

SteelCentral™ Network Performance Management Deployment Guide

NetProfiler
NetShark
Flow Gateway
Packet Analyzer
NetExpress

May 2016

riverbed®

© 2016 Riverbed Technology, Inc. All rights reserved.

Riverbed and any Riverbed product or service name or logo used herein are trademarks of Riverbed. All other trademarks used herein belong to their respective owners. The trademarks and logos displayed herein cannot be used without the prior written consent of Riverbed or their respective owners.

F5, the F5 logo, iControl, iRules and BIG-IP are registered trademarks or trademarks of F5 Networks, Inc. in the U.S. and certain other countries. Linux is a trademark of Linus Torvalds in the United States and in other countries. VMware, ESX, ESXi are trademarks or registered trademarks of VMware, Incorporated in the United States and in other countries.

Portions of SteelCentral™ products contain copyrighted information of third parties. Title thereto is retained, and all rights therein are reserved, by the respective copyright owner. PostgreSQL is (1) Copyright © 1996-2009 The PostgreSQL Development Group, and (2) Copyright © 1994-1996 the Regents of the University of California; PHP is Copyright © 1999-2009 The PHP Group; gnuplot is Copyright © 1986-1993, 1998, 2004 Thomas Williams, Colin Kelley; ChartDirector is Copyright © 2007 Advanced Software Engineering; Net-SNMP is (1) Copyright © 1989, 1991, 1992 Carnegie Mellon University, Derivative Work 1996, 1998-2000 Copyright © 1996, 1998-2000 The Regents of The University of California, (2) Copyright © 2001-2003 Network Associates Technology, Inc., (3) Copyright © 2001-2003 Cambridge Broadband Ltd., (4) Copyright © 2003 Sun Microsystems, Inc., (5) Copyright © 2003-2008 Sparta, Inc. and (6) Copyright © 2004 Cisco, Inc. and Information Network Center of Beijing University of Posts and Telecommunications, (7) Copyright © Fabasoft R&D Software; Apache is Copyright © 1999-2005 by The Apache Software Foundation; Tom Sawyer Layout is Copyright © 1992 - 2007 Tom Sawyer Software; Click is (1) Copyright © 1999-2007 Massachusetts Institute of Technology, (2) Copyright © 2000-2007 Riverbed Technology, Inc., (3) Copyright © 2001-2007 International Computer Science Institute, and (4) Copyright © 2004-2007 Regents of the University of California; OpenSSL is (1) Copyright © 1998-2005 The OpenSSL Project and (2) Copyright © 1995-1998 Eric Young (eay@cryptsoft.com); Netdisco is (1) Copyright © 2003, 2004 Max Baker and (2) Copyright © 2002, 2003 The Regents of The University of California; SNMP::Info is (1) Copyright © 2003-2008 Max Baker and (2) Copyright © 2002, 2003 The Regents of The University of California; mm is (1) Copyright © 1999-2006 Ralf S. Engelschall and (2) Copyright © 1999-2006 The OSSP Project; ares is Copyright © 1998 Massachusetts Institute of Technology; libpq++ is (1) Copyright © 1996-2004 The PostgreSQL Global Development Group, and (2) Copyright © 1994 the Regents of the University of California; Yahoo is Copyright © 2006 Yahoo! Inc.; pd4ml is Copyright © 2004-2008 zefer.org; Rapid7 is Copyright © 2001-2008 Rapid7 LLC; CmdTool2 is Copyright © 2008 Intel Corporation; QLogic is Copyright © 2003-2006 QLogic Corporation; Tarari is Copyright © 2008 LSI Corporation; Crypt_CHAP is Copyright © 2002-2003, Michael Bretterklieber; Auth_SASL is Copyright © 2002-2003 Richard Heyes; Net_SMTP is Copyright © 1997-2003 The PHP Group; XML_RPC is (1) Copyright © 1999-2001 Edd Dumbill, (2) Copyright © 2001-2006 The PHP Group; Crypt_HMAC is Copyright © 1997-2005 The PHP Group; Net_Socket is Copyright © 1997-2003 The PHP Group; PEAR::Mail is Copyright © 1997-2003 The PHP Group; libradius is Copyright © 1998 Juniper Networks. This software is based in part on the work of the Independent JPEG Group the work of the FreeType team.

This documentation is furnished "AS IS" and is subject to change without notice and should not be construed as a commitment by Riverbed Technology. This documentation may not be copied, modified or distributed without the express authorization of Riverbed Technology and may be used only in connection with Riverbed products and services. Use, duplication, reproduction, release, modification, disclosure or transfer of this documentation is restricted in accordance with the Federal Acquisition Regulations as applied to civilian agencies and the Defense Federal Acquisition Regulation Supplement as applied to military agencies. This documentation qualifies as "commercial computer software documentation" and any use by the government shall be governed solely by these terms. All other use is prohibited. Riverbed Technology assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.



Riverbed Technology
680 Folsom Street
San Francisco, CA 94107

Phone: 415-247-8800
Fax: 415-247-8801
Web: <http://www.riverbed.com>

Part Number
712-00113-08

Contents

Preface	7
About This Guide	7
Audience	8
Document Conventions	8
Documentation and Release Notes	9
Contacting Riverbed	9
What Is New	9
Chapter 1 - SteelCentral NPM Overview	11
What Is SteelCentral NPM?	11
Overview of the SteelCentral NPM Products	12
NetProfiler and NetProfiler-v Overview	12
Flow Gateway and Flow Gateway-v Overview	13
NetShark and NetShark-v Overview	13
Packet Analyzer Overview	14
Chapter 2 - SteelCentral NPM Deployment Scenarios	15
Choosing the Right Equipment	15
Choosing a NetProfiler Model	18
Choosing a Flow Gateway Model	21
Choosing a NetShark Model	22
Choosing a NetShark Module on AppResponse	23
Choosing NetShark-v on SteelHead EX	23
Choosing Packet Analyzer	24
Deployment Scenarios	24
Choosing How to Deploy the NetShark and Packet Analyzer	24
Deploying the NetShark-v on SteelHead EX and Packet Analyzer	26
Deploying the NetExpress	27
Deploying the Standard NetProfiler and Flow Gateway	29
Deploying the Enterprise NetProfiler and Flow Gateway	31
Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer	32
Deploying the NetProfiler, Flow Gateway, and NetShark on AppResponse	34
Deploying the NetProfiler, Flow Gateway, NetShark, NetShark-v, and Packet Analyzer	35
Deploying the NetProfiler, Flow Gateway, NetShark, and NetShark-v on the SteelHead EX	36
Port and Protocol Dependencies	37
NetShark and Packet Analyzer Port Dependencies	38
NetProfiler and Flow Gateway Port Dependencies	39
SteelCentral Appliance Full-Solution Port Dependencies	40

SteelCentral Appliance Enterprise Solution Port Dependencies.....	41
NetProfiler and NetExpress Flow Storage	41
Types of Flow Storage.....	41
Flow Rate Estimation	42
Flow Storage Size Estimation.....	43
Flow Redundancy with SteelCentral	44
Chapter 3 - Flow Collection for SteelCentral	49
Base Requirements.....	49
Flow Data Fields Consumed by NetProfiler.....	52
Flow Type Considerations.....	53
Flow Collection Considerations.....	54
Flow Collection in Virtual Environments	54
Validating Flow Collection.....	54
Sample Third-Party Configurations.....	56
Configuring VMware ESXi 5.5 Using vSphere	56
Configuring Cisco 6500 Series Switches Running Native Cisco IOS CLI	57
Configuring Cisco 6500 Series Switches in Hybrid Mode.....	58
Configuring Cisco 7500 Series Router	59
Configuring Cisco 7600 Series Router	59
Configuring Cisco 3560 and 3750 Flexible NetFlow.....	59
Configuring the Cisco Nexus 7000 Flexible NetFlow	60
Configuring NetFlow Export for Cisco Nexus 1000V	61
Configuring IPFIX for Avaya (Nortel) 8300 and 8600.....	62
Configuring sFlow for HP Procurve 3500, 5400, and 6200	62
Chapter 4 - Packet Collection for SteelCentral	65
SteelCentral for Packet Collection	65
Port Mirroring and SPAN	66
Port Mirroring	66
Remote SPAN and Encapsulated Remote SPAN	67
Sample Port Mirror Configurations	69
VMware ESXi Distributed vSwitch Port Mirroring Versus Promiscuous Mode	75
Network TAP Instrumentation.....	76
VACL Configuration Examples.....	78
VACL Port Mirroring Configuration on Cisco 6500 Running CatOS.....	78
VACL Port Mirroring Configuration on Cisco Catalyst 6500 Running Cisco IOS Software.....	78
NetShark Passthru	79
Packet Deduplication	80
Snaplen	80

Chapter 5 - SteelCentral and SteelHead Integration	81
SteelHead and SteelCentral Overview	81
NetFlow Versus SteelFlow Net.....	82
SNMP Interface Persistence (ifindex)	83
SteelCentral Appliance Deployment Considerations.....	84
Enabling SNMP Polling.....	84
In-Path Deployments	84
Virtual In-Path Deployments.....	85
Server-Side Out-Of-Path Deployments.....	87
Configuring SteelHead for Flow Data Export	87
NetProfiler and SteelHead Integration.....	89
Configuring Riverbed QoS Integration.....	90
SteelHead Flow Integration	93
Chapter 6 - NetProfiler and SteelCentral Portal Integration	95
Chapter 7 - Additional SteelCentral Integration.....	97
SNMP Integration.....	97
SNMP Integration for Flow Sources	97
SNMP Integration for Device Management of SteelCentral Components	98
SNMP Integration for Sending Traps	98
SNMP for Switch Port Discovery	99
Switch Port Discovery Supported Routers and Switches	99
Active Directory	99
Integration for Active Directory 2008.....	101
Integration for Active Directory 2003.....	101
REST API.....	101
Chapter 8 - NetProfiler Analytics and Service Monitoring.....	103
Analytic License Limits per NetProfiler Platform	103
Understanding the Analytics License Limits.....	104
Reducing the Number of Metrics	105
Conditions Required for Baseline Establishment	107
Determining Which Metrics to Use.....	107
Chapter 9 - Troubleshooting the NetProfiler.....	111
RTT Values Not Available.....	111
Not Receiving Reports by Email.....	112
DNS Names Not Being Resolved in Reports.....	112
Reports Are Not DNS-Resolving All Addresses.....	113

Data in Reports Seems Inconsistent	113
Sensor Protocol Violations.....	114
Communication Issues.....	114
Switch Port Discovery Troubleshooting.....	115
Appendix A - Licensing.....	117
Licensing Overview.....	117
Runtime Licenses.....	117
Capacity Licenses	118
Option Licenses.....	118
License Installation	118
Current-Generation Hardware License Installation.....	118
Packet Analyzer Licensing	119
Assigning Licenses	119
Manual License Installation	120
Automatic License Upgrades.....	120
Evaluation Licenses	120
Licenses Available	120
Licensing the NetExpress and NetExpress-v	120
Licensing the NetProfiler and NetProfiler-v	121
Licensing the Enterprise NetProfiler Cluster	122
Licensing Flow Gateway and Flow Gateway-v	122
Licensing the NetShark.....	122

Preface

Welcome to the *SteelCentral Network Performance Management Deployment Guide*. Read this preface for an overview of the information provided in this guide, the documentation conventions used throughout, and contact information. This preface includes the following sections:

- [“About This Guide” on page 7](#)
- [“Documentation and Release Notes” on page 9](#)
- [“Contacting Riverbed” on page 9](#)
- [“What Is New” on page 9](#)

About This Guide

The *SteelCentral Network Performance Management Deployment Guide* describes best practices for configuring and deploying the network performance management (NPM) subset of the SteelCentral products, and it includes deployment information about SteelHeads, SteelHead Interceptors, SteelCentral AppResponse, and third-party appliances. The guide includes information about flow collection, packet collection, metrics, analysis, troubleshooting, and licensing.

This guide includes information relevant to the following products:

- Riverbed SteelCentral NetProfiler (NetProfiler)
- Riverbed SteelCentral NetProfiler (Enterprise NetProfiler)
- Riverbed SteelCentral NetProfiler (virtual edition) (NetProfiler-v)
- Riverbed SteelCentral NetExpress (NetExpress)
- Riverbed SteelCentral NetExpress (virtual edition) (NetExpress-v)
- Riverbed SteelCentral Flow Gateway (Flow Gateway)
- Riverbed SteelCentral Flow Gateway (virtual edition) (Flow Gateway-v)
- Riverbed SteelCentral NetShark (NetShark), including Embedded SteelCentral NetShark, which is NetShark functionality embedded in a Riverbed SteelHead.
- Riverbed SteelCentral NetShark (virtual edition) (NetShark-v)
- Riverbed SteelCentral Packet Analyzer (Packet Analyzer)
- Riverbed SteelHead (SteelHead)

- Riverbed SteelHead EX (SteelHead EX)
- Riverbed SteelHead Interceptor (Interceptor)
- Virtual Services Platform (VSP)
- Riverbed SteelCentral AppResponse (AppResponse)
- Riverbed SteelCentral Portal (Portal)

Audience

This guide is written for network administrators, operators, and engineers familiar with WANs, LANs, and the data center environment.

You must also be familiar with the user documentation posted on the Riverbed Support site for the products mentioned in [“About This Guide” on page 7](#).

Document Conventions

This guide uses the following standard set of typographical conventions.

Convention	Meaning
<i>italics</i>	Within text, new terms and emphasized words appear in <i>italic</i> typeface.
boldface	Within text, CLI commands, CLI parameters, and REST API properties appear in bold typeface.
Courier	Code examples appear in Courier font: <pre>amnesiac > enable amnesiac # configure terminal</pre>
< >	Values that you specify appear in angle brackets: interface <ip-address>
[]	Optional keywords or variables appear in brackets: ntp peer <ip-address> [version <number>]
{ }	Elements that are part of a required choice appear in braces: {<interface-name> ascii <string> hex <string>}
	The pipe symbol separates alternative, mutually exclusive elements of a choice. The pipe symbol is used in conjunction with braces or brackets; the braces or brackets group the choices and identify them as required or optional: {delete <filename> upload <filename>}

Documentation and Release Notes

To obtain the most current version of all Riverbed documentation, go to the Riverbed Support site at <https://support.riverbed.com>.

If you need more information, see the Riverbed Knowledge Base for any known issues, how-to documents, system requirements, and common error messages. You can browse titles or search for keywords and strings. To access the Riverbed Knowledge Base, log in to the Riverbed Support site at <https://support.riverbed.com>.

Each software release includes release notes. The release notes identify new features in the software as well as known and fixed problems. To obtain the most current version of the release notes, go to the Software and Documentation section of the Riverbed Support site at <https://support.riverbed.com>.

Examine the release notes before you begin the installation and configuration process.

Contacting Riverbed

This section describes how to contact departments within Riverbed.

- **Technical support** - If you have problems installing, using, or replacing Riverbed products, contact Riverbed Support or your channel partner who provides support. To contact Riverbed Support, open a trouble ticket by calling 1-888-RVBD-TAC (1-888-782-3822) in the United States and Canada or +1 415-247-7381 outside the United States. You can also go to <https://support.riverbed.com>.
- **Professional services** - Riverbed has a staff of professionals who can help you with installation, provisioning, network redesign, project management, custom designs, consolidation project design, and custom coded solutions. To contact Riverbed Professional Services, email proserve@riverbed.com or go to <http://www.riverbed.com/services-training/Services-Training.html>.
- **Documentation** - The Riverbed Technical Publications team continually strives to improve the quality and usability of Riverbed documentation. Riverbed appreciates any suggestions you might have about its online documentation or printed materials. Send documentation comments to techpubs@riverbed.com.

What Is New

Since the last release of the *SteelCentral Network Performance Management Deployment Guide* (March 2015), the following information has been updated:

- SteelCentral RPM Dashboards has been updated to SteelCentral Portal
- General refresh of [Chapter 2, “SteelCentral NPM Deployment Scenarios”](#) on page 15
- **Updated** - [“Base Requirements”](#) on page 49
- **Updated** - [“Cisco Nexus 1000V Virtual Switch SPAN”](#) on page 69
- **Updated** - [“Switch Port Discovery Supported Routers and Switches”](#) on page 99
- **Updated** - [“REST API”](#) on page 101

CHAPTER 1 SteelCentral NPM Overview

This chapter provides an overview of the SteelCentral NPM products discussed in this guide. It includes the following sections:

- [“What Is SteelCentral NPM?” on page 11](#)
- [“Overview of the SteelCentral NPM Products” on page 12](#)

What Is SteelCentral NPM?

SteelCentral NPM is an enterprise-wide network performance management solution that provides visibility into your data centers, offices, and users. SteelCentral NPM uses network flow data and packet-based performance metrics to provide in-depth visibility into network operations, discover applications, and monitor performance. SteelCentral NPM uses advanced behavioral analytics to track performance over time, and alerts you to any deviations from normal behavior, and enables you to identify and resolve problems before there is an impact on end users.

This deployment guide covers specific appliances of the SteelCentral NPM product line. For more details, see [“About This Guide” on page 7](#) and [“Overview of the SteelCentral NPM Products” on page 12](#).

When you deploy SteelCentral NPM in your network infrastructure, you gain the following advantages:

- Behavior analytics for proactive network and system response monitoring
- Dependency mapping for an always-accurate view of your applications and their dependencies
- Executive-level dashboards for a quick summary of service performance
- Network-level and application-level dashboards for quick identification and analysis of network issues
- Cost-effective visibility into remote sites by leveraging existing SteelHead deployments
- Ability to plan for and understand optimized WANs
- Ability for true end-to-end coverage of the enterprise with dashboard-to-flow-to-packet analysis, providing scalability and flexibility

SteelCentral NPM provides a full set of application-centric, site-centric, and business-centric views so that you can discover your critical services, optimize performance, and continuously monitor service levels. SteelCentral NPM provides a consistent view of the network by breaking down the silos between network, infrastructure, and security teams. In addition, built-in integration with the SteelHead WAN optimization products provides full visibility and control over application delivery.

For more details about SteelHead integration, see [“SteelCentral and SteelHead Integration”](#) on page 81.

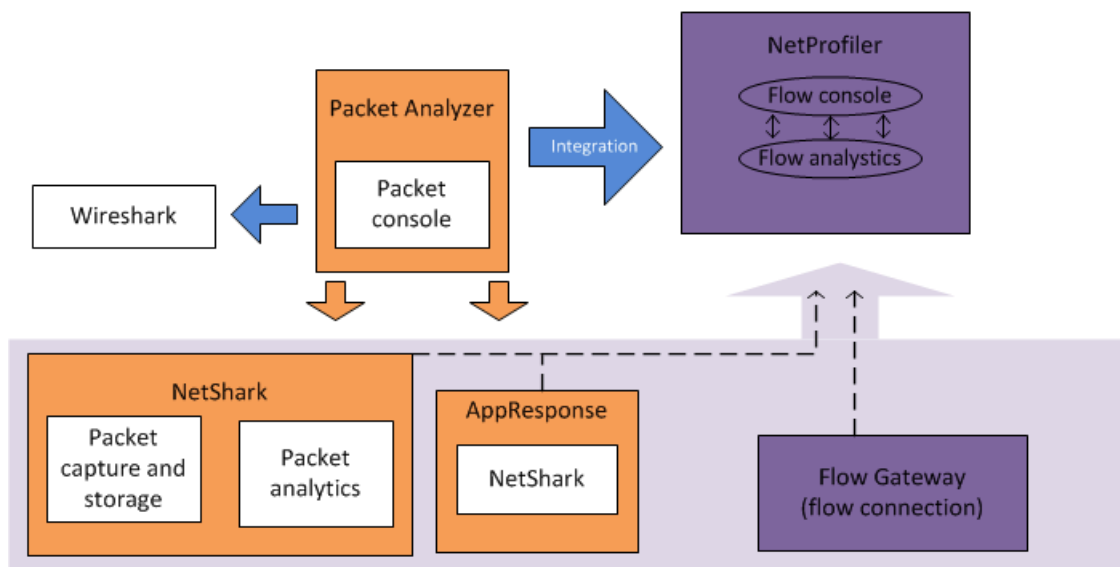
For more information about NetProfiler and the SteelCentral Portal, see [“NetProfiler and SteelCentral Portal Integration”](#) on page 95.

Overview of the SteelCentral NPM Products

This section describes the following SteelCentral products:

- [“NetProfiler and NetProfiler-v Overview”](#) on page 12
- [“Flow Gateway and Flow Gateway-v Overview”](#) on page 13
- [“NetShark and NetShark-v Overview”](#) on page 13
- [“Packet Analyzer Overview”](#) on page 14

Figure 1-1. SteelCentral Appliance Architecture



NetProfiler and NetProfiler-v Overview

The NetProfiler and NetProfiler-v provide, on a consistent-user interface, centralized reporting and analysis of the data collected by other SteelCentral products, SteelHeads, and flow exporting routers, switches, and other network devices. The NetProfiler offers performance analytics, security analytics, and proactive alerts for delivering application-aware monitoring and troubleshooting of your network. It combines all network data into a single data set with in-depth views that support flexible analysis and high-speed presentation of the information.

Members of the NetProfiler family of products include:

- **Standard NetProfiler** - Standard model designed for mid-size organizations supporting up to approximately 1,000,000 flows per minute (fpm).
- **Enterprise NetProfiler** - Designed to be expandable, supporting environments larger than the Standard NetProfiler up to 10,000,000 fpm.

- **NetProfiler-v** - Designed to allow easy deployment as part of a virtualized environment, supporting up to 2,000,000 fpm. You can deploy NetProfiler-v on VMware ESXi 5.0, 5.1, and 5.5.
- **NetExpress 470** - An entry-level appliance designed for small organizations. It includes NetProfiler, NetShark, and Flow Gateway functionality in one appliance and supports up to 120,000 fpm.

Note: For information about NetExpress 360 and 460, see the previous version of this deployment guide.

- **NetExpress-v** - An entry-level virtual appliance designed for small organizations. It includes NetProfiler, NetShark, and Flow Gateway functionality in one virtual appliance and supports up to 120,000 fpm. You can deploy NetExpress-v on VMware ESXi 5.0, 5.1 and 5.5.

For more information about the NetProfiler, see [“Choosing a NetProfiler Model” on page 18](#).

Flow Gateway and Flow Gateway-v Overview

The Flow Gateway and Flow Gateway-v collect flow data from routers, switches, and other network devices. These appliances support most standard flow types (NetFlow, sFlow, J-Flow, IPFIX, and so on). The Flow Gateway aggregates the data, deduplicates it, encrypts it, and sends it to the NetProfiler in a compressed format over a reliable and secure connection. The Flow Gateway can transmit data to up to five Standard NetProfilers or NetExpresses and supports up to 2,000,000 fpm.

You can deploy the Flow Gateway in the same location as the NetProfiler or regionally if you have multiple data centers. You can deploy the Flow Gateway-v on VMware ESXi 5.5.

For more information about the Flow Gateway, see [“Choosing a Flow Gateway Model” on page 21](#).

NetShark and NetShark-v Overview

The NetShark includes high-performance (multiple 1 GbE or 10 GbE interfaces) continuous packet capture, storage, and analysis. You can:

- use the NetShark for fast indexing and in-depth analysis of multiterabyte network traffic recordings.
- drill down to deliver micro-level flow resolution for analysis.

The NetShark:

- sends flow information, including performance metrics, to the NetProfiler and standard NetFlow v9 to other flow collectors.
- delivers real-time or historical deep-packet inspection (DPI) and analysis.

You can access the NetShark using Packet Analyzer. The NetShark uses the REST API for transferring data to Packet Analyzer.

NetShark-v is available in version 9.5 and later. NetShark-v operates similarly to the NetShark, but it is intended for use in virtual environments in which you want packet capture and continuous monitoring between virtual hosts or need to deploy packet capture in an environment in which deploying a full-NetShark does not make sense.

You can deploy NetShark-v on VMware ESXi v5.0, v5.1, and v5.5 and Microsoft Hyper-V Server 2012 R2.

This section contains the following topics:

- [“NetShark-v on SteelHead EX” on page 14](#)
- [“NetShark on AppResponse” on page 14](#)
- [“Embedded SteelCentral NetShark Overview” on page 14](#)

For more information about the NetShark, see [“Choosing a NetShark Model” on page 22](#).

NetShark-v on SteelHead EX

In RiOS 8.5 or later, SteelHead EX supports NetShark-v 10.5 using VSP. Deploying NetShark-v in VSP provides most of the functionality available from a full NetShark-v deployment.

For more information, see [“Choosing NetShark-v on SteelHead EX” on page 23](#).

NetShark on AppResponse

AppResponse 8.6.8 or later contains an embedded a NetShark-v engine (based on NetShark-v 10.0 code). This deployment of NetShark-v code provides most of the functionality available in the full NetShark and other NetShark-v deployments, except that it cannot perform Layer-7 DPI.

You must manually install the NetShark module with AppResponse 8.6.8. In AppResponse 9.0 or later, the NetShark module is included with the base software release, however depending on the functionality you want, you might need additional licenses.

For more information, see [“Choosing a NetShark Module on AppResponse” on page 23](#).

Embedded SteelCentral NetShark Overview

In RiOS 7.0 or later, the SteelHead includes limited NetShark functionality as Embedded SteelCentral NetShark. Embedded SteelCentral NetShark software enables on-demand packet capture on SteelHeads at remote sites, and it provides control and analysis of packet captures on remote SteelHeads directly from Packet Analyzer. As with the NetShark, you can use Embedded SteelCentral NetShark to drill down to deliver microlevel flow resolution for analysis using Riverbed XML-based protocol on top of an HTTPS connection for transferring data to Packet Analyzer. You do not need to transfer full packets until you need them.

Packet Analyzer Overview

Packet Analyzer seamlessly and securely integrates with a remote NetShark to deliver a complete and feature-rich distributed network analysis. Packet Analyzer is the only tool on the market to be fully integrated with Wireshark software, an open-source network protocol analyzer. While the NetProfiler provides visibility on all flows across the network, Packet Analyzer provides an in-depth view into problems requiring deep packet analysis.

For more information about Packet Analyzer, see [“Choosing Packet Analyzer” on page 24](#).

CHAPTER 2 SteelCentral NPM Deployment Scenarios

Deployment of SteelCentral NPM products requires advanced planning to ensure that you install the appliances to capture critical traffic. This chapter describes best practices that enable you to deploy your appliances efficiently, without wasting resources on unnecessary coverage. This chapter includes the following deployment sections:

- [“Choosing the Right Equipment” on page 15](#)
- [“Deployment Scenarios” on page 24](#)
- [“Port and Protocol Dependencies” on page 37](#)
- [“NetProfiler and NetExpress Flow Storage” on page 41](#)
- [“Flow Redundancy with SteelCentral” on page 44](#)

Choosing the Right Equipment

This section describes the equipment choices available to you. It includes the following sections:

- [“Choosing a NetProfiler Model” on page 18](#)
- [“Choosing a Flow Gateway Model” on page 21](#)
- [“Choosing a NetShark Model” on page 22](#)
- [“Choosing a NetShark Module on AppResponse” on page 23](#)
- [“Choosing NetShark-v on SteelHead EX” on page 23](#)
- [“Choosing Packet Analyzer” on page 24](#)

When determining what kind of equipment you need at each site—whether that site is a data center, a branch office, or a large building on a corporate campus—answer the following questions:

- What kind of information do I want to know about this location? Do I need response-time information, optimization information, WAN link bandwidth information, and application usage and performance information?
- Do I have an extensive virtualized environment already in place?
- How many users and how much traffic am I expecting at this location, now and several years into the future?

- What kind of physical resources do I have at this location? Are there technicians that can help maintain the hardware?
- What kind of network resources do I have at this location? Can my switch provide Switched Port Analyzer (SPAN) or mirror traffic? Do I have a packet broker that I can use to aggregate and process packets. Can my switches and routers provide flow information?
- Do I have sufficient bandwidth to transfer flow data between this location and the NetProfiler?
- How much visibility do I need at this location?
- Do I need packet-level visibility to view object calls and individual transactions within the application?

The following table shows additional NetProfiler solutions for several reporting attributes you might want to capture.

Environment Types	NetProfiler Solution
Small-size environments needing a single appliance solution with packet capture.	NetExpress 470
Medium-size to large-size environments.	Standard NetProfiler
Large-size to enterprise-size environments.	Enterprise NetProfiler
Virtualized environments with limited flow requirements.	NetExpress-v 470
Virtualized environment in medium-size to large-size environments with large capacity VMware hosts.	NetProfiler-v
Calculate basic NetFlow v9 data from packet data.	NetShark or NetShark-v

The following table shows additional SteelCentral solutions for several reporting attributes you might want to capture.

Tasks	SteelCentral Appliance Solutions
Accurately calculate response-time information for nonoptimized flows.	NetShark or NetExpress
Accurately calculate response-time information for optimized flows.	SteelHead with either a NetShark or NetExpress providing packet-level visibility on the server side
Report on and monitor link bandwidth information: for example, to monitor percent utilization.	Flow Gateway, NetExpress, and one of either Enterprise NetProfiler or Standard NetProfiler
Obtain detailed packet information: for example, to analyze network traffic in case of a security violation.	Physical or virtual NetShark or NetExpress 470
Gain visibility into virtualized environments.	NetShark-v or NetExpress-v 470

In choosing the right equipment, you want to make sure that the data you receive is the data you need. The following table describes some of the different flow formats supported by SteelCentral and specifies the features available within these formats.

Flow Format	Flow Data	SteelHead	NetShark
Source of data for the NetProfiler	x	x	x
Source and destination IP number, IP protocol ingress interface, IP type of service, number of bytes and packets, start and end times of flow, and TCP flags	x	x	x
Connection throughput for nonoptimized connections			x
Monitor traffic from a SPAN or TAP			x
Performance metrics		x	x
Throughput of 1 Gbps and 10 Gbps			x
Deep-packet inspection (DPI)	x	x	x
Layer-7 application fingerprinting	x	x	x
VoIP metrics	x	x	x
Web transaction timing (object load times)			x
Full and continuous packet capture, TB storage			x
Remote management module		x	x

In most environments different sites have varying numbers of users and volume of network traffic. A site with 10 users transferring large files all day generates many fewer flows than a site with 200 users making extensive use of web-based applications. For calculation purposes, Riverbed recommends that you use 40 to 60 fpm as the estimated average flows per minute per IP endpoint. Exact flows per minute depend on the traffic characteristics in your environment and can vary beyond the estimates.

Use multiplication to estimate the maximum number of flows per minute. For example, 100 users that each generate 50 fpm produce an approximate flow rate of 5,000 fpm. However, if the site has servers that are accessed from remote locations, the overall flow rate is likely to increase, potentially by a large amount. If you have used flow tools in the past, you might already have some flow estimates. You can also look at session counts on SteelHeads, firewalls, or load balancers to assist in obtaining flow estimates.

You must have the appropriate number of technical staff on site. In a small remote office of only a few non-technical people, deploying a virtual version of an appliance might make more sense than installing one or more physical appliances.

Consider other network equipment that is already installed at a site when you decide which SteelCentral product to install. If an office site contains multiple large switches capable of generating NetFlow or SPAN and port mirror traffic, a NetShark or Flow Gateway might make sense. Conversely, if a small office contains only a wireless internet router with no other network infrastructure, you have limited options for deploying visibility solutions; it is possible you might not need a SteelCentral solution in this location.

If you have a site that reports significant quantities of data across a WAN, consider the bandwidth used by SteelCentral for the transfers. Typical WAN bandwidth use is approximately 1.5 percent of monitored traffic. SteelCentral products report flows once per minute. If reporting multiple megabytes of traffic per minute seriously impedes the performance of WAN links, you might need a different solution: for example, restricting the amount of data being monitored.

Choosing a NetProfiler Model

The NetProfiler is licensed on a flow-per-minute basis, with the number of flows counted after all flows have been deduplicated. You want to choose the right NetProfiler model so that you do not receive inaccurate results as a result of receiving excessive numbers of flows. Consider the following factors when you are deciding which type and model of NetProfiler to install:

- The size of the current network you want to monitor
- The planned expansion of coverage expected during the lifetime of the product

NetExpress

There are two models of the NetExpress: NetExpress 470 and NetExpress-v 470. The functionality of the appliance and virtual editions is the same. Each version shares the same licensed capacities and maximum packet storage capabilities.

After deduplication, NetExpress 470 has flow rates ranging from 15,000 to 120,000 fpm. The NetExpress is best suited for smaller organizations or a small subsection of a larger network. Because the NetExpress can forward the flows it receives and processes directly to a different model of NetProfiler, this deployment can make sense for sites in which there is a need for local visibility and broader enterprise-wide visibility.

The NetExpress includes functionality similar to that of the NetProfiler, NetShark, and Flow Gateway. You can install a more compact deployment by combining the functionality into a single physical or virtual appliance.

The following table shows the NetExpress model options. Most features are available in the base unit. The NetExpress is 1U high, and you can upgrade in the field to the next-higher flow-rate version. If you want to analyze traffic within a software-defined network, you can add a software-defined network license.

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
NetExpress 470 F1	Up to 15,000 fpm		
NetExpress 470 F2	Up to 30,000 fpm	Primary 10/100/1000 for management	4 x 1 Gbps Copper
NetExpress 470 F3	Up to 60,000 fpm	Provides up to 4Tb of packet storage	4 x 1 Gbps Fiber (SFP)
NetExpress 470 F4	Up to 90,000 fpm		2 x 10 Gbps Fiber (SFP)
NetExpress 470 F5	Up to 120,000 fpm	Provides up the 3Tb of flow storage	
NetExpress-v 470 F1	Up to 15,000 fpm		
NetExpress-v 470 F2	Up to 30,000 fpm		
NetExpress-v 470 F3	Up to 60,000 fpm	—	—
NetExpress-v 470 F4	Up to 90,000 fpm		
NetExpress-v 470 F5	Up to 120,000 fpm		

The Standard NetProfiler

The Standard NetProfiler is available as both a physical and virtual appliance. The primary difference between the two models is the amount of deduplicated fpm and available flow record storage that is available. The physical appliance limit is 10 Tb, and the for virtual appliance this limit is 50 Tb. The NetProfiler-v provides a broader range of flow limits.

The Standard NetProfiler has flow limits ranging from 200,000 to 1,000,000 fpm, and the NetProfiler-v has flow limits ranging from 15,000 to 2,000,000 fpm. The physical models are suited for medium-size organizations with between 4,000 and 20,000 hosts, and the virtual models with between 3,000 and 40,000 hosts assuming an average of 50 fpm per host.

The Standard NetProfiler and the NetProfiler-v cannot forward flows to other NetProfilers, nor can the Standard NetProfiler and the NetProfiler-v receive flows directly from flow sources (with the exception of NetShark). Because the Flow Gateways and NetSharks can forward flows to five distinct NetProfilers, you can use the Standard NetProfiler to monitor a small subset of a larger network. You can send the flows from the NetSharks and Flow Gateways to the local Standard NetProfiler monitoring a network subset and up to four additional NetProfiler or NetExpress systems.

The following table shows the Standard NetProfiler model options. Most features are available in the base unit. Each physical appliance is 2U high, and you can upgrade in the field to the next-higher flow-rate version through a software license. If you want to analyze traffic within a software-defined network, you can add a software-defined network license.

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
NetProfiler-F1	Up to 200,000 fpm	Primary 10/100/1000 for management	Storage area network (SAN) card (two fiber host bus adapter [HBA] ports)
NetProfiler-F2	Up to 400,000 fpm		
NetProfiler-F3	Up to 600,000 fpm		
NetProfiler-F4	Up to 800,000 fpm		
NetProfiler-F5	Up to 1,000,000 fpm		
NetProfiler-v-F1	Up to 15,000 fpm	—	—
NetProfiler-v-F2	Up to 30,000 fpm		
NetProfiler-v-F3	Up to 60,000 fpm		
NetProfiler-v-F4	Up to 90,000 fpm		
NetProfiler-v-F5	Up to 200,000 fpm		
NetProfiler-v-F6	Up to 400,000 fpm		
NetProfiler-v-F7	Up to 600,000 fpm		
NetProfiler-v-F8	Up to 800,000 fpm		
NetProfiler-v-F9	Up to 1,000,000 fpm		
NetProfiler-v-F10	Up to 2,000,000 fpm		

Enterprise NetProfiler

The Enterprise NetProfiler has a minimum flow limit of 1,000,000 fpm; you can increase the flow limit with Expansion modules up to a maximum of 10,000,000 fpm. Each module provides support for an additional 1,000,000 fpm (until you reach the maximum of 10,000,000 fpm supported). In terms of hosts, an Enterprise NetProfiler can support at least 20,000 hosts, assuming an average of 50 fpm per host. Each additional Expansion module adds support for another 20,000 hosts using the same assumptions.

The following table shows the Enterprise NetProfiler model options. Most features are available in the base unit. The base Enterprise NetProfiler is composed of two 1U units and one 2U unit. Each Expansion module is an additional 2U unit. If you want to analyze traffic within a software-defined network, you can add a software-defined network license.

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
Enterprise NetProfiler-UI (required)	Part of base system	Primary 10/100/1000 for management Out-of-band management	—
Enterprise NetProfiler-DB (required)	Part of base system	Primary 10/100/1000 for management Out-of-band management	—
Enterprise NetProfiler-AN (required)	Base system, up to 1,000 K fpm	Primary 10/100/1000 for management Out-of-band management	SAN card (two fiber HBA ports)
Enterprise NetProfiler-EX (optional; add 0-9)	Expansion unit, each one adding 1,000 K FPM	Primary 10/100/1000 for management Out-of-band management	SAN card (two fiber HBA ports)
Enterprise NetProfiler-DP (required only when you have five or more Expansion modules)	N/A: used to balance flows on larger systems	A Primary 10/100/1000 interface for management and a secondary 10/100/1000 interface that can be bonded to the primary for high performance environments. Out-of-band management	—

Note: If you are using a storage area network (SAN) with an Enterprise NetProfiler cluster, you must have the SAN option for each Analyzer (AN) and Expansion (EX) modules. You cannot mix SAN and non-SAN storage in a single Enterprise NetProfiler cluster.

To plan for future expansion, you must know the current estimated number of flows per minute and the expected flows per minute in the timeframe being planned for. For example, if the network currently has 6,000 hosts and is expected to grow to 14,000 hosts in the next two years, a Standard NetProfiler is sufficient to handle the growth. A software license and hardware upgrade enable a Standard NetProfiler licensed for 400,000 fpm to be upgraded to 800,000 fpm.

Another example is a network currently providing service to 14,000 hosts and expected to grow to 25,000 in the next year. In this situation, an Enterprise NetProfiler is a better choice. The Standard NetProfiler is sufficient in a 14,000-host network, but it is unlikely to provide adequate performance and capacity to manage the network when it grows to 25,000 hosts.

Take the same considerations when you plan for virtual deployments and physical appliances. While the NetProfiler-v might have higher overall flow capacities, it still has upper limits. When you deploy virtual solutions, make sure you do not overrun both the maximum licensed capacities and the maximum available hardware on the virtual platform.

To provide visibility into a subset of the network and a full overview of the entire network, you can install a smaller capacity NetProfiler (such as NetProfiler-v or a NetExpress, or a Standard NetProfiler sufficient for local traffic flow) in combination with a larger capacity NetProfiler (such as a Standard or Enterprise NetProfiler sufficient for overall traffic flow). This combination ensures that the system meets both the immediate and planned growth needs.

While all the NetProfiler and NetExpress devices including the standard, enterprise cluster, and virtual appliances can be supplied with dual network interfaces, those interfaces are not designed for high availability (HA) or other forms of redundancy. You might be able to add a second IP address to the interface; however this is not a supported function and there is no guarantee that the address will remain through a system reboot or upgrade. Changes from version to version can change the functionality of the secondary interface. The only exception to this change is the Dispatcher module on the Enterprise Cluster where the two interfaces may be bonded to provide improved performance on large capacity systems.

Choosing a Flow Gateway Model

The Flow Gateway is available as both a physical and virtual model. The primary difference between the two options is ease of deployment. A single physical Flow Gateway is sufficient to manage a large number of devices. Both Flow Gateway and Flow Gateway-v provide up to 2,000,000 deduplicated flows per minute.

The upgrade from the smallest to the largest Flow Gateway within a class family (virtual or physical) requires only a license change.

The following table shows the available Flow Gateway models.

Flow Gateway Model	Deduplicate Flow Rate
Flow Gateway-F1	Up to 200,000 fpm
Flow Gateway-F2	Up to 400,000 fpm
Flow Gateway-F3	Up to 600,000 fpm
Flow Gateway-F4	Up to 800,000 fpm
Flow Gateway-F5	Up to 1,000,000 fpm
Flow Gateway-F6	Up to 2,000,000 fpm
Flow Gateway-v-F1	Up to 15,000 fpm
Flow Gateway-v-F2	Up to 30,000 fpm
Flow Gateway-v-F3	Up to 60,000 fpm
Flow Gateway-v-F4	Up to 90,000 fpm
Flow Gateway-v-F5	Up to 200,000 fpm
Flow Gateway-v-F6	Up to 400,000 fpm
Flow Gateway-v-F7	Up to 600,000 fpm
Flow Gateway-v-F8	Up to 800,000 fpm
Flow Gateway-v-F9	Up to 1,000,000 fpm
Flow Gateway-v-F10	Up to 2,000,000 fpm

Consider the following requirements when you deploy the Flow Gateway:

- The flows sent to the Flow Gateway are within the licensed limits.
- The geographic coverage is appropriate.
- The flow capacity of the NetProfiler is not exceeded by the total deduplicated capacity of the devices sending data.

- There are sufficient VMware resources available, if you choose to deploy Flow Gateway-v.
- The amount of deduplication in flows being sent from different sources.

For a medium-size organization with multiple disparate geographic locations, a single 2,000,000 fpm Flow Gateway might be able to manage the overall load; however, this configuration might not make sense if a significant amount of your corporate WAN bandwidth is consumed with the transmission of flow data: for example, from remote sites to a centrally located Flow Gateway. Because flow data is usually transmitted through UDP, an unreliable protocol, the likelihood of packet loss is increased the further a packet travels. In this sort of environment, Riverbed recommends that you deploy multiple smaller Flow Gateways or Flow Gateway-vs at major locations.

All the Flow Gateways (both physical and virtual) are supplied with dual network interfaces that are not designed for high availability (HA) or other forms of redundancy. You can assign an address to the secondary (AUX) interface. You can only use the AUX interface for receiving flows from flow sources. You cannot manage the device through this interface.

Choosing a NetShark Model

The NetShark is available as both a physical and virtual model. The primary difference between the two is packet storage capacity and packet processing speed (including write to disk speed). You must choose the appropriate NetShark platform to ensure that you can store the desired quantity of packets and that they are available for analysis. Capture cards are separately ordered items and you must make sure the cards ordered meet the needs of your network.

Deploy the NetShark with one or more capture cards. With the NetShark 6170 series you can have one or more 48 TB or 72 TB storage modules. If you have multiple modules, they must be identical capacities.

Note: In NetShark 10.9 and later, support for more than 72 TB of packet storage is achieved by daisy-chaining multiple storage units to a single NetShark 6170.

If you store every packet traversing a heavily loaded 10-Gbps network, the 72 TB of space on a NetShark 6170 with a 72-TB storage unit allows approximately 23 hours of storage (assuming a throughput of 7 Gbps). The 576 TB of space allows approximately 184 hours of storage (assuming a throughput of 7 Gbps).

If you store every packet traversing a heavily loaded 1-Gbps network, the 32 TB of space on a NetShark 4170 allows approximately 163 hours of storage (assuming a throughput of 1 Gbps).

The following table shows the NetShark models.

NetShark Base	Form Factor	Storage	Maximum Capture Cards	FPM Exported
NetShark 2170	1U	8 TB	1 (4 x 1 Gbps Cu, 4 x 1 Gbps Fiber, 2 x 10 Gbps Fiber)	400,000
NetShark 4170	2U	32 TB	2 (4 x 1 Gbps Cu, 4 x 1Gbps Fiber, 2 x 10Gbps Fiber, 4 x 10Gbps Fiber)	800,000
NetShark 6170	2U	48 to 576 TB	2 (4 x 1Gbps Cu, 4 x1 Gbps Fiber, 2 x 10 Gbps Fiber, 4 x 10 Gbps Fiber)	1,200,000

Note: Only the licensed maximum flow rate is guaranteed.

The following table shows the NetShark-v models.

Model	Storage	Capture Interfaces	Export to the NetProfiler
VSK-00050	50 GB	4	50,000 fpm
VSK-00200	1 TB	4	50,000 fpm
VSK-00400	2 TB	4	50,000 fpm

Note: Only the licensed maximum flow rate is guaranteed.

For more information about installing and configuring NetShark-v, see the appropriate documentation on the Riverbed Support site.

Choosing a NetShark Module on AppResponse

If you have deployed an AppResponse in your network, you can deploy a NetShark-based module directly on that appliance. This module provides most of the NetShark appliance analysis functionality to packets that are detected by the AppResponse. The NetShark module on AppResponse can build and forward traffic flows to a NetProfiler or NetExpress with no additional licensing requirements. For more advanced analysis and access to packet data (for example, using Packet Analyzer), you must purchase a NetShark license separately from any existing AppResponse licenses. With a full NetShark license you can access the entire packet store on the AppResponse through Packet Analyzer and perform most Packet Analyzer functions against those packets. Currently, access using the REST API is not supported on the NetShark module running on AppResponse.

The NetShark module running on AppResponse through version 9.0.3 is based on NetShark 10.0.6 code and is missing any of the newer functionality available with more recent versions of NetShark.

Choosing NetShark-v on SteelHead EX

In RiOS 8.5 or later, you can deploy NetShark-v directly on the SteelHead EX platform on VSP. The NetShark-v has most of the functionality of a physical NetShark except that it does not perform Layer-7 DPI. The NetShark performs Layer-7 DPI using the same engine as the SteelHead EX and therefore performing Layer-7 DPI analysis with the NetShark-v on the SteelHead EX is redundant. To make sure you are receiving Layer-7 DPI information, export SteelFlow Net from the SteelHead EX in addition to deploying the NetShark.

Riverbed recommends that you deploy the NetShark-v on the SteelHead EX in branch environments in which you want additional visibility but deploying additional hardware is challenging. For NetShark-v on the SteelHead EX to receive packets, you must configure the packets to flow through either the primary or auxiliary interface on the SteelHead EX.

Choosing Packet Analyzer

When you deploy Packet Analyzer, you must consider only how many users must perform deep-packet analysis. Packet Analyzer is Windows client software and can be licensed as follows:

- **Per installed machine** - Requires a license for each system on which you install Packet Analyzer. That license is permanently associated with that system and can only be used by the Packet Analyzer installed on that system. You must purchase a different license for each system on which Packet Analyzer is installed.
- **Concurrent licensing** - Provides a pool of licenses from which a Packet Analyzer instance can draw a license. For concurrent licensing to work properly you must have a license server—a NetProfiler of any flavor (NetExpress/Standard/Enterprise Cluster, physical appliance or virtual) or NetShark of any flavor (physical appliance or virtual) and network access between the Packet Analyzer installation and license server. A license is checked out and is associated with a particular Packet Analyzer installation for 24 hours or until it is released by the client software. you cannot release a license from the server.

Packet Analyzer can analyze traffic from either a virtual or physical NetShark, an Embedded SteelCentral NetShark probe, the NetShark module on AppResponse, or any standard packet capture files. You do not need Packet Analyzer for NetProfiler-level analysis and troubleshooting.

Deployment Scenarios

This section describes the following deployment scenarios:

- [“Choosing How to Deploy the NetShark and Packet Analyzer” on page 24](#)
- [“Deploying the NetShark-v on SteelHead EX and Packet Analyzer” on page 26](#)
- [“Deploying the NetExpress” on page 27](#)
- [“Deploying the Standard NetProfiler and Flow Gateway” on page 29](#)
- [“Deploying the Enterprise NetProfiler and Flow Gateway” on page 31](#)
- [“Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer” on page 32](#)
- [“Deploying the NetProfiler, Flow Gateway, and NetShark on AppResponse” on page 34](#)
- [“Deploying the NetProfiler, Flow Gateway, NetShark, NetShark-v, and Packet Analyzer” on page 35](#)
- [“Deploying the NetProfiler, Flow Gateway, NetShark, and NetShark-v on the SteelHead EX” on page 36](#)

Choosing How to Deploy the NetShark and Packet Analyzer

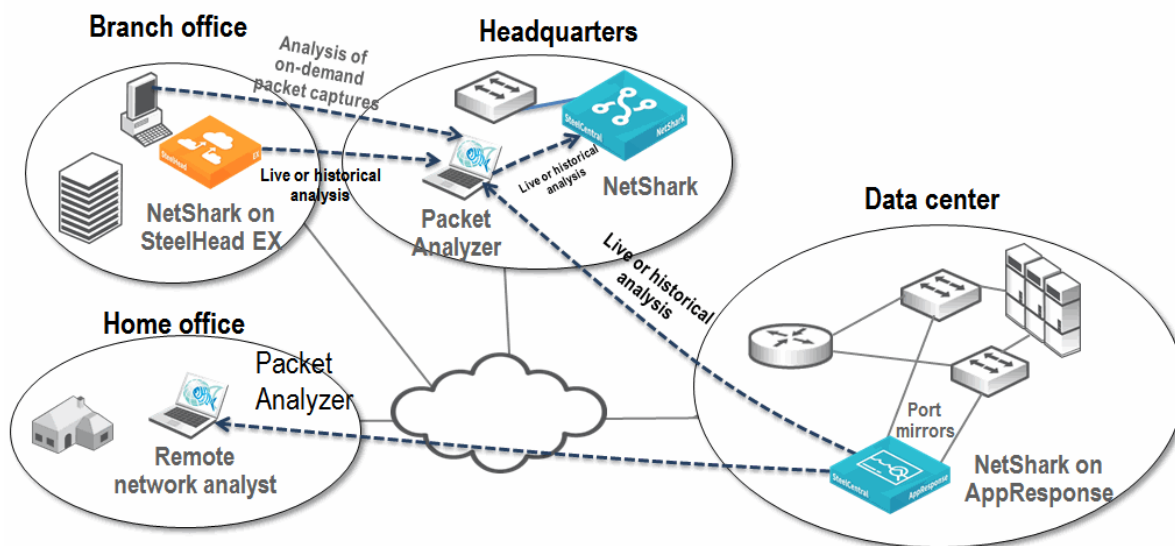
You can deploy the NetShark and Packet Analyzer in the following scenarios:

- [“Deploying NetShark and Packet Analyzer” on page 26](#)
- [“Deploying the NetShark on AppResponse and Packet Analyzer” on page 26](#)

No matter which deployment scenario you use, the following information applies. You deploy the NetShark and Packet Analyzer if you want extremely detailed views of network traffic and to improve troubleshooting of your network-related issues. Because the NetShark can only receive packets, you must install the NetShark in a location close to the source of the majority of the packets you are interested in collecting and analyzing. Remember that this deployment does not include the higher-level analysis and reporting that is included with the NetProfiler.

Figure 2-1 shows an example NetShark (any variation) and Packet Analyzer deployment. Although this example shows only a single NetShark appliance, you might need additional NetShark appliances for large data centers or to monitor additional locations.

Figure 2-1. Example NetShark and Packet Analyzer Deployment



For most deployment scenarios, you want to deploy the NetShark in the data center. The majority of network traffic in most corporate environments is focused in the data center, making it an ideal place to put a traffic-monitoring solution. You can investigate and analyze the conversations flowing between end users and the servers in the data center—this information is invaluable when troubleshooting a network issue.

When you install the NetShark in the data center, you do not always catch all traffic. However, this uncaptured traffic is often not of great interest or significant volume: for example, local print traffic to a floor or building. If you want to monitor traffic that does not go through the data center, you can place additional NetSharks at strategic wiring closets or deployed in the branch office on the SteelHead EX. Because there are many NetShark sizes, you can choose one solution that is appropriate for the data center and a smaller NetShark for a remote wiring closet if appropriate. The availability of the NetShark-v enables you to leverage the power of NetShark in conjunction with an existing VMware ESXi, or Microsoft Hyper-V, or SteelHead EX VSP environment and extend visibility into parts of your network that were not previously practical. For available NetShark models, see [“Choosing a NetShark Model” on page 22](#).

If you connect the NetShark to a network mirror port or TAP, you can detect network activity that you want to monitor but not necessarily store: for example, you can create watches (configurable alerts based on the behavior of certain traffic) to detect certain conditions without capturing the traffic. You can preserve only the desired information, thereby reducing the amount of disk space you use compared to a capture job that captures all traffic.

For more details about watches, see the *SteelCentral Packet Analyzer Reference Manual*.

To store network traffic, you must define capture jobs on the NetShark. Capture jobs tell the NetShark what sort of packets to store to disk. The capture job can store packets based on a wide variety of criteria, including physical source, medium, or various aspects of packet-header information. Each NetShark network interface can support 15 running, non-indexing capture jobs. If any of these jobs include indexing, the number of running capture jobs supported on the interface is 14. You can create an unlimited number of stopped capture jobs. For example, you can define one capture job to capture all traffic on specific VoIP ports, and define another capture job to capture all traffic destined for specific critical servers. The different capture jobs operate independently of each other and can capture overlapping data.

Use Packet Analyzer to analyze the data stored on a NetShark and to look at real-time and historical traffic. You can use a variety of views, filters and drill-down techniques to analyze traffic. Packet Analyzer has numerous views available to assist with the analysis. Because Packet Analyzer can connect to multiple NetSharks, the location of Packet Analyzer is not as critical as the location of the NetShark. When you use Packet Analyzer to apply views to packets captured on NetShark, only a small amount of information is actually passed between Packet Analyzer and NetShark. This small amount of data helps to optimize the performance and limit the impact of manipulating the data on the network.

Packet Analyzer does not pull down all of the packets unless prompted to open the packets in Wireshark or save the packets to a local file. Typically, when you save packets or pull packets into Wireshark, you have already filtered the data so that you move only specific packets.

When looking at real-time (or near real-time) data, increasing the physical distance between Packet Analyzer and the NetShark can force more traffic to travel on the WAN. Retrieving full-packet data from the NetShark across a WAN, can use significant amounts of bandwidth and possibly have an impact on other traffic.

To prevent these issues, place Packet Analyzer as close as possible to the NetShark. If you place Packet Analyzer across a slower WAN from the NetShark, you must apply appropriate views and filtering before viewing raw packets or saving these packets locally to minimize the impact these actions have on overall network performance.

Deploying NetShark and Packet Analyzer

The NetShark provides up to 72 TB for packet storage and supports up to eight 1-Gbps and 10-Gbps network interfaces in various combinations depending on what model of NetShark you deploy. This combination provides a powerful solution for packet capture and analysis, while also allowing you to forward flow data to a NetProfiler.

Deploying the NetShark on AppResponse and Packet Analyzer

The NetShark on AppResponse provides a number of benefits over a traditional NetShark, NetShark-v, or NetShark on SteelHead EX deployment. If you have the NetShark on AppResponse, you can analyze packets with both the AppResponse processes and the NetShark processes. Having both processes enables the AppResponse to feed flow data to the NetProfiler. Packet Analyzer can access the entire packet store of the AppResponse, enabling for hundreds of terabytes of packet storage and the use of high-speed 10-Gbps NICs.

When you deploy the NetShark on AppResponse you can configure the NetShark to forward flow data to NetProfilers without any additional license requirements. If you want to access the NetShark with Packet Analyzer, you need an additional NetShark license and Packet Analyzer license.

The NetShark-v on AppResponse does not support capture jobs. Instead, all traffic captured is in a single job that uses the same packet store as AppResponse. This single job enables access to all packets detected and stored by AppResponse without having to use space to double store packets.

Deploying the NetShark-v on SteelHead EX and Packet Analyzer

The NetShark-v is a perfect solution for an enterprise looking for better visibility into smaller branch offices with SteelHead EX. The NetShark-v on the SteelHead EX provides enables you to gather traffic from external sources (such as SPANs off switches, TAPs off links, or aggregators) and generate advanced metrics, including VoIP metrics (MOS, Jitter, and so on) and end-user experience information (response time, client delay, and so on).

The NetShark-v on the SteelHead EX has the following limitations:

- Packet storage space is limited to the space available in VSP.
- There is no access to packets on internal interfaces on the SteelHead. You must use a SPAN or TAP (or other method of aggregating packets) to feed packets to the NetShark-v through the auxiliary or primary network interface.
- Other VSP appliances may impact the performance of the NetShark-v.

If you are deploying the NetShark-v on the SteelHead EX, make sure that the SteelHead in question has sufficient resources (disk, memory, CPU) in the VSP environment to support the NetShark-v. Also make sure that the volume of data coming to the NetShark-v does not exceed its capacity.

The NetShark-v on the SteelHead EX does not have access to the packets traversing within the SteelHead itself. Instead you must use the auxiliary port as a capture port (running in promiscuous mode) and connect an external source of packets to that port. The external source of packets can be a SPAN from another device, a TAP from a link, or some other source.

For more details, see [“Port Mirroring and SPAN” on page 66](#).

Deploying the NetExpress

The NetExpress is the ideal solution for a small enterprise looking for an advanced monitoring solution. The NetExpress has the following primary deployment scenarios:

- Acting as a standalone system for smaller network environments
- Integrated as part of a broader system that provides narrower views of portions of a larger network

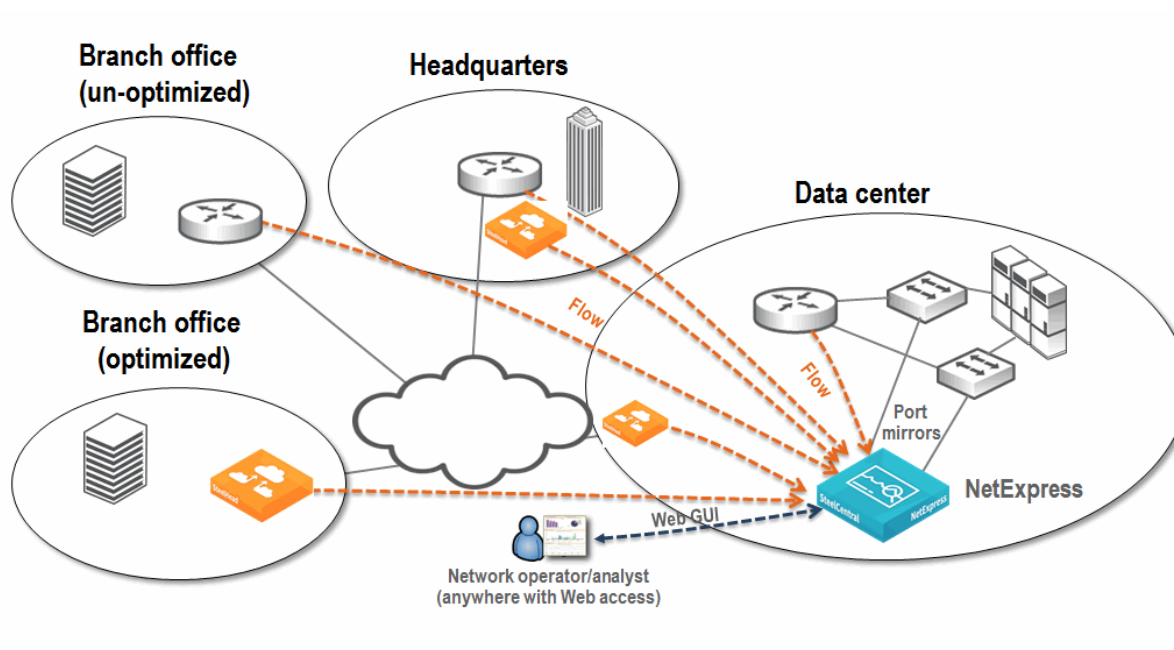
Standalone Deployment

The NetExpress, both the physical appliance and virtual edition, is designed to act as a standalone system capable of both receiving network traffic information and providing all the reporting and monitoring functionality expected of NetProfiler solutions. You can achieve this with four 1-Gbps or two 10-Gbps network monitor interfaces for receiving packets and a single 1-Gbps network interface for receiving flow data directly from the flow sources.

The physical location of the NetExpress is extremely important. Generally, you install the NetExpress at a data center, close to the core switching and routing infrastructure. This location creates the shortest connection paths between devices, and provides the most flexible monitoring.

Figure 2-2 shows an example of a standalone NetExpress deployment. Flow is collected locally at the data center from routers and SteelHeads, and additional flow is collected from remote sites. There is port mirroring of traffic for critical applications, sent directly to the NetExpress monitoring ports.

Figure 2-2. Example Standalone NetExpress Deployment



Integrated Deployment

When you deploy the NetExpress as an integrated deployment, consider what portions of the network you want visible. There are a few ways to restrict visibility directly within the NetProfiler, but you can use the NetExpress to more effectively simulate restricted visibility.

The NetExpress receives data directly from certain sources—through built-in or virtual monitoring ports and direct flow reception—and can forward certain data to other NetProfilers. Because of this, you can use the NetExpress to limit the view to a subset of the network and still feed a larger system to provide a full view.

When you deploy a NetExpress with a limited view, make sure that the NetExpress supports the required flow rate for the covered area. The largest NetExpress has a limit of 120,000 deduplicated flows per minute, and it can be quickly overwhelmed by larger deployments. For example, using an estimate of 40 fpm per host, 120,000 deduplicated flows should provide coverage for a network consisting of roughly 2,400 hosts; flows over 120,000 are dropped and you cannot specify which flows are kept and which are dropped.

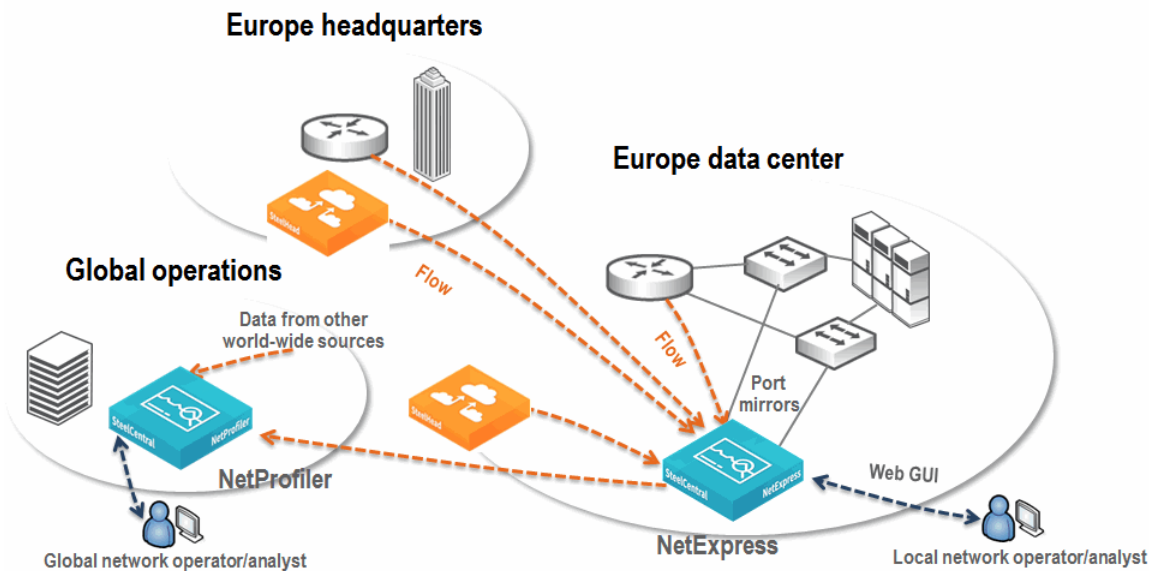
An approximation of how much traffic a host can generate on a per-minute basis for internal communications is 50 fpm per host. If the host is also communicating with resources on the internet, then the number of flows generated can be higher.

Consider physical location when you deploy the NetExpress in an integrated environment. The NetExpress provides all the same advanced network metrics as other NetProfilers when deployed correctly. For example, to provide accurate optimized response time, you must deploy the monitoring interfaces between the server-side SteelHead and the servers. If you place the NetExpress in the wrong location (for example, between the client-side SteelHead and the clients), you can prevent accurate collection and calculation of these values.

If you install the NetExpress in close proximity to the data center, and do not plan appropriately, you can lose visibility into other important areas of the network. If the data center is in one physical location but you are going to use the NetExpress to monitor a separate physical location, you have to choose between not having local visibility or not having certain information available. You can avoid this issue by deploying an additional component, such as the NetShark.

Figure 2-3 shows the NetExpress as part of a larger deployment that includes a Standard NetProfiler. This example shows that the local network operator monitors all traffic on the NetExpress and can configure local policies and local service dashboards. The data received by the NetExpress is also sent to a global NetProfiler. Collection from other sources by the global NetProfiler is not shown.

Figure 2-3. Example NetExpress in a Larger Deployment Environment



Deploying the Standard NetProfiler and Flow Gateway

This deployment is useful if you have an environment that encompasses from several thousand to tens of thousands of hosts in several different physical locations. The primary objective of this deployment is to provide visibility into what is happening on the network from a very high level.

Because the primary objective is to provide basic network visibility, you do not need the NetShark to provide network performance information and Layer-7 application identification information. You can deploy only the NetProfiler and the Flow Gateway to detect what hosts communicate with what other hosts during what time periods.

This scenario requires you to collect data from routers, switches, SteelHeads, and other sources of flow information. Flow data is forwarded to a NetProfiler through one or more Flow Gateways for analysis and reporting. While there are no limitations on the distance between physical devices, network and bandwidth requirements might impact placement decisions.

When deploying the Flow Gateway and Standard NetProfiler, you have a choice between physical and virtual appliances. The Standard NetProfiler physical version supports up to 1,000,000 fpm and the virtual version supports up to 2,000,000 fpm. Because the virtual NetProfiler supports significantly more flows per minute than the physical appliance, there are many scenarios in which deploying the virtual version can make more sense. Be sure to confirm you have sufficient resources on your VMware ESXi host for the level of NetProfiler you choose. Running a NetProfiler-v with insufficient resources might cause a variety of different issues.

The virtual and physical Flow Gateway both support the same number of flows, up to 2,000,000 fpm, so the choice between the virtual and physical versions is less important. As with the NetProfiler-v, if you choose a Flow Gateway-v, make sure you have sufficient resources on the ESXi host to support the required flows per minute.

Answer the following questions to decide between physical and virtual appliances:

- Do you have sufficient ESXi 5.0, 5.1, or 5.5 infrastructure available to support the NetProfiler and Flow Gateway deployments?
- Do you expect to need more than 1,000,000 fpm on the Standard NetProfiler or 2,000,000 fpm on the virtual NetProfiler in the near future?
- Is your virtual infrastructure located close enough to the flow sources so that you will not send excess data across the WAN?

If you have multiple data centers and are using server virtualization, you can use the NetProfiler and Flow Gateway as tools to show you which clients are using which servers and where those clients are located. You can optionally look within a software defined network carrying traffic between your virtual servers and data centers. If you are using QoS on SteelHeads, you can configure bandwidth to be properly allocated and used in the most efficient manner. You can also monitor your WAN links to ensure that you have sufficient bandwidth for high-traffic times of the day by looking at real-time performance and historical information.

For more information about QoS, see [“Configuring Riverbed QoS Integration” on page 90](#).

If your organization has multiple physical locations that are not connected with high amounts of bandwidth (such as a branch connected to a main office through a DSL line), Riverbed recommends that you deploy multiple Flow Gateways throughout the enterprise, with one at each location you want to monitor. Due to concerns about native flow data being sent across a WAN, placing a Flow Gateway at locations with sufficient traffic makes sense if the location warrants monitoring in the first place. A small Flow Gateway or Flow Gateway-v can support up to thousands of hosts (reporting devices).

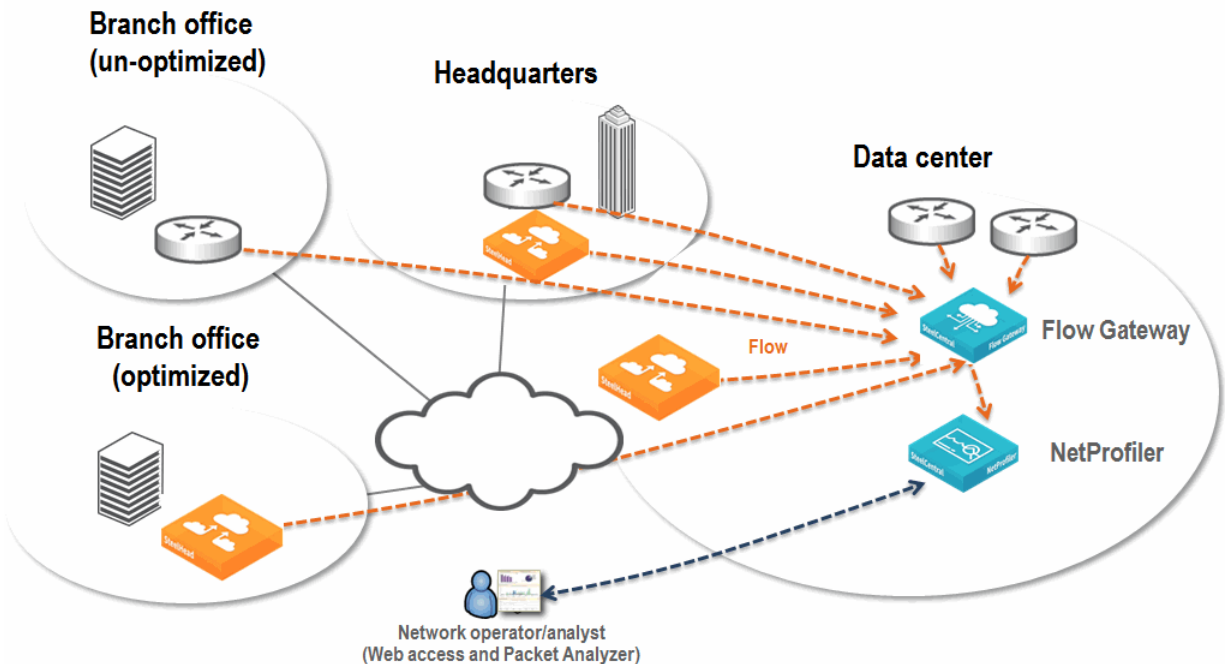
If you host your data center in a location that is not populated with employees, it might make sense to place your Flow Gateway at this site. The data you capture at the data center, even though there are few local employees, is often critical from a monitoring perspective. If there is small amounts of additional data that you want to monitor that does not warrant an additional Flow Gateway, you can send that data across the WAN.

When deciding where to deploy the Flow Gateway, consider the location of the two sides of the conversations you want to monitor. If the traffic is between remote clients and servers in a single data center, then you might not need to place the Flow Gateway at the remote office or send flow data from the remote office to a Flow Gateway in the data center. Because all critical traffic traverses the data center, a single Flow Gateway monitoring all the traffic in, out, and within the data center might be sufficient. If you want accurate per-interface statistics, then you might have to deploy additional Flow Gateways at smaller sites to properly gather interface statistics.

You do not have to install the NetProfiler in the data center because the physical location of the NetProfiler is much less important than the position of the Flow Gateway. Because there can still be considerable network traffic generated even after flows are compressed, Riverbed recommends that you install the NetProfiler as close as possible to the largest sources of flow data.

Figure 2-4 shows an example deployment that includes the NetProfiler and Flow Gateway. All SteelHeads and routers at remote sites, and routers within the data center, send flow data. There are no data flows from smaller sites (not shown in Figure 2-4). Because these much smaller sites primarily communicate back to the data center, traffic detection is based upon collection from the data center routers and SteelHead.

Figure 2-4. Example NetProfiler and Flow Gateway Deployment



Deploying the Enterprise NetProfiler and Flow Gateway

For very large environments, the Enterprise NetProfiler provides an expandable solution that can process flows from hundreds of thousands of hosts. The Enterprise NetProfiler provides a robust solution, allowing visibility ranging from a high-level overview, down to the flow level. The Enterprise NetProfiler has a base flow rate of 1,000,000 fpm, and you can add additional modules to support 1,000,000 fpm per module for a maximum supported flow capacity of 10,000,000 fpm, after deduplication. Because flows are stored across multiple expansion units, the amount of disk space for flow storage is increased as capacity increases.

If you have a very large organization, the physical location of the Enterprise NetProfiler becomes less critical. When you have enough traffic for this solution, you have multiple locations that are sending data. If there is a heavy enough concentration of flow in one specific area, it makes sense to install the Enterprise NetProfiler close to this source. In any case, you must locate the Enterprise NetProfiler in an appropriate facility with sufficient bandwidth, power, and cooling.

If you are deploying a very large Enterprise Cluster (consisting of a combination of more than five Analyzer and Expansion modules) you must deploy a separate Dispatcher module. On the Dispatcher module, you have the option of bonding the primary and AUX interfaces into a single channel. This bonding provides double the bandwidth between the Dispatcher module and associated switch (2 Gbps inbound and 2 Gbps outbound). Because of the high volume of data that must be moved between devices (Flow Gateway and NetShark to Enterprise Cluster and between modules on the Enterprise Cluster) Riverbed recommends that you perform this channel bonding.

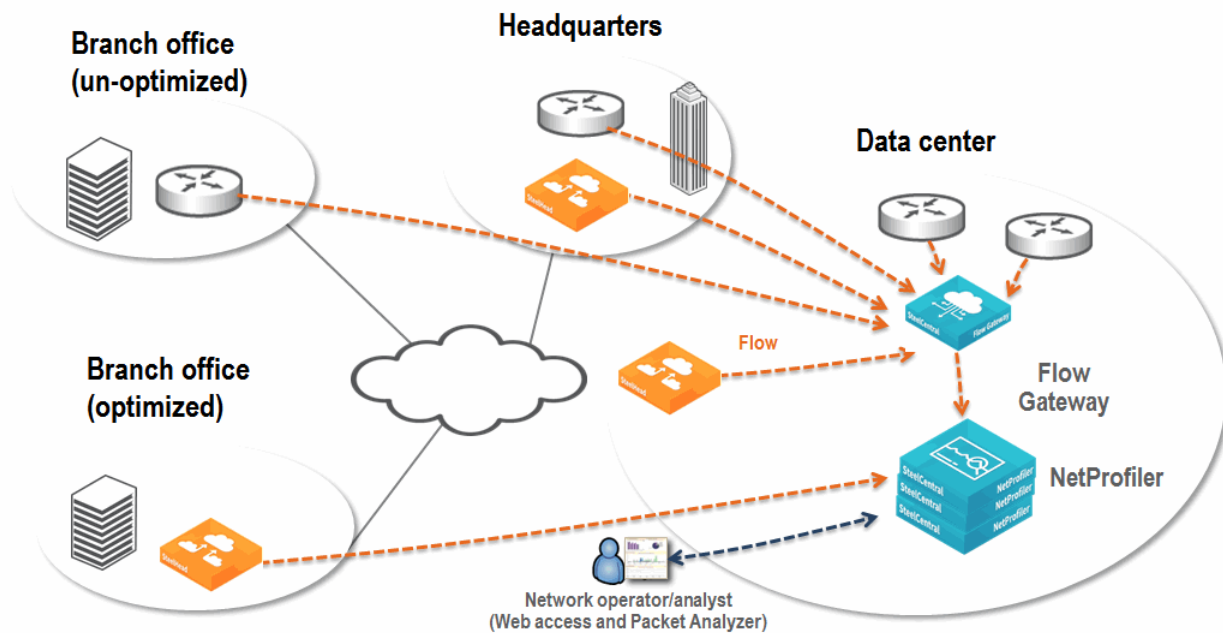
The most important factor in this deployment is to install the Flow Gateways in the correct locations. The Flow Gateway needs to collect data effectively. Depending on the size of your organization, there are several ways to deploy the Flow Gateway that make sense.

You can deploy fewer, larger-capacity Flow Gateways if the number of data collection points is relatively low and the flow rate at those locations is very high. Fewer, larger-capacity Flow Gateways is a good choice if your organization has one or two large data centers to which all clients connect and where all the collected information is concentrated. If you place one or more large Flow Gateways in a data center, you can collect all the data necessary. With proper placement, the Flow Gateway can detect conversations from the clients in the remote office to servers in the data centers.

When choosing between a physical and virtual Flow Gateway to deploy with an Enterprise NetProfiler, make sure you have sufficient ESXi infrastructure to support the size of the virtual device you deploy—include considerations for future expansion and current needs.

Figure 2-5 shows a Flow Gateway collecting and deduplicating the data flow, then forwarding the flow to the Enterprise NetProfiler. Because this deployment does not require network performance and deep packet analysis, you do not need to install the NetShark. This solution enables you to report, analyze, and troubleshoot traffic across the entire large enterprise network.

Figure 2-5. Example Enterprise NetProfiler and Flow Gateway Deployment



Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer

This scenario expands a standard NetProfiler and Flow Gateway deployment to include a NetShark and Packet Analyzer. The NetShark and Packet Analyzer enable you to:

- see network performance data (response time, server delay, and so on) and TCP health information (TCP retransmission).
- detect Layer-7 DPI application information identifying the applications running on the network, independent of ports and protocols in use.
- drill-down from the high-level view provided by the NetProfiler to successively lower-level views until you reach the packet-level view.

When you deploy NetShark you can choose between a physical and virtual appliance. The following questions can help you decide:

- Do you have sufficient ESXi or Hyper-V infrastructure to properly deploy a NetShark-v?
- Do you need more than 2 TB of packet storage?

The physical location of the NetShark is extremely important. The NetShark provides extensive packet capture and analysis capabilities. You must place the device in a location where it receives the maximum amount of critical traffic.

You must decide what information you want to monitor before you decide where to place the NetShark. If you have a single data center and the traffic to and from that data center is the most critical, you should place the NetShark so it can monitor the critical links or VLANs in the data center. However, if your servers contain critical data and are located in a special area (outside the traditional corporate data center), then you might want to place the NetShark in this area. For more information about various methods of collecting packet data, see [“Packet Collection for SteelCentral” on page 65](#).

As a best practice, Riverbed does not recommend that you use a NetShark to monitor a WAN, unless you want packet-level visibility into the WAN link. The routers or SteelHeads on either end of the link are likely to provide flow data that includes link use information down to the level of the individual conversations.

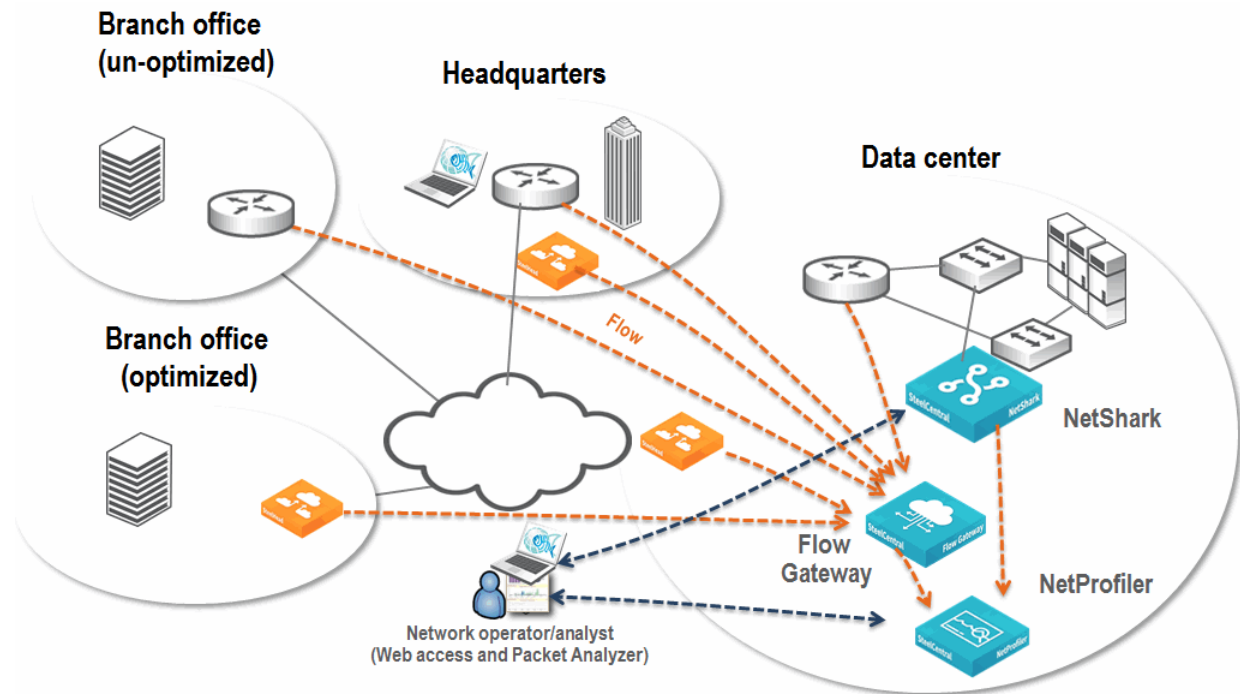
For performance reasons, you might need to limit the amount of data sent to a NetShark. A data center at a medium-size organization can easily exceed 10 Gpbs, which exceeds the write-to-disk capabilities of the NetShark.

The NetShark provides the following ways to monitor the appropriate traffic:

- **Physical** - Collecting packets by using SPAN, port mirroring, and TAP on only the desired links
- **Virtual** - Selecting only those specific packets that you want to monitor using the built-in filtering capabilities

Figure 2-6 shows an example deployment that includes a NetProfiler, Flow Gateway, a NetShark, and Packet Analyzer. Routers and SteelHeads send flow data across the network to the Flow Gateway and provide wide visibility into the network. A NetShark sits off of switches in the data center and collects packets for deeper visibility. Flow data from the NetShark merges with all other flow data collected by the NetProfiler. You can log in to the NetProfiler to view applications flowing across the entire network. When troubleshooting, if you need deeper packet-level analysis, the NetProfiler Management Console automatically launches Packet Analyzer. This configuration takes you from the NetProfiler view of flow data directly into Packet Analyzer views of packet data.

Figure 2-6. Example NetProfiler, Flow Gateway, NetShark, and Packet Analyzer Deployment



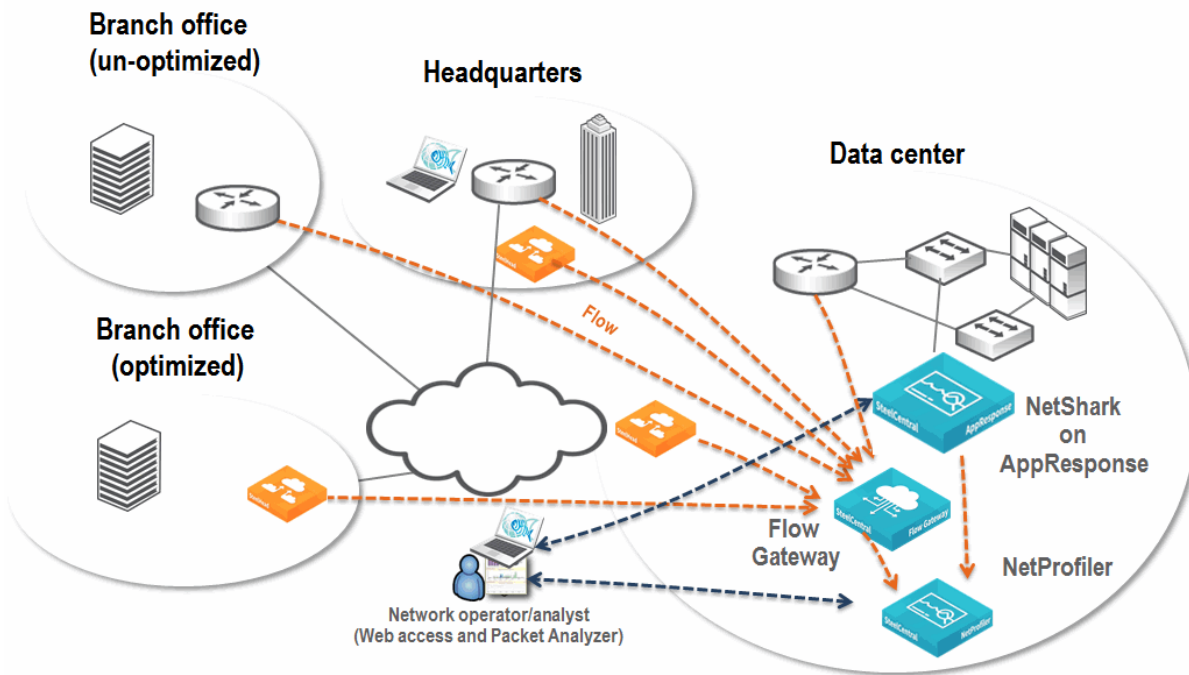
Deploying the NetProfiler, Flow Gateway, and NetShark on AppResponse

This deployment expands upon the NetProfiler deployment described in [“Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer” on page 32](#). When you add an AppResponse with a NetShark to the deployment, you gain a number of benefits. These include:

- Deep packet level visibility the NetShark provides to the NetProfiler
- Access to the packets detected by the AppResponse including the expansive storage options
- Ability to drill down from the NetProfiler to the AppResponse
- Functionality of the AppResponse

When deploying the AppResponse with the NetShark and the NetProfiler, you have an extremely powerful network and application troubleshooting tool.

Figure 2-7. Example NetProfiler, Flow Gateway, and NetShark on AppResponse Deployment

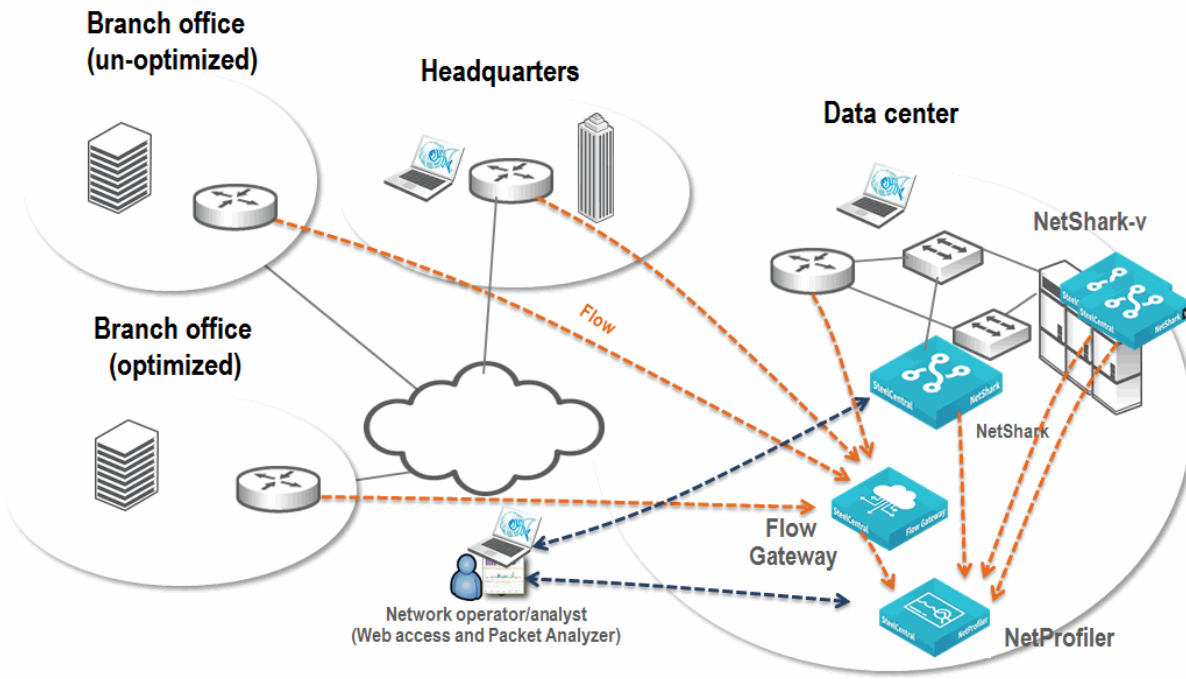


Deploying the NetProfiler, Flow Gateway, NetShark, NetShark-v, and Packet Analyzer

This deployment expands upon the NetProfiler deployment described in [“Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer” on page 32](#). By adding the NetShark-v, you add visibility into the physical network, and you have visibility into the relationship between virtual machines hosted on an ESXi platform in the virtual environment.

Figure 2-8 shows an example deployment that includes the NetProfiler, Flow Gateway, NetShark, NetShark-v, and Packet Analyzer. Deploy the NetShark-v on each ESXi platform in which you want visibility. Metrics are sent from within the virtual environment to the NetProfiler. Using Packet Analyzer, you can also perform packet analysis.

Figure 2-8. Example NetProfiler, Flow Gateway, NetShark, NetShark-v, and Packet Analyzer Deployment



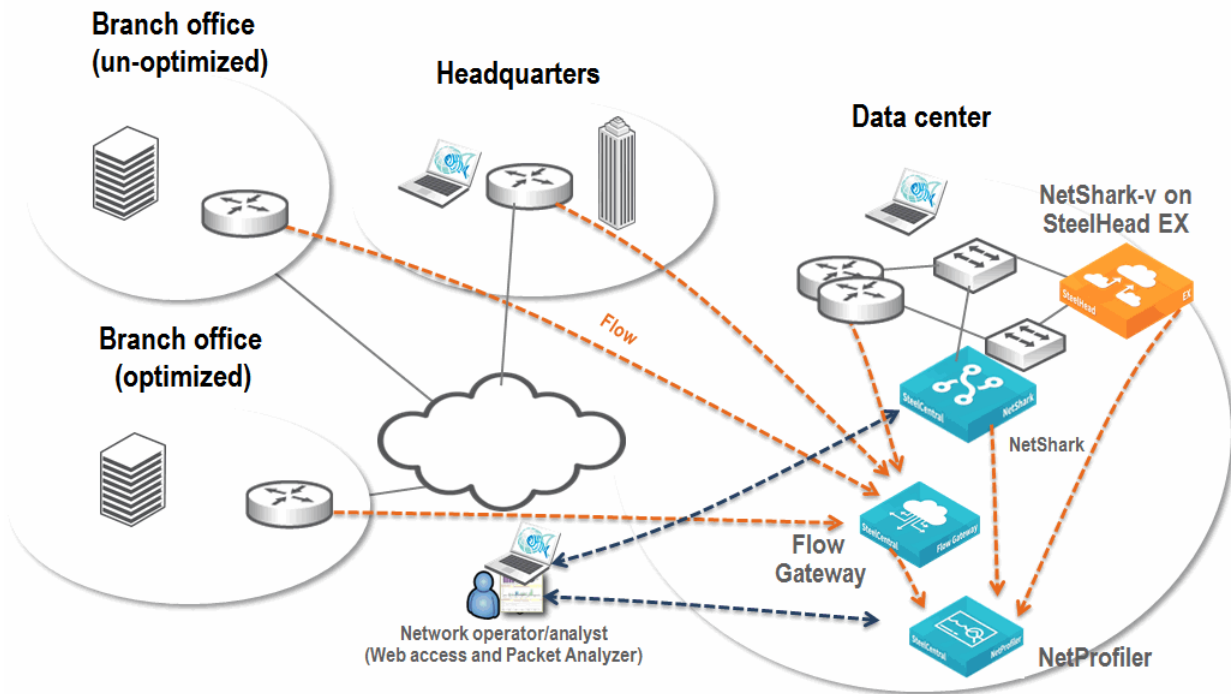
Deploying the NetProfiler, Flow Gateway, NetShark, and NetShark-v on the SteelHead EX

This deployment expands upon the NetProfiler deployment described in [“Deploying the NetProfiler, Flow Gateway, NetShark, and Packet Analyzer” on page 32](#). When deploying the NetShark on the SteelHead EX as part of an overall deployment, you might want to deploy the NetShark-v if any of the following conditions are true:

- You do not have an existing VMware ESXi infrastructure.
- The location does not warrant a full NetShark appliance.
- You need visibility at a packet level.
- You have, or are planning to, deploy the SteelHead EX

By deploying NetShark on the SteelHead EX, you gain the flexibility and visibility that NetShark provides while leveraging the benefit of existing hardware and avoiding the expense and complexity of deploying a physical appliance.

Figure 2-9. Example NetProfiler, Flow Gateway, NetShark, NetShark-v, and SteelHead EX Deployment



Port and Protocol Dependencies

To ensure that the SteelCentral products can communicate correctly, you must open up certain ports across any existing firewalls residing between the various SteelCentral components including NetProfiler, Flow Gateway, NetShark, NetShark on AppResponse, NetExpress, routers, switches, and other devices. The figures in this section show which ports and protocols are necessary for different deployment scenarios. The figures also show external port dependencies for various integrations. For more details about external integrations, see [“Additional SteelCentral Integration” on page 97](#).

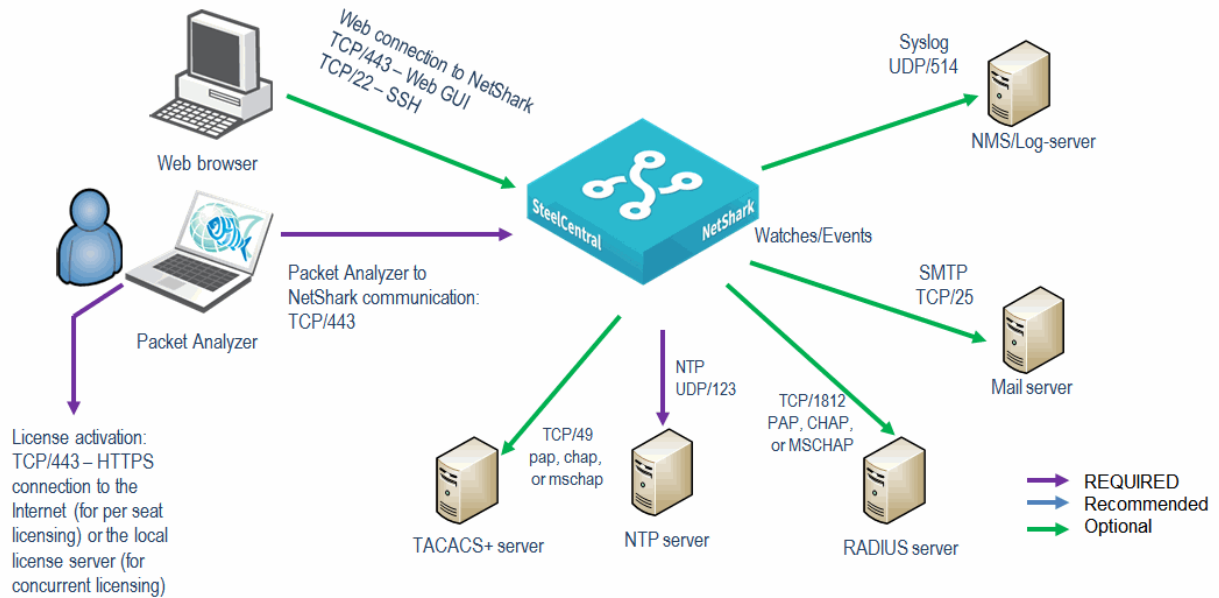
This section describes the following dependencies:

- [“NetShark and Packet Analyzer Port Dependencies” on page 38](#)
- [“NetProfiler and Flow Gateway Port Dependencies” on page 39](#)
- [“SteelCentral Appliance Full-Solution Port Dependencies” on page 40](#)
- [“SteelCentral Appliance Enterprise Solution Port Dependencies” on page 41](#)

NetShark and Packet Analyzer Port Dependencies

Figure 2-10 shows which ports and protocols you must have open for communications within a Packet Analyzer and a NetShark deployment. External connections are optional, depending on which integrations you use.

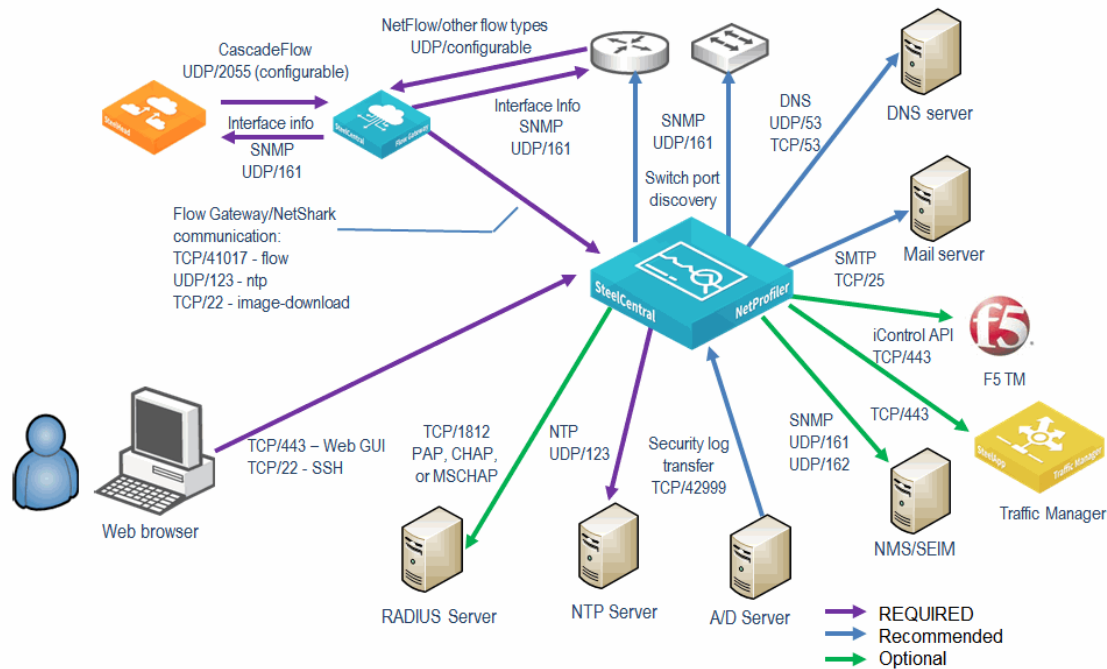
Figure 2-10. Packet Analyzer and NetShark Dependencies



NetProfiler and Flow Gateway Port Dependencies

Figure 2-11 shows which ports and protocols you must have open for communication within a NetProfiler and Flow Gateway deployment. External connections are optional, depending on which integrations you use.

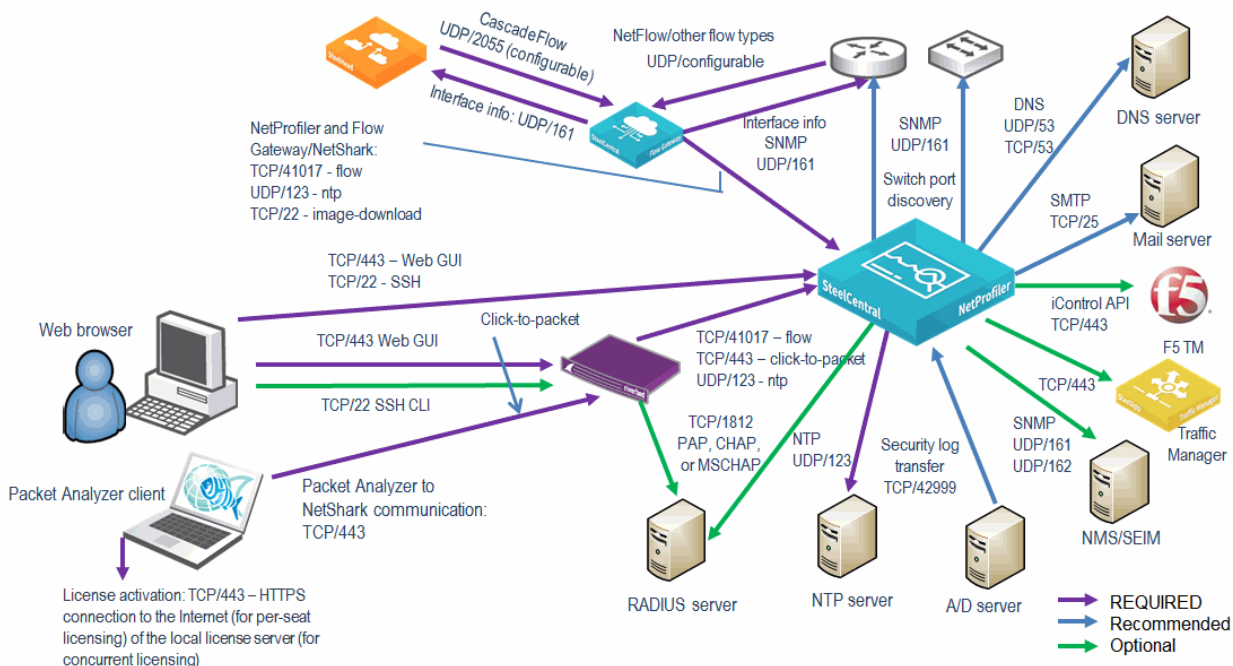
Figure 2-11. NetProfiler and Port Dependencies



SteelCentral Appliance Full-Solution Port Dependencies

Figure 2-12 shows which ports and protocols you must have open for communications within a NetProfiler, Flow Gateway, and NetShark deployment. External connections are optional, depending on which integrations you use.

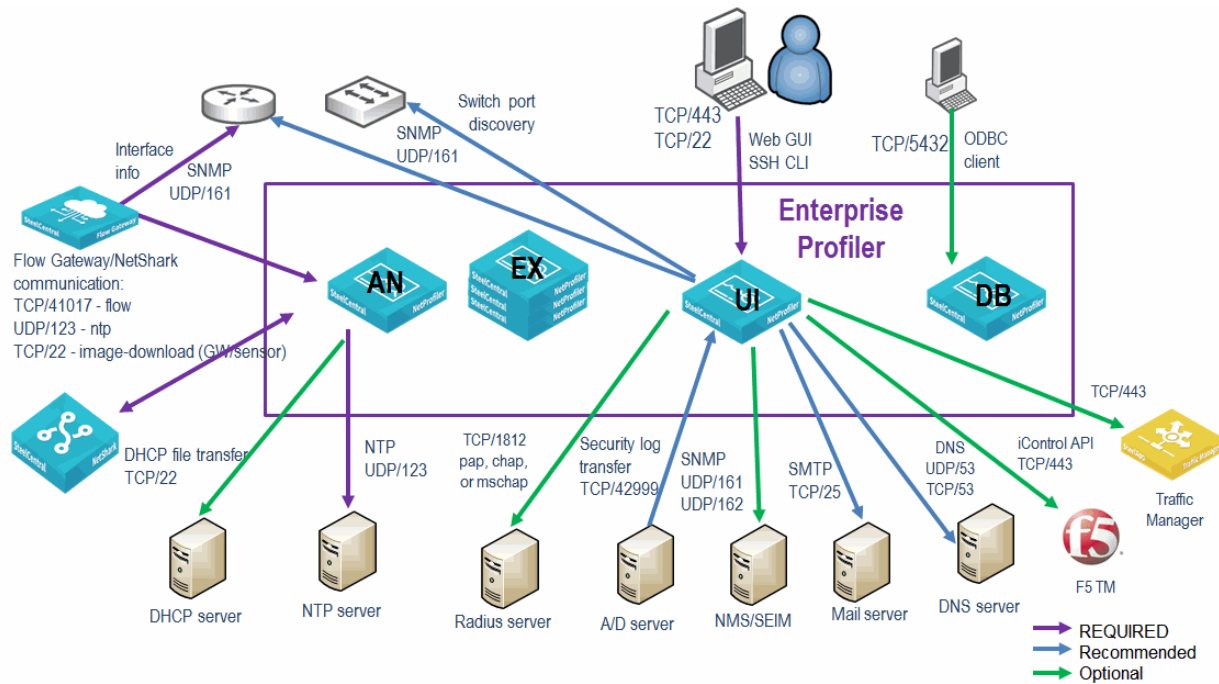
Figure 2-12. Full Solution Port Dependencies



SteelCentral Appliance Enterprise Solution Port Dependencies

Figure 2-13 shows which ports and protocols you must have for communications within an Enterprise NetProfiler, Flow Gateway, and NetShark deployment. Many external connections are optional, depending upon which integrations you use. You must have all components that compose the Enterprise NetProfiler on the same network subnet, using 1-Gb ports and connected to the same switch. If you include a dedicated dispatcher in the deployment, the switch it is connected to must have two available 1-Gb ports and a support channel bonding to allow for the higher bandwidth required.

Figure 2-13. Enterprise Solution Port Dependencies



NetProfiler and NetExpress Flow Storage

This section describes estimating and sizing your flow rate and flow storage requirements. It includes the following sections:

- “Types of Flow Storage” on page 41
- “Flow Rate Estimation” on page 42
- “Flow Storage Size Estimation” on page 43

Types of Flow Storage

The Standard and Enterprise NetProfiler support the following types of flow storage:

- **Local storage** - Disk internal to the SteelCentral appliance.
- **Storage area network (SAN)** - Near-line storage.

These storage mechanisms operate differently and provide their own benefits and disadvantages.

Local Storage

Every NetExpress, NetProfiler, and Enterprise cluster includes some amount of local storage. The following table lists the system storage type.

System Type	Physical Disks	Storage (Raw)
NetExpress (all physical models)	2	3 TB flow / 4 TB packet
NetProfiler (all physical models)	12	10 TB
Enterprise NetProfiler per Analyzer, or Expansion module	12	20 TB

The primary advantages to local storage are performance and cost. Because the storage is located within the system, it is extremely fast when performing reads and writes, limited only by the physical disks and bus connecting the disks to the core system. You cannot order or operate the NetProfiler without local storage.

The biggest disadvantage of local storage is that it limits expansion. There is no option to increase the capacity or number of internal disks.

Storage Area Network

Storage area network (SAN) provides the most robust external storage solution for Standard and Enterprise NetProfiler. SAN is a reliable solution that enables the NetProfiler to very quickly access an extremely large amount of storage—up to 50 TB.

When you connect a SAN to the NetProfiler, the SAN functions as a new disk partition. The NetProfiler treats the SAN as another disk that is part of the appliance, enabling it to easily access the data. You can also offload much of the disk-related processing to the add-in card you use to connect the SAN.

You must use a separate logical unit number (LUN) and fiber connection between each Analyzer or Expansion module and the SAN when you connect a SAN to an Enterprise NetProfiler cluster. This configuration increases the performance drastically, because each analyzer is given its own dedicated channel to the SAN and potentially its own set of drives on the SAN (depending on the SAN configuration). Also, when you deploy a SAN with an Enterprise Cluster you must either deploy the SAN with all Analyzer and Expansion modules or no Analyzer and Expansion modules. You cannot deploy a SAN with part of an Enterprise Cluster.

Flow Rate Estimation

The most important aspect of system sizing is the number of flows your system receives every minute. To accurately estimate the number of flows, you must know if a host accesses only internal resources or internal and external resources.

You must take into account the following considerations when estimating flows:

- Estimate your flows per minute.

In general, you can expect a minimum of 40 flows per host per minute across all internal systems. Because different systems send different numbers of flows, it is probable that some hosts send much more than 40 fpm, while some systems send fewer.

Riverbed recommends that you consider large appliances when the estimate for flow rate approaches the limits of current or planned appliance system: for example, if you estimate 49,000 fpm, consider if a system capable of supporting 50,000 flows is too restrictive.

- Consider how much of the traffic the NetProfiler detects.

A network with 10,000 hosts, only 25 percent of which forwards traffic to the NetProfiler, has a very different requirement than a network with 5,000 hosts, that all forward traffic. The network with 10,000 hosts has an estimated 100,000 fpm, and the network with 5,000 hosts has an estimated 200,000 fpm. You must know the total number of hosts on a network and the number of hosts that have their flows forwarded to a collector for accurate sizing.

- Estimate the current traffic level and what is expected to occur in the future.

Sizing a system for 2,000 hosts when an additional 500 seats will be added over the next year (or conversely 500 seats will be removed) might result in a significant missizing of the system. While you might not know detailed plans for an enterprise over the lifetime of the deployed system, Riverbed recommends that you make estimates based on as much information as is possible. Additionally, if you are performing a proof of concept (PoC) on a small portion of a network, ensure that the final solution is sized for the actual implementation and not the PoC implementation.

Flow Storage Size Estimation

Another primary sizing decision factor is how long you want to keep the stored flows available for retrieval. The amount of space occupied by each flow is dynamic depending on the number of hops being stored. Theoretically, you can store up to tens of billions of individual unique flows if each flow represents only data from a single hop. In a more realistic scenario, flows can represent data from up to ten hops, which is an eight times increase in the flow size (and a corresponding decrease in space available to store flows).

Note: This estimation assumes that you have 50 percent of available flow storage dedicated to storing flow records. On normal implementations, half of the storage is dedicated to storing the minute-by-minute flows and half is used to store presummarized data, called *rollups*.

If a NetExpress receives one flow per minute, it can store flows covering a time span of more than 4,500 years. On a more realistic network sending 50,000 fpm, the same NetExpress can store more than 60 days worth of flow data.

If you must store flows (at the realistic flow rate) for at least 180 days, then the storage included on the NetExpress is inadequate. If you have not yet purchased your appliances, make sure you choose a larger system (for example, a Standard NetProfiler) that has the amount of storage or supports a SAN storage option that meets your needs.

If you have already installed the system, you might need to upgrade to a larger capacity appliance to meet the desired storage capacity. In the example above, because NetExpress does not support a SAN alternative, you must upgrade to a larger system (such as a Standard NetProfiler) with additional storage capacity. When you determine the upgrade path to take, ensure that the path you choose is the most logical approach. You can move up from one system to another but not backwards. For example, you can move from a NetExpress to Standard NetProfiler or from Standard NetProfiler to Enterprise NetProfiler cluster. You cannot move from Enterprise NetProfiler cluster to a Standard NetProfiler.

When you choose the appropriate upgrade path, take into account the desired flow storage capacity and the cost of the system. While the Enterprise NetProfiler cluster might meet the storage needs of the network in the example above, unless you expect significant growth in the near future, it is unlikely to be as cost effective as a Standard NetProfiler with the SAN option.

Flow Redundancy with SteelCentral

This section describes how to send and store redundant flow using SteelCentral. SteelCentral support the following options to obtain redundant flows:

- **Option 1** - All unique flows go to one Flow Gateway, and the Flow Gateway sends the flows to two NetProfilers.
- **Option 2** - All unique flows go to two Flow Gateways, and each Flow Gateway sends the flows to one NetProfiler.
- **Option 3** - A blend of option 1 and option 2.

Figure 2-14 shows Option 1: all unique flows go to one Flow Gateway, and the Flow Gateway sends the flows to two NetProfilers.

Figure 2-14. Option 1: Sending Flows to a Single Flow Gateway




Figure 2-15 shows how you can scale Option 1 to the size of your organization. You can use this design if you run out of flow volume and have multiple Flow Gateways. Keep in mind that each flow source can send to only a single Flow Gateway. If the Flow Gateway goes offline, you lose the flows for that Flow Gateway for the period of time it is not operational.

Figure 2-15. Scalable Option 1: Sending Flows to a Single Flow Gateway




Figure 2-16 shows Option 2: all flows go to two Flow Gateways. Each Flow Gateway sends its flows to only one NetProfiler. This design provides Flow Gateway redundancy as well as NetProfiler redundancy.

Figure 2-16. Option 2: Sending Flows to Two Flow Gateways




Figure 2-17 shows how you can scale Option 2 to the size of your organization. You can scale this design as long as each Flow Gateway sends to only one NetProfiler.

Figure 2-17. Scalable Option Two: Sending Flows to Two Flow Gateways





Figure 2-18 shows Option 3: flow sources that go to only one Flow Gateway are sent to two NetProfilers, and the flows that go to both Flow Gateways are sent to a single NetProfiler. Riverbed recommends this option for cases that absolutely require this level of complexity. Flow Gateways must have a per-flow-source configuration, which means that you must specify the flow sources on the Flow Gateway for export. Failure to specify the sources to export can result in reporting and performance issues on the NetProfiler and communications issues between the NetProfiler and Flow Gateway.

Figure 2-18. Option 3: Blending Options 1 and 2



The NetShark can send flows directly to two NetProfilers. When configuring flow redundancy, you want to plan how the configuration is synchronized between the two NetProfilers. You can synchronize your configuration by using the SteelCentral backup and restore mechanism. One of the NetProfilers acts as the primary appliance on which configuration changes are made. The second NetProfiler acts as the secondary appliance, which receives a copy of the configuration from the primary appliance.

For more information about backup and restore, see the *SteelCentral NetProfiler and NetExpress User's Guide*.

CHAPTER 3 Flow Collection for SteelCentral

This chapter describes the flow collection for SteelCentral. It includes the following sections:

- [“Base Requirements” on page 49](#)
- [“Flow Data Fields Consumed by NetProfiler” on page 52](#)
- [“Flow Type Considerations” on page 53](#)
- [“Flow Collection Considerations” on page 54](#)
- [“Flow Collection in Virtual Environments” on page 54](#)
- [“Validating Flow Collection” on page 54](#)
- [“Sample Third-Party Configurations” on page 56](#)

In this chapter, the term *flow* refers to standard flow types, including NetFlow, sFlow, NetStream, IPFIX, and jFlow. The term *device* refers to a router, switch, or another networking device that supports standard flow export.

For details about collecting data from NetSharks, see [“Packet Collection for SteelCentral” on page 65](#). For details about collecting SteelFlow Net from SteelHeads, see [“SteelCentral and SteelHead Integration” on page 81](#).

Base Requirements

Your device configuration must meet the following requirements:

- Set the active time-out setting for flows to 60 seconds.
- Configure devices that support NetFlow v5, v7, or v9 with no aggregation. Riverbed recommends that you use v9 with no sampling if possible.
- Configure devices that support sFlow v2, v4, or v5 with the lowest possible sampling rate. Riverbed recommends that you use v5.
- Configure devices to export flow to the NetExpress management interface or the Flow Gateway management or auxiliary interface (but not both).

- Synchronize devices with an NTP source. Riverbed recommends that you synchronize devices with the same NTP source used by the NetProfiler. For proper operation and reporting, you must synchronize the time stamps on the network equipment and the NetExpress or Flow Gateway.
- Riverbed recommends that you do not adjust the inactive time-out setting from the default setting of 15 seconds. If you want to adjust the setting, the timeout must be less than 60 seconds.
- When you use NetFlow v5, make sure to add the **ip route-cache flow** (or appropriate) command for all active interfaces and VLANs in addition to the ones you regularly use. Because NetFlow v5 is typically ingress only, you can calculate egress only by aggregating ingress from the opposing interfaces.
- If NetFlow v9 is available, you can selectively control which interfaces to use and specify both ingress and egress. Additionally, with NetFlow v9, you can configure the device to include the time to live (TTL) metric in the flow record using the CLI. This enables ordered-path reporting in the NetProfiler. To enable TTL export, enter one of the following commands:
 - If using standard NetFlow configuration, the command syntax from global configuration mode is **ip flow-capture ttl**.
 - If using flexible NetFlow configuration, the command syntax within the flow record template is **match ipv4 ttl maximum**.
- Because by default flow data is nondeterministic (in that the flows do not specify client/server by default), Riverbed recommends that you enable the flow initiator indicator in NetFlow v9. Use the **collect connection initiator** command on Cisco routers and switches running the correct version of Cisco IOS software.
- Riverbed recommends that you configure SNMP access to any devices sending flow to the NetProfiler. Standard flow export provides information with only SNMP ifindex values. By enabling SNMP on these devices, the NetExpress or Flow Gateway can look up the actual names, descriptions, speeds, and other information about the interfaces. For more information about SNMP integration, see [“SNMP Integration for Flow Sources” on page 97](#).

Additional requirements and considerations for Cisco equipment:

- If you use NetFlow on a Cisco 4500 switch, the Supervisor Engine IV or V must be equipped with a NetFlow Services daughter card and the required software versions.
- If you use NetFlow on a Cisco 6500 switch equipped with both Multilayer Switch Feature Card (MSFC) and Supervisor 1 (SUP1) modules, you must enable NetFlow on the router level and the switch level. The route-once-switch-many concept applies to this hardware configuration. A new flow is first routed by the MSFC module before it is placed in the MLS cache and is switched. The NetProfiler must receive NetFlow data from both modules to avoid missing any data. A similar concept applies to a chassis with SUP2 or 720 modules.
- If you use NetFlow with the Cisco Nexus 7000 series, and you are using NX-OS v4, you must have a minimum version of NX-OS v4.2(8). If you are using NX-OS v5, you must have a minimum version of NX-OS v5.2(1). Earlier NX-OS releases have incorrect packets-per-second and bits-per-second statistics.

If you are using a Cisco Nexus 5000 series, you cannot export NetFlow from the device to a Flow Gateway for the purposes of forwarding that data to a NetProfiler. The Cisco Nexus 5000 is a Layer-2 switch and does not contain the required information in the flow record for NetProfiler to take advantage of the record. You can forward Layer-2 NetFlow to a Flow Gateway with the purpose of forwarding that flow on to other devices using the native flow forwarding function of the Flow Gateway.

- NetFlow export from the Cisco Adaptive Security Appliance (ASA) does not include standard NetFlow records. Cisco ASA exports NetFlow Secure Event Logging (NSEL) in a NetFlow wrapper. NSEL is event driven, exporting bytes only for the first and last packet in the flow. With early versions on Cisco ASA, there was no concept of an active timer, so you did not get updates on a reliable schedule. NSEL 9.1(2) introduced the ability to send scheduled updates on NSEL records. The NSEL records combined with NetProfiler 10.7 or later enable you to send flow data from a Cisco ASA to a NetProfiler and leverage the information available in the NSEL data. In addition to the usual information (source and destination IP, protocol, source and destination port, ingress and egress ifindex values), NetProfiler uses the following fields:

- ICMP type
- ICMP code
- High-level event code
- Milliseconds since UNIX Epoch that the event occurred
- Milliseconds since the UNIX Epoch that the flow was created
- Delta number of bytes from source to destination
- Delta number of bytes from destination to source

Compared with standard NetFlow v5 the following fields are missing from NSEL:

- Packets in flow
 - Total number of Layer-3 bytes in packets in the flow
 - SysUptime at the start of the flow
 - SysUptime at the time the last packet of the flow was received
 - Cumulative TCP flags
 - IP Type of Service
- Some Cisco devices support NetFlow export for Layer-2 switched traffic in addition to Layer-3 traffic. Generally, Layer-2 switched NetFlow is available for forwarding ASICs PFC3B, PFC3BXL, or PFC3C. For verification on whether your hardware or software supports Layer-2 NetFlow, see Cisco documentation. Use the following command to enable NetFlow export for Layer-2 (if your hardware or software supports Layer-2 traffic export):

```
Router(config)# ip flow export layer2-switched vlan <vlan-list>
```

While Flow Gateway and NetProfiler cannot consume Layer-2 NetFlow for the purposes of analyzing and reporting on network traffic, Flow Gateway can forward Layer-2 NetFlow as part of its native flow forwarding functionality. For more information, see the *SteelCentral Flow Gateway User's Guide*.

Flow Data Fields Consumed by NetProfiler

The following table shows the flow fields that are consumed by NetProfiler. When you configure third-party devices to export flow, you must include as many of the following fields as supported by the third-party device.

Field Name	Description	Flow Versions with Support
Connection initiator	Indicates which host initiated the conversation. Used for proper client/server determination.	NetFlow v9
Source IP address	Source IP address of conversation.	All standard versions
Destination IP address	Destination IP address of conversation.	All standard versions
Inbound SNMP ifindex	SNMP ifindex that identifies the interface through which the conversation is received for the device.	All standard versions
Outbound SNMP ifindex	SNMP ifindex that identifies the interface through which the conversation is transmitted out of the device.	All standard versions
Packet count	Number of packets sent during the conversation.	All standard versions
Byte count	Number of bytes sent during the conversation.	All standard versions
Timestamps	Time stamps for the beginning and end of the conversation.	All standard versions
Source port	Source port being used.	All standard versions
Destination port	Destination port being used.	All standard versions
TCP flags	Set TCP flags	NetFlow v5 and v9 on most devices, sFlow v5
Layer-4 protocol	Layer-4 protocol identifier.	All standard versions
QoS information	Type of service (TOS), differentiated services code point (DSCP).	All standard versions
Time-to-live (TTL)	Time-to-live value observed when the packet traversed the reporting device.	NetFlow v9
Application identifier	Layer-7 application identifier.	NBAR through NetFlow v9 with specific hardware (also available from Packeteer through FDR records), Citrix AppFlow, NBAR v2, SteelHead (8.5 and later), NetShark 10.5, Palo Alto
Retransmitted bytes and retransmitted packets	TCP transmission counters.	SteelFlow Net from SteelHeads, and NetSharks
Network round-trip time	Measurement of round-trip time across the network.	SteelFlow Net from SteelHeads and NetSharks

Field Name	Description	Flow Versions with Support
Total response time, server delay, client delay	Measurement of response time metrics across the network.	SteelFlow Net from NetSharks
VoIP metrics: <ul style="list-style-type: none"> • MOS • R-Factor score • Jitter • RTP packet loss 	Voice-over-IP-metrics computed by the NetShark.	NetShark 9.5 and later export
Loss	A count of the number of lost packets from the sequencing information.	MediaNet
Jitter	Mean jitter for the RTP stream.	MediaNet
ICMP Type	ICMP type.	ASA NSEL
ICMP Code	ICMP code.	ASA NSEL
Event	High-level event code.	ASA NSEL
Event Time	Time since the UNIX epoch when the event occurred.	ASA NSEL
Forward Flow Delta Bytes	Source to destination specific traffic counts.	ASA NSEL
Reverse Flow Delta Bytes	Destination to source specific traffic counts.	ASA NSEL

Flow Type Considerations

A Bluecoat Packeteer shaper supports flow-detail-records v2 (FDR) for application identifier collection. FDR records include information beyond that included in basic NetFlow including Layer-7 application information. The NetProfiler understands these records.

If the device is a SteelHead, Riverbed recommends that you use SteelFlow Net instead of NetFlow to collect additional metrics available as a part of SteelFlow Net, such as Layer-7 DPI, retransmit, and response time information. If the device supports multiple types of flow formats, NetFlow is preferable to sFlow.

Riverbed recommends that you use NetFlow v9 with TTL and that you export the TTL field so that network segment graphs are available in the NetProfiler. If NetFlow v9 is not available, use NetFlow v5.

Flow coalescing (also known as *deduplication*) is the process in which flows reported by different devices merge together into a single record. The NetProfiler and NetExpress merge records from the NetSharks and SteelHeads, NetFlow sources (routers, switches, and so on), and IPFIX sources and then tracks all the interfaces the flow crosses and any changes in traffic volume. A maximum of either 8 or 10 individual sources are supported depending on the appliance in use.

Flow coalescing does not occur between sFlow or sampled NetFlow sources and other flow types. These flow records do not merge with other record types (SteelFlow Net, IPFIX, NetFlow, and so on) because this data is sampled. You can have data inconsistencies if sFlow is merged with nonsampled sources. When you want to see sFlow or NetFlow reported conversations, make sure that your traffic query includes the source device of the flow to ensure the information is reported as accurately as possible.

Flow Collection Considerations

The NetProfiler combines all flows that it receives for the same five pieces of information: source IP address, destination IP address, source port, destination port, and protocol. The NetProfiler and NetExpress track every source device that sent the flow, for up to a total of 8 to 10 in-path devices (up to 16 to 20 total interfaces) per flow. You should gather flows from as many sources as possible, because the NetProfiler and NetExpress licenses are based on total deduplicated flow volume, versus number of devices or number of network interfaces for which flow is received. Do not forget the 8 to 10 device limit.

If you exceed the supported-device limit, the NetProfiler and NetExpress prioritize the devices kept per flow in the following order:

- SteelHeads
- NetSharks
- All other flow sources

If you exceed the supported-device limit within the same category, the devices with the lowest IP address are included. There is no way to change the prioritization of devices. Any data from devices over the 8 or 10 device limit is dropped.

Flow Collection in Virtual Environments

The vSphere v5 distributed vSwitch and the Cisco Nexus 1000V support flow export in a virtual environment. Either solution provides visibility into the virtualized environment, including:

- intrahost virtual machine traffic (virtual machine-to-virtual machine traffic on the same host)
- interhost virtual machine traffic (virtual machine-to-virtual machine traffic on different hosts)
- virtual machine-physical infrastructure traffic

Note: In VMware ESXi 5.5 and earlier, only the distributed vSwitch supports the export of NetFlow data.

Validating Flow Collection

You can validate flow collection from the NetProfiler activity displayed on the Devices/Interface page.

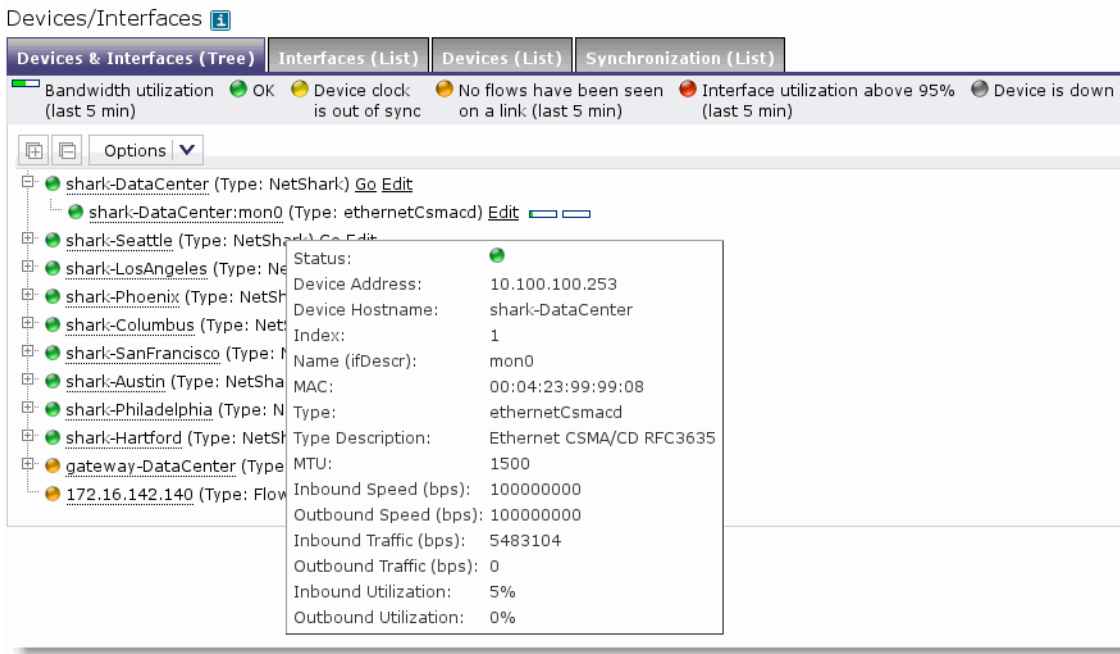
To validate flow data

1. Choose System > Devices/Interface.
2. Select the Devices & Interfaces (Tree) tab to view appliances that send data to the NetProfiler.
3. Expand the display for each Flow Gateway to view which devices the Flow Gateway is receiving flow data from.

4. Further expand the display for each flow-sending device listed in the tree to see specific interfaces.
5. Hover your cursor over the name of each interface to see details about the interface.

Figure 3-1 shows an expanded NetShark-DataCenter, with further expansion of WAN-RTR-Hartford. The pop-up window shows the details about the interface WAN-RTR-Hartford:wan, including inbound and outbound speed and utilization.

Figure 3-1. NetProfiler Device/Interfaces Page



When you validate flow collection on the Devices/Interface page, you can encounter the following display issues:

- If you do not see interface names and speeds, it is likely because you have not configured SNMP polling to the devices.

For details about how to configure SNMP polling for the flow-sending devices, see the *SteelCentral NetProfiler and NetExpress User's Guide*.

- If you see only outbound or only inbound traffic, it can be that you are not exporting bidirectional traffic for that particular interface.

All interfaces for which a flow record is received are in the list, even though you might not be currently exporting flow for that interface. You might see the data if the device is exporting data for the opposing interface and the flow outbound interface is the one in question. For example, you are exporting flows for Interface 1, but the flow is destined for Interface 2. When the flow is received on Interface 1, the record indicates that it is destined for Interface 2. Therefore, Interface 2 is in the list, even though you might not be exporting for Interface 2.

Sample Third-Party Configurations

This section has several third-party configuration examples that show you how to enable NetFlow export to the NetExpress or NetProfiler. Refer to vendor documentation specific to your device and version software. Commands complete various actions, depending upon device software version.

This section includes the following tasks:

- [“Configuring VMware ESXi 5.5 Using vSphere” on page 56](#)
- [“Configuring Cisco 6500 Series Switches Running Native Cisco IOS CLI” on page 57](#)
- [“Configuring Cisco 6500 Series Switches in Hybrid Mode” on page 58](#)
- [“Configuring Cisco 7500 Series Router” on page 59](#)
- [“Configuring Cisco 7600 Series Router” on page 59](#)
- [“Configuring Cisco 3560 and 3750 Flexible NetFlow” on page 59](#)
- [“Configuring the Cisco Nexus 7000 Flexible NetFlow” on page 60](#)
- [“Configuring NetFlow Export for Cisco Nexus 1000V” on page 61](#)
- [“Configuring IPFIX for Avaya \(Nortel\) 8300 and 8600” on page 62](#)
- [“Configuring sFlow for HP Procurve 3500, 5400, and 6200” on page 62](#)

Configuring VMware ESXi 5.5 Using vSphere

The following example uses VMware vSphere to configure an ESXi v5.5 distributed vSwitch to export flow data.

To configure flow on the ESXi 5.5 distributed vSwitch through vSphere

1. Log in to the vSphere Client and select the Networking inventory view.
2. Right-click the vSphere distributed switch in the inventory pane, and select Edit Settings.
3. Select the NetFlow tab.
4. Specify the IP address and port of the NetFlow collector.
5. Specify the vSphere distributed switch (VDS) IP address.

With an IP address to the vSphere distributed switch, the NetFlow collector can interact with the vSphere distributed switch as a single switch rather than interacting with a separate, unrelated switch for each associated host.

6. (Optional) Use the up and down menu arrows to set the Active flow export time-out and Idle flow export time-out.
7. (Optional) Use the up and down menu arrows to set the Sampling rate.

The sampling rate determines what portion of data NetFlow collects, with the sampling rate number determining how often NetFlow collects the packets. A collector with a sampling rate of 2 collects data from every other packet. A collector with a sampling rate of 5 collects data from every fifth packet.

8. (Optional) Select Process internal flows to collect data only on network activity between virtual machines on the same host.
9. Click OK.

Configuring Cisco 6500 Series Switches Running Native Cisco IOS CLI

The following example uses the native Cisco IOS CLI to configure the SUP and MSFC modules of a 6500 series switch. The following commands generally work with Cisco IOS Release 12.2 or later, except where specified. For further information, refer to the documentation for your Cisco IOS software release.

To configure the SUP and MSFC modules of a 6500 series switch

1. At the switch level (SUP2), enter the following commands to turn on NetFlow and set version, flow mask, and timing:

```
Router(config)# mls netflow
Router(config)# mls nde sender version 5
Router(config)# mls flow ip interface-full
Router(config)# mls nde interface
Router(config)# mls aging normal 32
Router(config)# mls aging long 64
```

2. At the routing module (MSFC), enter the following commands to set the device source interface, version, destination, and timeouts:

```
Router(config)# ip flow-export source loopback 0
Router(config)# ip flow-export version 9
Router(config)# ip flow-export destination <flow-gateway-or-netexpress_ip> <udp-port-number>
Router(config)# ip flow-cache timeout inactive 15 (this might be the default depending upon code version)
Router(config)# ip flow-cache timeout active 1
```

If you are running Cisco IOS Release 12.2(18) or later, use NetFlow v9. If NetFlow v9 is not available, use NetFlow v5.

If you are running Cisco IOS Release 12.3(14) or later and are exporting NetFlow v9, you can include export of the TTL, enabling the NetProfiler and NetExpress to show network segment diagrams:

```
Router(config)# ip flow-capture ttl
```

If you are running Cisco IOS Release 12.3(14) or later, running NetFlow v9, and have hardware that supports export of NBAR Layer-7 information, include the following command:

```
Router(config)# ip flow-capture nbar
```

3. To enable NetFlow on your interfaces, enter the following commands, where applicable, for each interface or interface grouping where you require NetFlow accounting (three types of interfaces):

```
interface <type> <slot>/<port>
```

For example:

```
Router(config)# interface fastethernet 0/1
Router(config-if)# ip route-cache flow
```

or

```
interface vlan <vlan-id>
```

For example:

```
Router(config)# interface vlan 3
Router(config-if)# ip route-cache flow
```

or

```
interface port-channel <channel-id>
```

For example:

```
Router(config)# interface port-channel 3
Router(config-if)# ip route-cache flow
```

4. Optionally, if you want to export Layer-2 switched flows (and your switch supports Layer-2 NetFlow export), enter the following command for the set of VLANs where you want the Layer-2 flows exported:

```
Router(config)# ip flow export layer2-switched vlan <vlan-list>
```

Configuring Cisco 6500 Series Switches in Hybrid Mode

The following example configures the SUP and MSFC modules of a Cisco 6500 series switch running in the hybrid mode.

To configure the SUP and MSFC modules of a 6500 series switch in hybrid mode

1. At the switch level (SUP), enter the following commands to enable NetFlow data export (NDE) and to set destination of flow, timers, and full flow:

```
Router(config)# set mls nde enable
Router(config)# set mls nde enable <flow-gateway-or-netexpress_ip> <udp-port-number>
Router(config)# set mls agingtime 16
Router(config)# set mls agingtime fast 32 0
Router(config)# set mls agingtime long-duration 64
Router(config)# set mls flow full
```

2. At the routing module (MSFC), enter the following command to configure NDE and set the destination of flow:

```
Router(config)# ip flow-export <ip-address> <udp-port> <version>
```

3. At the interface level, enter the following commands to enable NetFlow on each interface on which you want to collect statistics and set timers:

```
Router(config)# interface <type> <slot>/<port-adapter>
```

For example:

```
Router(config)# interface fastethernet 0/1
Router(config-if)# ip route-cache flow
Router(config)# ip flow-cache timeout active 1
Router(config)# ip flow-cache timeout inactive 15
```

Configuring Cisco 7500 Series Router

The following example uses the Cisco IOS CLI to configure a Cisco 7500 series router.

To configure a Cisco 7500 series router using the Cisco IOS CLI

1. Enter the following commands to configure NDE:

```
Router# configure terminal
Router(config)# ip flow-export <flow-gateway-or-netexpress_ip> <udp-port-number> <version>
```

2. Enter the following command to enable NetFlow at the interface level on each interface on which you want to collect statistics:

```
Router(config)# interface <type> <slot>/<port-adapter>
```

For example:

```
Router(config)# interface fastethernet 0/1
```

For 7500:

```
Router(config-if)# ip route-cache flow
```

3. Enter the following commands to set the NetFlow timers:

```
Router(config)# ip flow-cache timeout active 1
Router(config)# ip flow-cache timeout inactive 15
```

Configuring Cisco 7600 Series Router

The following example uses the Cisco IOS CLI to configure a Cisco 7600 series router.

To configure a Cisco 7600 series router using the Cisco IOS CLI

1. Enter the following commands to configure NetFlow Data Export (NDE):

```
Router(config)# ip flow-export <flow-gateway-or-netexpress_ip> <udp-port-number>
Router(config)# ip flow-export <version>
Router(config)# mls nde sender <version>
```

2. Enter the following command to enable NetFlow at the interface level on each interface on which you want to collect statistics:

```
interface <type> <slot>/<port-adapter>
```

For example:

```
Router(config)# interface fastethernet 0/1
Router(config-if)# ip flow ingress
```

3. Enter the following commands to set the NetFlow timers:

```
Router(config)# ip flow-cache timeout active 1
Router(config)# ip flow-cache timeout inactive 15
```

Configuring Cisco 3560 and 3750 Flexible NetFlow

The following example shows an example Flexible NetFlow configuration for the Cisco 3750 and 3560 series switches with NetFlow service module C3KX-SM-10G.

To configure Flexible NetFlow for a Cisco 3750 or 3560 switch

1. Enter the following commands to create the flow record:

```
Switch# flow record cascade-record
Switch# match ipv4 tos
Switch# match ipv4 protocol
Switch# match ipv4 source address
Switch# match ipv4 destination address
Switch# match ipv4 ttl
Switch# match transport source-port
Switch# match transport destination-port
Switch# collect counter bytes
Switch# collect counter packets
Switch# collect timestamp sys-uptime first
Switch# collect timestamp sys-uptime last
```

2. Enter the following commands to create the flow exporter and monitor:

```
Switch# flow exporter Cascade
Switch# destination <ip address of flow-gateway or netexpress>
Switch# transport udp <ip address of flow-gateway or netexpress>
Switch# flow monitor Cascade
Switch# record Cascade-record
Switch# exporter Cascade
Switch# cache timeout active 60
Switch# cache timeout inactive 15
```

3. Enter the following commands to enable export on a specific port:

```
Switch# interface TenGigabitEthernet1/1/1
Switch# ip flow monitor Cascade input
Switch# ip flow monitor Cascade output
```

Configuring the Cisco Nexus 7000 Flexible NetFlow

The following example uses Cisco Nexus OS 5.2.1 to configure NetFlow export. You must complete the set of commands in Step 5 for each Layer-3 interface.

To configure a NetFlow export using a Cisco Nexus 7000 Flexible NetFlow

1. Enter the following commands to configure a record to include all necessary fields for the NetProfiler, NetExpress, or Flow Gateway:

```
Switch# configure terminal
Switch(config)# flow record cascade-record
Switch(config-flow-record)# match interface input
Switch(config-flow-record)# match interface output
Switch(config-flow-record)# match ipv4 source address
Switch(config-flow-record)# match ipv4 destination address
Switch(config-flow-record)# match protocol
Switch(config-flow-record)# match transport source-port
Switch(config-flow-record)# match transport destination-port
Switch(config-flow-record)# collect flow direction
Switch(config-flow-record)# collect ipv4 tos
Switch(config-flow-record)# collect ipv4 ttl max
Switch(config-flow-record)# collect transport tcp flags
Switch(config-flow-record)# collect counter bytes
Switch(config-flow-record)# collect counter packets
Switch(config-flow-record)# collect routing next-hop address ipv4
Switch(config-flow-record)# collect timestamp sys-uptime first
Switch(config-flow-record)# collect timestamp sys-uptime last
```

- At the global level, enter the following commands to configure required timeout settings:

```
Switch# configure terminal
Switch(config)# feature netflow
Switch(config-netflow)# flow timeout active 60
Switch(config-netflow)# flow timeout inactive 15
Switch(config-netflow)# flow timeout session
```

- Enter the following commands to configure NetFlow export:

```
Switch# configure terminal
Switch(config)# flow exporter cascade-export
Switch(config-flow-exporter)# destination <ip address of flow-gateway or netexpress>
Switch(config-flow-exporter)# source ethernet 2/1
Switch(config-flow-exporter)# transport udp 2055
!--- Listening port configured on Flow Gateway
Switch(config-flow-exporter)# version 9
```

- Enter the following commands to configure flow monitor:

```
Switch# configure terminal
Switch(config)# flow monitor cascade-monitor
Switch(config-flow-monitor)# record netflow ipv4 cascade-record
Switch(config-flow-monitor)# exporter cascade-export
```

- Enter the following commands to apply a flow monitor to a VLAN or interface (one time for each Layer-3 interface):

```
Switch# configure terminal
Switch(config)# vlan 30
Switch(config-vlan)# ip flow monitor cascade-monitor input
```

Configuring NetFlow Export for Cisco Nexus 1000V

Configuring NetFlow export of the Cisco 1000V is similar to the physical Nexus switches running NX-OS (for example, Cisco Nexus 7000), with some variation in commands. The primary difference is that the Riverbed recommended configuration parameters are for the Cisco Nexus 7000 TTL export. Use the template shown in this example (TTL export is not an option on the Cisco Nexus 1000V).

To configure NetFlow export for a Cisco Nexus 1000V

- Enter the following commands to configure NetFlow Exporter and timing parameters:

```
n1000v# configure terminal
n1000v(config)# flow exporter cascade-export
n1000v(config-flow-exporter)# destination <ip address of flow-gateway or netexpress>
n1000v(config-flow-exporter)# source mgmt 0
n1000v(config-flow-exporter)# transport udp 2055
!--- Listening port configured on Flow Gateway
n1000v(config-flow-exporter)# version 9
n1000v(config-flow-exporter-version-9)# option exporter-stats timeout 60
n1000v(config-flow-exporter-version-9)# template data timeout 1200
n1000v(config-flow-exporter-version-9)# option interface-table timeout 3600
```

- Enter the following commands to configure flow monitor:

```
n1000v(config)# flow monitor cascade-monitor
n1000v(config-flow-monitor)# record netflow-original
n1000v(config-flow-monitor)# exporter cascade-export
n1000v(config-flow-monitor)# timeout active 60
n1000v(config-flow-monitor)# timeout inactive 15
```

3. Enter the following commands to apply the flow monitor to either each virtual interface or each port profile:

- For an interface:

```
n1000v(config)# interface vethernet 2
n1000v(config-if)# ip flow monitor cascade-monitor input
n1000v(config-if)# ip flow monitor cascade-monitor output
```

- For a port profile (the port profile must be configured with other appropriate parameters and inherited on the appropriate interfaces or port groups):

```
n1000v(config)# port-profile type vethernet <profile-name>
n1000v(config-port-prof)# ip flow monitor cascade-monitor input
n1000v(config-port-prof)# ip flow monitor cascade-monitor output
```

Configuring IPFIX for Avaya (Nortel) 8300 and 8600

The following example uses Nortel ERS 8300 and ERS 8600 to configure flow export. You use similar commands to configure other Nortel routers.

To configure IPFIX for Avaya (Nortel) 8300 and 8600

1. Enter the following command to enable IPFIX globally:

```
ERS# config ip ipfix state enable
```

2. Enter the following command to enable IPFIX at a port level, for each port where you want each export:

```
ERS# config ip ipfix port 5/2, 5/3, 5/4, 5/5, 5/6 all-traffic enable
```

3. Enter the following commands to set the timing parameters for the SteelCentral compatibility (active time-out is in minutes, export interval in seconds):

```
ERS# config ip ipfix active-timeout 1
ERS# config ip ipfix aging-interval 15
ERS# config ip ipfix export-interval 60
```

Depending on your router and software version, you might need to specify slot numbers in the previous commands. The following example shows the commands with slot numbers:

```
ERS# config ip ipfix slot 5 active-timeout 1
ERS# config ip ipfix slot 5 aging-interval 15
ERS# config ip ipfix slot 5 export-interval 60
```

4. Enter the following commands to enable export and to export to the NetExpress and Flow Gateway:

```
ERS# config ip ipfix exporter-state enable
ERS# config ip ipfix collector add <ip address of flow-gateway or netexpress> dest-port
<listening of flow-gateway or netexpress> enable true
```

or

```
ERS# config ip ipfix slot 5 exporter-state enable
ERS# config ip ipfix slot 5 collector add <ip address of flow-gateway or netexpress> dest-port
<listening of flow-gateway or netexpress> enable true
```

Configuring sFlow for HP Procurve 3500, 5400, and 6200

The following example uses Procurve 3500, 5400, and 6200 to configure flow export. You use similar commands to configure other HP Procurve devices.

To configure sFlow for HP Procurve 3500, 5400, and 6200

1. Enter configuration mode to configure the NetExpress or Flow Gateway as a flow destination:

```
ProCurve# configure
ProCurve(config)# sflow 1 destination <ip address of flow-gateway or netexpress> dest-port
<listening of flow-gateway or netexpress>
```

In this example, 1 is the sFlow instance. If this instance ID is already in use, then enter either 2 or 3 in the previous and the following commands.

2. Enter the following command to activate sampling:

```
ProCurve(config)# sflow 1 sampling all 500
```

The example shows a sampling rate of one out of every 500 packets. Riverbed recommends that you set the sampling rate to the lowest value recommended by HP; the lowest value recommended depends on device and link speed. In the example, all results use this HP-recommended sampling rate for all ports.

3. Enter the following commands to activate polling:

```
ProCurve(config)# sflow 1 polling all 60
```

In the example, all results are using this polling rate for all ports, and 60 indicates the polling and export interval.

4. Enter the following command to save the configuration:

```
ProCurve(config)# write memory
```


CHAPTER 4 Packet Collection for SteelCentral

This chapter describes the different methods for SteelCentral packet capture. You use packet capture to monitor traffic and analyze packets. This chapter includes the following sections:

- [“SteelCentral for Packet Collection” on page 65](#)
- [“Port Mirroring and SPAN” on page 66](#)
- [“Network TAP Instrumentation” on page 76](#)
- [“VACL Configuration Examples” on page 78](#)
- [“NetShark Passthru” on page 79](#)
- [“Packet Deduplication” on page 80](#)
- [“Snaplen” on page 80](#)

SteelCentral for Packet Collection

SteelCentral collects packets using one of the following devices:

- NetShark
- NetExpress 470 with built-in NetShark capability
- AppResponse with NetShark
- SteelHead EX with NetShark-v
- NetShark-v

All the above devices capture packets at up to line rate, write those packets to disk, perform analysis on those packets, allow easy retrieval of stored packets, and forward packets to capture devices.

You can forward flows from these devices to the NetProfiler based upon the packets that the appliance collects. In addition to standard NetFlow-type fields, these components send TCP health, performance, and end-user experience information. For more information about these components, see [“SteelCentral NPM Overview” on page 11](#) and [“SteelCentral NPM Deployment Scenarios” on page 15](#).

Port Mirroring and SPAN

This section contains the following topics:

- [“Port Mirroring” on page 66](#)
- [“Remote SPAN and Encapsulated Remote SPAN” on page 67](#)
- [“Sample Port Mirror Configurations” on page 69](#)
- [“Cisco Nexus 1000V Virtual Switch SPAN” on page 69](#)
- [“VMware ESXi Distributed vSwitch Port Mirroring Versus Promiscuous Mode” on page 75](#)

Port Mirroring


Port mirroring is the most popular method for collecting packets. Port mirroring is commonly referred to as switched port analyzer (SPAN). You can use the terms SPAN and port mirroring interchangeably. When you configure port mirroring, depending upon your hardware, you can mirror:

- select ports or select VLANs from a device to a monitoring port.
- all ports or all VLANs from the device to a monitoring port.

You can also, depending upon your hardware, configure mirroring on ingress, egress, or both, on the interfaces or VLANs you are monitoring.

[Figure 4-1](#) shows a monitoring configuration in which you detect traffic among all local servers. By monitoring an uplink port or VLAN, in addition to the local ports or VLANs, you can also detect traffic between all external hosts to the local hosts. The NetShark supports up to eight monitoring ports that enable you to duplicate this configuration multiple times using the same NetShark.

Figure 4-1. SPAN Connectivity



Best practices for port mirroring:

- For most monitoring and troubleshooting, you want to collect both sides of the conversation on a single monitor interface. This means that if you are capturing only a single port, you must mirror both directions—ingress and egress. If you are monitoring all ports or all communicating VLANs, you can capture ingress and egress only. Capturing ingress and egress on all ports or all VLANs is redundant, and the duplicate traffic can be deduplicated on the NetShark.
- When you set up port mirroring, you must follow best practices according to your switch vendor. Because many architectures use nonblocking methods that drop overages if you overrun a port mirror (for example, by sending multiple gigabits per second worth of packets from a single gigabit port), depending on the switch you use, there can be an adverse effect on traffic or switch performance.
- For large applications across numerous switches, you can use third-party port monitor aggregators for flexible configurations. Vendors that supply port monitor aggregators include, but are not limited to, Anue Systems, NetOptics, Gigamon, cPacket Networks, and VSS Monitoring.
- Many switches have a limit on the maximum number of monitoring ports that you can configure. This limit is often a maximum of two monitoring ports. If the limit is a problem in your environment, you can add a TAP to an existing monitoring port (essentially making a copy of the traffic already being monitored by another device), or you can use VLAN access control lists (VACLs) to configure what amounts to an additional SPAN port, provided that your equipment supports VACLs. For more information, see [“VACL Configuration Examples” on page 78](#).

Remote SPAN and Encapsulated Remote SPAN

This section describes the following SPAN variations:

- [“RSPAN” on page 67](#)
- [“ERSPAN” on page 68](#)


Riverbed recommends Remote SPAN (RSPAN) and Encapsulated Remote SPAN (ERSPAN) techniques in special circumstances only. With some routers and switches, an adverse impact on performance can occur with configuration of RSPAN or ERSPAN. Read the appropriate documentation and release notes for the hardware and software of your switch or router.

RSPAN

RSPAN enables an extension of a SPAN over the network to another switch on a Layer-2 nonroutable RSPAN VLAN. You can use RSPAN when you have one or more access switches and you want to configure a SPAN to a single NetShark or NetExpress monitoring port at a distribution switch. To ensure that network traffic is not impeded, dedicate a trunk port to carry the traffic from the access switches to the distribution switch.

Figure 4-2 shows a monitoring configuration in which you detect traffic to and from local servers on two different switches. The monitoring port is on an upstream switch. The NetShark and NetExpress have two or more monitoring ports that enable you to duplicate this configuration multiple times using the same NetShark or NetExpress.

Figure 4-2. RSPAN Connectivity




ERSPAN

ERSPAN enables an extension of a SPAN over the network to another switch through a routed GRE-encapsulated tunnel. You can use ERSPAN when a NetShark or NetExpress is monitoring from a distant switch. In this case, you must have adequate bandwidth over the routed path that carries the mirrored traffic so that mirroring does not adversely affect production network traffic.

Note: Riverbed recommends that you use ERSPAN with extreme caution because it can seriously impact overall network performance.

Figure 4-3 shows a monitoring configuration that enables you to detect traffic to and from local servers on two different switches when the monitoring port is on an upstream switch over a routed network. The NetShark and NetExpress have two or more monitoring ports that enable you to duplicate this configuration multiple times using the same NetShark or NetExpress.

Figure 4-3. ERSPAN Connectivity



You must use ERSPAN in a virtualized environment that uses the Cisco Nexus 1000V. The Cisco Nexus 1000V mirrors traffic sent between virtual machines by sending ERSPAN to an external Cisco Catalyst 6500 switch. Take caution when using ERSPAN because it can use a significant amount of bandwidth and resources that send remote traffic.

Sample Port Mirror Configurations

This section includes the following SPAN port configuration examples:

- “Cisco Nexus 1000V Virtual Switch SPAN” on page 69
- “Cisco Catalyst 6500 SPAN” on page 73
- “Cisco Nexus 5000 SPAN” on page 74
- “Cisco Nexus 1000V ERSPAN to Cisco Catalyst 6500” on page 75

SPAN port configurations vary depending upon device and software version. For more information, see the documentation that came with your device.

For details about Cisco switch configuration examples, go to http://www.cisco.com/en/US/products/hw/switches/ps708/products_tech_note09186a008015c612.shtml.

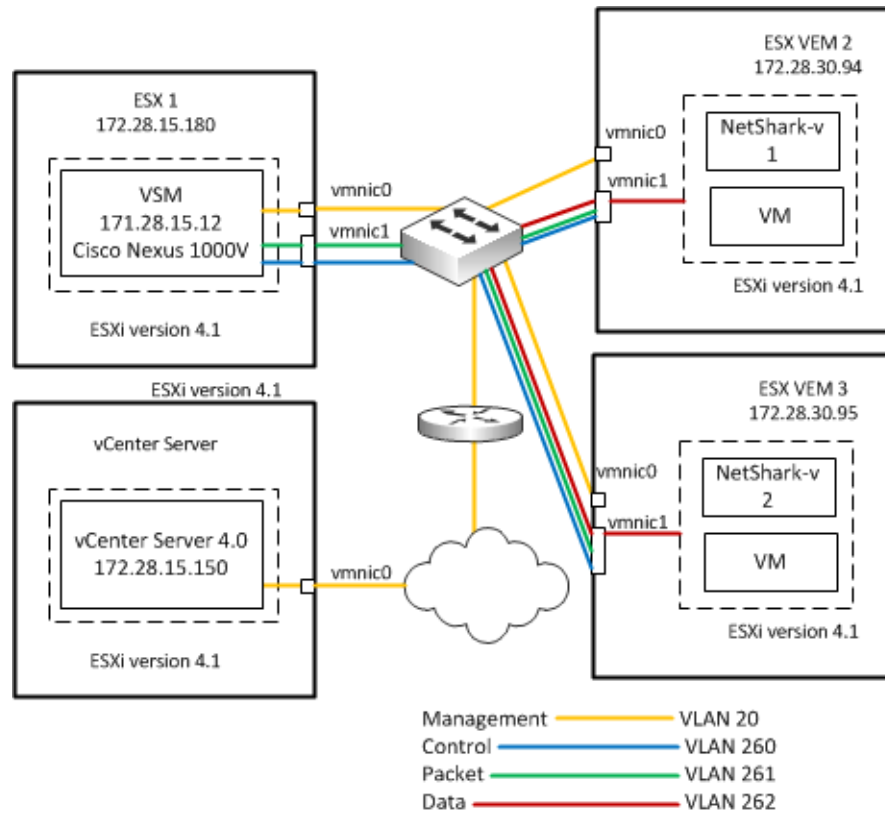
Cisco Nexus 1000V Virtual Switch SPAN

This section describes how to deploy NetShark-v with the Cisco Nexus 1000V switch in a virtualized environment. It includes the following sections:

- “Cisco Nexus 1000V Series Switch Overview” on page 70
- “Enabling a SPAN Session on Cisco Nexus 1000V Switch” on page 72

The NetShark-v software is an application version of the NetShark appliance that provides deep visibility into virtualized environments by monitoring all applications traversing the hypervisor. Cisco Nexus 1000V series switches deliver highly secure, multitenant services by adding virtualization intelligence to the data center network. [Figure 4-4](#) shows a typical deployment scenario.

Figure 4-4. Cisco Nexus 1000V Virtual Switch SPAN



Note: This section describes deploying NetShark-v with the Cisco Nexus 1000V switch. For more information about NetShark-v, see <https://support.riverbed.com/content/support/software/steelcentral-npm/net-shark.html>.

You must have the following requirements:

- VMware ESXi 4.1 or later
- NetShark-v ova package
- Cisco Nexus 1000V VSM and VE package
- VMware vSphere client

Cisco Nexus 1000V Series Switch Overview

The Cisco Nexus 1000V series switches are virtual machine access switches running the Cisco NX-OS operating system for VMware vSphere environments. Operating inside the VMware ESX or ESXi hypervisors, the Cisco Nexus 1000V Series provides:

- policy-based virtual machine connectivity.
- mobile virtual machine security and network policy.

- nondisruptive operational model for your server virtualization and networking teams.

The Cisco Nexus 1000V series switch has two major components:

- **Virtual Ethernet Module (VEM)** - runs inside the hypervisor. The VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware virtual switch functionality. The VEM uses the VMware vNetwork Distributed Switch API to provide advanced networking capability to the virtual machine (VM). The VEM takes configuration information from the VSM and performs Layer-2 switching and advanced networking functions:
 - Port channels
 - QoS
 - Security: private VLAN, access control lists (ACLs), and port security
 - Monitoring: NetFlow, SPAN, and ERSPAN
- **External Virtual Supervisor Module (VSM)** - manages the VEMs. The VSM is the supervisor for the virtual switch chassis. The VEM can exist as a virtual machine running on an ESX or ESXi host, or as a separate physical appliance. Instead of physical line-card modules, the VSM supports multiple VEMs running in software inside the physical server. Configuration of the VEM is performed through the VSM and is automatically propagated to the VEMs. Instead of configuring soft switches inside the hypervisor on a host-by-host basis, you can define configurations for immediate use on all VEMs being managed by the VSM, from a single interface.

To complement the ease of creating and provisioning virtual machines, the Nexus 1000V Series includes the Port Profile feature to address the dynamic nature of server virtualization from the perspective of the network. Port Profile is a new NX-OS feature that provides configuration of ports by category, enabling the solution to scale to a large number of ports.

Port Profiles enable you to define VM network policies for different types or classes of VMs from the VSM. You apply the profiles to individual VM virtual NICs (vNICs) through the VMware's vCenter UI for transparent provisioning of network resources. Port Profiles are a scalable mechanism for configuring networks with large numbers of VMs.

Port Profile follows the VM throughout its lifecycle, whether the VM is being migrated from one server to another, suspended, hibernated, or restarted. In addition to migrating the policy, the VSM also moves the VM network state, such as the port counters and flow statistics. VMs participating in traffic-monitoring activities, such as Cisco NetFlow or ERSPAN, can continue these activities uninterrupted by VMware vMotion operations. When a specific Port Profile is updated, the Nexus 1000V Series automatically provides live updates to all of the virtual ports using that same Port Profile.

Note: NetShark-v with the Cisco Nexus 1000V switch is not officially supported.

Enabling a SPAN Session on Cisco Nexus 1000V Switch

This section describes how to configure a local SPAN session.

Note: The Nexus 1000V switch also supports ERSPAN. For more information about the ERSPAN capability, see the Cisco page at <http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/span.pdf>.

To configure a local SPAN session, each local SPAN session must have at least one destination port (also called a monitoring port) that receives a copy of traffic from the source ports or VLANs. A destination port has these characteristics:

- Can be any physical or virtual Ethernet port or a port channel
- Cannot be a source port
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session
- Receives copies of transmitted and received traffic for all monitored source ports. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.
- Must be on the same host (line card) as the source port
- On Local SPAN, the source interface and destination interface are on the same device.

Consider the following requirements before you begin SPAN configuration:

- You can configure a maximum of 64 SPAN sessions (Local SPAN plus ERSPAN) on the VSM.
- A maximum of 32 source VLANs are allowed in a session.
- A maximum of 128 source interfaces are allowed in a session
- You can configure a port in a maximum of four SPAN sessions.
- You cannot use the destination port in one SPAN session as the destination port for another SPAN session.
- You cannot configure a port as both a source and destination port.
- For VLAN SPAN sessions switched on the same VLAN with both receive and transmit configured, two packets (one from receive and one from transmit) are forwarded from the destination port.

Important Notes:

- SPAN sessions are created in the shut state by default.
- When you create a SPAN session that already exists, any additional configuration is added to that session. To make sure the session is cleared of any previous configuration, you can delete the session first.

To configure a local SPAN session

- On the Nexus 1000V enter the following CLI commands?

```
configure terminal
no monitor session session-number monitor session session-number description description
source {interface type | vlan | port-profile} {number | range} [rx | tx | both] (Optional)
Repeat above line to configure additional SPAN sources
(Optional) filter vlan {number | range}
```



```
(Optional) Repeat above line to configure all source VLANs to filter. destination {interface
type | port-profile} {number | range} (Optional) Repeat above line to configure all SPAN
destination ports. no shut
(Optional) exit
(Optional) interface ethernet slot/port[-port]
(Optional) switchport trunk allowed vlan {vlan-range | add vlan-range | except vlan-range |
remove vlan-range | all | none}
(Optional) Repeat above line to configure the allowed VLANs on each destination port.
(Optional) show interface ethernet slot/port[-port] trunk
(Optional) copy running-config startup-config
```

For more information, see:

- http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0/install/software/guide/install_n1000v.html#wp114347
- http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/data_sheet_c78-492971.html
- http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0_4_s_v_1_3/system_management/configuration/guide/n1000v_system_9span.html
- <http://shouldhavegonewithcisco.com/2009/08/10/integrating-the-nexus-1000v-vsm-with-vcenter/>

Cisco Catalyst 6500 SPAN

The following steps describe how to configure a SPAN for all traffic for VLANs 1 through 100 using a Cisco Catalyst 6500 SPAN. You must only capture ingress on the VLANs to monitor all traffic.

To configure a SPAN for all traffic for VLANs 1 through 100 using a Cisco Catalyst 6500 SPAN

1. From the switch CLI, enter configuration mode to set up a monitor session and configure the source traffic you want to monitor:

```
Switch# configure terminal
Switch(config)# monitor session 1 source vlan 1-100 rx
```

2. Enter the following command to configure the destination port where the NetShark or NetExpress monitoring port is connected:

```
Switch(config)# monitor session 1 destination gigabitethernet 4/3
```

The following example shows how to capture all traffic to and from sources on the downstream port 5/1 and send the collected traffic to port 5/3.

To configure a SPAN for all traffic to and from a downstream switch on port 5/1 using a Cisco Catalyst 6500 SPAN

1. From the switch CLI, enter configuration mode to set up a monitor session and configure the source traffic you want to monitor:

```
Switch# configure terminal
Switch(config)# monitor session 1 source gigabitethernet 5/1 both
```

2. Enter the following command to configure the destination port where the NetShark or NetExpress monitoring port is connected:

```
Switch(config)# monitor session 1 destination gigabitethernet 5/3
```

Cisco Nexus 5000 SPAN

The following example shows how to configure a SPAN for all traffic for VLANs 1 to 100. The Cisco Nexus 5000 collects all traffic ingress to the VLANs. The example shows that using a SPAN on ingress works as well as VLANs 1 to 100.

To configure a SPAN for all traffic for VLANs 1 to 100 using a Cisco Nexus 5000 SPAN

1. From the switch CLI, enter configuration mode to set up a monitor session:

```
Switch# configure terminal
Switch(config)# monitor session 1
Switch(config-monitor)# exit
Switch(config)#
```

2. Enter the following commands to configure the destination port to which the NetShark or NetExpress monitoring port is connected (first set the port as a monitoring port, and next place it into the created session):

```
Switch(config)# interface ethernet 5/4
Switch(config-if)# switchport monitor
Switch(config-if)# exit
Switch(config-if)# monitor session 1
Switch(config-monitor)# destination interface ethernet 5/4
```

3. While still in configuration mode, enter the following command to configure the source traffic you want to monitor:

```
Switch(config-monitor)# source vlan 1-100
```

The following example shows all traffic using SPAN to and from a downstream switch on port 5/2. You want to make sure that you are capturing all traffic to and from sources on the downstream port. Capture traffic in both directions on the port (default if unspecified).

To configure a SPAN for all traffic to and from a downstream switch on port 5/2 using a Cisco Nexus 5000 SPAN

1. From the switch CLI, enter configuration mode to set up a monitor session:

```
Switch# configure terminal
Switch(config)# monitor session 1
Switch(config-monitor)# exit
Switch(config)#
```

2. Enter the following commands to configure the destination port to which the NetShark or NetExpress monitoring port is connected (first, mark the port as a monitoring port, and next place it into the created session):

```
Switch(config)# interface ethernet 5/5
Switch(config-if)# switchport monitor
Switch(config-if)# exit
Switch(config-if)# monitor session 1
Switch(config-monitor)# destination interface ethernet 5/5
```

3. While still in configuration mode, enter the following command to configure the source traffic you want to monitor:

```
Switch(config-monitor)# source interface ethernet 5/2 both
```

For additional information about Cisco Nexus 5000 and NetShark, see <http://supportkb.riverbed.com/support/index?page=content&id=S24538>.

Cisco Nexus 1000V ERSPAN to Cisco Catalyst 6500

The following example shows how to configure an ERSPAN for Cisco Nexus 1000V to a Catalyst 6500. You must configure both the Cisco Nexus 1000V and the Catalyst 6500. This example shows data collection from VLANs 1 through 10 on the Cisco Nexus 1000V switch. The example uses a ERSPAN identifier of 100 for the configuration.

To configure the Cisco Nexus 1000V to collect data on VLANs 1 to 10

1. From the switch CLI, enter configuration mode to set up a monitor session and provide a description:

```
Switch# configure terminal
Switch(config)# monitor session 1 type erspan-source
Switch(config-monitor)# desc cascadeerspansource
```

2. Enter the following command to select which ports or VLANs to monitor:

```
switch (config-monitor)# source vlan 1-10
```

3. Enter the following commands to provide the destination IP address of the 6500 switch (use any reachable IP address on the 6500) and an identifier:

```
Switch (config-monitor)# destination ip <6500-ip-address>
Switch (config-monitor)# erspan-id 100
Switch (config-monitor)# no shut
```

To configure the Cisco Catalyst 6500 to ERSPAN

1. From the switch CLI, enter configuration mode to set up a monitor session and provide a description:

```
Switch# configure terminal
Switch(config)# monitor session 1 type erspan-destination
Switch(config-monitor)# desc cascadeerspansource
```

2. Enter the following commands to configure the specific destination interface, identifier, and receiving IP address:

```
Switch (config-monitor)# destination interface gix/y/z
Switch (config-monitor)# source
Switch (config-monitor)# erspan-id 100
Switch (config-monitor)# ip address <6500-ip-address>
Switch (config-monitor)# no shut
```

VMware ESXi Distributed vSwitch Port Mirroring Versus Promiscuous Mode

Port mirroring can mirror all the traffic coming in or going out of particular virtual ports on a virtual distributed switch. Promiscuous mode repeats the traffic it receives to any virtual adapter that has entered promiscuous mode. Promiscuous mode cannot forward traffic to a particular port on the virtual switch. In other words, any virtual machine connected to the port group that is in promiscuous mode can capture the traffic. This behavior makes using promiscuous mode a potential security risk. Riverbed recommends that you consult your network and security teams before you configure promiscuous mode.

Time Stamping

NetShark provides software-based time stamping of incoming flows. For some applications, such as certain financial transactions, performing time stamping in software does not provide the level of detail needed. To provide support for the additional granularity needed, the NetShark (but not the NetShark-v) supports external time stamping of incoming packets. NetShark supports time stamps from the following appliances:

- Gigamon (Header/Trailer/Trailer X12-TS)
- Anue (requires Advanced Packet Processing module)
- cPacket
- VSS (Time stamp Only and Port ID & Time stamp)
- Arista Packet Broker (Series 7150)

Packet Slicing

Packet slicing is the process of selectively forwarding packets or portions of packets from the packet aggregator to the collector. When a packet is sliced, only a portion of that packet can be forwarded: for example, only the headers are forwarded. When performing packet slicing on a Gigamon 2404 and forwarding the sliced packets to the NetShark, the packet lengths continue to appear correct in both Packet Analyzer views and during packet capture (PCAP) export. The payload (or whatever portion of the packet that is sliced off) is not available on the NetShark even though the packet length is displayed correctly. There is nothing to configure on the NetShark for proper support of packet slicing from the Gigamon 2404.

Packet slicing performed on other packet aggregators results in incorrect packet lengths being reported in NetShark and other devices.


Network TAP Instrumentation

You can insert passive network taps as another method for collecting packet data. This device sits inline on a physical link and makes a copy of all traffic passing through to a monitoring device. You can classify taps as follows:

- **Basic TAPs** - Make a copy of the signal on the wire to a secondary port for monitoring. When you use a passive TAP, you must use two monitoring ports on the NetShark for each single link that you monitor, because the TAP uses a separate port to copy the traffic in each direction.

Figure 4-5 shows a TAP on a link between Device A and Device B. The TAP copies traffic in the direction from Device A to Device B on one port and the direction from Device B back to Device A on a second port.

Figure 4-5. Basic TAP Connectivity



- **Regeneration taps** - Enables you to send the same traffic for the same monitored link to multiple devices. These taps are useful if you want to send traffic from one link to both a NetShark, AppResponse, or NetExpress and another device (such as an IDS).
- **Aggregation taps** - Enables you to aggregate both directions of traffic on a monitored link through a single port so that you need only a single port on the NetShark, AppResponse, or NetExpress for a link you want to monitor. If you use this method, you can potentially miss some packets if the full-duplex link exceeds the available single duplex line rate of the capture point.

Some aggregation taps can regenerate and send traffic from a monitored link to multiple monitoring devices (sometimes referred to as port aggregation). Some aggregation taps can combine multiple monitored links to one or more monitoring devices, sometimes referred to as link aggregation.

Other aggregation taps can split traffic and spread the incoming packets among various different collectors allowing for load balancing and packet slicing.

If you are using an aggregation tap, check with you tap vendor for specific capabilities and options.

- **Advanced/Intelligent taps** - Many of the same vendors that offer intelligent SPAN or port-mirror solutions also offer solutions you can use for taps.

Best practices for TAP deployment:

- Ensure that you understand which type of TAP you are using, keeping in mind that basic taps require two monitoring ports per monitored link, one for traffic in each direction.
- You can use taps on existing SPAN and port-monitoring ports. Using taps is useful if there are no longer SPAN and monitoring ports available on the switch you want to monitor.
- You can chain taps. For example, if you already have a TAP deployed to a monitoring device such as an IDS, you can TAP into the feed to the IDS for monitoring with the NetShark or NetExpress.
- Ensure all packets from a single conversation are sent to the same capture port. NetShark is unable to calculate certain metrics (such as cRTT) if the packets are seen on different capture ports.

VACL Configuration Examples

You can use a VLAN access control lists (VACLs), which are used to mirror ports, for cases when your switch supports only a limited number of in-use SPAN ports. This section includes the following examples:

- [“VACL Port Mirroring Configuration on Cisco 6500 Running CatOS” on page 78](#)
- [“VACL Port Mirroring Configuration on Cisco Catalyst 6500 Running Cisco IOS Software” on page 78](#)

VACL configuration varies based upon device and software version number. For details, see the documentation specific to your device and software version.

VACL Port Mirroring Configuration on Cisco 6500 Running CatOS

The following example shows VACL port mirroring configuration for a Cisco Catalyst 6500 running CatOS. Apply the configuration to the switch only; there is no MSFC component. Connect the capture port where the NetShark or the NetExpress are monitoring interfaces to trunk ports.

To configure VACL port mirroring on a Cisco Catalyst 6500 running CatOS

1. Enter the following commands to create the VACL and specify it as a capture VACL:

```
> set security acl ip SteelCentralMonitor permit any any capture
> show security acl info SteelCentralMonitor editbuffer
```

2. Enter the following command to commit the VACL to NVRAM:

```
> commit security acl SteelCentralMonitor
```

3. Enter the following command to map the VACL to all VLANs you want to monitor:

```
> set security acl map SteelCentralMonitor vlan1,vlan2,vlan3
```

4. Enter the following commands to specify the capture port on which you have connected the NetShark or NetExpress monitoring port (enables for normal switching and creates a copy on the capture port):

```
> set security acl capture-ports 5/3
> show security acl capture-ports
```

VACL Port Mirroring Configuration on Cisco Catalyst 6500 Running Cisco IOS Software

The following example shows VACL port mirroring configuration for Cisco Catalyst 6500 running Cisco IOS software. Apply the configuration to the switch only; there is no MSFC component.

To configure VACL port mirroring on a Cisco Catalyst 6500 running Cisco IOS software

1. From the switch CLI, enter the following commands to create the VACL:

```
Switch# configure terminal
Switch(config)# ip access-list SteelCentralMonitor
Switch(config-access-list)# permit ip any any
Switch(config-access-list)# exit
Switch(config)#
```

2. Enter the following commands to configure the assigned capture or monitoring port as a trunk port (interface 5/3):

```
Switch(config)# interface GE5/3
Switch(config-if)# no ip address
Switch(config-if)# switchport
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation dot1q
```

3. Enter the following commands to define the VLAN access map:

```
Switch(conf)# vlan access-map <map-name-seq#>
Switch(conf-map_name)#
```

4. Enter the following commands to configure the action clause as capture for the access map:

```
Switch(conf-map_name)# match ip address SteelCentralMonitor
Switch(conf-map_name)# action forward
```

OR

```
Switch \ (conf-map_name)# action forward capture
Depending on Cisco IOS rev
Switch(conf-map_name)# exit
```

5. Enter the following commands to apply the access map to all VLANs that you want to monitor:

```
Switch (conf)# vlan filter map_name vlan-list 1-10,15,16...
```

6. Enter the following commands to specify the capture port (previously configured trunk port):

```
Switch (conf)# interface GE5/3
Switch (config-if)# switchport capture
```

NetShark Passthru

Riverbed recommends that you do not use NetShark in passthru mode. Passthru mode is not available on xx70 platforms.

The only reason to use passthru mode is to TAP a SPAN port that another device is already using. For example, an IDS. In passthru mode, you do not have to configure an additional SPAN on the device. With this solution, you are using two ports on the NetShark to monitor a single SPAN port.

When you place the NetShark in passthru mode, it acts as a TAP on a live interface. In passthru mode, the NetShark passes traffic between two physical interfaces on the same card.

Note: The passthru mode does not fail-to-wire. If the NetShark loses power or stops operating for whatever reason, the link does not pass traffic.

Packet Deduplication

Depending upon the packet capture method you use, you might send multiple copies of the same packet to the NetShark. This can occur when you are:

- port mirroring multiple VLANs from the same monitoring port and the packet is routed on the device from one VLAN to another. Even if you are mirroring only ingress to the VLAN, the switch can mirror a copy of the packet when it enters the first VLAN, and mirror a second copy when it enters the second VLAN.
- port mirroring both ingress and egress on the port or VLAN and the packet is routed into and out of the same port or VLAN.
- using an aggregating TAP and the packets are detected on both ports being aggregated.
- using an intelligent monitoring solution that is capturing the same packet from multiple ports and is not performing deduplication.

If any of these actions apply to your environment, Riverbed recommends that you use packet deduplication on the NetShark, AppResponse, and NetExpress. You must enable this feature on the NetShark for each capture port. The NetShark deduplicates packets by evaluating the packet identifier along with other information in the packet. Deduplicated packets can properly capture TCP retransmissions while duplicate packets, due to instrumentation, are dropped. The NetShark performs deduplication on a per-port basis, meaning if duplicate packets are received on separate interfaces, they are not deduplicated.

Some network devices might retransmit TCP packets as part of their normal operation. If you are collecting packets on both sides of such devices to the same port on the NetShark or the NetExpress, enabling packet deduplication does not remove retransmission counts as these are true retransmits. The following are two examples:

- If you capture traffic between an Interceptor and SteelHead and are using full transparency, the packets appear as retransmissions to the NetProfiler and NetExpress if the packets are captured on the same port as the originating packets. To avoid this, do not capture traffic between the Interceptor and SteelHead if configured with full transparency.
- If you capture traffic on either side of a Cisco ASA firewall to the same port on the NetShark, the ASA has a security feature that is enabled by default to help protect against TCP session hijacking. This feature causes the ASA to rewrite sequence numbers for packets traversing it, resulting in observed retransmitted packets if the packets are captured on the same NetShark or NetExpress monitoring port. To avoid this, you can disable the connection random-sequence-number feature on ASA, or you can change your instrumentation so that you do not capture traffic from both sides of the firewall to the same monitoring port.

Snaplen

Snaplen is an abbreviation for snapshot length. Snaplen equals the number of bytes captured for each packet. Having a snaplen smaller than the maximum packet size on the network might allow you to store more packets. However, you are likely limited in your ability to inspect and extract the full packet content.

In most cases, you want to configure the NetShark to capture full packets. If this is your case, do not adjust the default snaplen parameter. Full packet analysis within Packet Analyzer can require the entire packet be captured. If you adjust the snaplen to be smaller, some Packet Analyzer views might not be able to fully analyze the data. For example, volumes of data represented can appear to be smaller than they actually are, and more detailed views do not contain all the data necessary for full analysis.

CHAPTER 5 SteelCentral and SteelHead Integration

This chapter describes how to configure SteelCentral and the SteelHead into an integrated solution for traffic monitoring and troubleshooting. When you integrate SteelCentral and the SteelHead into your environment, you can successfully analyze the traffic on your network for capacity planning and troubleshooting, and thus realize the benefits of optimization. This chapter includes the following sections:

- [“SteelHead and SteelCentral Overview” on page 81](#)
- [“SteelCentral Appliance Deployment Considerations” on page 84](#)
- [“Configuring SteelHead for Flow Data Export” on page 87](#)
- [“NetProfiler and SteelHead Integration” on page 89](#)

SteelHead and SteelCentral Overview

This section describes a summary of SteelCentral and includes the following sections:

- [“NetFlow Versus SteelFlow Net” on page 82](#)
- [“SNMP Interface Persistence \(ifindex\)” on page 83](#)

The primary integration point of the SteelHead and SteelCentral is the SteelFlow Net export from the SteelHead. The SteelHead sends the NetExpress or Flow Gateway an enhanced version of NetFlow called SteelFlow Net. SteelFlow Net includes:

- NetFlow v9 extensions for round-trip time measurements that enable you to understand volumes of traffic across your WAN and end-to-end response time.
- extensions that enable the NetProfiler and NetExpress to properly measure and report on the benefits of optimization.

For more information about flow collection, see [“SteelHead Flow Integration” on page 93](#).

Note: In RiOS 7.0.1 and later, RSP was replaced with Virtual Services Platform (VSP). VSP comes preinstalled in the SteelHead EX. For more information about VSP, see the *Steelhead EX Management Console User’s Guide*.

NetFlow Versus SteelFlow Net

NetFlow provides detailed records of conversations in the network. A basic NetFlow record includes the IP addresses of the endpoints (workstations, servers, printers, and so on), the port or protocol used, traffic volume in bits and packets, TCP flags ingress and egress interface, and so on. These records are sent to a flow collector such as the Flow Gateway. When the records are intelligently combined and stored, you can read them to understand traffic throughout the network for reporting, troubleshooting, and automatic detection of network and application issues.

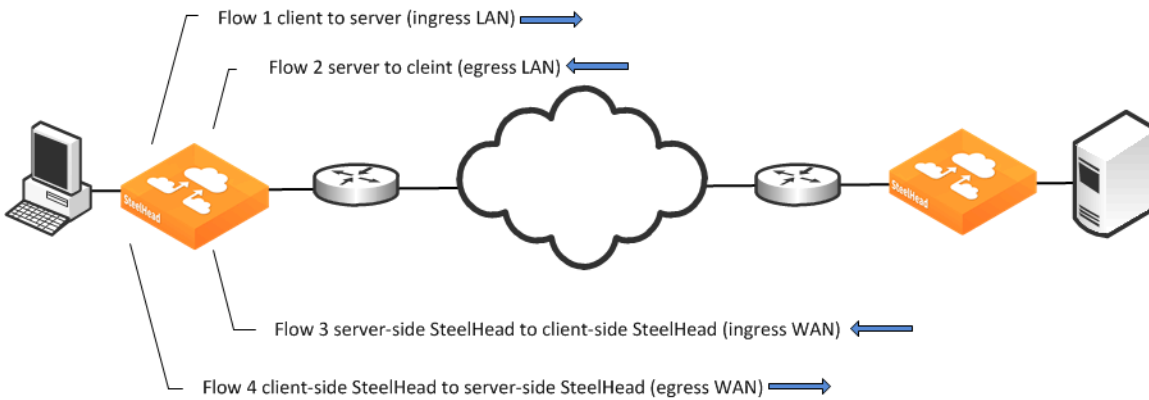
The SteelHead supports the export of flow data in a variety of different formats. Among these formats are standard NetFlow v5 and v9. Use only these standard versions when exporting flow to non-SteelCentral flow collectors or to an earlier version of the NetProfiler, NetExpress, or Flow Gateway. SteelFlow Net was created by Riverbed. It extends NetFlow v9 with multiple custom extensions that provide additional details specific to flows passing through the SteelHead.

The following table shows NetProfiler feature compatibility in RiOS 6.0 and later.

Feature	Flow Version	NetProfiler Version
Basic traffic reporting	NetFlow v5 and v9	NetProfiler 8.0 and later
Enhanced reporting: <ul style="list-style-type: none"> Automatic identification of SteelHead pairs Automatic identification of LAN-WAN interfaces WAN optimization reporting 	SteelFlow Net compatible (v5-based)	NetProfiler 8.3 and later
Enhanced reporting: <ul style="list-style-type: none"> Automatic identification of SteelHead pairs Automatic identification of LAN-WAN interfaces WAN optimization reporting End-to-end response time measures for optimized session QoS DPI 	SteelFlow Net (v9-based)	NetProfiler 8.4 and later

When you use SteelFlow Net, the SteelHead sends four flow records for each optimized TCP session to the NetFlow collector: ingress and egress for the inner channel connection, and ingress and egress for the outer channel. A pass-through connection still sends four flow records even though there are no separate inner and outer channel connections. In either case, the NetProfiler, NetExpress, and Flow Gateway merges these flow records together with flow data collected for the same flow from other devices.

Figure 5-1. Optimized Flow Using NetFlow v9



SNMP Interface Persistence (ifindex)

One of the more common NetProfiler and NetExpress reporting issues happens when interface names and index number associations change across reboots of the SteelHead. This is a direct result of nonpersistent SNMP interface index names, an issue that was the default configuration in RiOS versions prior to 7.0. While you do not need to manually enable ifindex persistence (beginning with new systems running RiOS 7.0 or later; you might need to enable ifindex persistence in upgraded systems), if issues occur with index names and numbers changing on NetProfiler and NetExpress, check to see if ifindex persistence is enabled.

Note: When enabling index persistence on the SteelHead there can be a one time reassignment of index values resulting in changes to index-name mappings. This change can result in inaccurate historical information (for flows received prior to the reassignment).

SNMP uses index values to reference interface names. Only index values are included in the flow record when the SteelHead sends flow data to the NetExpress or Flow Gateway. An SNMP poll from the NetExpress and Flow Gateway to the SteelHead is used to map the index number to an interface name. The SNMP interface index value-to-name mapping can potentially change across reboot and upgrades, causing the NetProfiler to incorrectly map the SteelHead interfaces. Use the ifindex-persistence command on the SteelHead to permanently pin SNMP interface names to SNMP index values.

To enable ifindex persistence

- On the SteelHead, connect to the CLI and enter the following commands:

```
sh > enable
sh # configure terminal
sh (config) # snmp ifindex-persist
#-- You must restart the optimization service, so that NetFlow can use the configured ifindex.
sh (config) # exit
sh # write memory
sh # restart
sh # show service
```

To verify that SNMP persistence is enabled

- On the SteelHead, connect to the CLI, enter the following commands, and look for *Persistent ifindex*:
yes:

```
sh # show snmp
SNMP enabled: yes
System location:
System contact:
Engine ID: 0x8000430b805dc6257f4b328d15
Read-only community: riverbed
Traps enabled: yes
Interface listen enabled: no
Trap interface: primary
Persistent ifindex: yes
No Listen Interfaces.
No trap sinks configured.
```

For additional details about ifindex values, see [“SNMP Integration for Flow Sources” on page 97](#).

SteelCentral Appliance Deployment Considerations

This section provides recommendations on how to improve the accuracy of the exported flow data when you deploy SteelCentral in specific SteelHead deployment scenarios, and it describes certain SteelHead features. It includes the following sections:

- [“Enabling SNMP Polling” on page 84](#)
- [“In-Path Deployments” on page 84](#)
- [“Virtual In-Path Deployments” on page 85](#)
- [“Server-Side Out-Of-Path Deployments” on page 87](#)

Enabling SNMP Polling

The Flow Gateway appliance uses SNMPv3 to poll the SteelHead. This polling requires the SteelHead access control list (ACL) to allow the Flow Gateway to poll OID 1.3.6.1.2. One of the steps in creating an ACL on the SteelHead is creating a View that lists OIDs that are included or excluded from access. If OID 1.3.6.1.2 is not in the *included* section, the Flow Gateway cannot retrieve data when it polls the SteelHead.

In-Path Deployments

In an in-path configuration, you deploy SteelHeads in the physical path of the client and server, where they detect all traffic. You can source flow export from either the SteelHead primary or auxiliary interface to the Flow Gateway or the NetExpress. Enable flow export for all traffic on all SteelHead interfaces that have traffic traversing them: for example, lan0_0 and wan0_0 for a single in-path SteelHead deployment.

Riverbed recommends that you enable simplified routing when using SteelCentral and SteelHead in-path configurations. Simplified routing avoids situations where traffic can potentially run through the SteelHead more than once—commonly known as ricochet. When packet ricochet occurs, the same traffic is reported by the SteelHead multiple times, which causes an unexpected increase in bandwidth, packets, and other traffic statistics in various NetProfiler reports. Ricochet can happen when you install the SteelHead in a different subnet from the client or server and you do not configure the appropriate static routes to avoid traffic passing through the SteelHead multiple times on the way to and from the default gateway.

For more details about simplified routing and in-path deployments, see the *SteelHead Management Console User's Guide* and the *SteelHead Deployment Guide*.

Virtual In-Path Deployments

This section describes SteelHead virtual in-path deployments. It contains the following topics:

- [“Interceptor Considerations” on page 86](#)
- [“Subnet Side Rules” on page 86](#)

In a virtual in-path deployment, the SteelHeads are placed physically out of the path but virtually in the path between the clients and servers. This deployment differs from a physical in-path deployment in that a packet redirection mechanism is used to direct packets to the SteelHead.

Redirection mechanisms include policy-based routing (PBR), Web Cache Communication Protocol (WCCP), the Interceptor, and a Layer-4 load balancer.

In a virtual in-path deployment the SteelHead most likely detects only the traffic that has been redirected to it based on the configured ACL on the router, and not all the destined for the WAN. As a result, you cannot report on the actual WAN-link use based on the SteelHeads reports only. Therefore, you must enable NetFlow export on the router to capture information about the traffic that is not redirected to the SteelHead.

Enabling NetFlow on the router enables reporting traffic on the actual WAN link. If the SteelHead is using correct addressing, the optimized connections are reported showing the SteelHeads IP addresses as the end points of the flow and not the original client/server IP addresses. This can cause some double counting in reports under certain circumstances. However, you must not include the WAN interface of the router in the WAN optimized group, as it is not an endpoint in optimization.

In a virtual in-path configuration, the SteelHeads are connected with their WAN interface only, so that they do not have sufficient information to determine the flow direction of pass-through traffic. Therefore the IP addresses of the subnets passing through the SteelHead need to be manually configured belong to either the WAN or LAN side of the network from a SteelHead perspective. You can specify this configuration in the subnet side rules table.

For more information about configuring subnet side rules, [“Subnet Side Rules” on page 86](#) and the *SteelHead Deployment Guide*.

As a best practice in virtual in-path deployments, enable NetFlow on the primary and auxiliary interfaces and export flow data for the optimized traffic only from the SteelHead. Use the router to export the pass-through flow data. Additionally, configure LAN- and WAN-side subnets in the subnet side rules table.

Prior to RiOS 6.0, you must run the following commands on the SteelHead that is running virtually in-path to enable the SteelHead to determine the direction of flows of optimized traffic on the WAN interface:

```
enable
config terminal
ip flow-export destination <ip-address> <port> interface wan0_0 fakeindex on
```

In RiOS 6.0 and later, the fake index feature is enabled automatically, and the direction of flows of optimized traffic is determined automatically by the SteelHead. Subnet Side Rules still must be configured.

If reports are showing abnormally large bandwidth on the SteelHead WAN interface, it is an indication that either fakeindex was not enabled or LAN subnets were not properly configured. A flow list on such traffic would show flows with both the ingress and egress interfaces of the SteelHead as the WAN, such as wan0_0.

To get information about only the nonoptimized traffic, create a report using a host subnet (or host address) with the SteelHead client IP address.

For more details about virtual in-path deployments, see the *SteelHead Deployment Guide*.

Interceptor Considerations

When you deploy SteelHeads virtually in-path and load-balanced by an Interceptor, in addition to ensuring that you have fake index enabled, make sure that you avoid capturing inner-channel traffic. When you deploy a NetShark or NetExpress, do not place either appliance so it captures traffic between the SteelHeads and Interceptor, because this configuration does not detect any information about the nonoptimized network traffic. You only detect communication between the SteelHeads and Interceptor.

If you configure the SteelHead with full-transparency, you must exclude traffic between the SteelHeads and Interceptor when you configure port mirroring for collection by the NetShark or NetExpress; otherwise, you might detect excessive retransmission and incorrect or missing round-trip-time (RTT) metrics.

Subnet Side Rules

In virtual in-path deployments, the SteelHeads are connected with their WAN interfaces only, so that they do not have sufficient information to determine the flow direction of pass-through traffic. Therefore, you must manually configure the IP addresses of the subnets passing through the SteelHead to belong to either the WAN or LAN side of the network (from a SteelHead perspective).

If you do not have any LAN subnets configured, the SteelHead does not discern whether the traffic is passing from the WAN to the LAN or in the opposite direction. This lack of configuration can result in over reporting traffic in a particular direction or for a particular interface.

To configure subnet side rules on a SteelHead

1. On the SteelHead, choose Networking > Network Services: Subnet Side Rules.
2. Select Add a Subnet Side Rule (Figure 5-2).

Figure 5-2. Subnet Side Rules Page

▼ Add a Subnet Side Rule ✕ Remove Subnet Rules ⇅ Move Subnet Rules...

Insert Rule At: Start ▾

Subnet: 0.0.0.0/0

☒ Subnet is on the **LAN** side of this appliance

☐ Subnet is on the **WAN** side of this appliance

Add

Rule	Source
default	all

3. From the drop-down list, select one of the following:

- Start
- End
- Rule number

The rules are evaluated in order; evaluation stops when a rule is matched. Rules must be in proper order for your network.

4. Specify a subnet using a valid CIDR notation.

5. Select whether the subnet is on the LAN or WAN side of the appliance.

6. Click **Add** to save the rule.

7. Continue to add rules until you have mapped all subnets.

8. Click **Save** to save your changes permanently.

Server-Side Out-Of-Path Deployments

An out-of-path deployment is a network configuration in which the SteelHead is not in the direct physical path between the client and the server. In an out-of-path deployment, the SteelHead acts as a proxy. This configuration is suitable for data center locations in which physical in-path or virtual in-path configurations are not possible. The out-of-path solution uses Network Address Translation (NAT); therefore there is no direct correlation between the client and server conversation and the traffic over the WAN. You can still create valuable reports with this configuration. However, you cannot view the benefit of optimization.

In a server-side out-of-path (SSOOP) deployment, you enable NetFlow on the primary and auxiliary interface and export flow data for only the optimized traffic from the SteelHead. Similar to the virtual in-path deployment, you configure the router to export the pass-through flow data, as the SteelHead detects only optimized data in this configuration.

Prior to RiOS 6.0, and similar to the virtual in-path deployment, you must enable fake index to properly report on the direction of the optimized traffic through the SteelHead. You do not need to configure subnet side rules because the out-of-path SteelHead detects only optimized traffic and never passes traffic.

In RiOS 6.0 and later, fake index is enabled automatically.

For more details about SSOOP, see the *SteelHead Deployment Guide*.

Configuring SteelHead for Flow Data Export

This section explains how to configure the SteelHead for flow data export.

To enable flow data export on the SteelHead

1. On the SteelHead, choose Networking > Network Services: Flow Statistics.
2. Select Enable Flow Export.

3. Clear Enable Top Talkers.

Riverbed recommends that you disable top talkers unless you have it enabled for reporting on the SteelHead. Enabling top talkers forces the active flow time-out to 30 seconds. This time-out results in flow updates being sent to the collector twice per minute. While Riverbed recommends a time-out of 60 seconds in most cases, there are no issues with sending updates more frequently. The collector correctly collates the incoming information and sends a single, complete, update to the NetProfiler once per minute.

Riverbed recommends that you check Export QoS and Application Statistics to CascadeFlow Collectors in all cases.

4. Specify 60 seconds in the Active Flow Timeout field.

5. Specify 15 seconds in the Inactive Flow Timeout field.

6. Click **Apply**.

7. Click **Save** to save your changes permanently.

Figure 5-3. Flow Export Page

The screenshot shows a web interface with two main sections: 'Flow Statistics Settings' and 'Flow Export Settings'. In the 'Flow Statistics Settings' section, there are three checkboxes: 'Enable Application Visibility' (checked), 'Enable WAN Throughput Statistics' (checked), and 'Enable Top Talkers' (unchecked). Below the 'Enable Top Talkers' checkbox are two radio buttons: '24-hour Report Period (Higher Granularity)' (selected) and '48-hour Report Period (Lower Granularity)' (unselected). The 'Flow Export Settings' section contains 'Enable Flow Export' (checked) and 'Export QoS and Application Statistics to CascadeFlow Collectors' (checked). Below these are two input fields: 'Active Flow Timeout:' with a value of '60' and 'Inactive Flow Timeout:' with a value of '15', both followed by the unit 'seconds'. At the bottom left of the form is a dark blue button labeled 'Apply'.

To configure the Flow Collector and Exporting interfaces

1. On the Flow Export page, select the Add a New Flow Collector tab.
2. Specify the IP address and listening UDP port number of the NetExpress or the Flow Gateway.

3. Select the version of flow data to be exported:
 - For the NetProfiler 8.4 and later, use SteelFlow Net.
 - For the NetProfiler 8.3.2, select SteelFlow Net compatible and select the LAN Address check box.
4. For the desired interfaces, select All to export both optimized and nonoptimized flow information.
5. Click **Add** to add the NetExpress or Flow Gateway to the collector list.
6. Click **Save** to save your changes permanently.

Figure 5-4. Flow Collector and Exporting Interfaces

Flow Collectors:

Collector IP Address: Port:

Version:

Packet Source Interface: (Interface used for the source IP of the flow packets.)

LAN Address: ☒ Show

Capture Interface **primary**:

Capture Interface **lan0_0**:

Capture Interface **lan1_0**:

Capture Interface **wan0_0**:

Capture Interface **wan1_0**:

Capture Interface **rios_lan0_0**:

Capture Interface **rios_lan1_0**:

Capture Interface **rios_wan0_0**:

Capture Interface **rios_wan1_0**:

Enable Filter: ☐ Enable

Filter:

Note: The filter is applicable to CascadeFlow and Netflow v9 only. Flow reports will only be sent for IP/Subnets included in this list. If the filter is not enabled, all will be reported.

The filter should be of the form "IP/Subnet" or "IP:Port", one entry per line.

When you click **Apply** and **Add**, the Management Console updates the running configuration. Your changes are written to disk only when you save your configuration.

NetProfiler and SteelHead Integration

The NetProfiler integrates with the SteelHead in two ways: Riverbed QoS Integration and Flow Integration. This section contains the following topics:

- [“Configuring Riverbed QoS Integration” on page 90](#)
- [“SteelHead Flow Integration” on page 93](#)

Configuring Riverbed QoS Integration

Starting with NetProfiler 10.5 and RiOS 8.5, and newer, Quality of Service (QoS) reporting is enhanced by integrating the NetProfiler to collect QoS configuration from the SteelHeads and provide a centralized monitoring and reporting location for QoS historical and real-time traffic as it traverses the WAN.

Note: You must be running RiOS 8.5.x or later.

Integrating Riverbed QoS with NetProfiler enables the NetProfiler to act as a single view into an entire Riverbed QoS infrastructure. NetProfiler acts as a single point of contact, enabling visibility into performance, utilization, and availability across multiple SteelHeads without having to connect to multiple devices or do significant amounts of manual work to correlate data. With Riverbed QoS and NetProfiler integration you can:

- view how much bandwidth is being consumed between sites for specific Riverbed QoS classes.
- ensure bandwidth is assigned appropriately for the needs of the class.
- identify classes that are under-used or over-used.
- confirm the correct applications are running within each class.
- collect QoS configuration summaries
- conduct drill-in investigations and troubleshoot QoS related issues.

When properly configured with REST access codes, the NetProfiler is able to query against individual SteelHeads and retrieve information about how individual classes are configured, minimum and maximum bandwidth assignments per class, and other information. You can run reports to show information for individual SteelHeads and aggregate class details.

Integrating Riverbed QoS with the NetProfiler is composed of the following steps, which require administrator-level access on the SteelHead and NetProfiler:

- [“To configure flow export” on page 90](#)
- [“To configure REST API access” on page 91](#)
- [“To configure the NetProfiler” on page 91](#)

This section requires you be familiar with the configuring QoS on the SteelHead. For details, see the *SteelHead Deployment Guide*.

Note: Outbound QoS (basic or advanced) is assumed to be enabled and configured on the SteelHead. Only Outbound QoS is supported.

The first steps of the configuration are completed on the SteelHead.

To configure flow export

1. On the SteelHead, choose Configure > Networking > Flow Statistics to display the Flow Statistics page.
2. In the Flow Export Settings box, select Enable Flow Export and Export QoS and Application Statistics to SteelFlow Net Collectors.

3. Select the Add a New Flow Collector tab, specify the Collector IP address and port on the NetProfiler, and select the appropriate interfaces.

To configure REST API access

1. On the SteelHead, choose Configure > Security > REST API Access.
2. Select Enable REST API Access and click **Apply**.
3. Select the Add Access Code tab.
4. Specify a useful name in the Description of Use field.
5. Select Generate New Access Code and click Add.
A new code is generated.
6. Expand the new entry and copy the access code.
7. Save the copied access code for the following procedure.

If enabling REST API on multiple SteelHeads, you can copy and paste this code into other SteelHeads. You must perform these steps on every SteelHead on which you want to receive application Flow statistics and share the QoS configuration with NetProfiler.

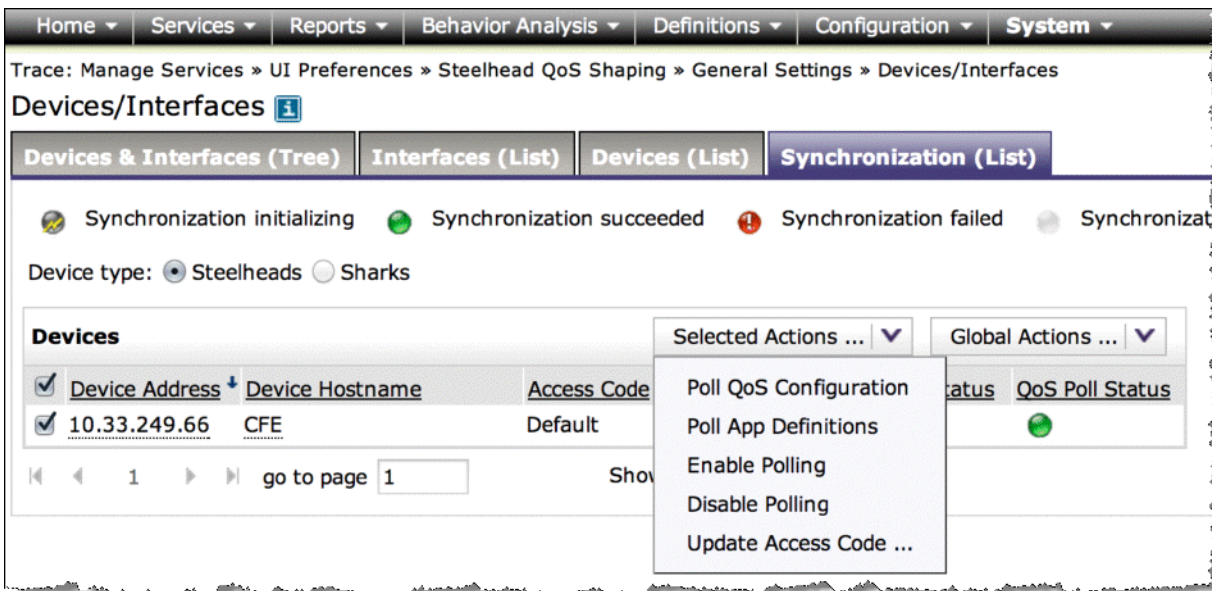
The rest of the configuration is performed on the NetProfiler.

To configure the NetProfiler

1. On the NetProfiler, choose Configuration > General Settings and Data Sources.
2. Select Use NetFlow/IPFIX and specify the port you entered on the SteelHead.
3. Choose System > Devices/Interfaces.
4. Select Synchronization (List) tab.
5. Select the Steelheads radio button and select the SteelHead in the Devices list.

Depending on network conditions and traffic flows, the SteelHead might not immediately appear in the device list.

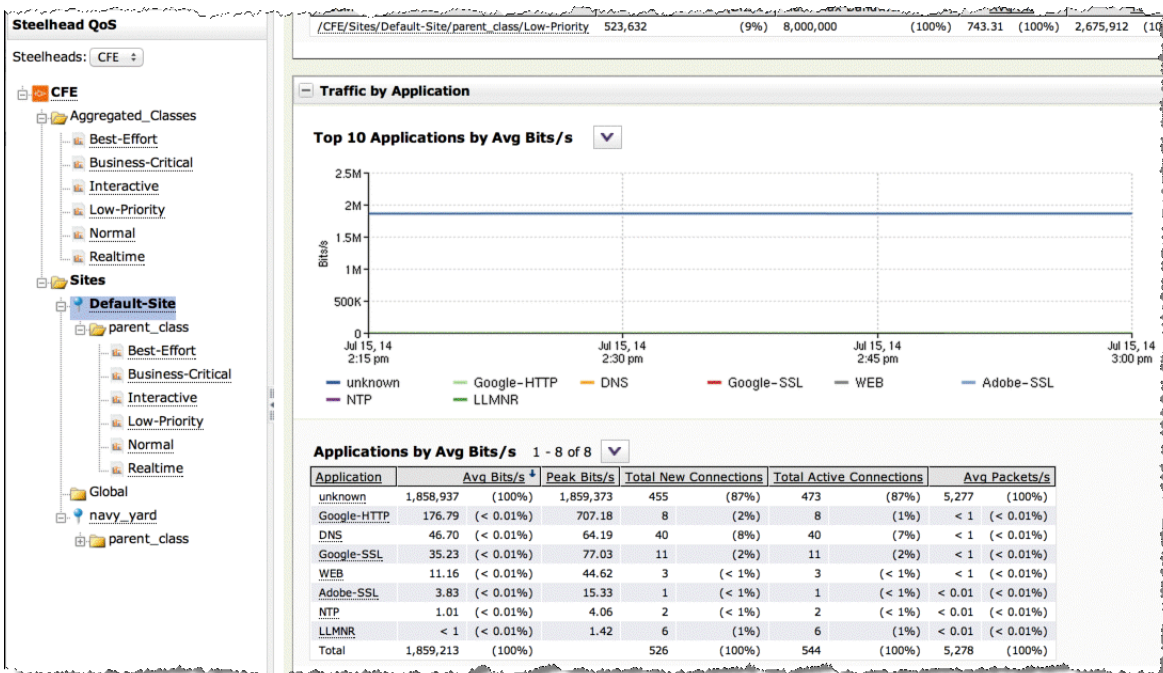
Figure 5-5. Devices/Interfaces Page



6. From the Selected Actions drop-down list, select Update Access Code.
7. In the Access Code window, specify the same access code you previously used in the SteelHead configuration step and click OK.
8. From the Devices/Interfaces page, make sure the SteelHead is selected. From the Selected Actions drop-down list, select Enable Polling
The NetProfiler starts polling the SteelHead QoS configuration and generates reports according to your configuration.
9. To verify the operation, Choose on Reports > SteelHead QoS Shaping.

10. Expand the SteelHead to view the polled QoS details.

Figure 5-6. QoS Details



SteelHead Flow Integration

Because you can deploy SteelHeads in multiple places in your network, including small and large branch offices, they are ideal collection points for data on network and application performance. The ability to leverage the SteelHead to forward data to a Flow Gateway collector and then on to a NetProfiler enables you to increase the number of locations data is collected from without having to deploy additional collectors around the network.

The data collected from SteelHeads provides significantly more information than what is generally available from routers and switches. SteelFlow Net leverages the extensibility of NetFlow v9 to include significant additional information. The additional information available from SteelFlow Net includes Layer-7 application information, packet and byte retransmit counts, Round Trip Time, and more.

For more information about SteelFlow Net, see [“SteelHead and SteelCentral Overview” on page 81](#).

Keep the following in mind when you integrate SteelHead with NetProfiler:

1. Set the active flow time-out to 60 seconds.

You can set the active flow time-out to less than 60 seconds; however, that results in additional record updates being sent to the collector result in additional load on the SteelHead, Flow Gateway, and network. There is no benefit to sending updates more frequently than once per minute because the information is consolidated and forwarded to the NetProfiler once per minute.

2. The inactive flow time-out must be set to a value less than 60 seconds. Riverbed recommends an inactive flow time-out of 15 seconds.

The inactive flow time-out sends a final update for flows that have stopped transmitting data but have not sent a FIN or RST ending the conversation. This time-out is suggested because this will remove the flow from the flow tracking table on the SteelHead. Because this table has a limited number of entries freeing up entries from a completed flow ensures there is sufficient room for active flows.

3. The flow format should be SteelFlow Net.

There are four formats supported: NetFlow v5, NetFlow v9, SteelFlow Net-compatible, and SteelFlow Net. SteelFlow Net is the most recent flow type and provides the most complete support for the most metrics. Choosing a different flow format continues to provide data; however, the amount of information forwarded from the SteelHead is more limited than what is available with SteelFlow Net.

CHAPTER 6 NetProfiler and SteelCentral Portal Integration

The SteelCentral Portal provides a single place to take data from multiple application performance management and network performance management data sources and present that view on a single screen. The NetProfiler is one of a variety of data sources for the SteelCentral Portal (other sources include but are not limited to SteelCentral AppInternals and AppResponse). The data from NetProfiler is presented in conjunction with data from other sources. Intelligent drill-down menus enable easy access to underlying and more detailed data available from different sources. Because the SteelCentral Portal uses built-in intelligence to ensure the correct data source is associated with the correct data, you can easily pivot from data from one source (such as the NetProfiler) to a view from a different source (such as AppResponse).

The following caveats apply when using the NetProfiler as a data source for dashboards. These caveats include limits on what information is presented and performance related issues:

- Currently, you cannot display all data available in the NetProfiler in a SteelCentral Portal view. For example, Service Maps, NetProfiler, the SteelCentral Portal widgets, WAN Optimization Reports, and user information are not available directly on the SteelCentral Portal. However, this information is available after you have pivoted from the SteelCentral Portal to a report directly in the NetProfiler.
- Riverbed recommends that as a best practice, you keep the data refresh rates to the lowest practical value possible to enable the SteelCentral Portal integration to have the least impact on the NetProfiler performance. The SteelCentral Portal uses REST API calls to access information about the NetProfiler. Because the information is retrieved in the same manner as any other report that has a significant number of information with rapid refresh rates, this can have an impact on the NetProfiler overall performance.

The closer the SteelCentral Portal server is physically located to the NetProfiler, the less impact network latency has on the overall performance of the integration. Having significant latency between the SteelCentral Portal server and NetProfiler can introduce performance issues or prevent proper and expedient updating of the SteelCentral Portal pages.

You can configure all or one of your NetProfilers as data sources for the SteelCentral Portal. There is no support for configuring the same NetProfiler multiple times as this has an adverse impact on the performance of the NetProfiler.

To minimize the chances of performance issues and maximize the available functionality between the NetProfiler and the SteelCentral Portal, Riverbed recommends that you use the most recent versions available.

CHAPTER 7 Additional SteelCentral Integration

SteelCentral includes a number of additional integrations that enable you to complete your deployment by evaluating additional data or integrating with other management and reporting systems. This chapter includes information about the most commonly used SteelCentral integrations:

- [“SNMP Integration” on page 97](#)
- [“SNMP for Switch Port Discovery” on page 99](#)
- [“Switch Port Discovery Supported Routers and Switches” on page 99](#)
- [“Active Directory” on page 99](#)
- [“REST API” on page 101](#)

For additional assistance, contact Riverbed Professional Services.

SNMP Integration

This section describes the following SNMP integrations. It includes the following sections:

- [“SNMP Integration for Flow Sources” on page 97](#)
- [“SNMP Integration for Device Management of SteelCentral Components” on page 98](#)
- [“SNMP Integration for Sending Traps” on page 98](#)

SNMP Integration for Flow Sources

When devices send data flow, standard SNMP interface identifiers (ifindex values) indicate which interfaces the data flow traverses. You must map the interface identifiers to names and descriptions to identify them on the NetProfiler. You must also obtain the link speed information so you can convert raw bandwidth numbers to link utilization percentages. The NetProfiler, NetExpress, and Flow Gateway gather the following information using standard SNMP from all devices sending standard flow or SteelFlow Net:

- Device name
- Interface names
- Interface descriptions
- Interface capacities

Ensure that you configure firewalls between the NetExpress or Flow Gateway and flow-reporting devices to enable SNMP access between the NetExpress or Flow Gateway and remote device. If there are any access rules on the flow-reporting devices, you must enable these access lists to allow SNMP access from the NetExpress or Flow Gateway.

For more information about configuring of the NetProfiler and NetExpress for SNMP collection of these items, see the *SteelCentral NetProfiler and NetExpress User's Guide*.

SNMP Integration for Device Management of SteelCentral Components

You can monitor the status of SteelCentral through SNMP from an external SNMP manager. Currently, the NetProfiler, Flow Gateway, and NetExpress generally support the industry-standard UCD-SNMP-MIB. For more information, see <http://www.oidview.com/mibs/2021/UCD-SNMP-MIB.html>.

When you use SNMP, it is normal to detect high CPU and memory use. This does not mean that the appliance is experiencing a problem. Because SteelCentral are appliances and not standard servers, processes tend to hold the entire CPU for normal use (100 percent CPU utilization is normal) and make efficient use of available memory resources.

For the Enterprise NetProfiler, you can poll each physical component separately.

Because the NetProfiler reports a broken connection with a Flow Gateway or NetShark, a best practice is to configure your SNMP manager to send health traps and only poll the NetProfiler Event Manager module.

SNMP Integration for Sending Traps

The NetProfiler can send traps through SNMPv1 or SNMPv3 to third-party trap receivers. You can customize which types of traps to send to which devices within the notifications pages of the NetProfiler UI. Some of the use cases for sending SNMP traps are as follows:

- Sending the NetProfiler or NetExpress health messages to a third-party network manager or SNMP device manager
- Sending service and performance and availability events to a third-party network manager
- Sending security events to a security event manager (SEM)

You must configure the third-party device receiving the trap with the NetProfiler MIB (labeled Mazu-MIB). This MIB is available from either the Riverbed Support site or NetProfiler help downloads page.

You can route the appropriate events to the appropriate devices by first configuring recipient groups within the NetProfiler and then configure which events are sent to which recipients. Recipient groups can contain email recipients and SNMPv1 or SNMPv3 recipients. For more information about how to configure these notifications, see the *SteelCentral NetProfiler and NetExpress User's Guide*.

SNMP for Switch Port Discovery

Standard flow records identify hosts by their IP addresses. The NetProfiler supports discovery of MAC addresses and switch port locations of individual hosts based upon their IP addresses. This enables the NetProfiler to display complete information, as shown in [Figure 7-1](#). This information appears in multiple places throughout the NetProfiler UI.

Figure 7-1. MAC and Host Switch Information Displayed as a Result of Switch Port Discovery



Switch Port Discovery Supported Routers and Switches

Switch Port Discovery is based on the NetDisco open source library version 2. For further details on supported devices, see the appropriate NetDisco documentation available on the NetDisco development website.

Active Directory

The NetProfiler provides a user identity feature that maps active directory (AD) usernames with IP addresses. This feature enables you to view:

- the users associated with an end station on the network ([Figure 7-2](#)).

- all the end stations that a user has logged into (Figure 7-3).

Figure 7-2. Users Logged into a Given Host for a Selected Time Period

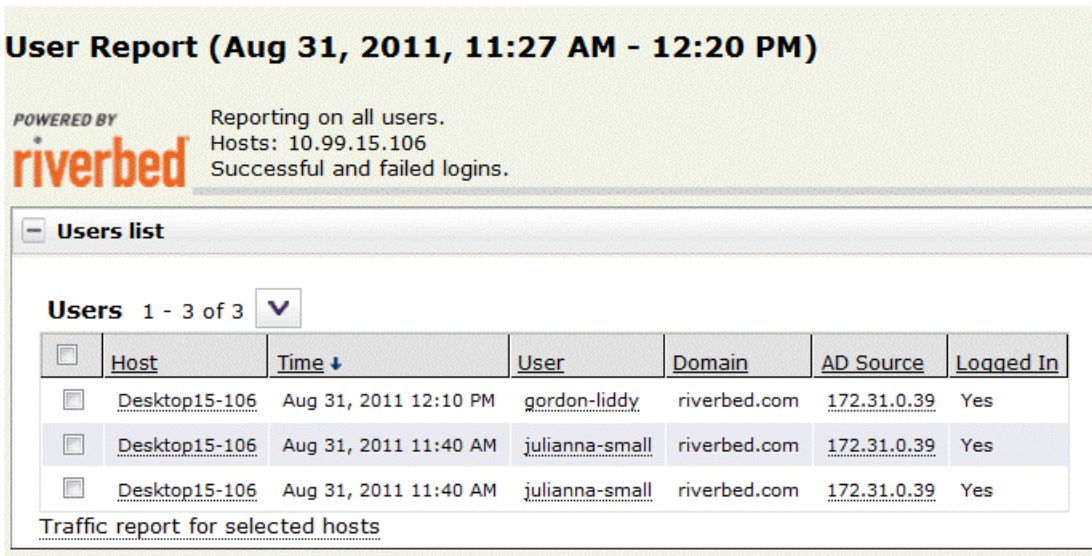
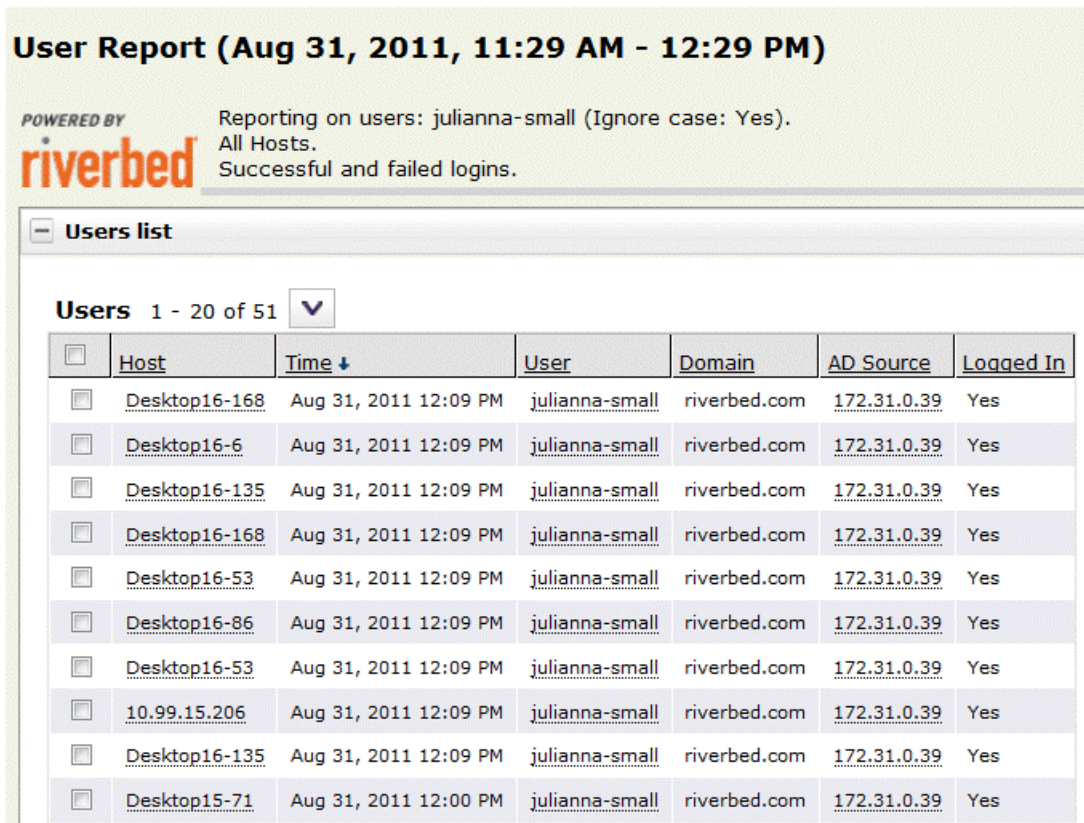


Figure 7-3. Hosts That a Given User Logged into for a Selected Time Period



This feature relies on the security audit events obtained from one or more Microsoft active directory domain controllers. You can send this event data directly to the NetProfiler from a domain controller, or for AD-2008, an event collector host. Riverbed provides a service application named SteelCentral AD Connector that forwards the appropriate events from a domain controller, or event collector, to the NetProfiler.

Integration for Active Directory 2008

For AD-2008, you must use the SteelCentral AD Connector 2.0. You can install the connector on either the domain controllers or an event collector, but Riverbed recommends that you install the AD Connector on the event collector. Even though installation on a domain controller is easier, you must install it as many times as the number of domain controllers. The installation on an event collector requires a few more steps, including planning of event-collecting topology (if not already implemented in the environment), but it requires no additional product installed on domain controllers and provides more flexible delivery paths.

For more information about configuring integration with AD-2008, see the documentation provided on the Riverbed Support site.

You can download the connector and the document from either the Riverbed support site or directly from the NetProfiler help downloads page.

Integration for Active Directory 2003

For AD-2003, you must use the SteelCentral AD Connector 1.5. You can install the connector on the domain controllers or another Windows server acting as a collector within the same domain controller, but Riverbed recommends that you install the connector on the domain controller. Installing the connector directly on the domain controller requires no messaging between the domain controllers and an external collector, whereas if you use the external collector, you need significant inter-system communications.

For more information about configuration integration with AD-2000 and AD-2003 environments, see Technical Note #29: Microsoft AD Integration for User Identity.

You can download the connector and the document from Riverbed Support or directly from the NetProfiler help downloads page.

REST API

SteelCentral 10.0.5 and later include a REST API for the NetShark and NetProfiler. Using REST API, you can perform such activities as:

- packet captures on NetShark
- traffic reports on NetProfiler
- WAN optimization reports on NetProfiler
- appliance configuration

You can see more information about specific functions of the REST API directly on your NetProfiler at <https://{NetProfiler}/rest/api/index.php>, and the following links on the Riverbed Support site:

- <https://support.riverbed.com/apis/>
- <https://support.riverbed.com/apis/steelscript/>

CHAPTER 8 NetProfiler Analytics and Service Monitoring

SteelCentral service monitoring simplifies discovering, modeling, and configuring monitoring for enterprise applications. This chapter describes how you can account for the analytic license limit when mapping services. This chapter also provides best practices for how to stay within these limits and which specific metrics to use.

This chapter includes the following sections:

- [“Analytic License Limits per NetProfiler Platform” on page 103](#)
- [“Understanding the Analytics License Limits” on page 104](#)
- [“Reducing the Number of Metrics” on page 105](#)
- [“Conditions Required for Baseline Establishment” on page 107](#)
- [“Determining Which Metrics to Use” on page 107](#)

This chapter assumes you are familiar with the NetProfiler and NetExpress user documentation on the Riverbed Support site.

Analytic License Limits per NetProfiler Platform

If you are running the NetProfiler 9.0.5 or later with a deployed analytics license, the following table shows the number of analytics policy metrics allowed system wide.

NetProfiler Model	Available Analytic Policy Metrics	Supported Locations
NetExpress	Up to 5000	1,800
Standard NetProfiler	Up to 7500	1,800
Enterprise NetProfiler	Up to 100,000	1,800

The available analytic policy metrics number includes the total number of analytic policy metrics to be tracked system wide. It includes the following:

- All metrics tracked through service configuration, up to eight metrics per segment monitored
- All metrics tracked through performance and availability policy configuration:
 - Link congestion, up to two metrics per policy configured
 - Link availability, up to two metrics per policy configured

- Application performance, up to seven metrics per policy configured
- Application availability, up to two metrics per policy configured

The following policy types do not count toward the analytics limit:

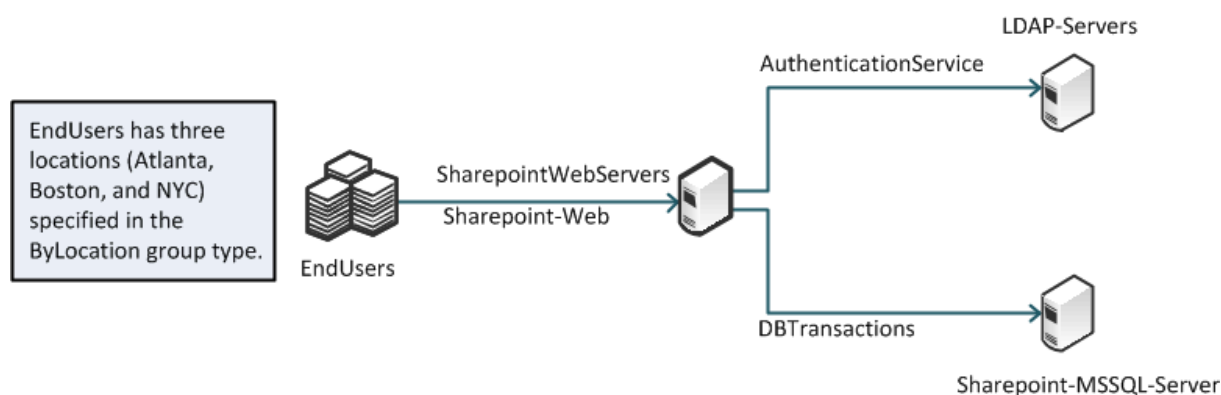
- Security policies
- User-defined policies

Understanding the Analytics License Limits

You can configure hundreds of metrics on the NetProfiler by clicking a few checkboxes. If you understand the concept behind the check boxes, you can make intelligent decisions about how to configure services for optimal performance.

Figure 8-1 shows a simple SharePoint service configuration. The SharePoint web servers are connected to three end-user locations (Atlanta, Boston, and New York), and they are supported by LDAP and MS-SQL on the back end.

Figure 8-1. Simple SharePoint Service Configuration



Using the example shown in Figure 8-1, five metrics are selected (Figure 8-2) for each of the segments during service discovery for the SharePoint service.

Figure 8-2. Five Metrics for SharePoint Service

Metrics Per Segment						
Segment	Conn BW	Conn New Conns	Effncy # TCP Rsts	Effncy TCP Retrans	BW	UserExp Rsp Time
DBTransactions	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
AuthenticationService	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>
Sharepoint-Web	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		<input checked="" type="checkbox"/>

For the backend segments, DBTransactions and AuthenticationService, there are a total of 10 metrics used (five metrics times two segments). However, for SharePoint-Web, there are three end-user locations. Because it is necessary to monitor each location individually, this yields a total of 15 metrics (5 metrics x 3 locations). The total number of metrics used to monitor the SharePoint service is 25.

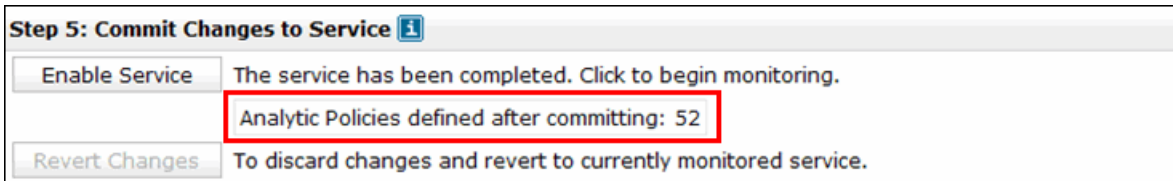
A location with 100 end users is more representative of a larger enterprise network. If there are 100 end-user locations, the SharePoint service requires more than 500 metrics (510 metrics to be exact) for monitoring.

You can view the number of metrics on a running system:

- during the process of service discovery using the wizard.

On the Commit Changes to Service page of the wizard, you have to commit changes to the service before the analytics are configured. The page shows a count of how many metrics to be added when you save the service.

Figure 8-3. Metrics Shown on the Commit Changes to Service Page




The Commit Changes to Service Page only shows how many metrics are added when the service is committed. It does not display the total used or how many are still available.

- on the Information page.

The System > Information page includes the Policies Usage display, which shows the total number of metrics currently in use and the overall limit for the system. The difference between the number of metrics currently used and the limit (5000, 7500, or 10,000, depending upon the platform) is the number remaining. The remaining number is available for you to configure.

Figure 8-4. Metrics Shown on the Information Page



Reducing the Number of Metrics

You can reduce the number of metrics used with services in the NetProfiler in the following ways:

- **Reducing locations** - You can reduce the number of locations you monitor to reduce the number of metrics in use. The following methods reduce the number of locations:
 - You can reduce the number of groups in the group type used for identifying end-user locations, for a reduction across all services.

- You can reduce the number of locations you use for a specific service. [Figure 8-5](#) shows how to reduce the number of locations during discovery for any service by choosing a subset of the end-user locations to monitor.

Figure 8-5. Reducing the Number of Locations on the Customize Monitoring on the Locations Page

Step 4: Customize Monitoring -- Locations ⓘ

Segments that face end-users can be broken out by location. Select which locations you want to monitor for each front end segment.

Front End Segments	Locations
Web	<input type="radio"/> All <input checked="" type="radio"/> Selected (3 out of 10) done

Austin
 Cambridge
 Columbus
 DataCenter
 Hartford
 LosAngeles
 Philadelphia
 Phoenix
 SanFrancisco
 Seattle

[Save & Next>](#)

Both methods cause a reduction in the number of configured policies.

- **Not tracking by groups** - You can turn off by-group tracking and track all front-end metrics as one item instead of a detailed breakdown by groups. This might make sense if all sites or groups have similar characteristics (for example, response time) when attaching to the front end of the application. To turn off by-group tracking, edit and clear the Track By End-User checkbox on the front-end end-user component.

You continue to receive a list of clients that have an impact if an analytic has an error.

- **Reducing monitoring** - You can reduce the number of metrics being monitored by eliminating a metric per end-user location and a metric for each back-end segment in the service. [Figure 8-6](#) shows four metrics configured by default and four additional metrics you can configure. Turning off only one of these metrics, particularly if there is a large number of end-user locations, can result in a large reduction in metrics you use.

Figure 8-6. Default and Configurable Metrics on the Customize Monitoring on the Metrics Page

Step 4: Customize Monitoring -- Metrics ⓘ

Determine which metrics you want to have collected

[Add Metrics](#)

Metrics Per Segment									
Segment	Conn Active	Conn BW	Conn New Conns	Effncy # TCP Rsts	Effncy TCP Retrans	UserExp Avg App Thr/Conn	UserExp Avg Conn Dur	UserExp Rsp Time	
Web	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	

[Save & Next>](#)

As mentioned earlier, you can aggregate locations into regions. You do not reduce the number of metrics you have, but you can more easily manage the Services Dashboard. If you have hundreds of locations in the ByLocation group, you can create a ByRegion group type instead. Specify the group type to represent the end-user locations on the General Settings page. The default setting specifies the ByLocation group type.

Conditions Required for Baseline Establishment

A segment must collect a certain amount of data before it can establish daily and weekly baselines. The daily baseline requires three days and one hour of collected data. Until the system has been running and collecting data for a segment for this period of time, the analytic metrics related to the segment metrics remains in an initializing state.

The weekly baseline requires three weeks and one day of data. Additionally, for each metric to initialize, there must be some data for the metric in 50 percent of each 15-minute time period. This means that if the segment is a backup that runs only once per day or is active for only a few hours a day, the segment does not initialize.

Determining Which Metrics to Use

When you configure service monitoring within the NetProfiler, Step 4 of the configuration wizard enables you to decide which metrics to monitor for each segment within the service. The metrics are organized according to three different categories. The following table summarizes all metrics.

Category	Metric Name	Enabled	Dips	Spikes
Connectivity (Conn)	New connections			Yes
	Active connections	Yes		
	Bandwidth		Yes	
User experience (UserExp)	Response time	Yes		Yes
	Average connection duration			Yes
	Average application throughput per connection		Yes	
Efficiency (Effncy)	TCP retransmissions	Yes		Yes
	Number of TCP resets	Yes		Yes

The following metrics are enabled by default:

- Active connection rate (detecting spikes)
- Response time (detecting spikes)
- TCP resets (detecting spikes)
- TCP retransmissions (detecting spikes)

When you consider which metrics to use for an application, take into account what each segment is responsive for versus the service as a whole. Front-end segments might have very different characteristics than back-end segments.

For example, if you have a web-connected front end, you might detect numerous brief connections versus a back-end segment for the same service, which might have only continuous database interactions over only a handful of connections. For another service, you might have a front-end segment that uses Citrix, which might keep connections open throughout very long periods of time, while back-end connections to application servers might be shorter in duration but greater in number. For details about characteristics to consider per metric, see [“Reducing the Number of Metrics” on page 105](#).

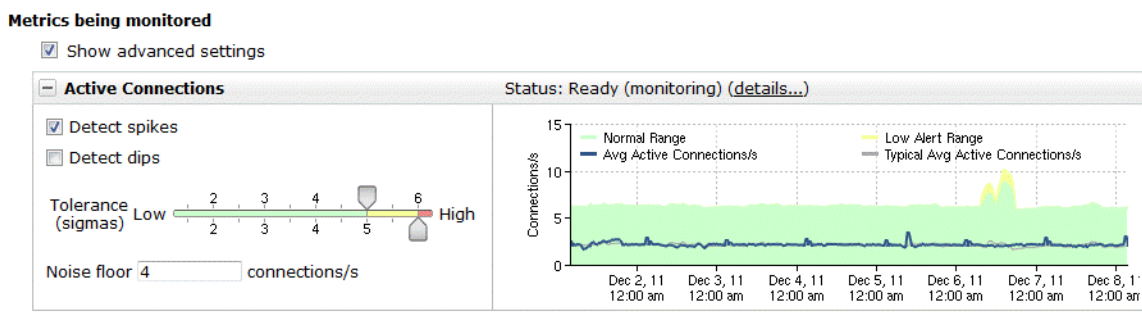
The four default, enabled metrics satisfy a majority of TCP-based application segments, although for segments with a low number of connections, you might want to disable or change the settings on the active connections metric.

When you choose which metrics to include, use the following best practice guidelines:

- **Applications with low connectivity rates** - For back-end segments for which applications have connectivity between only a few servers, or front-end segments for which only a few clients are connected at a time, the active connection rate is very low, and the tolerance band might be very tight. You might detect alerts when there is a very minor change in connectivity (one new session connects longer than what is normal).

For these situations, you can disable this metric, or you can increase the tolerance band and add an appropriate noise floor to the metric. The noise floor can help control minor fluctuations. [Figure 8-7](#) shows a segment that has only a few connections active per second, with a raised tolerance to 5 for low and 6 for high, and an added noise floor of four connections per second.

Figure 8-7. Metric for a Low Connectivity Rate Segment



- **Active connection rate metric consideration without weekly seasonality** - If you are trying to keep the number of alerts low, Riverbed recommends that you not disable the active connections metric until after the weekly baseline is set. The baseline is three weeks and one day of data.
- **UDP applications** - For UDP applications, the TCP health and TCP performance measurement-based metrics do not work. You can disable TCP resets, TCP retransmissions, and response time. For UDP segments that have periodic bandwidth, you can enable the bandwidth metric.
- **Back-end segments with continuous communications** - For many back-end segments, you can enable the average application throughput per connection metric. This metric tracks the bandwidth that is consumed during the active parts of the session. You are alerted when the baselined value dips below the threshold. This dip can indicate that less data is transferred, which can indicate that the application efficiency has dropped, and this can have an impact on user experience.
- **Single-transaction-oriented TCP sessions** - For application segments that tend to set up a new TCP session for each transaction, you can enable the average connection duration metric. This metric tracks the duration of the connections and alerts you if it dips below that baseline. For this type of segment, tracking new connections in addition to active connections can also be beneficial.
- **Revisit metrics and tuning after three weeks of data** - Although three days and one hour of data are required for the analytic metrics to initialize, it takes three weeks and one day for the analytics to begin using a weekly baseline. This baseline becomes more predictable when you monitor weekly seasonality is monitored (for example, lower traffic volumes on the weekend). Tuning and final decisions on which metrics might not be best for the segment are made after this time period.

- **Understanding the characteristics of your application** - To better understand the characteristics of the segments on your application, you can run service-level-objective (SLO) reports after the segments have initialized. The SLO reports enable you to see the baselined periodicity of each metric. If the segment has not yet initialized, you can run reports to gain a better understanding of the segment characteristics. Running reports in this manner helps you to choose which metrics to use per segment and to fine-tune after initialization.

CHAPTER 9 Troubleshooting the NetProfiler

Troubleshooting SteelCentral can be a complex process involving looking at multiple different areas of the product.

In its simplest form, the NetProfiler is divided into two distinct areas: the web-based UI and the supporting infrastructure (system processes, scripts, and so on). Problems can occur in one area only or across multiple areas. For example, identifying why specific data is missing from a report might reveal a problem with the way the report is being run (a UI-based problem), a problem with the underlying query processing engine (an infrastructure-related problem), or a combination of the two.

This chapter includes troubleshooting information in the following areas:

- [“RTT Values Not Available” on page 111](#)
- [“Not Receiving Reports by Email” on page 112](#)
- [“DNS Names Not Being Resolved in Reports” on page 112](#)
- [“Reports Are Not DNS-Resolving All Addresses” on page 113](#)
- [“Data in Reports Seems Inconsistent” on page 113](#)
- [“Sensor Protocol Violations” on page 114](#)
- [“Communication Issues” on page 114](#)
- [“Switch Port Discovery Troubleshooting” on page 115](#)

RTT Values Not Available

When you run a report with RTT columns, sometimes the included rows do not show data. Some of the reasons data might not be available include the following:

- **Non-TCP flow** - The NetProfiler calculates RTT information only for TCP-based flows. Any flow that is not TCP-based (ICMP, UDP, and so on) does not display RTT information.
- **Flow not seen by the NetShark or NetExpress** - The flow must be detected by a NetShark or NetExpress to calculate RTT information. If a flow is reported only from a network router through NetFlow, there is not any RTT information available.

If the flow is only partially seen (for example, due to asymmetric routing) by the NetShark, or NetExpress, the RTT information is not valid and is discarded.

- **Retransmits during the initial flow setup** - The NetProfiler does not show RTT information if there is retransmitted information during the initial flow setup. If the NetProfiler does not know which packet is the correct packet to use, there is no way to accurately calculate RTT information.
- **Delay between the packets** - If the delay between the SYN, SYN-ACK, ACK, and DATA packets is too great (multiple minutes), the RTT timer might expire and RTT information is not calculated.
- **Drops questionable values** - The NetProfiler takes a conservative approach and drops values that are questionable.
- **Flows that start before the initial startup time** - Because RTT information is calculated only during the initial flow setup, any flows that started before the beginning of the report time frame do not include RTT information: for example, a report from 14:00:00-15:00:00 that includes flows that started prior to 14:00:00 do not show RTT information.

Not Receiving Reports by Email

Information and errors related to the emailing of reports are stored in the NetProfiler audit log. Anytime an email is sent, an entry is placed in the audit log, as are any error messages that are received during the transmission process.

When you configure the SMTP server you must use an appropriate host and required account information. You can choose to use a username and password for authentication. Ensure that you configure the SMTP server to enable connections from the NetProfiler and that it is able to relay messages to the appropriate destinations.

Note: The NetProfiler does not currently support encrypted SMTP.

DNS Names Not Being Resolved in Reports

You must complete the following steps for the NetProfiler to resolve IP addresses to DNS names:

1. Configure the DNS servers in the General Settings page.
2. In its UI Preferences settings, you must configure each user account to resolve IP addresses to DNS names.

The following are options for DNS resolution in UI Preferences:

- **Resolve hostnames using DNS** - Turns on the use of DNS to resolve an IP address into a host name.
- **Resolve hostnames for hosts managed by DHCP** - Uses the optional SteelCentral DHCP integration to use the information provided during DHCP imports for name resolution.
- **Suppress DHCP/DNS search domains from resolved hostnames** - Suppresses the display of the listed search domains (for example, riverbed.com) when showing names.

Reports Are Not DNS-Resolving All Addresses

To prevent the NetProfiler from overwhelming DNS servers, you can place the following limits on resolving hosts. You can control the:

- number of DNS resolution requests the NetProfiler sends to a DNS server at one time.
- maximum number of addresses the NetProfiler attempts to resolve addresses for.

Riverbed recommends that you set these values to 500 each to prevent overwhelming the DNS server. This means that the NetProfiler does not attempt to resolve more than 500 rows from any one table and sends up to 500 requests at one time. You can increase these values to allow more addresses to be processed.

Be aware that the NetProfiler waits only one second for responses from DNS servers, to prevent slow DNS servers from delaying the return of reports. Any addresses that are not resolved before the time expires do not display the associated DNS name.

Data in Reports Seems Inconsistent

This issue is one of the hardest issues to troubleshoot. Consider the following possible causes:

- **Mismatched report types** - If you run a host-centric NetProfiler report versus an interface-based report.

A host-centric report is any report you run from the Reports > Traffic Reports page, and select the Host, Application, or Advanced tab. These reports always query the database based upon the perspective of the hosts in the hosts field. When you run these queries, you look for all hosts matching the query conditions, and all output is from the perspective of the hosts.

An interface-centric report is any report you run from the Reports > Traffic Reports page, and select the Interface tab. These reports always run from the perspective of the interfaces in the interface field. When you run these queries, you are querying for all interfaces matching the query conditions, and all output is from the perspective of the interfaces.

- **Missing data** - The primary causes of missing data are as follows:
 - **No coverage of the desired data** - With a large network, you can have pockets of data that are not covered by devices reporting to the NetProfiler: for example, a branch office might not have a device reporting traffic internal to the branch. You cannot report on traffic for which there is no coverage.
 - **Too many devices reporting data** - The NetProfiler currently is limited to storing data from a five devices. If flow is reported from more than five devices, the data is not retained. This results in the reported traffic rates for those devices being inaccurately low.
 - **Missing directional data (ingress or egress)** - When you export NetFlow, depending on the version and device type, you might not be able to export both ingress and egress data for each interface. Consider a common example in which only ingress data is received with NetFlow v5. When the flow records are sent, the egress interface is indicated in the record, even though the statistics are counted on an ingress interface. When NetProfiler receives these ingress records, it assumes that the data is preserved as it passes through the device. This means that if the record is received on Interface 1, and the record indicates the egress interface is Interface 2, NetProfiler assumes that this amount of data leaves the device on Interface 2.

In some cases where this assumption is not valid: for example, on a router with a 10-Mbps interface on one side, and a 1.5-Mbps interface on the other side, and the router is forced to dropped data due to the 1.5-Mbps interface being oversubscribed. Because the NetProfiler has no way of knowing how much data actually went out the other interface, the numbers can be incorrect. If the 10-Mbps interface is receiving 5 Mbps, the NetProfiler reports 5 Mbps leaving the 1.5Mbps-interface because this is the best data NetProfiler received.

- **Routing issues** - The NetProfiler must detect all the details of each flow that it reports on to provide accurate information. When you have asymmetric routing (where traffic takes one path from client-to-server and a different path from server-to-client), the NetProfiler can miss one side of the conversation. This results in inaccurate information from reports.
- **Different data resolution** - While the NetProfiler provides several different data resolutions, many other devices do not provide multiple resolutions or do not allow detailed control over which resolution you use. Comparing a report from two different resolutions (for example, one hour on the NetProfiler and five minutes on a router) is very likely to result in differences in reported values.
- **Reporting device settings** - Reports are not correct when the NetFlow export settings on the reporting device (for example, the router) are not as per the Riverbed recommendation.

Sensor Protocol Violations

The Sensor protocol violation (SPV) error is one of the more common errors that you can receive when using the NetProfiler. The error appears as a system event indicating that the violation occurred between the NetProfiler and one or more reporting devices (Flow Gateway or NetShark). Two potential causes of the SPV error are slow transfers and lack of time synchronization:

- **Slow transfers** - Because the NetProfiler must receive data from remote devices and analyze it all within one 60-second period, there is very little leeway for slow transfers. The NetProfiler allows 48 seconds in each minute for all remote devices to send their data (transfers happen in parallel). Any data sent to the NetProfiler, any data after the 48 seconds is ignored and SPV errors are generated.

If you detect SPV errors, send a file transfer between the NetProfiler and remote device reported in the SPV and calculate how fast data is transferring. If the transfer rate is extremely low and the number of flows is high, this is the likely issue.

If the issue is time sensitive—it only occurs when backups are also running—you must perform the test around the same time that the issue occurred.

- **Lack of time synchronization** - If one NetShark or Flow Gateway is unable to NTP synchronize with the NetProfiler, the time on that device can drift sufficiently from the NetProfiler time. The NetProfiler then has difficulty ensuring that the data being sent is for the same time slice the NetProfiler is currently processing. You must ensure that no firewall or ACL are blocking NTP.

Communication Issues

NetProfilers communicate with each other on select ports. These ports must be open between the NetProfiler and the remote device for all functions to work correctly. The following primary ports are used for communications among devices:

- **TCP/41017** - Used to pass data back and forth between the NetProfiler and remote devices such as the NetShark and Flow Gateway. You must open this port bidirectionally between devices.

- **UDP/123** - Used for NTP synchronization between the NetProfiler and the remote devices such as the NetShark and Flow Gateway. The NetProfiler acts as the NTP server for the remote devices.

If anything is blocking communications on the specified ports, that portion of the NetProfiler system does not work correctly. For example, if UDP/123 (NTP) is blocked between the NetProfiler and NetShark, the time on the NetShark is likely to drift, resulting in inaccurate reports.

To ensure that a port is open, use the following telnet and CLI commands:

```
[mazu@cascade-gateway etc]$ telnet 10.38.7.8 41017
Trying 10.38.7.8...
Connected to 10.38.7.8.
Escape character is '^]'.
^]
telnet> quit
Connection closed.
```

In this example, the connection was successful, and the CLI output shows port TCP/41017 is open between the host and 10.38.7.8. Had the port been closed, the conversation might have looked like this:

```
[mazu@cascade-gateway etc]$ telnet 10.38.7.2 41017
Trying 10.38.7.2...
telnet: connect to address 10.38.7.2: connection refused
```

The connection was rejected by 10.38.7.2 on port TCP/41017.

Switch Port Discovery Troubleshooting

If device polling fails, ensure that rules on the device allow polling from the NetProfiler or NetExpress. Also check the device is in the list of supported devices in [“Switch Port Discovery Supported Routers and Switches” on page 99](#). If polling still fails, you can enter the following CLI commands to verify SNMP communication and support.

For a lookup router, enter the command:

```
NetProfiler# snmpinfo <router-ip-address> --dev --fw --version <1|2> --comm <read-only-
community-string>
```

For a switch, enter the command:

```
NetProfiler# snmpinfo router-ip-address> --dev --fw --version <1|2> --comm <read-only-
community-string>
```

Command output includes information about whether or not the device is reachable and supported. You can give Riverbed Support the output of the command for additional information.

For more details, see [“Additional SteelCentral Integration” on page 97](#).

APPENDIX A Licensing

This appendix explains various aspects of licensing the SteelCentral. It includes the following sections:

- [“Licensing Overview” on page 117](#)
- [“License Installation” on page 118](#)
- [“Packet Analyzer Licensing” on page 119](#)
- [“Assigning Licenses” on page 119](#)
- [“Manual License Installation” on page 120](#)
- [“Automatic License Upgrades” on page 120](#)
- [“Evaluation Licenses” on page 120](#)
- [“Licenses Available” on page 120](#)

Note: Flow-per-minute limits changed in February 2013. For information about the SteelCentral prior to February 2013, consult documentation for that software release.

Licensing Overview

The NetProfiler and Flow Gateway systems have the following types of licenses:

- [“Runtime Licenses” on page 117](#)
- [“Capacity Licenses” on page 118](#)
- [“Option Licenses” on page 118](#)

Runtime Licenses

All devices require a runtime license. Runtime licenses permit basic system operation. If you do not have a runtime license installed, your system operates in a basic mode, which enables basic UI interaction, but does not process new flow data.

Capacity Licenses

Capacity licenses control the maximum amount of data the device can process. Capacities are set based on the maximum number of flows that the device can process each minute (for example, an F1 Flow Gateway can accept no more than 150,000 fpm).

Different types of SteelCentral devices have different flow capacities:

- **Flow Gateway** - 150,000, 400,000, 600,000, 800,000, 1,000,000, and 2,000,000 fpm
- **Flow Gateway-v** - 15,000, 30,000, 60,000, 90,000, 200,000, 400, 0000, 600,000, 800,000, 1,000,000, 2,000,000 fpm
- **NetExpress 470** - 15,000, 30,000, 60,000, 90,000, and 120,000 fpm
- **NetProfiler** - 200,000, 400,000, 600,000, 800,000, and 1,000,000 fpm
- **NetProfiler-v** - 15,000, 30,000, 60,000, 90,000, 200,000, 400, 0000, 600,000, 800,000, 1,000,000, and 2,000,000 fpm
- **Enterprise NetProfiler cluster** - 1,000,000 or more, depending on the number of Analyzer or Expansion modules you have installed

Some SteelCentral products do not require a capacity license. The NetShark enables data to be processed up to the capacity of the system without requiring additional capacity licenses. However, the NetShark ability to export flows to a NetProfiler is limited by the license. Depending on the NetShark model, the maximum flows exported ranges from 400,000 to 1,200,000 fpm. (These limits are the maximum licensed limits, but not guaranteed limits. Actual limits vary depending on the load on the NetShark.)

Option Licenses

Option licenses enable you to add functionality to the NetExpress, NetProfiler, NetProfiler-v, or Enterprise NetProfiler cluster. Software defined network is an example of functionality that requires an option license.

The security module does not require you to purchase an additional part number. The NetExpress, NetProfiler, and Enterprise NetProfiler cluster appliances support the security module through the runtime license. On the NetExpress, you do not need to purchase an additional license for operational analytics or service monitoring. That license is included with the NetExpress runtime license.

License Installation

License installation varies depending on the type of hardware the license is installed on. SteelCentral systems based on current-model hardware use an automatic license server that provides license keys directly to the SteelCentral appliance or through a web portal that enables you to manually install the license. Older hardware relies on manually installed licenses, using instructions provided with the license.

Current-Generation Hardware License Installation

When you purchase a license for your SteelCentral system for xx60 hardware, a license key is generated and automatically assigned to the system in the Riverbed licensing portal.

If the SteelCentral appliance system has internet access, you can automatically download the new license to the device. If the SteelCentral appliance system does not have internet access, you can manually add the new license to the system by entering the provided key into the SteelCentral Management Console.

Packet Analyzer Licensing

Packet Analyzer has two licensing methods:

- **Per Seat Licensing** - You use per seat licensing when you have a set of users who are using Packet Analyzer on a regular basis and cannot tolerate limited access. Network administrators responsible for analyzing and troubleshooting a network are often given per seat licensing. Per seat licensing assigns an individual license to each installation of Packet Analyzer. The assigned license is permanently associated with the Packet Analyzer installation. Packet Analyzer launches without needing access to any sort of license server.
- **Concurrent Licensing** - You use concurrent licensing when you have a set of users who need access to Packet Analyzer on an occasional basis. Users who might analyze packet traces or access NetShark data infrequently are good candidates for concurrent licensing. Also if you have users around the world who do not need to access Packet Analyzer at the same time concurrent licensing is a good option.

For concurrent licensing to work you must have a license server with valid Packet Analyzer licenses assigned to it. NetProfiler, NetExpress, and NetShark can all act as valid license servers. When Packet Analyzer launches the user specifies the appropriate license server and Packet Analyzer attempts to retrieve a license from that system. Licenses are assigned to a Packet Analyzer installation for up to 24 hours, with licenses expiring at midnight on the license server. Once a license is assigned to a Packet Analyzer instance there is no way to release the license and additional communication between the Packet Analyzer instance and the license server is not required. You can use Packet Analyzer offline with no issues.

Assigning Licenses

When you purchase SteelCentral appliance hardware, all appropriate licenses are assigned to each unit. The build-time license (indicating what type of system the unit is: for example, the Flow Gateway) is installed during the manufacturing process and you cannot change it. Additional licenses (capacity and option licenses) are also generated during the manufacturing process but are not installed directly on the system at that time.

You can retrieve the additional licenses on the Riverbed licensing portal. For example, if you purchase a single Flow Gateway with an F1 capacity license, a license is generated on the Flow Gateway that enables it to support 200,000 fpm. This license is not activated, but it is available to you in the Riverbed licensing portal. The license automatically downloads to the Flow Gateway when the Flow Gateway licenses are synchronized with the portal.

Sometimes you must manually assign licenses to specific hardware. This helps ease the process of installing different models of hardware in different locations.

For example, if you want to deploy 10 Flow Gateways at sites around the world, each site requires a different capacity level. You purchase four F4 Flow Gateways, two F3 Flow Gateways, and four F1 Flow Gateways. If each Flow Gateway is automatically assigned an operational and a capacity license, you must ensure that the appropriate units are shipped to the appropriate destinations. However, with the flexible licensing model, all you have to do is ship a Flow Gateway to each location. When the Flow Gateway is at the location and installed, you can use the licensing portal to assign capacity licenses to the Flow Gateways installed at the other locations. You can deploy different capacity licenses (or option licenses) to different systems without interacting directly with each system prior to deployment.

Manual License Installation

The Riverbed licensing portal provides an easy-to-use, automatic system to ensure that all your devices are up to date. However, there are times when the SteelCentral appliance cannot access the portal: for example, on secure networks with no internet access. In these cases, you must manually access the licensing portal, retrieve the desired license keys, and manually add them to the SteelCentral.

Automatic License Upgrades

When you purchase an upgrade for a specific SteelCentral appliance, the licensing process is automatic. After the purchase is approved, the licensing portal receives a new, upgraded license. The next time that appliance checks with the license server, the new license is downloaded and installed. You do not need to remove old licenses: installed license with the highest capacity always take precedence.

Evaluation Licenses

You can evaluate all of the licenses SteelCentral support. Evaluation licenses are provided with specific expiration dates, after which the license no longer works. You can test out some functionality, such as monitoring analytics, for a specific period of time. The expiration date is displayed on the license in the licensing portion of the UI.

Because all licenses are installed with expirations, you can install the runtime license as a temporary license. When the runtime license expires, the system stops functioning correctly. You then must install a license with a later expiration date or no expiration.

Licenses Available

This section describes the licenses available for different SteelCentral.

Licensing the NetExpress and NetExpress-v

The NetExpress 470 has five capacity licenses and one optional license. The capacity licenses provide:

- 15,000 (F1) fpm
- 30,000 (F2) fpm
- 60,000 (F3) fpm
- 90,000 (F4) fpm
- 120,000 (F5) fpm

The optional license provides VXLAN support.

The NetExpress-v 470 has five capacity licenses and one optional license. The capacity licenses provide:

- 15,000 (F1) fpm
- 30,000 (F2) fpm
- 60,000 (F3) fpm
- 90,000 (F4) fpm
- 120,000 (F5) fpm

The optional license provides:

The optional license provides VXLAN support.

Licensing the NetProfiler and NetProfiler-v

The NetProfiler has five capacity licenses and two optional licenses. The capacity licenses provide:

- 200,000 (F1) fpm
- 400,000 (F2) fpm
- 600,000 (F3) fpm
- 800,000 (F4) fpm
- 1,000,000 (F5) fpm

The optional licenses provide:

- SAN support.
- VXLAN support.

NetProfiler-v has ten capacity licenses and a single optional license. The capacity licenses provides:

- 15,000 (F1) fpm
- 30,000 (F2) fpm
- 60,000 (F3) fpm
- 90,000 (F4) fpm
- 200,000 (F5) fpm
- 400,000 (F6) fpm
- 600,000 (F7) fpm
- 800,000 (F8) fpm
- 1,000,000 (F9) fpm
- 2,000,000 (F10) fpm

The optional license provides VXLAN support.

Licensing the Enterprise NetProfiler Cluster

The Enterprise cluster provides a base capacity license of 1,000,000 fpm. You can expand up to nine additional Expansion modules, each providing 1,000,000 fpm of capacity. No additional capacity licenses are required for the Enterprise cluster. When you deploy an Enterprise Cluster with four or more EX modules, you must also deploy a dispatcher.

The optional license provides VXLAN support.

Licensing Flow Gateway and Flow Gateway-v

The Flow Gateway has six different capacity licenses and no optional licenses. The capacity licenses provide:

- 200,000 (F1) fpm
- 400,000 (F2) fpm
- 600,000 (F3) fpm
- 800,000 (F4) fpm
- 1,000,000 (F5) fpm
- 2,000,000 (F6) fpm

Flow Gateway-v has 10 capacity licenses and no optional licenses. The capacity licenses provide:

- 15,000 (F1) fpm
- 30,000 (F2) fpm
- 60,000 (F3) fpm
- 90,000 (F4) fpm
- 200,000 (F5) fpm
- 400,000 (F6) fpm
- 600,000 (F7) fpm
- 800,000 (F8) fpm
- 1,000,000 (F9) fpm
- 2,000,000 (F10) fpm

Licensing the NetShark

The NetShark does not have any capacity or optional licenses available.

NetShark Model	Flow Export Capacity
2179	400 k
4170	800 k
6170	1,200 k