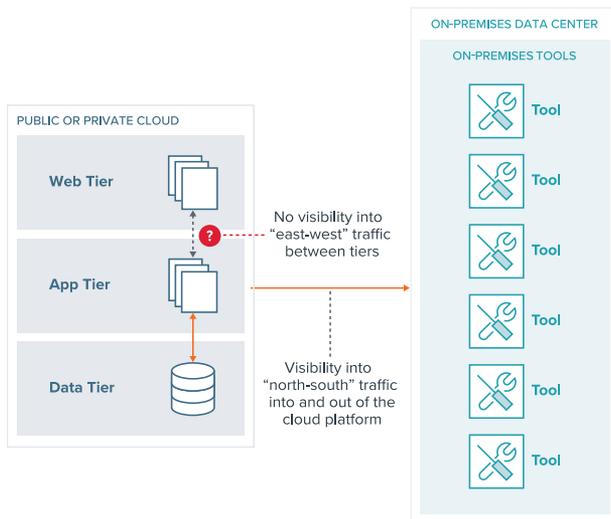


Figure 4-4: The enterprise can't put TAPs or SPAN ports in the cloud platform to monitor traffic between the application tiers.

If you want to run applications on more than one cloud platform you have to deal with even more tool sets, making end-to-end visibility that much harder to achieve.

Monitoring Applications on Cloud Platforms

Let's look at how NGNPBs provide pervasive visibility across cloud platforms.

Visibility modules and vTAPs in the cloud

A NGNPB deploys a visibility node in each cloud-based virtual private network, such as an Amazon VPC or an Azure Virtual Network (VNet) (Figure 4-5). The visibility node uses vTAP agents to monitor network traffic to each instance (VM) in that private network. The vTAPs send copies of the network packets to the visibility node.

Some providers of IaaS platforms on public clouds have started to introduce their own native, agentless vTAPs that can mirror traffic to the visibility node.

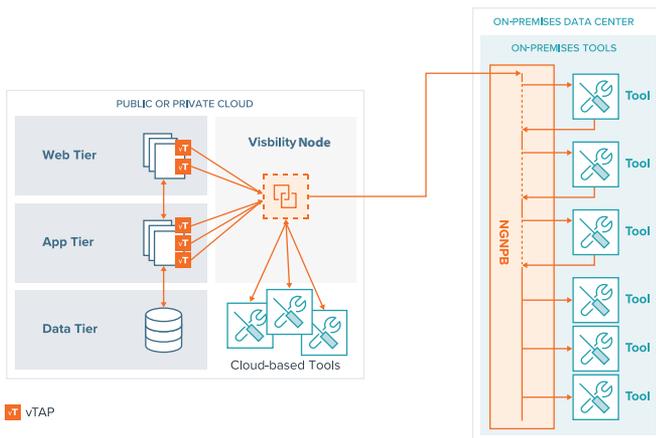


Figure 4-5: Visibility nodes use vTAPs to monitor traffic between application instances, which is then sent to tools on the cloud platform or in the datacenter.

The visibility node aggregates the traffic from the vTAPs, applies targeting rules created by an administrator, and forwards the packets and metadata, typically via a Generic Routing Encapsulation (GRE) or VXLAN tunnel. The traffic can be routed to:

- ✓ A “tool tier” of security, analytics, and performance monitoring software running on the cloud platform.
- ✓ The central NGNPB platform in the enterprise data-center, which in turn distributes the packets and metadata to the tools running there.

Figure 4-5 also illustrates a point about flexibility. The NGNPB gives enterprises options to:

- ✓ Use security and performance monitoring tools that run on the cloud platform if these provide unique capabilities, or...
- ✓ Use existing tools in the datacenter so they can have one tool set across all environments, or...
- ✓ Do some of each based on the pros and cons for each type of tool.

Orchestration

NGNPBs typically include a central orchestration and management module. This module is integrated via REST APIs with monitoring and management tools from the platform vendors, such as AWS CloudWatch and Azure Network Watcher.

The integration allows the NGNPB to be notified when changes take place in the cloud-based virtual private network. It can then take appropriate actions, such as scaling up to monitor new VMs, or deploying a visibility node when application instances are created on a new virtual private network.

The orchestration and management module can also be a central point of control to create and manage policies for collecting and filtering data from the workloads in the cloud (Figure 4-6).

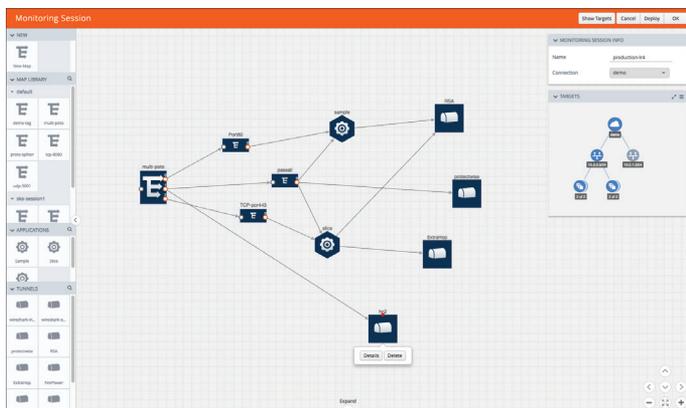


Figure 4-6: The NGNPB orchestration and management module can manage policies for monitoring network traffic on the cloud platform. (Source: Gigamon)

One tool for physical, virtual, and cloud environments

An NGNPB can acquire, aggregate, and distribute network traffic and metadata for physical, virtual, *and* cloud environments. This means that you can collect traffic from any point in your extended network infrastructure and route that traffic to any tool connected to the NGNPB. It doesn't matter whether the tool is on premises or in the cloud. You don't have to migrate or reconfigure security and performance tools when workloads are migrated.

Making Security a Technology Enabler Instead of a Blocker

In this chapter we covered a lot of technical detail about virtual environments and IaaS cloud platforms. The underlying point, though, is that a NGNPB can simplify the secure adoption of new technologies.

In this situation, NGNPBs mitigate or eliminate trade-offs between visibility and cost, and between performance and security.

With an NGNPB, an enterprise can move applications to virtual environments and cloud platforms without:

- ✓ Impairing visibility into network traffic and jeopardizing security, compliance, or performance monitoring
- ✓ Needing to acquire, learn, integrate, and manage new sets of network monitoring tools

NGNPBs can support CloudOps teams by:

- ✓ Helping the organization overcome concerns about visibility in the cloud
- ✓ Easing the migration of applications from private to public clouds
- ✓ Enabling consistent visibility across multi-platform environments

NGNPBs can help service providers gain visibility into subscriber activities. By reducing the cost of monitoring subscribers they can also increase average profit per user.

Chapter 5

Use Case: Out-of-Band Security and Performance Monitoring Tools

In this chapter

- Review types of out-of-band tools
 - Learn the value of comprehensive visibility and unsampled data
 - See how NGNPBs overcome the limitations of SPAN ports
-

The Proliferation of Out-of-Band Tools

Out-of-band tools inspect copies of the packets that travel on the network. Unlike inline tools, they do not change or block the live traffic.

Out-of-band tools can be grouped into three categories:

1. **Security** tools, such as IDS, DLP, and sandboxing products, which inspect traffic to identify IOCs, malware, and confidential information.
2. **Analytics** tools, such as SIEM, UEBA, forensic, and big data analytics tools, that examine network traffic and user behavior to detect patterns and anomalies.

3. **Performance monitoring tools**, such as network performance monitoring, application performance monitoring, and customer experience management tools, which measure and analyze service levels of the computing infrastructure from different perspectives.

The variety and number of these tools have been growing, because they help enterprises:

- ✓ Identify zero-day attacks and advanced persistent threats (APTs) that cannot be detected with traditional security controls
- ✓ Leverage technologies like artificial intelligence, machine learning, and big data analytics that detect patterns in vast quantities of data
- ✓ Utilize micro-segmentation and other techniques to isolate and monitor smaller zones within the computing infrastructure
- ✓ Increase employee and customer satisfaction by measuring and improving network and application availability and performance

Comprehensive Visibility for Analytics Tools

Analytics tools typically are used to uncover imminent and ongoing attacks by:

- ✓ Finding patterns in data and correlating IOCs and actions that are typically part of a campaign by attackers
- ✓ Determining baseline levels of activity and then identifying outliers and anomalous behaviors

But you can't find patterns when some of the clues are hidden by blind spots, or establish baselines and identify anomalies with incomplete data.

That is where next-generation network packet brokers come in. NGNPBs ensure that out-of-band analytics tools have visibility into *all* network traffic, including east-west traffic in

datacenters, traffic in virtual environments, and traffic to and from cloud applications and platforms.

End-to-End Visibility for Performance Monitoring Tools

NetOps teams rely on a variety of performance monitoring and management tools to maintain service levels acceptable to employees and customers. These tools are used to:

- ✓ Precisely measure the time required to send network traffic between locations and complete application transactions
- ✓ Troubleshoot problems by identifying bottlenecks in applications, servers, and network equipment
- ✓ Measure network and application availability
- ✓ Document compliance with, and violations of, service level agreements (SLAs)
- ✓ Generate data that can be used to optimize network and application performance
- ✓ Measure trends to establish growth needs and set budgets

Major categories of these tools include:

- ✓ Network performance management (NPM) tools, which gather network-related metrics like network response times broken down by port, IP address, protocol, and other parameters
- ✓ Application performance management (APM) tools, which gather application-related metrics like application response times, throughput, and error rates
- ✓ Customer experience monitoring tools, which focus on performance as perceived by customers and employees

Because these tools need to pinpoint issues that may occur at any point in a transaction, there isn't much point in measuring only part of the path. NGNPBs can acquire and aggregate traffic from physical, virtual, and cloud environments and

deliver it in a single data stream so performance monitoring tools obtain visibility into complete transactions, from original source to final destination.

NGNPs also acquire east-west traffic within datacenters, virtual environments, and cloud platforms, including traffic between application tiers and between zones in datacenters. This information enables performance monitoring tools to see application and network performance broken down into small segments, which helps them perform detailed analyses of normal and anomalous behaviors.

Unsampled Data

“Unsampled data” sounds like a bad thing. In fact, unsampled data is critical for analytics and performance monitoring tools.

Many organizations obtain copies of network traffic for out-of-band tools from SPAN ports on routers and switches. When network traffic volumes rise, these devices often save processing cycles and bandwidth through strategies such as:

- ✓ Sending only a small percentage of the traffic (a “sample”) through SPAN ports
- ✓ Sending only metadata or a summary of the traffic
- ✓ Automatically dropping packets that are malformed, are unusually long or short, or have bad checksums

These techniques can impair the usefulness of analytics tools, because they don’t get all the data they need to achieve accuracy. Also, performance monitoring tools can be misled when they don’t have access to all the packets that have traversed the network.

NGNPs address these challenges by sending all the packets that each tool needs, unsampled, unsummarized, and including malformed packets.

Packets and Metadata

Depending on the type of analysis required, out-of-band tools may inspect full network packets, or just metadata that summarizes a few key parameters. NGNPBs can forward network traffic and generate metadata at the same time, then use L2-L7 filtering to send one or the other, or some subset, to each tool. The tools get exactly what they need, and no more.

De-duplication and VLAN Tagging

NGNPBs provide shared services that increase the accuracy and effectiveness of performance monitoring tools.

Duplicate packets can be caused by inter-VLAN communication, misconfigured network switches, and packets that traverse paths with multiple TAP and SPAN ports. Duplicates artificially increase packet and byte counts, preventing performance monitoring tools from uncovering the real causes of bottlenecks. NGNPBs eliminate these distortions with de-duplication.

NGNPBs can also add port and VLAN tags to packets, showing the port where the packets entered the network, and sometimes the source location (for example, a remote office). This information can help performance monitoring tools trace problems back to their source.

Overcoming the Limitations of SPAN Ports

Many enterprises feed copies of traffic to out-of-band tools through SPAN ports, which have significant limitations.

More access points

Most network devices have only two SPAN ports with limited bandwidth. Network and security teams often compete to attach tools to those access points. Adding devices to increase the number of SPAN ports is expensive.

SPAN ports and packet replication consume resources on network devices and can degrade the performance of routers

and switches. Also, when network devices are overloaded with production network traffic, they often de-prioritize the SPAN port function and drop packets. These situations compromise the integrity and completeness of the replicated feed.

NGNPB appliances typically offer from a few dozen to more than 100 ports, ranging from 100 Mb to 100 Gb. Organizations can scale up and feed as many out-of-band tools as they want with just one or a few units.

Increased flexibility and manageability

A shortage of SPAN ports also reduces the flexibility and manageability of the network and security infrastructure. In most organizations it takes time to replace one tool with another, or to attach a tool to perform troubleshooting.

Because NGNPB appliances have so many ports and are easy to manage, administrators can quickly provide access to network traffic to tools for troubleshooting and testing. Attaching additional tools involves simple configuration changes on the NGNPB, with no impact on the production network.

Also, as we discussed in Chapter 2, you can test inline tools in out-of-band mode without slowing down the network, then toggle those tools to inline mode in seconds.

Monitoring at Full Line Speeds

When you use out-of-band tools for troubleshooting and problem diagnosis, they must keep up with the traffic flow so they can give you data about the current state of the network. NGNPBs make it easier to monitor at full line speed by:

- ✓ Using targeting and filtering so each performance monitoring tool receives only the packets or metadata that it needs
- ✓ Balancing network traffic so tools of the same type share the load when activity in one area spikes
- ✓ Freeing up capacity on the tools by offloading tasks such as SSL/TLS decryption, metadata generation, and de-duplication

Chapter 6

Into the Future

In this chapter

- Review how network packet brokers evolved by adding new network monitoring and security features
- Explore how NGNPBs help NetOps and SecOps groups collaborate
- Learn how NGNPBs can work with data lakes and empower threat hunters

How We Got Here: The Evolution of NGNPBs

First-generation network packet brokers (NPBs) were designed primarily to help network administrators and network operations staffs better monitor and manage enterprise networks. These NPBs collected network packet and flow data to help administrators monitor network performance and troubleshoot network problems.

Since the NPBs were already using TAPs and SPAN ports to monitor network traffic, the next step was to add other network-related functionality, including:

- ✓ Network traffic aggregation to collect network traffic throughout the enterprise and direct it to network tools, eliminating blind spots in network monitoring

- ✓ Application filtering that goes beyond L-2 to L-4 attributes such as source and destination IP addresses and ports, to provide flexible filtering based on L-5 to L-7 application-related criteria
- ✓ A set of network-related shared services, including SSL/TLS decryption, packet de-duplication, and header stripping to reduce network traffic and improve traffic performance and accuracy
- ✓ Packet slicing and masking to enforce compliance with privacy regulations and reduce the volume of irrelevant traffic being sent to tools

NPBs also provided centralized management of network monitoring so NetOps teams could visualize network activity and performance data across physical and virtual environments in datacenters and remote offices.

Adding security functionality

As network packet brokers continued to evolve into NGNPBs they also added security-related functionality. Some of the major functionality additions included:

- ✓ Inline bypass and load balancing to improve the resiliency and performance of security tools
- ✓ Centralized management and automation tools to define and distribute policies for managing the flow of traffic to security and security analytics tools
- ✓ Application identification capabilities to make tools application aware, so they can apply policies and inspect traffic on an application-by-application basis

These additions have made NGNPBs valuable tools for security operations (SecOps) staffs responsible for keeping security and analytics tools fully operational, and for security analysts and architects who want to increase the use of security tools while controlling (or lowering) costs.

Addressing virtual and cloud platforms

Many of the most recent enhancements to NGNPBs involve providing visibility into applications and services on virtual and cloud platforms. We reviewed these in Chapter 4.

ON THE WEB



For insights into the evolution of network packet brokers, see the ZK Research white paper, “[How to Strengthen Security While Optimizing Network Performance](#).”

Where We Are Going: Enabling New Initiatives

Next-generation network packet brokers support best practices and new technologies that are strengthening network management and cybersecurity. Let’s look at how NGNPBs are enhancing SIEMs, helping NetOps and SecOps groups work together, preserving historical data, and supporting threat hunting.

SIEMs: Going Beyond Logs

SIEMs originally focused on correlating and analyzing log data from servers and security tools like firewalls, IPSs, and anti-malware products. However, they were often handicapped by the limitations of this data, difficulties in correlating data, lack of context, and lack of visibility into events in virtual and cloud environments. Many installations face performance issues related to the explosion in network traffic and server and tool logs.

Today, NGNPBs are helping make SIEMs more effective tools for security operations center (SOC) and incident response (IR) teams. Contributions include:

- ✓ Feeding network metadata to SIEMs that they can correlate with log data to develop deeper insights into threat behaviors based on context
- ✓ Providing pervasive visibility so SIEMs have access to encrypted traffic and traffic from throughout the enterprise, including virtual and public cloud environments

- ✓ Filtering and de-duplicating network traffic so SIEMs are not overwhelmed by data they can't use
- ✓ Masking sensitive information so SOC and IR team members do not violate privacy and security regulations



NGNPBs can also enrich the metadata being sent to SIEMs. For example, they can append HTTP and HTTPS return codes and DNS query and response information to NetFlow and IPFIX records. This contextual information helps analysts spot problems like potential command and control communications with external websites and rogue DNS services on the network.

Harmonizing NetOps and SecOps

Network operations (NetOps) and security operations (SecOps) teams share the goal of providing secure, fast, reliable networks. However, the teams often have conflicting priorities: optimizing network and application performance and employee productivity on the one hand and maximizing security on the other.

As a result of historical factors and institutional rivalries, the two teams often use different tool sets to perform the same tasks. Not only does this lead to duplication and extra costs, it often results in conflicting views on what is happening and what can be done to solve problems.

An NGNPB can help harmonize the interests and activities of NetOps and SecOps teams by:

- ✓ Giving both NetOps and SecOps a complete view of network traffic from a single source, with data that can be used for network monitoring and optimization *and* for threat detection and incident response
- ✓ Speeding up the evaluation and deployment of new tools without requiring lengthy change control processes

- ✓ Allowing NetOps and SecOps to set policies in their own realms without affecting the other team; for example, letting NetOps manage the flow of data to network and application performance monitoring tools, while SecOps manages data going to firewalls, IPSs, SIEMs and security analytics tools
- ✓ Filtering traffic to security devices and offloading tasks like decryption and metadata generation, so NetOps teams don't have to worry about security tools slowing down network performance when traffic spikes
- ✓ Protecting network performance by avoiding planned and unplanned outages caused by security tools



Gartner analysts Sanjit Ganguli and Lawrence Orans discuss the conflicts between NetOps and SecOps and how to reduce the tension in their research note: [Align NetOps and SecOps Tool Objectives with Shared Use Cases](#).

Data Persistence and Data Lakes

IT organizations are investing heavily in analytics and forensics tools so they can:

- ✓ Analyze and reverse-engineer APTs and other complex attacks
- ✓ Find all systems affected by attacks so they can be remediated
- ✓ Establish baselines of normal behavior so they can identify abnormal and anomalous behaviors that indicate malicious activities
- ✓ Chart trends that can be used to strengthen security controls and plan for network growth

To produce accurate results, these tools need historical data – lots of historical data – including detailed packet capture data as well as network metadata.

However, until recently data and metadata generated by NGNPs and network devices were used in real time or near-real time and then dropped, or summarized, or stored on

systems inaccessible to other tools. No data persistence mechanisms were available that could handle cost-effectively the volume and variety of data and metadata that was generated.

New storage technologies are making it practical to store and use much greater volumes of historical data and share it with a wide variety of analytics tools for many types of analysis.

Data lakes

Among the best examples of flexible, cost-effective persistent storage methods are data lakes. Data lakes are storage repositories that hold vast amounts of structured and unstructured data in native (raw) formats. Each data element has a unique ID and is tagged with multiple types of metadata.

The advantages of data lakes include:

- ✓ Rapid deployment and flexibility, because there is no need to define schemas or clean up or summarize information before storing data
- ✓ Accessibility of data to a wide variety of analytics, data mining, data visualization, and modeling tools, including those that utilize artificial intelligence, machine learning, and pattern recognition technologies
- ✓ Fidelity and data provenance, because original versions of the data are always retained and linked to transformed versions, so analysts can analyze the original data in new ways and auditors can use it to demonstrate compliance
- ✓ Centralization of storage across tool stacks

NGNPs are a good source of data for data lakes because they:

- ✓ Acquire and aggregate network traffic from across the enterprise
- ✓ Generate metadata in a consistent format for all network traffic
- ✓ Work with a wide variety of network traffic formats, so the data can be retained in its original form

Increasingly, IT organizations will be coupling NGNPBs with data lakes and other new storage methods to better support analytic and forensic activities.

Empowering Threat Hunters

Most IT security activities are either preventative or reactive. They aim to block threats at the network perimeter or respond to incidents and remediate breaches after they are detected.

Recently, some enterprises have added a new type of security activity: threat hunting. Threat hunters study the modus operandi or “tradecraft” of threat actors, including their tools, techniques, and procedures (TTPs). They use this knowledge to search for clues that attackers have been active on the network, employing techniques that do not depend on having known signatures or IOCs. These techniques allow threat hunters to detect APTs and other sophisticated attacks that use previously unknown methods or that have managed to evade existing controls.



The types of attacker activities and tools that threat hunters look for include in-memory malware, persistence techniques such as storing shellcode within registry keys, use of common administrative tools like Windows Management Instrumentation (WMI) and Windows Sysinternals PsExec to perform reconnaissance and malicious tasks on the network, and use of tools like KERBEROS to steal user credentials (“KERBEROASTING”).

Many threat hunting methods involve analyzing network metadata to find evidence of command and control communication into and out of the network and lateral movement within the network.

NGNPBs are an excellent source of network data and metadata that threat hunters require to answer questions like:

- ☑ Is anyone using protocols that enable remote authentication, such as SSH, SMB, and RDP?
- ☑ Do any persistent objects have a history of initiating network connections to remote sites?

- ☑ What is the distribution of certificate authorities (CAs) associated with persistent objects, and do any of those CAs have weak or poor reputations?
- ☑ What remote sites have a history of failed logon attempts?

Chapter 7

Selecting the Right NGNPB

In this chapter

- Explore criteria you can use to evaluate NGNPBs
 - Understand the importance of three types of comprehensiveness
 - Learn how to select an NGNPB supplier that will be a good long-term partner for your strategic investment
-

A Strategic Investment

So far in this guide we have talked about next-generation packet brokers in general terms, as if they were all the same. But not all products have equivalent capabilities. Because NGNPBs have the potential to become a central part of your network and security infrastructure, you want to find one that meets your needs today and will continue to do so in the future. In this chapter we highlight criteria you should use to compare NGNPBs and select the one that best fits your organization.

Comprehensiveness of the Solution

Throughout this guide we have stressed the importance of comprehensive visibility and the advantages of working with all types of data sources and security and network tools.

Physical, virtual, and cloud environments

To do their jobs, security tools need to be able to detect all the IOCs in the enterprise; analytics tools need to be able to baseline all the traffic and metadata; and performance monitoring tools need to track transactions from end to end. To meet these requirements, an NGNPB must be able to acquire and aggregate both north-south and east-west traffic in virtual environments and on public and private cloud platforms.

The NGNPB's management tools should provide "single pane of glass" visibility into activities in physical, virtual, and cloud environments and allow policies to be set consistently across them. This means applying the same policies on all platforms *and* being able to share traffic data across them so on-premises tools can analyze traffic from the cloud and vice-versa.

Inline and out-of-band tools

To simplify management and minimize duplication, the same NGNPB should be able to handle the needs of both inline tools such as firewalls, IPSs, ATPs, WAFs, and anti-malware products, and of out-of-band tools such as IDS, DLP, SIEM, security analytics, and network and application performance monitoring products.

The ability to toggle tools back and forth between out-of-band and inline modes can help balance performance and security for latency-sensitive applications and makes it easier to test and deploy new inline tools.

Traffic and metadata formats

Most enterprise environments include old communications equipment, specialized protocols (e.g., SIP, CDP, LLDP), and cutting-edge communications technologies (e.g. TLS 1.3 and IPv6). A NGNPB should be able to handle the complete gamut, so it can acquire traffic from all these sources, generate relevant metadata, and output the metadata in formats that can be used by a wide range of security, analytics, and performance monitoring tools.

Scalability and Capacity for Ultra-high-speed Networks

NGNPBs need to be able to scale up to support large numbers of inline and out-of-band tools with very high volumes of traffic.

Even if you are not using 40GB or ultra-high-speed 100GB networks today, you probably will be in the future. If your NGNPB can't handle traffic at these speeds you will be faced with unpleasant trade-offs between falling behind or having to sample traffic to keep up, both of which can undermine security and performance monitoring.

An NGNPB should offer equipment with high-capacity ports and a track record of handling high volumes of traffic in real-world production environments.

Integration with Cloud Platforms

In Chapter 4 we reviewed the challenges of acquiring network traffic and metadata in virtual environments and in public and private clouds. To address these challenges, an NGNPB should be integrated with popular virtualization and cloud platforms such as VMware NSX, Cisco ACI, Amazon AWS, Microsoft Azure, and OpenStack so that it can:

- ✓ Deploy visibility nodes on cloud-based virtual private networks like Amazon VPCs and Azure VNets
- ✓ Automate responses to dynamic changes in workloads by exchanging information with cloud platform monitoring and management tools like AWS CloudWatch and Azure Network Watcher
- ✓ Share network traffic and metadata with security, analytics, and performance monitoring tools running on the cloud platforms (as well as tools in enterprise datacenters)

The NGNPB vendor must perform this integration for each cloud platform. Check whether the NGNPB you are considering has the functionality you need on the cloud platform(s) you are using.

Also, if you are using more than one IaaS platform, see if the NGNPB can give you a single console for monitoring and managing policies consistently across all of them as well as your on-premises infrastructure.

Vendor Focus and Track Record

Because you are investing in a strategic infrastructure technology, you need to select a product that will continue to be enhanced and evolved to meet your needs in the future. These days that is not a given.

Attributes you might look for include:

- ✓ A history that demonstrates a long-term commitment to network packet brokers
- ✓ A track record addressing (or even better anticipating) customer needs and the evolution of the product category
- ✓ A road map that shows how the vendor will address the next big opportunities and challenges you will face in security and network operations
- ✓ Investments and acquisitions that demonstrate commitment to the road map and to technological innovation
- ✓ An active, growing ecosystem of technology partners
- ✓ A large, committed customer base, including enthusiastic reference customers

Closing Thoughts

There is no good one-sentence definition of next-generation network packet brokers. It takes a while to grasp what they are and everything they can do. Indeed, most people are surprised by the wide range of benefits they offer for security and network management.

We hope this guide has given you a good foundation in the features and uses of NGNPBs. The next step is to talk to other organizations like yours and see how they have used NGNPBs. Or if you are a hands-on type, try one in your own environment.

Appendix: Key Features of NGNPs

Application filtering winnows network traffic based on L-2 to L-7 characteristics so tools don't waste resources on packets they can't process and low risk traffic (such as streaming video from commercial sources), and so they don't inspect sensitive information in violation of privacy and security regulations. See pages 15 and 26.

De-duplication eliminates duplicate network packets to reduce network traffic and prevent tools from wasting resources by inspecting the same packets multiple times. See page 31.

Header stripping removes specified headers, tags, and encapsulation from packets so tools do not waste resources trying to interpret them, or drop packets with unreadable header types. See page 31.

Inline bypass sends network traffic around inline security tools when they fail, to prevent application outages. Inline bypass can also refer to a range of features that allow network traffic to be redirected to increase network and application availability. See Chapter 2, especially pages 15-16.

Load balancing distributes network traffic across two or more tools of the same type to smooth out network performance, reduce costs by making better use of available capacity, retain the use of existing devices when adding new ones, enhance redundancy and availability, and allow devices to be upgraded or replaced without causing outages. See pages 12-14.

Masking helps organizations comply with regulatory requirements by overwriting passwords, account numbers, Social Security numbers, and other specified fields with meaningless characters. See page 32.

Metadata generation provides tools with information about network traffic and packets, such as summaries of traffic flows and the sources, destinations, protocols and characteristics of packets, to help with incident response, threat hunting, forensics, performance monitoring, and historical and trend analysis. See pages 27-30.

Monitoring “east-west” traffic in cloud platforms uses virtual TAPs (vTAPs) to provide tools with visibility into network traffic flowing between application and service instances on public and private cloud platforms. See pages 38-40.

Monitoring “east-west” traffic in virtual environments uses virtual TAPs (vTAPs) to provide tools with visibility into network traffic flowing between application modules on virtual machines in virtual environments. See pages 34-37.

Packet slicing removes the payloads from network packets to reduce network traffic to tools that only analyze headers, cut storage requirements, and comply with regulations that prohibit the inspection and storage of sensitive information. See page 31-32.

SSL/TLS decryption creates a “decryption zone” so multiple tools can inspect SSL/TLS traffic without having to use their own resources to decrypt and re-encrypt the packets. See pages 21-26.

Toggling inline tools between out-of-band and inline modes allows organizations to change priority from minimizing latency to tightening security, and to deploy tools quickly as soon as testing and “learning” processes have been completed. See pages 17-18.

VLAN tagging, also called “source port labeling,” adds to packets information on where they entered the network to help performance monitoring and troubleshooting. See page 47.

The Market Leader in Network Visibility Solutions

We're changing the game on how you see, control and contextualize traffic across your physical, virtual and cloud networks. Reduce risk, complexity and costs to meet your business and innovation needs now, not later – with powerful visibility and insights into your organization's network activity.

WHAT WE DO

Gigamon delivers powerful visibility and insight into your organization's network activity, enabling you to reduce risk, complexity and costs.



Simplify your current security architecture



Realize the true ROI of your security tools



Gain real visibility into critical security threats

To learn more visit
www.gigamon.com



More than you ever thought possible: security... network availability...cost savings...visibility into virtual environments and cloud platforms!

Next-generation network packet brokers (NGNPBs) have been transformed over a few short years. Today, they ensure that IT security and performance monitoring tools have comprehensive visibility into network traffic across the enterprise and operate at peak efficiency and availability. They help security and network operations teams see into “blind spots,” detect attacks, and protect the performance of networks and applications.

- **Next-generation network packet brokers** — learn their four key capabilities
- **Inline bypass** — see how NGNPBs increase the efficiency and availability of security and performance monitoring tools
- **Shared services** — explore how NGNPBs take over decryption, metadata generation, and other tasks for hundreds of tools
- **Visibility into blind spots** — examine how NGNPBs provide visibility into traffic inside virtual environments and cloud platforms
- **Inline and out-of-band tools** — understand the ways NGNPBs support all types of security and performance monitoring tools

Find out how network visibility helps incident responders, NetOps and SecOps teams, data analysts and threat hunters do the important work of digital transformation and network security.

About the Author

Jon Friedman has over 20 years experience in industry analysis and marketing roles at software and IT services companies. He has described cutting-edge technologies and their business benefits for more than 40 high-tech companies. Jon has a BA from Yale and an MBA from Harvard.



CYBEREDGE
P R E S S

Not for resale

ISBN 978-1-948939-10-2



9 781948 939102 >