Cascade Deployment Guide

April 2012

riverbed®

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
199 Fremont Street
San Francisco, CA 94105

Phone: 415.247.8800
Fax: 415.247.8801
Web: http://www.riverbed.com

Part Number
712-00113-01

Contents

Preface

Welcome to the Cascade 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

The Cascade Deployment Guide describes best practices for configuring Cascade appliances deployment, including Steelhead appliances, Interceptor appliances, and third-party appliances. The guide includes information about flow collection, packet collection, metrics, analysis, troubleshooting, and licensing.

Audience

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

You must also be familiar with the following:

- *Cascade Profiler and Cascade Express User's Guide*
- *Cascade Sensor and Cascade Gateway User's Guide*
- *Shark Appliance User Manual*
- *Virtual Cascade Shark Appliance Quick Start Guide*
- *Cascade Pilot Reference Manual*
- *RSP User's Guide* (RSP) or *Steelhead EX Management Console User's Guide* (VSP)
- *Riverbed Deployment Guide*

Document Conventions

This guide uses the following standard set of typographical conventions.

Convention	Meaning
italics	Within text, new terms and emphasized words appear in italic typeface.
boldface	Within text, CLI commands and GUI controls appear in bold typeface.

Courier	Code examples appear in Courier font:
	`amnesiac > enable`
	`amnesiac # configure terminal`

< >	Values that you specify appear in angle brackets: **interface <ipaddress>**
[ ]	Optional keywords or variables appear in brackets: **ntp peer <addr> [version <number>]**
{ }	Required keywords or variables appear in braces: **{delete <filename> | upload <filename>}**
|	The pipe symbol represents a choice between the keyword or variable to the left or right of the symbol (the keyword or variable can be either optional or required): **{delete <filename> | upload <filename>}**

Additional Resources

This section describes resources that supplement the information in this guide. It includes the following information:

- "Release Notes" on page 2
- "Riverbed Documentation and Support Knowledge Base" on page 3
- "Online Documentation" on page 3

Release Notes

The following online file supplements the information in this guide. It is available on the Riverbed Support site at
https://support.riverbed.com.

Online File	Purpose
<product>_<version_number> <build_number>.pdf	Describes the product release and identifies fixed problems, known problems, and work-arounds. This file also provides documentation information not covered in the guides or that has been modified since publication.

Examine the release notes file before you begin installation and configuration. It contains important information.

Riverbed Documentation and Support Knowledge Base

For a complete list and the most current version of Riverbed documentation, log in to the Riverbed Support site at
https://support.riverbed.com.

The Riverbed Knowledge Base is a database of 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.

Online Documentation

The Riverbed documentation set is periodically updated with new information. To access the most current version of Riverbed documentation and other technical information, consult the Riverbed Support site located at
https://support.riverbed.com.

Contacting Riverbed

This section describes how to contact departments within Riverbed.

Internet

You can learn about Riverbed products at
http://www.riverbed.com.

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.

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

The following sections have been added since the *Cascade Deployment Guide February 2012* release:

- "Using Virtual Shark" on page 14
- "Deploying the Profiler, the Gateway, the Shark, Virtual Shark, and the Pilot" on page 25

CHAPTER 1

Cascade Overview

This chapter describes an overview of the Riverbed Cascade product suite. It contains the following sections:

- "What Is Cascade?" on page 5
- "The Cascade Components" on page 6

What Is Cascade?

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

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

- Behavior analytics for proactive monitoring
- Dependency mapping for an always-accurate view of your applications and their dependencies
- Executive-level dashboards for a quick summary of service performance
- Cost-effective visibility into remote sites by leveraging Steelhead appliances
- Ability to plan for and understand optimized WANs
- Ability to go broad (for true end-to-end coverage of the enterprise) and deep (with dashboard-to-flow-to-packet analysis), providing scalability and flexibility
- Easy to deploy, manage, and monitor

Cascade provides a full set of application-centric, site-centric, and business-centric views so you can discover your critical services, optimize performance, and continuously monitor service levels. Cascade's consistent view of the network breaks down the silos between network, application, infrastructure, and security teams and shortens the time to resolve problems. Built-in integration with Steelhead appliance WAN optimization products provides full visibility and control over application delivery. For more details on Steelhead appliance integration, see "Cascade Steelhead Appliance Integration" on page 59.

The Cascade Components

This section describes the following Cascade products:

- "The Cascade Profiler Overview" on page 6
- "The Cascade Gateway Overview" on page 7
- "The Cascade Sensor and Sensor-VE Overview" on page 7
- "The Cascade Shark and Virtual Cascade Shark Overview" on page 7
- "The Cascade Pilot Overview" on page 8

Figure 1-1 shows the integrated Cascade components.

Figure 1-1. Cascade Architecture

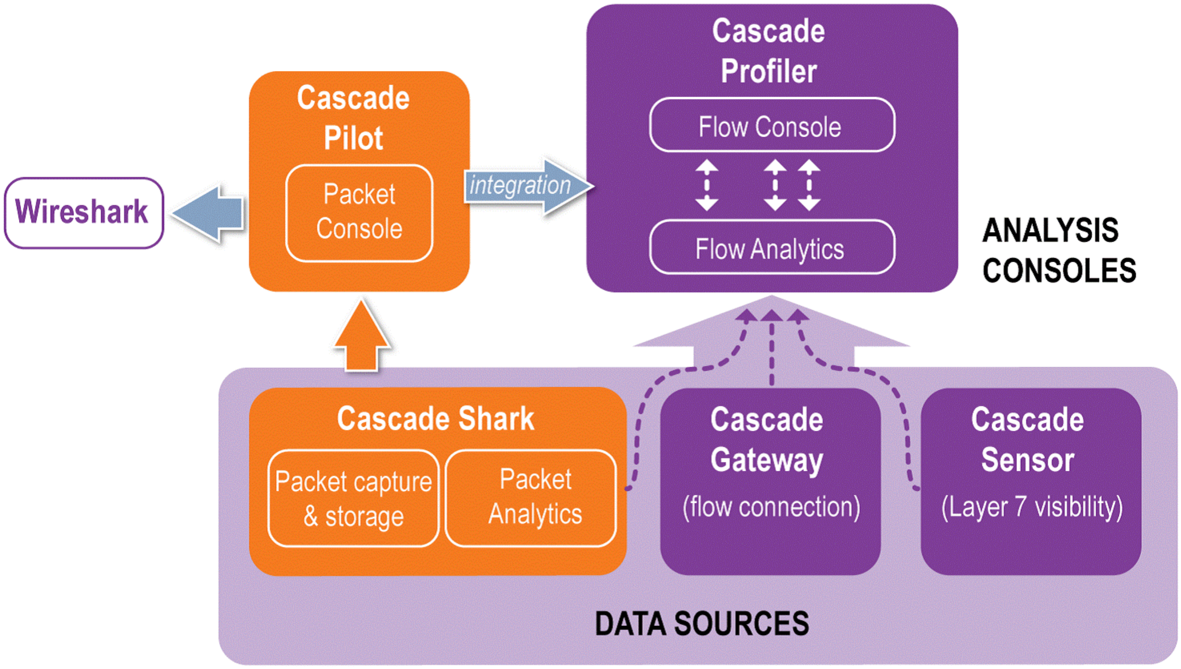

The Cascade Profiler Overview

The Cascade Profiler provides centralized reporting and analysis of the data collected by other Cascade appliances, Steelhead products, and flow exporting routers and switches, on a single user interface. The Profiler offers performance analytics, security analytics, and proactive alerts for delivering application-aware monitoring and troubleshooting to your network. It combines all network data into a single data set with in-depth views that support flexible analysis of the information.

The levels of the Profiler are as follows:

- **The Express** - Light-weight model designed for entry-level and small organizations (no more than 3,750 hosts), which includes some Sensor and Gateway functionality.
- **The Standard Profiler** - Standard model designed for mid-level organizations (up to around 25,000 hosts).
- **The Enterprise Profiler** - Designed to be expandable, supporting up to approximately 180,000 hosts.

For more information about the Profiler, see "Using the Profiler" on page 11.

The Cascade Gateway Overview

The Cascade Gateway collects flow data from routers, switches, and other network devices. It supports most standard flow types (NetFlow, sFlow, J-Flow, IPFIX, and so on). It aggregates the data, compresses it, encrypts it, and sends it to Cascade Profiler. It can transmit data to up to two Standard Profilers or Express appliances. You can deploy it in the same location as the Profiler, or regionally if you have multiple data centers.

For more information about the Gateway, see "Using the Gateway" on page 13.

The Cascade Sensor and Sensor-VE Overview

The Cascade Sensor passively inspects packets from port mirroring ports or taps. It provides Layer-7 application classification and sends performance information, including end-user experience metrics, to the Profiler. You can deploy a virtual edition, Sensor-VE, as a RiOS Services Platform (RSP) package on Steelhead appliances. Sensor-VE includes all the capabilities of Cascade Sensor with the exception of Layer-7 application classification.

For more information about the Sensor, see "Using the Sensor" on page 13.

The Cascade Shark and Virtual Cascade Shark Overview

The Cascade Shark is a high-performance (1 GbE or 10 GbE) continuous packet capture, storage, and analysis appliance. You can use it for fast indexing and in-depth analysis of multiterabyte network traffic recordings. With the Shark, you can drill down to deliver micro-level flow resolution for analysis. The Shark sends performance reports to the Profiler. The Shark delivers real-time or back-in-time deep packet inspection and analysis. You can access the Shark using the Pilot. The Shark uses Riverbed's unique XML-based protocol on top of an HTTPS connection for transferring data to the Pilot.

As of Cascade v9.5, you can choose to deploy the Virtual Cascade Shark.

For more information about the Shark, see "Using the Shark" on page 14.

Embedded Cascade Shark Overview

In RiOS v7.0 and later, the Steelhead appliance includes embedded Shark functionality. Embedded Cascade Shark allows on-demand packet capture on Steelhead appliances at remote sites, and control and analysis of packet captures on remote Steelhead appliances directly from the Pilot. As with the Shark, with Embedded Cascade Shark you can drill down to deliver microlevel flow resolution for analysis using Riverbed's XML-based protocol on top of an HTTPS connection for transferring data to the Pilot. You do not need to transfer full packets until you need them.

The Cascade Pilot Overview

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

For more information about the Pilot, see "Using the Pilot" on page 14.

CHAPTER 2

Cascade Component Deployment Scenarios

Deployment of the Cascade product line requires advanced planning to ensure that you install the appliances to capture critical traffic, and that you deploy your appliances efficiently without wasting resources on unnecessary coverage. This chapter includes the following deployment sections:

- "Choosing the Right Equipment" on page 9
- "Deployment Scenarios" on page 15
- "Port and Protocol Dependencies" on page 28
- "Cascade Flow Storage" on page 32

Choosing the Right Equipment

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

- "Using the Profiler" on page 11
- "Using the Gateway" on page 13
- "Using the Sensor" on page 13
- "Using the Shark" on page 14
- "Using the Pilot" on page 14

When determining 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—consider the following:

- What kind of information do I want to know about this location? Do I need response-time information? Optimization information? WAN link bandwidth information? Application usage information?
- How many users and how much traffic am I expecting at this location now and in 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 SPAN and mirror traffic? Can my routers provide flow information?

- Do I have sufficient bandwidth to transfer flow data between this location and the Profiler?
- How much visibility do I need at this location?
- Do I need packet-level visibility to view objects calls and individual transactions within the application?

The following table shows Cascade solutions for several scenarios.

Scenario	Cascade Solution
Accurately calculate response-time information for non-optimized flows.	The Shark or the Sensor
Accurately calculate response-time information for optimized flows.	Steelhead appliance and Shark or the Sensor on the server-side
Report on and monitor link bandwidth information: for example, to monitor percent use.	The Gateway
Monitor types of Layer-7 applications on the network: for example, to ensure that no one violates the network usage policy.	The Sensor
Obtain detailed packet information: for example, to analyze network traffic in case of a security violation.	The Shark preferred, the Sensor the acceptable
Gain visibility into virtualized environments.	Virtual Shark

Different sites have varying numbers of users and volume of network traffic. A site with 10 users transferring large files all day generates far fewer packets and flows then a site with 200 users viewing Web pages. For calculation purposes, Riverbed uses you 20 flows per minute as the estimated average flows per minute. Exact flows per minute depend on the traffic characteristics in your environment.

Use multiplication to estimate the maximum number of flows per minute. For example, 100 users each generating 20 flows per minute produce a flow rate of roughly 2,000 flows per minute. 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 firewalls or load balancers for flow estimates.

You must have the appropriate number technical staff on site. In a small remote office of only a few non-technical people, deploying a Sensor-VE might make more sense than installing a physical box.

Consider other network equipment that is already installed at a site when you decide what sort of Cascade appliances to install. If an office site contains multiple large switches capable of generating NetFlow or SPAN and port mirror traffic, a Shark, a Sensor, or a Gateway might make sense. Conversely, if a small office contains only a wireless Internet router with no other network infrastructure, the options for deploying visibility solutions are extremely limited or possibly not needed.

If you have a site that reports significant quantities of data across a WAN, consider the bandwidth used for the transfers. Typical WAN bandwidth use is 1.5% of monitored traffic. Cascade devices 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 (such as restricting the amount of data being monitored).

Using the Profiler

The Profiler is licensed on a flow-per-minute basis, after all flows have been deduplicated. You want to choose the right Profiler model, so that you do not receive inaccurate results and performance issues. Consider the following factors when you are deciding what type and model of Profiler to install:

- The size of the current network you want to profile
- The planned expansion of coverage

Each Profiler model—the Express, the Standard, and the Enterprise—is designed for a specific organization size.

The Express Profiler

The Express has flow rates ranging from 25,000 to 75,000 flows per minute (after deduplication), and is best suited for either smaller organizations or a small subsection of a larger network. Because the Express can forward the flows it receives directly to a different model Profiler, this deployment can make sense for sites where there is a need for local visibility and enterprise-wide visibility.

The Express includes functionality similar to that of the Profiler and the Gateway Cascade components. The additional capabilities enable a compact deployment.

The following table shows the Express model options. Most features are available in the base unit. The base unit is 2U high and you can upgrade in the field to the next higher flow-rate version. There is an additional platform analytics license (LIC-CAP-0360-CAA) if you want analytics and service dashboards.

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
CAX-000360 (L)	up to 25K FPM	Primary 10/100/1000 for management Four ports 10/100/1000 to monitor Sensor	Two port 1- Gbps SFP (LR and SR option for additional Sensor monitoring)
CAX-000360 (M)	up to 50K FPM,		
CAX-000360 (H)	up to 75K FPM		

The Standard Profiler

The Standard Profiler has flow limits ranging from 100,000 to 500,000 flows per minute, making it ideally suited for mid-size organizations with between 4,000 and 25,000 hosts (assuming an average of 20 flows per minute per host).

The Standard Profiler cannot forward flows to other Profilers, nor can the Standard Profiler receive flows directly from flow sources (for example, routers). Because the Sensors, Gateways, and Sharks can forward flows to two distinct Profilers, and the Shark can forward flows to two distinct Profilers, you can use the Standard Profiler to monitor a small subset of a larger network. You can send the flows from the Cascade devices (Sharks, Sensors, and Gateways) to the local Standard Profiler (monitoring a network subset) as well as to another Profiler.

The following table shows the Standard Profiler model options. Most features are available in the base unit. Each unit is 2U high and you can upgrade in the field to the next-higher flow-rate version. There is an additional platform analytics license (LIC-CAP-2260-CAA) if you want analytics and service dashboards (monitoring the entire network).

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
CAP-02260 (L)	up to 100K FPM	Primary 10/100/1000 for management	SAN card (two fiber HBA ports)
CAP-02260 (M)	up to 200K FPM		
CAP-02260 (H)	up to 500K FPM		

The Enterprise Profiler

The Enterprise Profiler has a minimum flow limit of 800,000 flows per minute; you can expand that with expansion modules—up to a maximum of 3.6 million flows per minute. Each module provides support for an additional 400,000 flows per minute. In terms of hosts, a standard Enterprise Profiler can support at least 40,000 hosts (assuming an average of 20 flows per minute per host) and potentially much more.

The following table shows the Enterprise Profiler model options. Most features are available in the base unit. The base Enterprise Profiler is composed of two 1U units or one 2U unit. Each expansion module is an additional 2U unit. There is an additional platform analytics license (LIC-CAP-4260-CAA) if you want analytics and service dashboards.

Base Unit and Flow License	Deduplicate Flow Rate	Included Ports	Optional Expansion Ports
CAP-04260-UI (required)	Part of base system	Primary 10/100/1000 for management	N/A
CAP-04260-DB (required)	Part of base system		
CAP-04260-AN (required)	Base system, up to 800K FPM		SAN card (two fiber HBA ports)
CAP 4260-EX (optional; add 0-7)	Expansion unit, each one adding 400K FPM		
CAP 4260-DP (required only when you have five or more 4260-EX)	N/A: used to balance flows on larger systems		N/A

To plan for future expansion, you must know the current estimated number of flows per minute and the expected number in the time frame being planned for. For example, if the network currently has 6,000 hosts and is expected to grow to 9,000 hosts in the next two years, a Standard Profiler is sufficient to handle the growth. A software license and hardware upgrade allows a Profiler licensed for 200,000 flows per minute to be upgraded to 500,000 flows.

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 Profiler is a better choice. The Standard Profiler is sufficient in a 14,000-host network, but it is unlikely to provide adequate performance and license limits to manage the network when it grows to 25,000 hosts.

If you want to provide visibility into a subset of the network and a full overview, you can install a Standard Profiler along with an Enterprise Profiler. This ensures that the system meets both the immediate and planned growth needs.

Using the Gateway

A single Gateway is often sufficient to manage a large number of devices, because Gateways provide a wide range of licensed flow limits (from 100,000 to 1.4M flows per minute). Additionally, the upgrade from the smallest to the largest Gateway requires only a license change. The following table shows the available Gateway models.

Gateway Model	Deduplicate Flow Rate
CAG-02260-F1	up to 100 K FPM
CAG-02260-F2	up to 200 K FPM
CAG-02260-F3	up to 500 K FPM
CAG-02260-F4	up to 800 K FPM
CAG-02260-F5	up to 1.4 M FPM

Consider the following when you deploy a Gateway:

- The flow limits are within the licensed limits.
- The geographic coverage is appropriate.
- The flow capacity of the Profiler is not exceeded by the devices sending data.

For a mid-size organization with multiple disparate geographic locations, a single 1.4 M flow Gateway can manage the overall load—you must have the appropriate-sized Profiler because the Gateway can send more flow than the standard Profiler can receive. However, this 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 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, it is a good idea to deploy multiple smaller Gateways at major locations.

Using the Sensor

When you choose Sensor options, consider the number and type of interfaces already installed and your overall system capability. The Sensor can identify Layer -7 applications. The Sensor can process no more than 5 Gbps of packets. If you need to process more packets than a single Sensor can effectively handle, you need multiple Sensors.

The following table shows the Sensor option. You start with the base unit and add appropriate media cards.

Base Unit	Included Ports	Optional Expansion Ports
CAS-02260	Primary 10/100/1000 for management Four ports - 10/100/1000 copper for monitoring Sensor	Two port 10 Gbps SFP (LR and SR option)

Using the Shark

You must choose the appropriate Shark appliance (4 TB, 8 TB, 16 TB, or 32 TB) to ensure that the appropriate quantity of packets are stored and are available to be analyzed. You deploy the Shark with one or more capture cards.

Consider the number of interfaces possible and the amount of disk space available for storage. Capture cards are separately ordered items. Model selection depends on the expected network packet rate and retention time you want. The largest Shark appliance comes with roughly 32 TB of disk space for packet storage. If you want to store every packet traversing a heavily loaded 1000 Mbps network, the 32 TB of space allows for approximately nine hours of storage.

The following table shows the Shark models.

Shark Base	Form Factor	Storage	Maximum Capture Cards
CSK-01100-BASE	1U	4 TB	1 (NIC-CSK-2TX or NIC-CSK-4TX-C)
CSK-02100-BASE	2U	8 TB	2 (supports all card types)
CSK-02200-BASE	2U	16 TB	2 (supports all card types)
CSK-03100-BASE	3U	16 TB	2 (supports all card types)
CSK-03200-BASE	3U	32 TB	2 (supports all card types)

Using Virtual Shark

You must choose the appropriate Virtual Shark model (50 GB, 1 TB, or 2 TB) to ensure that the appropriate quantity of packets are stored and are available to be analyzed. The available models are shown in the table below. You deploy Virtual Shark in VMware ESXi environments through connectivity to a standard vSwitch, distributed vSwitch, or a Cisco Nexus1000V Switch. You install Virtual Shark similarly to other virtual machine disk format (VMDKs) and uses a vSwitch promiscuous-mode port channel to obtain packets. If you use the Cisco Nexus 1000V, Virtual Shark uses a SPAN port. You can limit traffic to only inter-VM traffic by using the appropriate capture Berkeley packet filter (BPF).

The following table shows the Virtual Shark models.

Model	Storage	Capture Interfaces	Export to the Profiler
VSK-00050	50 GB	4	50K FPM
VSK-00200	1 TB	4	50K FPM
VSK-00400	2 TB	4	50K FPM

For more information on installing and configuring Virtual Shark, see the *Virtual Cascade Shark Appliance Quick Start Guide.*

Using the Pilot

When you deploy the Pilot, you must consider only how many users need to perform deep packet analysis. The Pilot is Windows client software and is licensed per installed machine. The Pilot can analyze traffic from a Shark appliance, an Embedded Cascade Shark probe, and any standard packet capture files. You do not need the Pilot for Profiler-level analysis and troubleshooting.

Deployment Scenarios

This section describes the following deployment scenarios:

- "Deploying the Shark and the Pilot" on page 15
- "Deploying the Express" on page 17
- "Deploying the Shark and the Sensor" on page 19
- "Deploying the Express, the Shark, and the Pilot" on page 19
- "Deploying the Express, the Shark, and the Pilot" on page 19
- "Deploying the Standard Profiler and the Gateway" on page 21
- "Deploying the Enterprise Profiler and the Gateway" on page 23
- "Deploying the Profiler, the Gateway, the Shark, and the Pilot" on page 24
- "Deploying the Profiler, the Gateway, the Shark, Virtual Shark, and the Pilot" on page 25
- "Deploying the Enterprise Profiler, the Shark, the Pilot, Sensor-VE, and the Gateway" on page 26

Deploying the Shark and the Pilot

You want to deploy the Shark and the Pilot to look at extremely detailed views of network traffic and to improve troubleshooting of your network-related issues. Because the Shark can only receive packets, you must install the Shark in a location close to the source of the majority of the packets. Remember that this deployment does not include the higher-level analysis and reporting that is included with the Profiler.

For most deployment scenarios, you can place the Shark in the data center. The majority of network traffic in most corporate environments is centered on 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 is invaluable when investigating a problem or security incident.

When you install the Shark in the data center, you do not always catch all traffic (such as printing traffic). However, this uncaptured traffic is usually not of great interest or significant volume. If you want to monitor traffic that does not go through the data center, you can place additional Sharks at strategic wiring closets. Because Cascade offers a variety of Shark sizes, you can choose one solution that is appropriate for the data center and a smaller Shark that works at a remote wiring closet. For available Shark models, see "Using the Shark" on page 14.

You can use other methods to obtain on-demand packet captures from remote locations, and then use the Pilot to analyze them. Although it is not feasible to analyze data center traffic without a packet capture appliance, you can analyze occasional smaller packet captures without an appliance. Additionally, any Steelhead appliance running RiOS v7 and later includes Embedded Cascade Shark. Embedded Cascade Shark allows you to control and analyze of packet captures on the Steelhead appliance directly from Pilot.

If you connect the Shark appliance 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* to detect certain conditions without capturing the traffic. You can preserve only the desired information, thereby reducing the amount of disk space used compared to a capture job that captures all traffic.

For more details on watches, see *Cascade Pilot Reference Manual.*

To store network traffic, you must define capture jobs on the Shark appliance. Capture jobs tell the Shark 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. You can define multiple capture jobs on each Shark appliance to capture packets that meet different, specific requirements. 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 might capture overlapping data.

Use the Pilot to analyze the data stored on a Shark appliance and to look at real-time and historical traffic. You can use a variety of filters and drill-down techniques to analyze traffic. There are numerous views available in the Pilot to assist with the analysis. Because the Pilot can connect to multiple Shark appliances, the location of the Pilot is not as critical as the location of the Shark. You can optimize the Pilot-to-Shark communication by applying views and filters to the data so that only the necessary information is sent between the Shark and the Pilot.

The Pilot 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 only the specific packets of interest are moved.

When looking at real-time (or near real-time) data, the more physical distance you put between the Pilot and the Shark, can create a situation where you use more bandwidth when you monitor additional data, and sources that provide the data. Full-packet data sent from the Shark across a WAN, can use significant amount of bandwidth and possibly have an impact on other traffic.

To prevent these issues, place the Pilot as close as possible to the Shark. If you place the Pilot across a slower WAN from the Shark, you must apply appropriate views and filtering before viewing raw packets or saving these packets locally.

Figure 2-1 shows an example Shark and Pilot deployment. Although this example shows only a single Shark appliance, you might need additional Shark appliances for large data centers or to monitor additional locations, such as remote data centers or closets.

Figure 2-1. Example Shark and Pilot Deployment

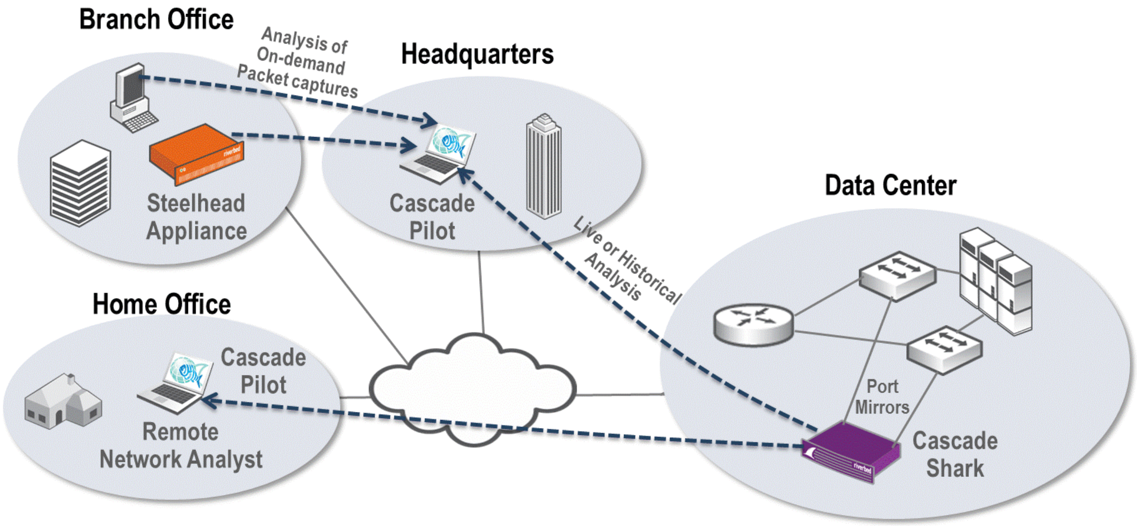

Deploying the Express

The Express is the ideal solution for a small enterprise looking for an advanced monitoring solution. The Express 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 Express 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 Cascade Profiler solutions. This is achieved with two 1-GB network monitor interfaces for receiving packets and flow data directly from the source.

The physical location of the Express is extremely important. Because there are only two monitoring interfaces for receiving and analyzing SPAN, mirror, and TAP traffic, you must place the device as close to the sources of that data as possible.

The Express is generally installed 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.

Additionally, because the Express can receive flow data directly from flow sources, it is important to place it close to sending devices (for example, core switches and routers) so there is no impact on the WAN.

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

Figure 2-2. Example Standalone Cascade Express Deployment

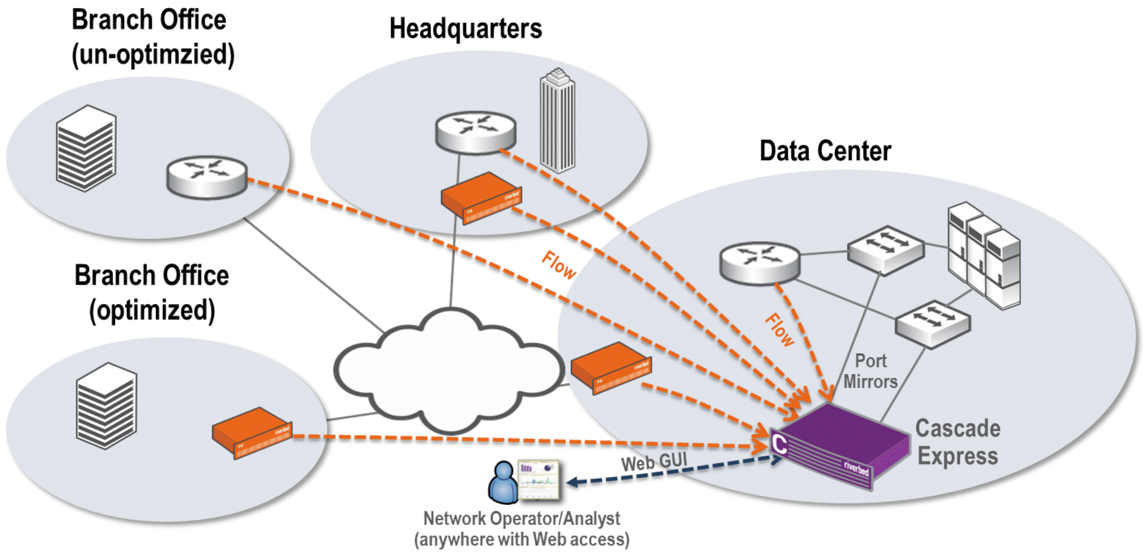

Integrated Deployment

When you deploy the Express as an integrated deployment, consider what portions of the network you want visible to users of the different devices being deployed. There is no way to restrict visibility directly within the Profiler, but you can use an Express to simulate limited visibility.

The Express receives data directly from the source (through built-in monitoring ports and direct flow reception) and can forward the data to other Profilers. Because of this, you can use the Express to limit the view to a subset of the network and still feed a larger system to provide a full view.

When you deploy an Express with a limited view, make sure that the Express supports the required flow rate for the covered area. The largest Express has a limit of 75,000 deduplicated flows per minute, and it can be quickly overwhelmed by larger deployments. For example, using an estimate of 20 flows per minute per host, 75,000 deduplicated flows should provide coverage for a network consisting of roughly 3,750 hosts; flows over the 75,000 are dropped.

Note: Twenty flows per minute per host is an approximation of how much traffic a host can generate on a per-minute basis for internal communications. If the host is also communicating with resources on the Internet, then the number of flows generated is higher.

Consider physical location when you deploy the Express in an integrated environment. The Express provides all the same advanced network metrics (such as standard and optimized response-time values) as any other Profiler when installed correctly. For example, to provide accurate optimized response time, the monitoring interfaces must be located between the server-side Steelhead appliance and the servers. If you place the Express in the wrong location (for example, between the client-side Steelhead appliance and the clients), you can prevent accurate collection and calculation of these values.

If you install the Express in close proximity to the data center, an 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 Express to monitor a separate physical location, you have to choose between either not having local visibility or not having certain information available (for example, optimized response-time information). You can avoid this issue by deploying an additional component (such as a the Shark). For details, see “Deploying the Express, the Shark, and the Pilot” on page 19.

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

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

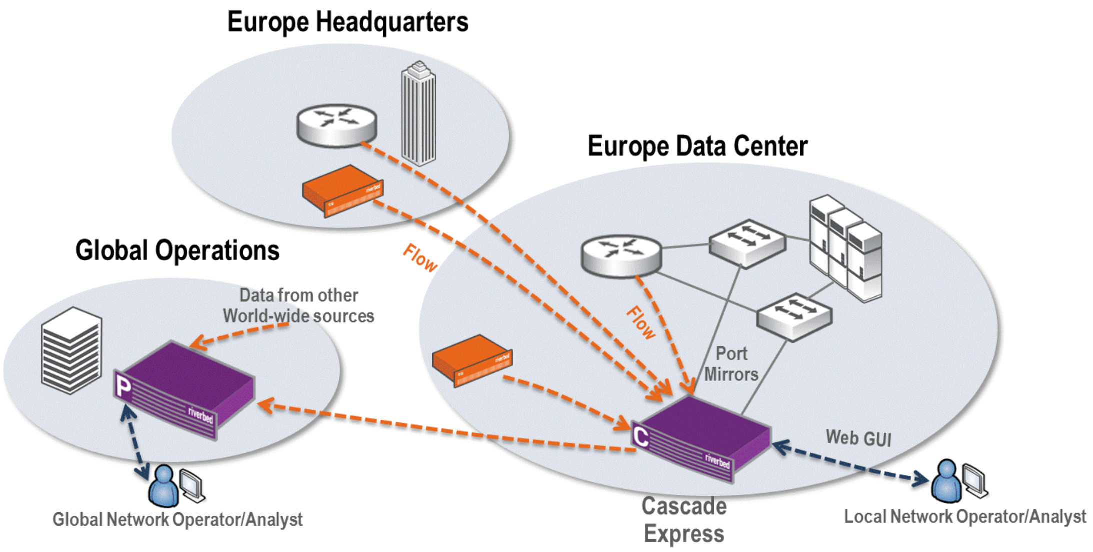

Deploying the Shark and the Sensor

This deployment scenario presents one approach to using the Shark and Sensor in the data center. If you need Layer-7 information exported to the Profiler, you must use the Sensor. Layer-7 identification is typically most important for data entering and leaving the data center. Layer-7 application activity can identify certain client-to-server applications that might use dynamic port assignments (for example, MS Exchange with dynamically assigned ports) and detect potential security issues (for example, downloads through bittorrent).

Because Layer-7 traffic detection is largely limited to traffic entering and leaving the data center, consider placing the Sensor at the data center edge or in the data center core. You can selectively capture data entering and leaving the data center through WAN-bound interfaces. However, traffic within the data center is typically more secure and better defined, so Layer-7 identification within the data center becomes less of a requirement.

The Shark appliance provides high-rate packet capture and performance information to the Profiler. The Shark is installed in the data center core or distribution tiers.

Alternatively, traffic within the data center is usually where packet capture becomes more of a requirement, because you might want packet analysis for client-to-server traffic and server-to-server traffic.

Figure 2-4 shows the Sensor deployed in the data center core, and SPANs ports or VLANs for traffic going to and from the WAN. This placement allows all traffic entering and leaving the data center to be identified at Layer-7, and also allows performance measurements and retransmission counts to be measured at the core.   One or more Sharks are deployed at the distribution tier. This allows packet collection and analysis for data destined for the data center infrastructure and for data remaining within the data center infrastructure. You can use Layer-4 mappings (naming of applications by IP and port) for all data collected, in addition to standard port name definitions. All necessary data is collected for service monitoring and analytics because both the Sensor and the Shark record and report response-time metrics, resets, retransmissions, connection counts, and traffic volumes per conversation.

Figure 2-4. Example Shark and Sensor Deployment

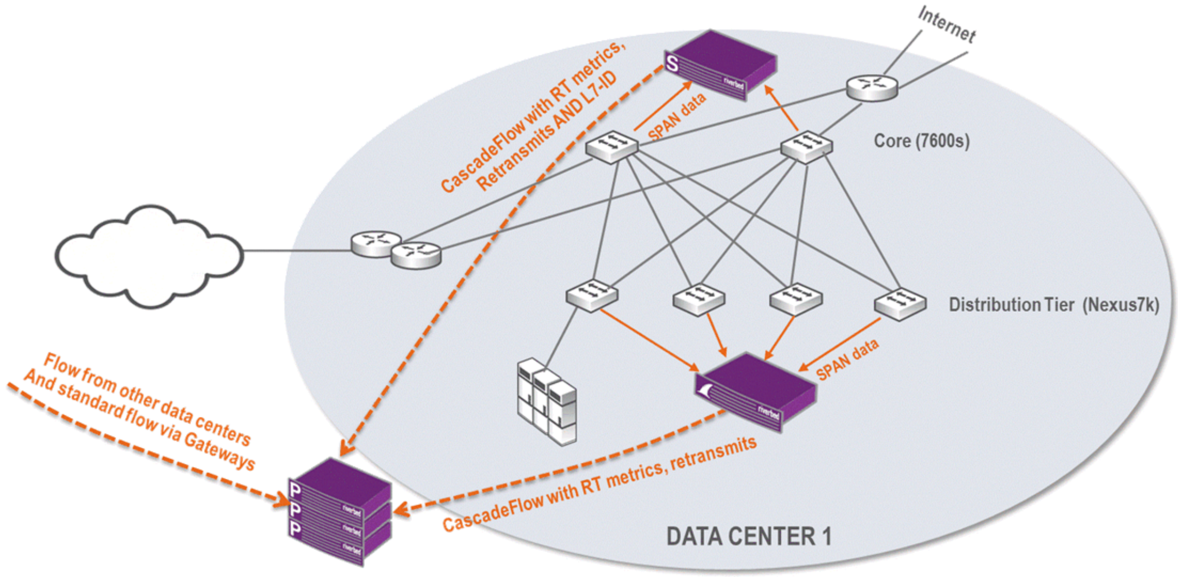
Internet
CascadeFlow with RT metrics,
Retransmits AND L7-ID
S
SPAN data
Core (7600s)
Distribution Tier (Nexus7k)
SPAN data
Flow from other data centers
And standard flow via Gateways
P
P
P
CascadeFlow with RT metrics, retransmits
DATA CENTER 1

Deploying the Express, the Shark, and the Pilot

This deployment is best for smaller environments where you want high visibility, dashboards, and examination, in addition to deep packet analysis.

In addition to acting as a standalone device, you can configure the Express to receive flow data from additional sources and integrate with other products. To provide both high- and low-level views of the data on the network, you can integrate the Express with the packet capturing abilities of the Shark appliance and the detailed analysis capabilities of the Pilot. This deployment has many advantages, including that the Express provides a high-level view of network activity but still allows easy drill-down access to the analysis engine on the Pilot.

When you plan to use the Express with the Shark and the Pilot, the considerations are similar to deploying the Express independently or as part of an integrated solution. Primarily, consider the flow limits. The Express is limited to no more than 75,000 deduplicated flows, and the Shark appliance can export up to 600,000 flows per minute (depending on the incoming traffic rate). Excess flows from the Shark (if any) are dropped by the Express and are not used in analysis or stored on the Express.

When you deploy the Express with the Shark, it is not critical that you place the Express in the data center, for the following reasons:

- You can leverage the Shark interfaces (instead of the Express interfaces) to monitor the traffic between the server-side Steelhead appliance and the data center.
- You can use Shark to augment reporting of non-optimized response-time information, rather than relying on only the Express to provide that data.

In an environment with multiple physical locations, you can remove the restriction of where you locate the Express. For example, you benefit from local visibility in a larger environment. You can place the Express at a remote office, monitoring local flow data and network traffic, and configure the server-side Steelhead appliance to forward only traffic not already sent from the Express. A server-side Shark appliance can also detect optimized flow response-time values.

One significant benefit of integrating the Express, the Shark, and the Pilot is the ability to use the Express high-level view of the network and automated monitoring functions (such as services and operational analytics), while retaining the ability to drill down to more detail when you use the Pilot, and even further, with Wireshark.

If you want to view packet data, you must define a capture job to capture the packets on the Shark appliance. Riverbed recommends that you set a capture job to store all packets being forwarded to the Express by the Shark, based on:

- your particular monitoring preferences.
- capacity.
- the network flow rate you want to monitor.
- other uses of the Shark appliance.

Figure 2-5 shows an example deployment that includes the Express, the Shark, and the Pilot. Routers and Steelhead appliances send flows from the network to the Express, providing broad visibility into the network. A Shark appliance monitors traffic from switches in the data center, collecting packets for deeper visibility. The data from the Shark merges with the flow data collected. You can log in to the Express and view applications flowing across the entire network. When troubleshooting, if you need deeper packet-level analysis, the Profiler UI automatically launches appropriate information into Pilot. This takes you from the Profiler view of network traffic directly into Pilot views of packet data.

Figure 2-5. Example The Express, Shark, and Pilot Deployment

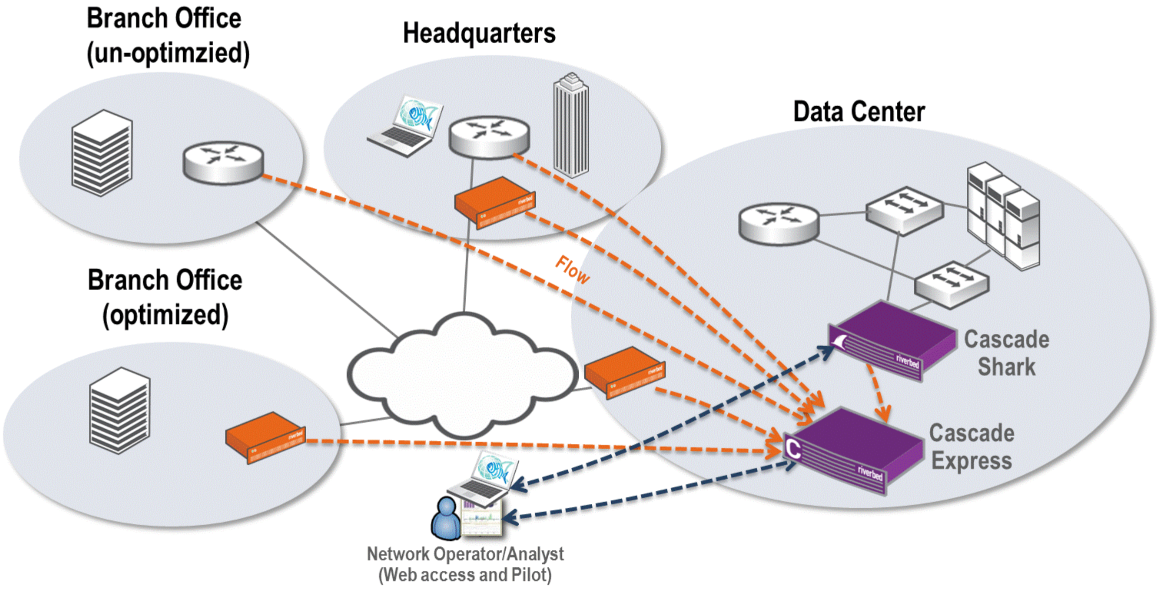

Deploying the Standard Profiler and the Gateway

This deployment is most common in an environment that encompasses a company with several thousand hosts and 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 Shark or the Sensor to provide network performance information and Layer-7 application identification information. You can deploy only the Profiler and the Gateway to detect what hosts communicate with what other hosts during what time periods.

This scenario requires you to collect data from switches and routers. The Steelhead appliances at each location forward the flows to a single Profiler for analysis and reporting. There is no restriction on the distance between physical locations for this implementation.

If you have multiple data centers and are interested in using virtualization, you can use the Profiler and the Gateway as tools to show you which servers are used by what clients and where those clients are located. You can also monitor WAN links to ensure that you have sufficient bandwidth for busy times by looking at real-time performance information and historical information.

If your organization has multiple physical locations that are not connected with high amounts of bandwidth (such as a gigabit MAN), Riverbed recommends that you deploy multiple Gateways throughout the enterprise, with one at each location you want monitored. Due to concerns with native flow data (for example, unreliable UDP transmission and no compression) being sent across a WAN, placing a Gateway at locations with sufficient traffic makes sense if the location warrants monitoring in the first place. A small Gateway (up to 100,000 flows per minute) can support up to thousands of hosts and offers the same benefits of deduplication, compression, and reliable encrypted transport of data to the Profiler.

If you have a small site that hosts a data center but has only a few employees, you might want to deploy a Gateway even if the number of hosts is relatively small, because of the nature of the data. If you want to monitor only data that is of low value (for example, minimal bandwidth or not-confidential), instead of deploying a Gateway at the site, you can flow data across the WAN.

When deciding where to deploy a 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 a Gateway at a remote office or send flow data from the remote office to a Gateway in the data center. Because all critical traffic is in the data center, a single Gateway monitoring all the traffic in, out, and within the data center is sufficient.

While you do not have to install the Profiler in the data center—the physical location of the Profiler is much less important position than the Gateway—Riverbed recommends that the Profiler be as close as possible to the largest sources of flow data. Because access to the Profiler is achieved primarily through the Web-based UI, placing the device physically close to the end users is not critical.

Figure 2-6 shows an example deployment that includes the Profiler and the Gateway. All Steelhead appliances 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 2-6). 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 appliance.

Figure 2-6. Example Profiler and Gateway Deployment

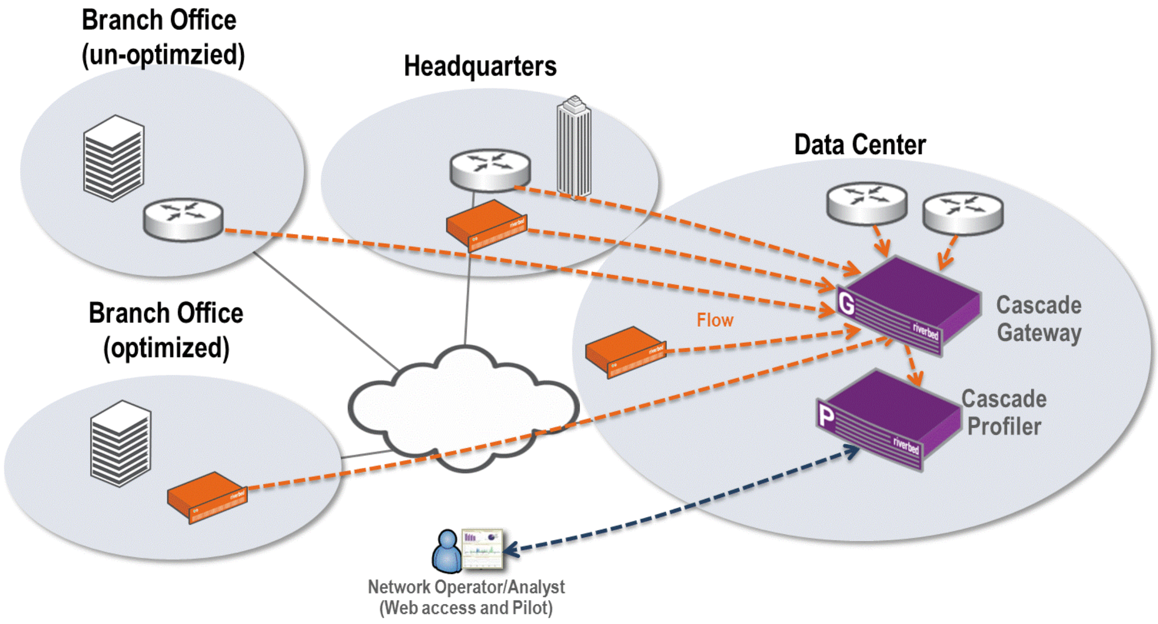

Deploying the Enterprise Profiler and the Gateway

For very large environments, the Enterprise Profiler provides an expandable solution that can process flows from tens of thousands of hosts. The Enterprise Profiler provides a robust solution, allowing visibility ranging from a high-level overview to the flow level. The Enterprise Profiler has a base flow rate of 800,000 flows per minute, and you can add additional modules to support 400,000 flows per minute per module for a maximum supported flow capacity of 3.6 million flows per minute (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 Profiler becomes less critical. When there is enough traffic to require this solution, you have multiple locations that are sending data. If there is a concentration of flow in one area (for example, a primary data center), it makes sense to locate the system close to this source. In any case, you must locate the Enterprise Profiler in an appropriate facility with sufficient bandwidth, power, and cooling.

The most important factor in this deployment is to install the Gateways in the correct locations. The Gateway needs to collect data effectively. Depending on the size of your organization (the number of remote offices, the information you want to monitor, and the amount of bandwidth available), there are several scenarios that make sense.

You can deploy fewer, larger-capacity Gateways if the number of data collection sites is relatively low and the flow rate at those locations is very high. This 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 Gateways in a data center, you can collect all the data necessary. With proper placement, the Gateway can detect conversations from the clients in the remote office to servers in the data centers.

Figure 2-7 shows two Gateways collecting and deduplicating the data flow, then forwarding the flow to the Enterprise Profiler. Because this deployment does not require network performance and deep packet analysis, you do not need to install the Shark or the Sensor. This solution enables you to report, analyze, and troubleshoot traffic across the entire large enterprise network.

Figure 2-7. Example Enterprise Profiler and Gateway Deployment

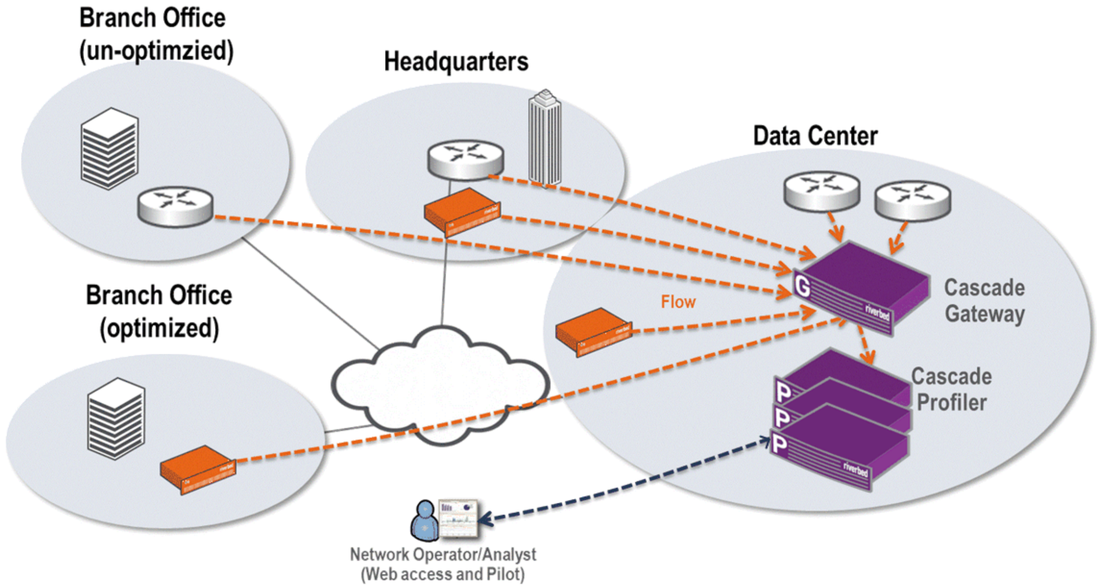

Deploying the Profiler, the Gateway, the Shark, and the Pilot

This scenario expands a standard Profiler and a Gateway deployment to include a Shark and a Pilot. The Shark and the Pilot allow you to:

- see network performance data (response time, server delay, and so on) and TCP health information (TCP retransmission).
- drill-down from the high-level view provided by the Profiler to successively lower-level views until you reach the packet-level view.

The physical location of the Shark is extremely important. The Shark appliance 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 Shark. If you have a single data center and the traffic to and from that data center is the most critical, you should place the Shark 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 Shark in this area. For more information about various methods of collecting packet data, see “Packet Collection for Cascade” on page 47.

It is not a best practice to use a Shark to monitor a WAN, unless you want packet-level visibility into the WAN link. The routers or Steelhead appliances 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 Shark. With a limit of eight 1-Gbps interfaces or two 10-Gbps interfaces, the Shark appliance has high capacities, but not unlimited packet-capture capacities. A data center at a medium-sized organization can easily exceed 20 Gb of traffic per second.

The Shark provides the following ways to ensure that the appropriate traffic is monitored:

- **Physical** - SPANing, mirroring, and TAPing only the desired links
- **Virtual** - Selecting only those specific packets that you want to monitor using the built-in filtering capabilities

Figure 2-8 shows an example deployment that includes a Profiler, a Gateway, a Shark, and a Pilot. Routers and Steelhead appliances send flows across the network to the Gateway and provide wide visibility into the network. A Shark appliance sits off of switches in the data center and collects packets for deeper visibility. Flow data from the Shark merges with all other flow data collected by the Profiler. You can log in to the Profiler to view applications flowing across the entire network. When troubleshooting, if you need deeper packet-level analysis, the Profiler UI automatically launches the Pilot. This takes you from the Profiler view of flow data directly into Pilot views of packet data.

Figure 2-8. Example Profiler, Gateway, Shark, and Pilot Deployment

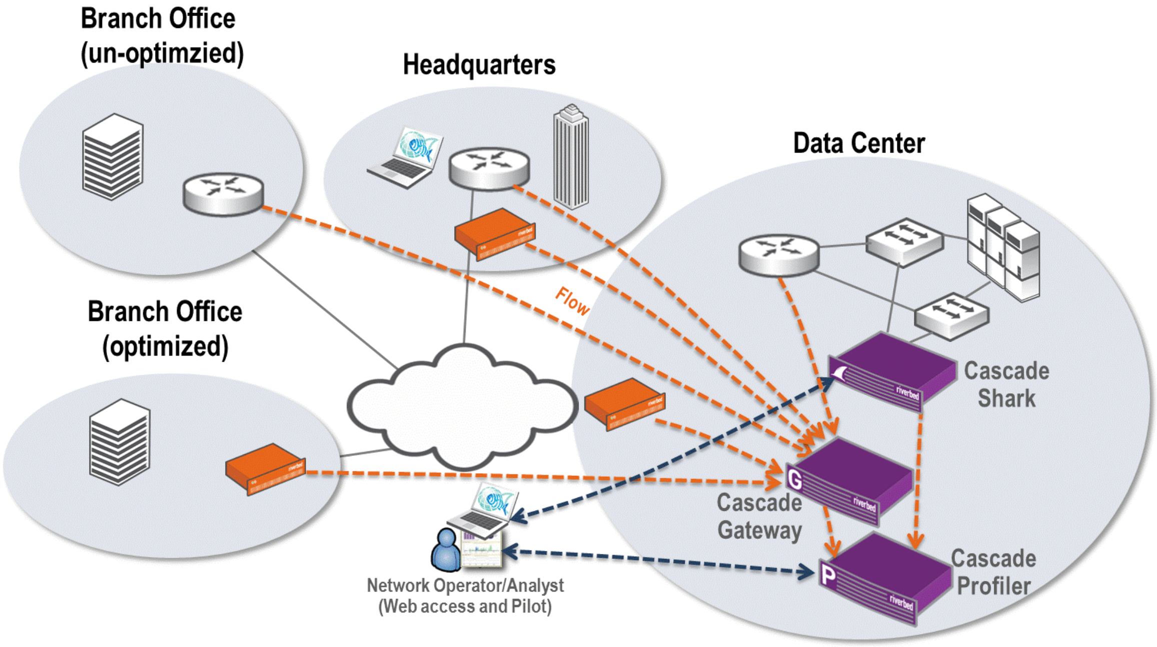

Deploying the Profiler, the Gateway, the Shark, Virtual Shark, and the Pilot

This deployment expands upon the Profiler deployment described in “Deploying the Profiler, the Gateway, the Shark, and the Pilot” on page 24. By adding Virtual Shark, you obtain visibility into the physical network, as well as visibility into the relationship between virtual machines hosted on an ESXi platform in the virtual environment.

Figure 2-9 shows an example deployment the includes the Profiler, the Gateway, the Shark, Virtual Shark, and the Pilot. Virtual Sharks are deployed on each ESXi platform in which you want visibility. Metrics are sent from within the virtual environment to the Profiler. Using the Pilot, you can also perform packet analysis.

Figure 2-9. Example Profiler, Gateway, Shark, Virtual Shark, and Pilot Deployment

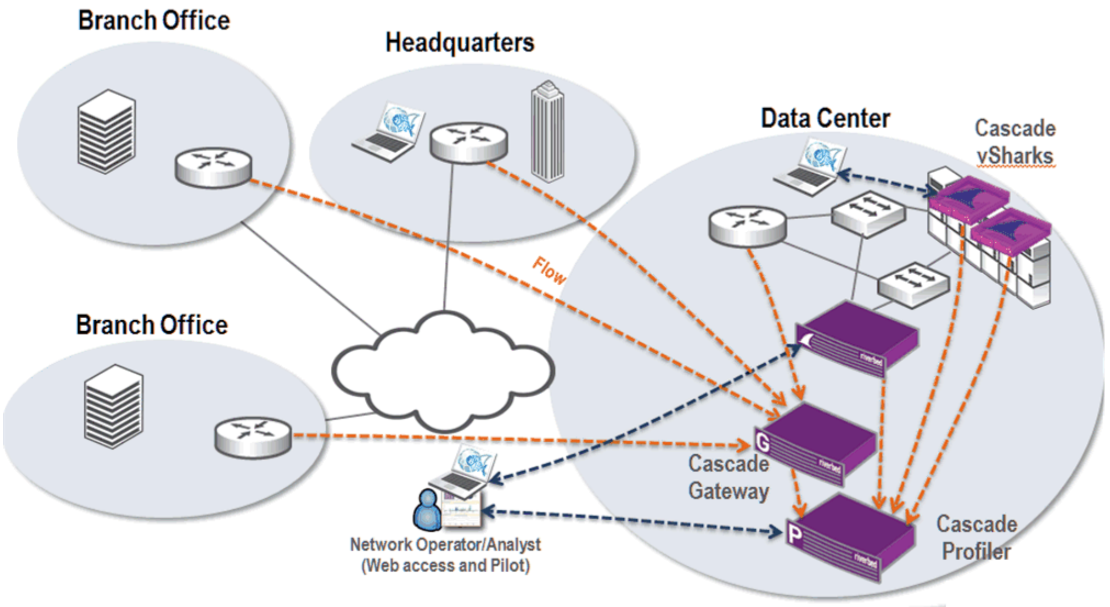

Deploying the Enterprise Profiler, the Shark, the Pilot, Sensor-VE, and the Gateway

This deployment expands the Enterprise Profiler deployment described in "Deploying the Enterprise Profiler, the Shark, the Pilot, Sensor-VE, and the Gateway" on page 26. The Enterprise Profiler provides the necessary capacity and processing power for large enterprise environments, and you can increase the functionality of the Enterprise Profiler by adding the Shark, the Pilot, the Sensor, the Sensor-VE, and the Gateway solutions.

When you deploy this solution, first determine the specific needs and areas you want to monitor. Because the Enterprise Profiler can provide many different types of information, you must decide what information you want for which portions of the corporate network.

If you want to provide all functionality in all areas, you must consider that this might not be practical from a cost or implementation standpoint. If you install a standalone Shark (or Sensor) at all locations, you can achieve an extremely high level of visibility into applications being run on the wire, but the cost might outweigh the benefit.

Because most enterprises have a limited number of data centers, the majority of the traffic of interest is between the remote hosts (clients) and the servers in the data center. If you place a Sensor or a Shark in the data center, it gathers the main traffic and eliminates the cost of placing a physical system at each remote site. In this deployment scenario, the Shark captures remote site packet traffic within the data center without remote packet capture.

You can also use the Sensor-VE solution, which provides most of the same functionality as the Sensor. Sensor-VE evaluates traffic based on packets (and provides response time calculation information) but without the added overhead and cost of another physical system. One benefit of using Sensor-VE versus collecting CascadeFlow from the Steelhead appliances is that with Sensor-VE you can use the Steelhead appliance auxiliary port to connect to a port mirror at a remote site. This enables collection of traffic that might not normally cross the Steelhead appliance or other flow sources.

This solution enables an in-depth view of the traffic on the network. The Shark can store packets at wire speeds in addition to forwarding flow information about the packets to the Profiler. Storing the packets enables you to easily drill-down from the high-level view in the Profiler, to the detailed packet analysis in the Pilot, and finally to the packet-level view in Wireshark.

The Shark and the Sensor have some similar functionality because they both analyze packets and forward enhanced flow-level detail to the Profiler. The Shark and the Sensor compute response times and report on TCP health (retransmissions). Whether you choose the Shark or the Sensor depends upon your needs.

The Sensor has an advanced application identification functionality exported in flow that the Shark does not have. The Sensor enables you to detect the application running on the wire (irrespective of port and protocol). This is highly useful, especially when you are monitoring specific traffic, such as monitoring for SSL performance issues or problems connecting to a Microsoft SQL server. However, the Shark, unlike Sensor, can write packets to disk at very high rates. This enables deeper analysis of more packet data and allows enables you to view full packets captured at wire speeds. If you require application identification and detailed packet level analysis, you might want to deploy both a Shark and a Sensor in the data center.

There is as much value in deploying a Gateway in this scenario versus a scenario without a Sensor or Shark. The Gateway can receive flow data from a variety of sources including traffic optimization information from Steelhead appliances. Additionally, you can configure the Gateway to receive data from multiple sources on unreliable UDP ports, and forward the data in an encrypted and compressed format using reliable TCP ports.

Figure 2-10 shows an example deployment that includes the Enterprise Profiler, the Shark, the Pilot, the Sensor-VE, and the Gateway. Two Shark appliances are placed in front of the critical server farms in two data centers. This enables you to measure response time and also to perform deeper analysis of application traffic for critical packet-level troubleshooting. All routers and Steelhead appliances across the entire enterprise network collect flow data, which enables broad visibility everywhere. Two Gateways are deployed for flow collection, and flow source devices are configured to send their flow to the Gateway that is geographically closest. Because deeper visibility is required at the remote sites, where you cannot obtain standard NetFlow, you can deploy a Sensor-VE to gain visibility into local traffic that never leaves the remote sites.

Figure 2-10. Enterprise Profiler, Shark, Pilot, Sensor-VE, and Gateway

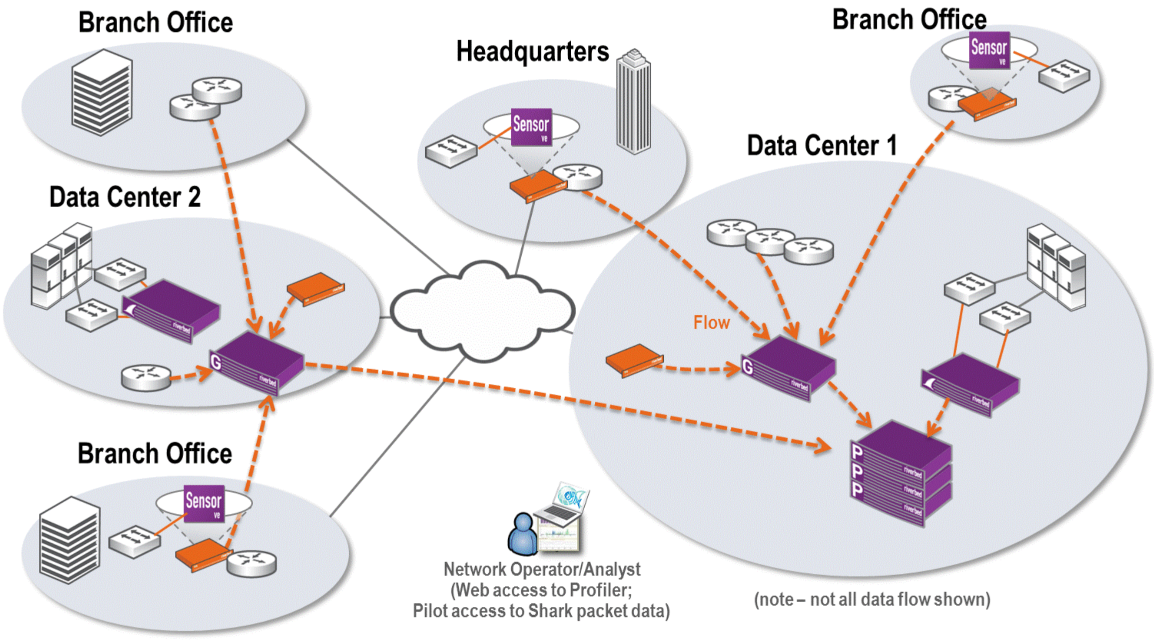

Port and Protocol Dependencies

To assure that the Cascade components communicate, you must open up ports across any existing firewalls. 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 on external integrations, see "Additional Cascade Integration" on page 73.

This section describes the following:

- "The Pilot and the Shark Port Dependencies" on page 29
- "The Profiler and the Gateway Port Dependencies" on page 30
- "Cascade Full-Solution Port Dependencies" on page 31
- "Cascade Enterprise Solution Port Dependencies" on page 32

The Pilot and the Shark Port Dependencies

Figure 2-11 shows which ports and protocols you must have open for communications within a Pilot and a Shark deployment. External connections are optional, depending which integrations you use.

Figure 2-11. Pilot and Shark Dependencies

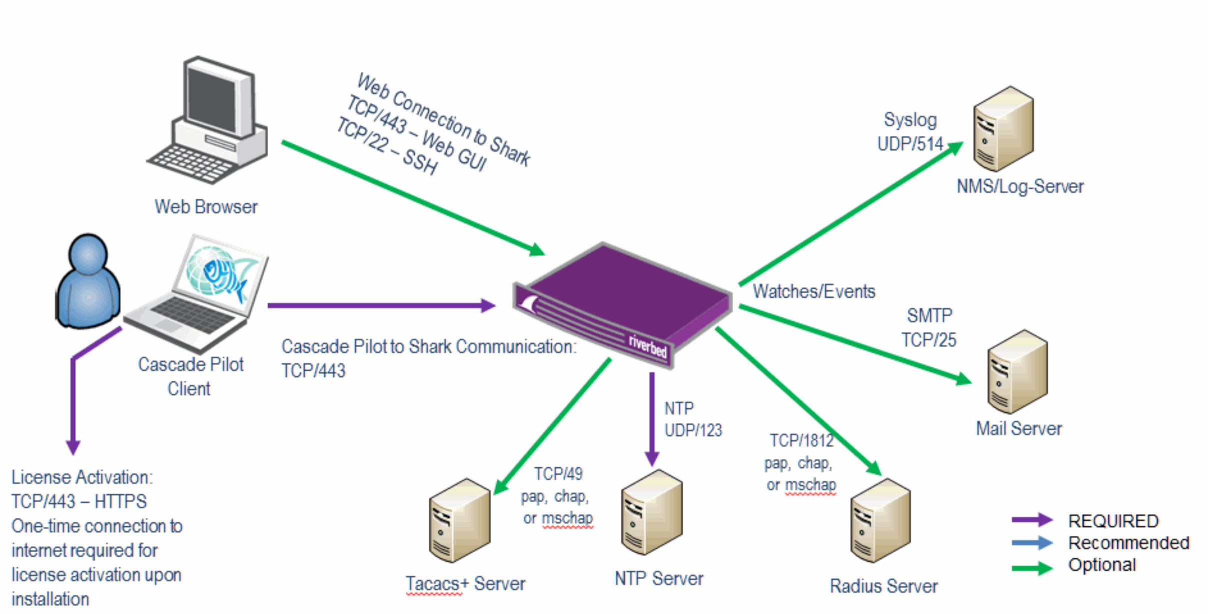

The Profiler and the Gateway Port Dependencies

Figure 2-12 shows which ports and protocols you must have open for communication within a Profiler and a Gateway deployment. External connections are optional, depending which integrations you use.

Figure 2-12. Profiler and Port Dependencies

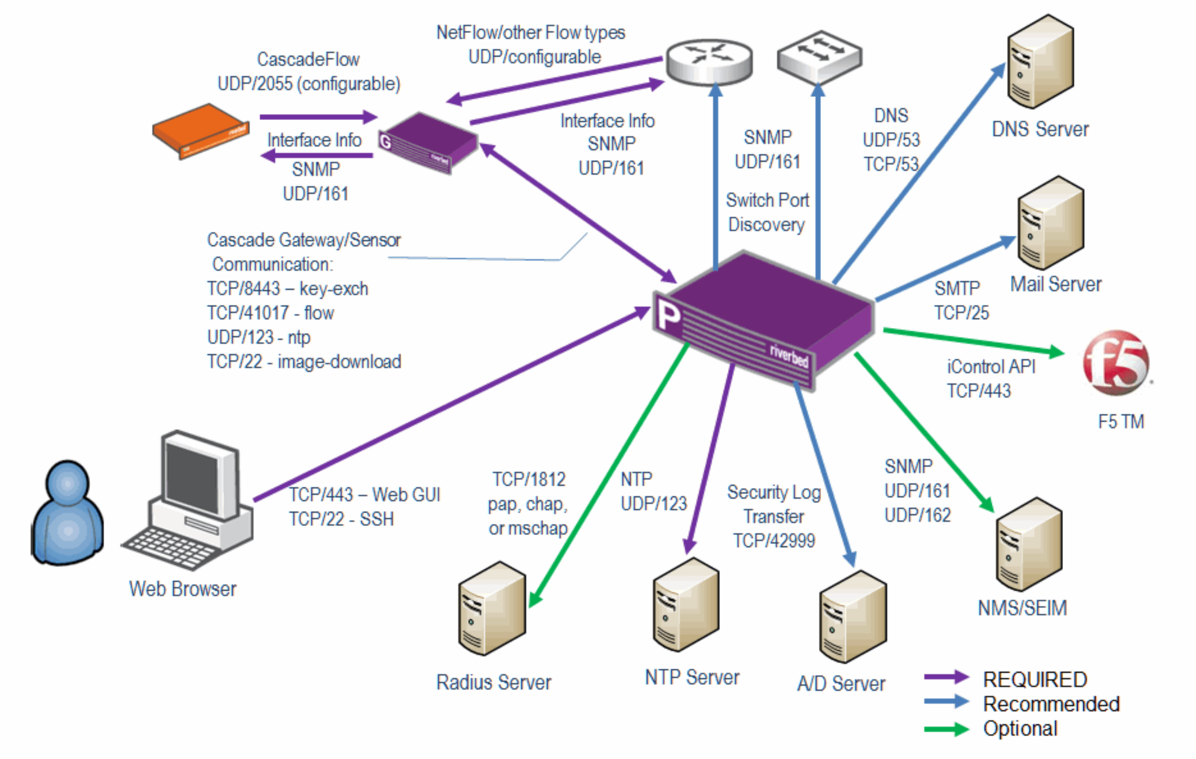

Cascade Full-Solution Port Dependencies

Figure 2-13 shows which ports and protocols you must have open for communications within a Profiler, a Gateway, a Sensor, and a Shark deployment. External connections are optional, depending which integrations you use.

Figure 2-13. Full Solution Port Dependencies

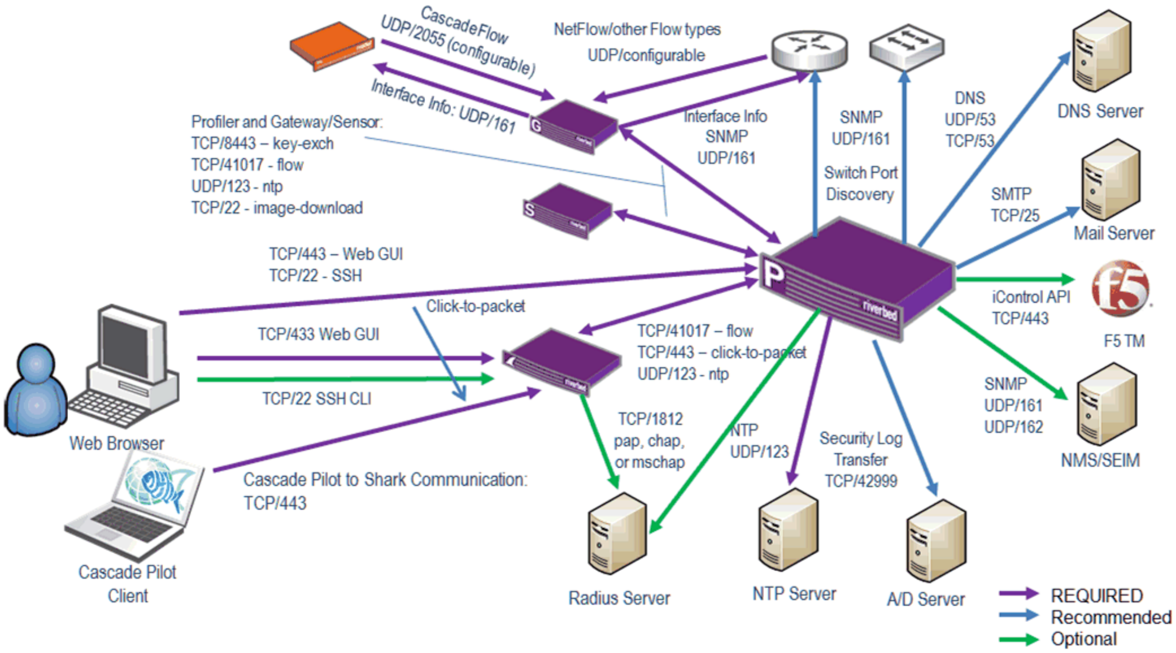

Cascade Enterprise Solution Port Dependencies

Figure 2-14 shows which ports and protocols you must have for communications within an Enterprise Profiler, a Gateway, a Sensor, and a Shark deployment. Many external connections are optional, depending upon integrations you use. You must have all components that compose the Enterprise Profiler on the same network subnet, using 1-Gb ports and preferably connected to the same switch.

Figure 2-14. Enterprise Solution Port Dependencies

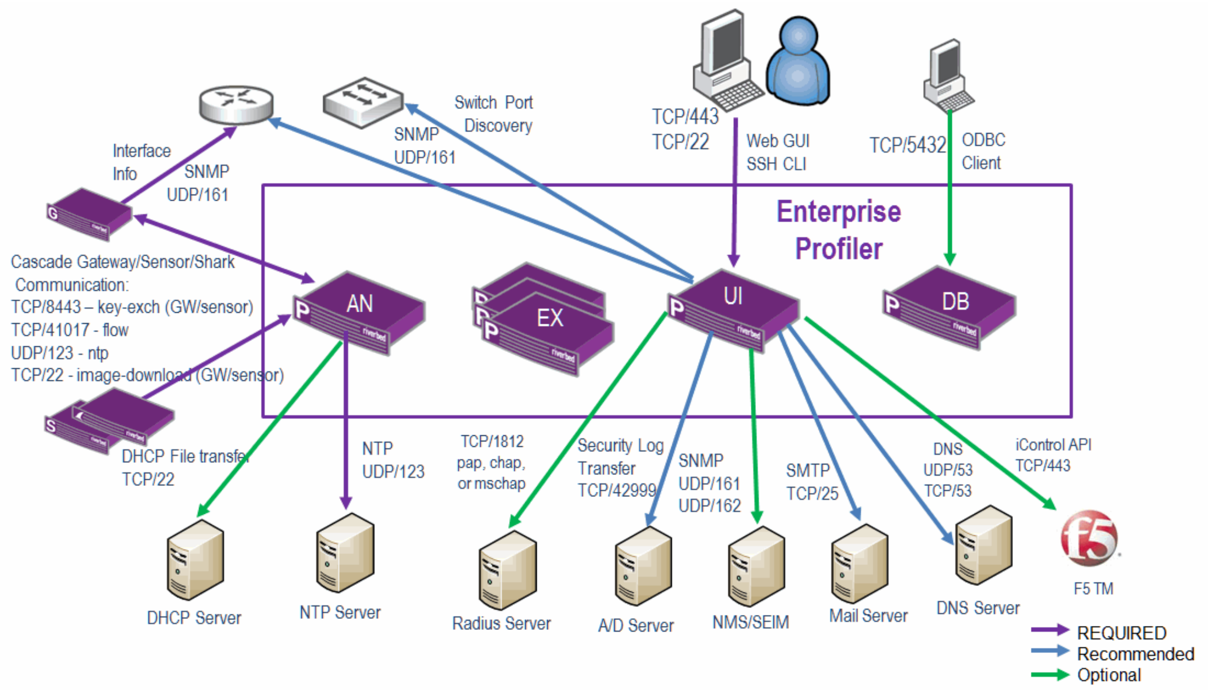

Cascade Flow Storage

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

- "Types of Flow Storage" on page 32
- "Flow Rate Estimation" on page 33
- "Flow Storage Size Estimation" on page 34

Types of Flow Storage

Cascade supports the following types of flow storage:

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

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

Local Storage

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

System Type	Physical Disks	Partitions	Flow Storage (Raw)
Express (CAX 100/200/300)	2	1	500 GB
CAP-02120	2	1	1000 GB
CAP-02140/02160	6	2	4000 GB
CAP-04110	6	2	4000 GB

Systems with a single partition use the partition to store the flow information, and the boot and configuration information for the Profiler. On systems with two partitions, one partition is used for the system software (configuration and so on) and the other is dedicated to flow storage.

System Type	Physical Disks	Flow Storage (Raw)
CAX-360 (L/M/H)	2	3.7 TB
CAP-2260 (L/M/H)	12	11.8 TB
CAP-4260	16	11.8 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. The cost of the storage is included in the cost of the Profiler system; no additional expenditure on external storage is required.

The biggest disadvantage of local storage is that is has expansion limits. There is no option for internal disk expansion.

Storage Area Network (SAN)

SAN provides the most robust external solution. SAN is a reliable solution that allows the Cascade device to very quickly access four or more petabytes of storage (limited by the JFS/2 file system used for the device).

When you connect a SAN to the Cascade system, the SAN functions as a new disk partition. Cascade treats the SAN as another disk that is part of the Profiler, enabling it to easily access the data. You can also offload much of the 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 and the SAN when you connect a SAN to an Enterprise cluster. This increases the performance drastically, because each analyzer is given its own dedicated channel to the SAN, and potentially has its own set of drives on the SAN (depending on the SAN configuration).

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 the host accesses internal hosts, or internal and external hosts.

You must take into account the following when estimating flows:

- Estimate your flows per minute.

  In general, you can expect 20 flows per host per minute across all internal systems (client systems, printers, servers, and so on). Because different system send different numbers of flows, it is probable that some hosts send much more then 20 flows per minute, while some systems send fewer.

  Riverbed recommends that you give serious thought when the estimate approaches the limits of a system; for example, if the estimate indicates 49,000 flows per minute you might want to consider if a system capable of supporting 50,000 flows is too restrictive.

- Consider how much of the traffic the Profiler detects.

  A network with 10,000 hosts, only 25% of which forwards traffic to the Profiler, has a very different requirement then a network with 5,000 hosts, that all forward traffic. The first example has an estimated 50,000 flows per minute, and the second example has an estimated 100,000 flows per minute. You must know the total number of hosts on a network and the number of hosts that have their flows forwarded.

- Estimates must include the current traffic level and what is expected to occur in the near 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 mis-sizing of the system. 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.

- Know which hosts are talking to each other.

  This is most difficult factor to quantify into an estimate. The 20 flows per host per minute number is based on conversations between only internal hosts (for example, internal clients and servers). The flow processing requirement changes drastically if a host is also accessing external systems (for example, Web sites), and those flows are being processed by the Profiler. An increase of 100% (to 40 flows per minute per host) is not unreasonable for hosts accessing Internet sites.

Flow Storage Size Estimation

Another primary sizing decision factor is how long you want to store flows. Because each flow currently occupies approximately 300 bytes, it is theoretically possible, depending on the platform, to store up to tens of billions of individual unique flows.

Note: This estimation assumes that you have 50% of available 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 pre-aggregated data, called *rollups*.

If an xx60 Express Profiler receives one flow per minute, it can store flows covering a time span of approximately 6,450 years. On a more realistic network sending 50,000 flows per minute, the same Express can store approximately 47 days worth of flow data.

If you need to store flows for 180 days (for example, for legal reasons), then the storage included on the Express is inadequate. You must upgrade to a larger system with additional storage capacity. When you determine the upgrade path to take, ensure that the path you choose is the most logical approach. There are two possible upgrades:

- Standard Profiler (with 11.8 TB of storage built in)
- Enterprise cluster (with 11.8 TB of storage in a base cluster and the ability to add an additional 82.6 TB of storage)

When you choose the appropriate upgrade, take into account the desired flow storage capacity and the cost of the system. While the Enterprise cluster might meet the needs of the network, unless you expect significant growth in the near future, it is unlikely to be cost effective.

CHAPTER 3 Flow Collection for Cascade

This chapter describes the flow collection for Cascade devices. It contains the following sections:

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.

A Bluecoat Packeteer shaper supports flow-detail-records v2 (FDR) for application identifier collection. If the device is a Steelhead appliance, Riverbed recommends that you use CascadeFlow instead of NetFlow to collect retransmit and response time information. If the device supports multiple types of flow formats, NetFlow is preferable to sFlow.

For details about collecting data from Cascade Sharks and Sensors, see "Packet Collection for Cascade" on page 47. For details about collecting CascadeFlow from Steelhead appliances, see "Cascade Steelhead Appliance Integration" on page 59.

Base Requirements

You must meet the following requirements to set up your router:

- Configure devices that support NetFlow for NetFlow v1, v5, v7, or v9 with no aggregation and no sampling. Riverbed recommends that you use v5 or v9.
- Configure devices that support sFlow for sFlow versions v2, v4, or v5 with the lowest possible sampling rate. Riverbed recommends that you use v5.
- Configure devices to export flow to the Express or Gateway management interface. Use default-configured destination port on the Cascade equipment.
- Synchronize devices with an NTP server. Riverbed recommends that you synchronize devices with the same NTP server used by the Profiler. For proper operation and reporting, the timestamps on the network equipment and Cascade equipment must be synchronized to the same time.

- Set the active timeout setting for flows to 60 seconds.

Note: Cisco IOS versions show this timeout in either minutes or seconds.

- Do not adjust the inactive timeout setting from the default setting of 15 seconds. If you must, the time 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 actively use. Because NetFlow v5 is typically ingress only, you can calculate egress only by aggregating ingress from the other 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 TTL command. This enables ordered-path reporting in the Profiler. 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**.
- You must configure SNMP access to any devices sending flow to the Profiler. Standard flow export provides information with only SNMP ifindex values. By enabling SNMP on these devices, Cascade 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 73.

Additional requirements and considerations for Cisco equipment include:

- 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 MSFC and 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 Profiler 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 5000 or Nexus 7000 series, and you are using NX-OS v4, you must have a minimum release 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 statistics and bits-per-second.
- NetFlow export from the Cisco ASA is does not include standard NetFlow records. Cisco ASA exports NSEL (NetFlow Event Log) in a NetFlow wrapper. NSEL is event driven, exporting bytes only for the first and last packet in the flow. There is no concept of an active timer, so you do not get regular updates. NSEL is considered a replacement for syslog that enables the ASA to scale. Cascade does not currently support reporting from NSEL.
- 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 <vlanlist>
```

Flow Data Fields Consumed by Cascade

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

Field Name	Description	Flow Versions with Support
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	Timestamps 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 Cascade Sensors and Packeteer through FDR records
Retransmitted bytes and retransmitted packets	TCP transmission counters	CascadeFlow from Steelhead appliances, Sensors, and Sharks
Network round-trip time	Measurement of round-trip time across the network	CascadeFlow from Steelhead appliances, Sensors and Sharks
Total response time, server delay, client delay	Measurement of response time metrics across the network	CascadeFlow from Sensors and Sharks

Flow Type Considerations

If you use Steelhead appliances, Riverbed recommends that you always use CascadeFlow. For all other devices, if you have a choice between getting data from devices supporting NetFlow versus devices supporting sFlow, Riverbed recommends that you use NetFlow. sFlow is typically supported only with sampling and is less accurate than NetFlow.

Riverbed recommends that you use NetFlow v9 with TTL. Riverbed recommends that you export the TTL field so that network segment graphs are available in the Profiler. If version 9 is not available, use NetFlow v5.

Flow Collection Considerations

The Profiler deduplicates all flows that it receives for the same five pieces of information: source IP address, destination IP address, source port, destination port, and protocol. Cascade tracks every source device that sent the flow, for up to a total of five in-path devices (up to 10 total interfaces) per flow. You should gather flow from as many sources as possible, because the Cascade license is based on total flow volume (versus number of devices or number of network interfaces for which flow is received). Do not forget the five-device (up to 10 interface) limit.

If you exceed the five-device limit, Cascade prioritizes the devices kept per flow in the following order:

- Steelhead appliances
- Sharks and Sensors
- Sensor-VEs
- All other flow sources

If you exceed the five-device limit within the same category, the device with the lowest IP address is included.

Flow Collection in Virtual Environments

The vSphere v5 vSwitch and the Cisco Nexus1000V 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.

Sample Third-Party Configurations

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

This section contains the following:

- "Configuring Cisco 6500 Series Switches Running Native IOS" on page 41
- "Configuring Cisco 6500 Series Switches in Hybrid Mode" on page 42
- "Configuring Cisco 7500 Series Router" on page 43
- "Configuring Cisco 7600 Series Router" on page 43
- "Configuring the Cisco Nexus 7000 Flexible NetFlow" on page 43
- "Configuring NetFlow Export for Cisco Nexus 1000V" on page 44
- "Configuring IPFIX for Avaya (Nortel) 8300 and 8600" on page 45
- "Configuring sFlow for HP Procurve 3500, 5400, and 6200" on page 46

Configuring Cisco 6500 Series Switches Running Native IOS

The following example uses the native IOS command line interface to configure the SUP and MSFC modules of a 6500 series switch. The following commands generally work with IOS v12.2 or later, except where specified. For further information, refer to the documentation for your version of IOS.

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 <cascade-gateway-or-express_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
```

Note: If you are running IOS v12.2(18) or later, use NetFlow v9. If NetFlow version 9 is not available, use version 5).

If you are running IOS v12.3(14) or later and are exporting NetFlow v9, you can include export of the TTL, enabling Cascade to show network segment diagrams:

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

If you are running IOS v12.3(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
```

Next, 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
```

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

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

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 <gateway-or-express_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 commands 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 IOS CLI to configure a Cisco 7500 series router.

To configure a Cisco 7500 series router using the IOS CLI

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

   ```
   Router# confg t
   Router(config) # ip flow-export <cascade-gateway-or-express_ip> <udp_port_number> <version>
   ```

2. Enter the following commands 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 IOS command line interface to configure a Cisco 7600 series router.

To configure a Cisco 7600 series router using the IOS CLI

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

   ```
   Router(config)# ip flow-export <cascade-gateway-or-express_ip> <udp_port_number>
   Router(config)# ip flow-export <version>
   Router(config)# mls nde sender <version>
   ```

2. Enter the following commands 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 the Cisco Nexus 7000 Flexible NetFlow

The following example uses Cisco Nexus OS v5.2.1 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 Cascade:

```
Switch# config t
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
```

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

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

3. Enter the following commands to configure NetFlow export:

```
Switch# config t
Switch(config)# flow exporter cascade-export
Switch(config-flow-exporter)# destination <cascade-gateway-or-express_ip>
Switch(config-flow-exporter)# source ethernet 2/1
Switch(config-flow-exporter)# transport udp 2055
!--- Listening port configured on Gateway
Switch(config-flow-exporter)# version 9
```

4. Enter the following commands to configure flow monitor:

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

5. Enter the following commands to apply a flow monitor to a VLAN or interface:

```
Switch# config t
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

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

   ```
   n1000v# config t
   n1000v(config)# flow exporter Cascade-export
   n1000v(config-flow-exporter)# destination <cascade-gateway-or-express_ip>
   n1000v(config-flow-exporter)# source mgmt 0
   n1000v(config-flow-exporter)# transport udp 2055     (Listening port configured on 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
   ```

2. 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 veth 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. Similar commands are used for configuring 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 Cascade compatibility (active timeout 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 Cascade:

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

or

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

Configuring sFlow for HP Procurve 3500, 5400, and 6200

The following example uses Procurve 3500, 5400, and 6200 to configure flow export. Similar commands are used for configuring other HP Procurve devices.

To configure IPFIX for Avaya (Nortel) 8300/8600

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

```
ProCurve# configure
ProCurve(config)# sflow 1 destination <gateway-or-express-ip-address> <gateway-or-express -
listening-port>
```

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 in using this 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 in 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 mem
```

CHAPTER 4

Packet Collection for Cascade

This chapter describes the different methods for Cascade packet capture. You use packet capture for traffic monitoring and packet analysis. This chapter contains the following sections:

Cascade Components for Packet Collection

Cascade collects packets using one of the following components:

- **The Shark** - Traffic capture and monitoring for high-rate packet capture and analysis
- **The Sensor** - Traffic monitoring with Layer-7 application identification
- **Sensor-VE** - Traffic monitoring running remotely on a Steelhead appliance Riverbed Services Platform(RSP)
- **The Express with built-in Sensor capability -** Port traffic monitoring with Layer-7 application identification

You can forward flows from these devices to the Profiler based upon the packets you collect. In addition to standard NetFlow-type fields, these components send TCP health information (TCP retransmission) and TCP performance information (response times). For more information about these components, see "Cascade Overview" on page 5 and "Cascade Component Deployment Scenarios" on page 9.

SPAN and 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 capture on ingress, egress, or both, on the interface or VLAN 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 Shark and the Sensor have two or more monitoring ports that enable you to duplicate this configuration multiple times using the same Shark or Sensor.

Figure 4-1. SPAN Connectivity

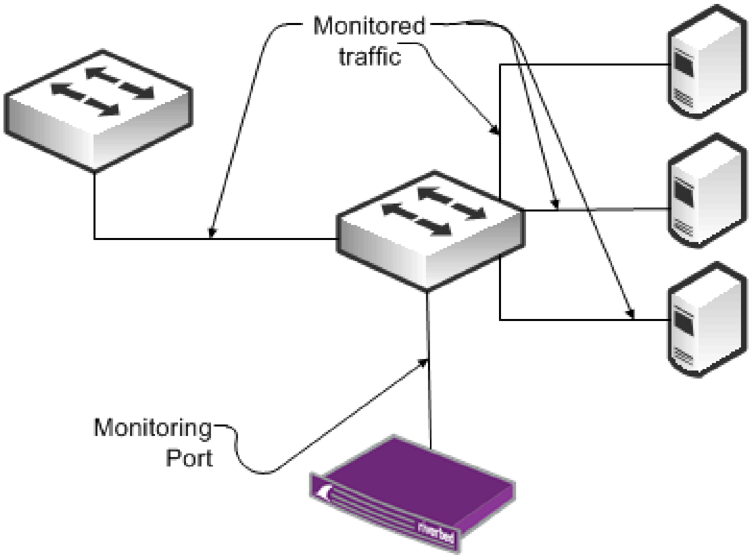

Best practices for port mirroring:

- For most monitoring and troubleshooting, you must collect both sides of the conversation. This means that if you are capturing only one port, you must mirror both directions (ingress and egress). If you are monitoring all ports or all VLANs that are communicating, you can capture ingress only. It is redundant to capture ingress and egress on all ports or all VLANs, and the duplicate traffic is deduplicated on the Sensor or at the Profiler level.
- When you set up port mirroring, you must follow best practices according to your switch vendor. Because many architectures use non-blocking 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 might 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 Anue Systems, NetOptics, Gigamon, cPacket Networks, and VSS Monitoring.

- Many switches have a limit on the number of monitoring ports that you can configure. This limit is often 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 the following sections in this chapter.

RSPAN and ERSPAN

This section describes the following SPAN variations:

- "RSPAN" on page 49
- "ERSPAN" on page 50

Riverbed recommends RSPAN and ERSPAN techniques in special circumstances. 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

Remote SPAN (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 Cascade 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 Shark and the Sensor have two or more monitoring ports that enable you to duplicate this configuration multiple times using the same Shark or Sensor.

Figure 4-2. RSPAN Connectivity

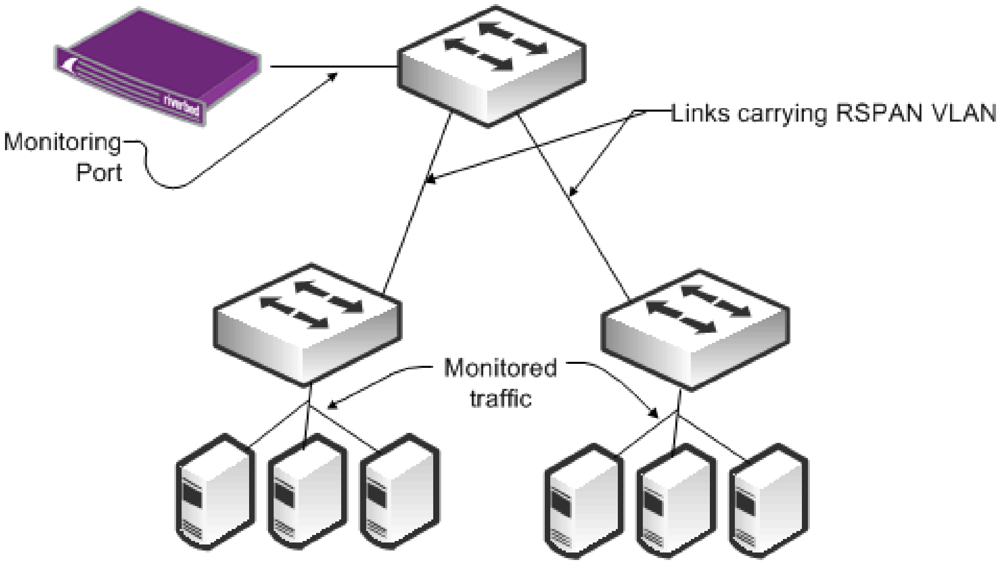

ERSPAN

Encapsulated remote SPAN (ERSPAN) enables an extension of a SPAN over the network to another switch through a routed GRE-encapsulated tunnel. You can use ERSPAN when Cascade 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.

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 Shark and the Sensor have two or more monitoring ports that enable you to duplicate this configuration multiple times using the same Shark or Sensor.

Figure 4-3. ERSPAN Connectivity

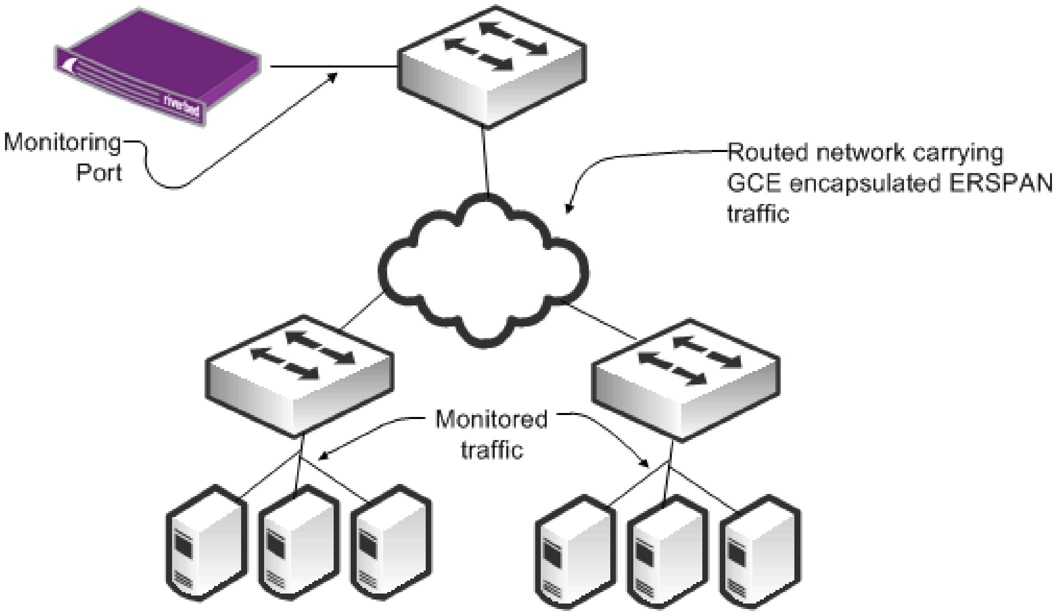

You must use ERSPAN in an virtualized environment that uses the Cisco Nexus 1000V. The Cisco Nexus 1000V mirrors traffic sent between virtual machines by sending ERSPAN to an external Cisco Catalyst 6500 switch.

Sample Port Mirror Configurations

This section contains the following SPAN port configuration examples:

- "Cisco Catalyst 6500 SPAN" on page 51
- "Cisco Nexus 5000 SPAN" on page 51
- "Cisco Nexus 1000V ERSPAN to Cisco Catalyst 6500" on page 52

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 Catalyst 6500 SPAN

The following steps illustrate how to configure a SPAN for all traffic for VLANs 1 through 100 using a Cisco Catalyst 6500 SPAN. You need only capture ingress (rx) 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 to be monitored:

   ```
   Switch# conf t
   Switch (config)# monitor session 1 source vlan 1-100 rx
   ```

2. Enter the following command to configure the destination port where the Cascade monitoring port is connected:

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

The following example shows capturing all traffic to and from sources on the downstream port 5/1 and sending 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 to be monitored:

   ```
   Switch# conf t
   Switch (config)# monitor session 1 source gigabitethernet 5/1 both
   ```

2. Enter the following command to configure the destination port where the Cascade monitoring port is connected:

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

Cisco Nexus 5000 SPAN

The following steps illustrate how to configure a SPAN for all traffic for VLANs 1 through 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 through 100.

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

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

   ```
   Switch# conf t
   Switch (config)# monitor session 1
   Switch (config-monitor)# exit
   Switch (config)#
   ```

2. Enter the following commands to configure the destination port to which the Cascade 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 to be monitored:

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

The following example shows all traffic SPANing 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# conf t
   Switch (config)# monitor session 1
   Switch (config-monitor)# exit
   Switch (config)#
   ```

2. Enter the following commands to configure the destination port to which the Cascade 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 to be monitored:

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

Cisco Nexus 1000V ERSPAN to Cisco Catalyst 6500

The following steps illustrate how to configure an ERSPAN for Cisco's virtual switch, 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 through 10

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

   ```
   Switch# conf t
   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# conf t
Switch (config)# monitor session 1 type erspan-destination
Switch (config-monitor)# desc CascadeErspanDest
```

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
```

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 Shark or the Sensor for one link that you monitor. This is because the tap uses a separate port to copy each direction's traffic.

  Figure 4-4 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 4-4. Basic Tap Connectivity

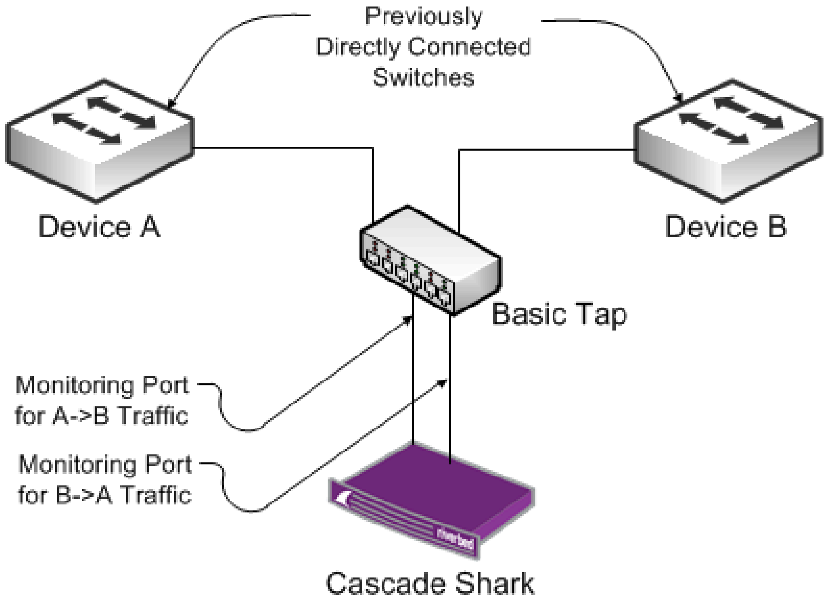
Previously Directly Connected Switches
Device A
Device B
Basic Tap
Monitoring Port for A->B Traffic
Monitoring Port for B->A Traffic
Cascade Shark

- **Regeneration taps** - Send the same traffic for the same link being monitored to multiple devices. These taps are useful if you want to send traffic from link to both Cascade and another device, for example, an IDS.
- **Aggregation taps** - Enable 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 Shark or the Sensor for a link you want to monitor. If you use this method, you can potentially miss some packets if the full-duplex link is running close to line rate in both directions.

  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*).
- **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.
- You can use taps on existing SPAN and port-monitoring ports. This is useful if there are no longer any more 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 Cascade.

VACL Configuration Examples

You can use a VLAN Access Control List (VACLs), which is used to mirror ports, for cases when your switch supports only a limited number of in-use SPAN ports. This section contains the following examples:

- "VACL Port Mirroring Configuration on Cisco 6500 Running CatOS" on page 54
- "VACL Port Mirroring Configuration on Cisco Catalyst 6500 Running IOS" on page 55

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 Shark or the Sensor 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 CascadeMonitor permit any any capture
> show security acl info CascadeMonitor editbuffer
```

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

```
> commit security acl CascadeMonitor|all
```

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

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

4. Enter the following commands to specify that the capture port on which Cascade monitoring is connected (allows 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 IOS

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

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

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

```
Switch# conf t
Switch(config)# ip access-list CascadeMon
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 CascadeMon
Switch (conf-map_name)# action forward
```

or

```
Switch (conf-map_name)# action forward capture // Depending on 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
```

Shark Passthru

The only reason to use passthru is to tap a SPAN port that another device is already using: for example, a Cascade Sensor. 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 Shark appliance to monitor a single SPAN port.

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

Note: The passthru function does not fail-to-wire. This means that if the Shark appliance loses power or stops operating for whatever reason, the link will not pass traffic.

Packet Deduplication

Depending upon the packet capture method you use, you might send multiple copies of the same packet to the Shark or the Sensor. 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 port deduplication on Cascade. Packets can be deduplicated by the Shark or the Sensor when necessary. You must enable this feature on the Shark, but it is enabled by default on the Sensor. Both appliances deduplicate packets by evaluating the packet identifier along with other information in the packet. Deduplicated packets enable TCP retransmissions to be captured while duplicate packets due to instrumentation are dropped. The Shark performs duplication on a per-port basis.

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 Shark or the Sensor, 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 appliance and Steelhead appliance and are using full transparency, the packets appear as retransmissions to Cascade if the packets are captured is on the same port as the originating packets. To avoid this, do not capture traffic between the Interceptor appliance and Steelhead appliance if configured with full transparency.
- If you are capture traffic on either side of a CISCO ASA Firewall to the same port on a the Shark or the Sensor, 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 Shark or Sensor 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 enables you to store of more packets, but you might not be able to inspect the full packet content.

In most cases, you want to configure the Shark to capture full packets. If this is your case, do not adjust the default snaplen parameter. Do not adjust this parameter. This enables full-packet analysis within the Pilot. If you adjust the snaplen to be smaller, some Pilot views cannot fully analyze the data. For example, volumes of data represented can appear to be smaller than they actually are, and more detailed views (for example, the Web Bandwidth by Object - Advanced view) do not contain all the data necessary for full analysis.

Because the Sensor is not typically used for full-packet analysis, the snaplen is set to 250 bytes by default. This enables TCP/IP headers and some application-level headers to be captured and viewed for a larger number of packets. However, because the Sensor performs packet sampling, it is not useful to view a full sequence of packets. In either case, setting snaplen on the Shark or the Sensor does not affect the accuracy of the flow data sent to the Profiler.

You cannot use snaplen with Express Sensor ports and Sensor-VE to view packet data.

CHAPTER 5

Cascade Steelhead Appliance Integration

This chapter describes how to configure Cascade components and the Steelhead appliance into an integrated solution for traffic monitoring and troubleshooting. When you integrate Profiler and the Steelhead appliance 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 Appliance and Cascade Overview" on page 59
- "Cascade Deployment Considerations" on page 62
- "Configuring Steelhead Appliance for Flow Data Export" on page 65
- "Configuring Cascade Sensor-VE on RSP" on page 67

Steelhead Appliance and Cascade Overview

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

- "NetFlow Versus CascadeFlow" on page 60
- "SNMP Interface Persistence (ifindex)" on page 61

The two primary integration points of the Steelhead appliance and Cascade are CascadeFlow export from the Steelhead appliance, and Sensor-VE (Virtual Edition):

- The Steelhead appliance sends the Gateway or the Express an enhanced version of NetFlow called Cascade*Flow*. CascadeFlow includes:
  - NetFlow v9 extensions for round-trip time measurements that allow you to understand volumes of traffic across your WAN and end-to-end response time.
  - extensions that allow Cascade to properly measure and report on the benefits of optimization.
- Sensor-VE collects packet level data and sends the resulting flow data to a standard Profiler or an Express. Sensor-VE monitors traffic in one of the following ways:
  - From the Steelhead appliance's auxiliary port
  - Traffic that passes through the Steelhead appliance internally

You can use SPAN or port mirror switches deployed in proximity to the Steelhead appliance when you monitor traffic from the Steelhead appliance's auxiliary port.

Note: In RiOS v7.0.1 and later, RSP was replaced with Virtual Services Platform (VSP). VSP comes preinstalled in the Steelhead EX appliance. For more information about VSP, see the *Steelhead EX Management Console User's Guide*.

NetFlow Versus CascadeFlow

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 Gateway. When the records are intelligently coalesced and stored they can be used to understand traffic throughout the network for reporting, troubleshooting, and automatic detection of network and application issues.

There are two versions of standard NetFlow supported by the Steelhead appliance: v5 and v9. Use only these standard versions when exporting flow to non-Cascade flow collectors or an earlier version of Cascade. CascadeFlow was created by Riverbed. It extends NetFlow v9 with 12 custom extensions that provide additional details specific to flows passing through the Steelhead appliance.

The following table shows Cascade feature compatibility in RiOS v6.0 and later.

Feature	Flow Version	Cascade Version
Basic traffic reporting	NetFlow v5 and v9	Profiler v8.0 and later
Enhanced reporting: • Automatic identification of Steelhead appliance pairs • Automatic identification of LAN-WAN interfaces • WAN optimization reporting	CascadeFlow compatible (v5-based)	Profiler v8.3 and later
Enhanced reporting: • Automatic identification of Steelhead appliance pairs • Automatic identification of LAN-WAN interfaces • WAN optimization reporting • End-to-end response time measures for both optimized and unoptimized sessions	CascadeFlow (v9-based)	Profiler v8.4 and later
Sensor-VE	Not applicable; the flow is sent directly from Sensor-VE	Profiler v8.4 and later

When you use CascadeFlow, the Steelhead appliance 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, Cascade merges these flow records together with flow data collected for the same flow from other devices.

Figure 5-1. NetFlow v9

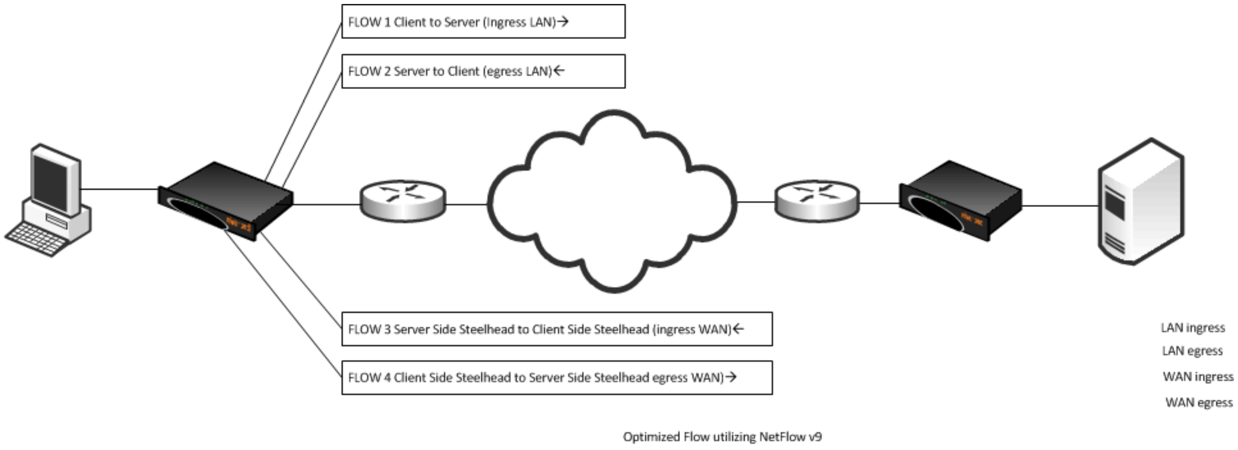
FLOW 1 Client to Server (Ingress LAN)→
FLOW 2 Server to Client (egress LAN)←
FLOW 3 Server Side Steelhead to Client Side Steelhead (ingress WAN)←
FLOW 4 Client Side Steelhead to Server Side Steelhead egress WAN)→
LAN ingress
LAN egress
WAN ingress
WAN egress
Optimized Flow utilizing NetFlow v9

SNMP Interface Persistence (ifindex)

One of the more common Cascade reporting issues happens when interface names change across reboots of the Steelhead appliance, which is a direct result of nonpersistent SNMP interface index names.

SNMP uses index values to reference interface names. Only index values are included in the flow record when the Steelhead appliance sends flow data to Cascade. An SNMP poll from Cascade to the Steelhead appliance 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 Profiler to incorrectly map the Steelhead appliance interfaces. Use the **ifindex-persistence** command on the Steelhead appliance to permanently pin SNMP interface names to SNMP index values.

To enable ifindex persistence

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

```
sh > en
sh # config t
snmp-server ifindex-persist
#--- You must restart the optimization service, so that netflow can use the configured ifindex
sh # exit
sh # write mem
sh # show service restart
```

To verify that SNMP persistence is enabled

- On the Steelhead appliance, connect to the CLI, enter the following command, 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 on ifindex values, see "SNMP Integration for Flow Sources" on page 73.

Cascade Deployment Considerations

This section provides recommendations on how to improve the accuracy of the exported flow data when you deploy Cascade in specific Steelhead appliance deployment scenarios, and describes certain Steelhead appliance features. It contains the following sections:

- "In-Path Deployments" on page 62
- "Virtual In-Path Deployments" on page 62
- "Out-of-Path Deployments" on page 64

In-Path Deployments

In an in-path configuration, Steelhead appliances are placed in the physical path of the client and server, where they detect all traffic. You can source flow export from either the Steelhead appliance primary or auxiliary interface to export flow data to the Profiler. Enable flow export for all traffic on all Steelhead appliance interfaces that have traffic traversing them (for example, lan0_0 and wan0_0 for a single in-path Steelhead appliance deployment).

Simplified Routing

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

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

Virtual In-Path Deployments

A virtual in-path deployment happens when traffic is directed to the Steelhead appliance through WCCP, PBR, the Interceptor appliance, or a Layer-4 load balancer. All traffic might not traverse the Steelhead appliance because of router redirection access control lists (ACLs). You must rely on NetFlow from the router, as well as the Steelhead appliance, to ensure that you are sending complete flow data to a Gateway or an Express. You only need to select optimized flows to export from the Steelhead appliance to a Gateway or an Express, but you must be sure to export NetFlow from the router for Cascade appliances to detect unoptimized and optimized traffic.

If you configure the virtual in-path environment to use correct addressing, when you export all flow from routers in-path between the Steelhead appliances, additional flows are exported (original IP and port numbering and inner-channel IP and port numbering). Selectively export NetFlow to avoid exporting the inner-channel traffic (unless you wan to track this traffic).

In a virtual in-path deployment, only the WAN interfaces are used for traffic interception. The Profiler optimization-benefit reporting relies on the separation of preoptimization traffic (known as LAN-side traffic), and post-optimization traffic (known as WAN-side traffic). To separately report LAN-side and WAN-side traffic with only the physical WAN interface connected, you must enable a pseudo-SNMP interface that represents the LAN-side traffic. Use the SNMP fake index command to enable the pseudo-SNMP interface. This command causes the Steelhead appliance to report LAN-side traffic as if it is on a LAN interface during the NetFlow export, even though you are not using the physical LAN interface.

Flow export can be sourced from either the primary or the auxiliary interface to export flow data to the Profiler.

To enable fake index

- If you have RiOS v6.0 or later and you want to export CascadeFlow or NetFlow v9, the fake index automatically starts with flow export, and no additional steps are required.

  or,

  If you have RiOS prior to v6.0, or when you want to export a flow type other than CascadeFlow or NetFlow v9, connect to the CLI and enter the following commands:

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

For more details on virtual in-path deployments, see the *Riverbed Deployment Guide*.

Interceptor Appliance Considerations

When you deploy Steelhead appliances virtually in-path and load-balanced by an Interceptor appliance, in addition to ensuring that fake index is enabled and avoiding capturing inner-channel traffic, you must also use care when capturing traffic with the Shark or the Sensor. If you configure the Steelhead appliance with full-transparency, you must exclude traffic between the Interceptor appliance and Steelhead appliances when you configure port mirroring for collection by the Shark or the Sensor.

When you configure port mirroring for collection by Shark or Sensor, you might detect excessive retransmissions and incorrect or missing RTT metrics if you do not exclude traffic between the Interceptor appliance and Steelhead appliances.

Subnet Side Rules

In virtual in-path deployments, configure VLAN subnets to allow the Steelhead appliance to properly determine the direction of non-optimized traffic that passes through the Steelhead appliance. If no LAN subnets are configured, the Steelhead appliance does not discern whether the traffic is passing from the WAN to the LAN or in the opposite direction. This can result in over-reporting traffic in a particular direction or for a particular interface.

In a virtual in-path deployment you must configure LAN subnets on each, because the LAN subnets behind each Steelhead appliance are typically unique.

To configure subnet side rules on a Steelhead appliance

1. On the Steelhead appliance, choose Configure > Networking > Subnet Side Rules.

2. Select Add a Subnet Side Rule.

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.

Figure 5-2. Subnet Side Rules Page

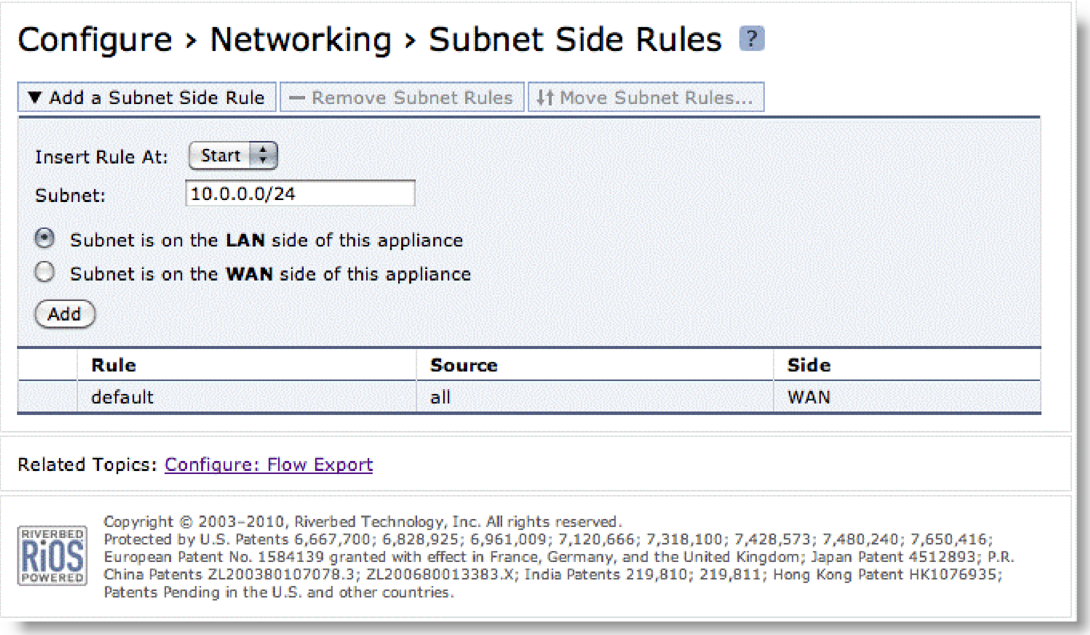

Configure > Networking > Subnet Side Rules ?

Insert Rule At: Start

Subnet: 10.0.0.0/24

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

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

	Rule	Source	Side
	default	all	WAN

Related Topics: Configure: Flow Export

Out-of-Path Deployments

In an out-of-path deployment, the Steelhead appliance is not directly in the data flow between the client and server. Because non-optimized traffic does not always pass through the Steelhead appliance, you must export NetFlow from the router for Cascade to detect non-optimized and optimized traffic. Export only optimized traffic from the Steelhead appliance. To export flow data to the Gateway or the Express, you can export flow from the primary or auxiliary interface.

As with virtual in-path configuration, you must enable SNMP fake index to properly report the direction of the optimized traffic through the Steelhead appliance. You do not need to configure subnet side rules, because the out-of-path Steelhead appliance detects only optimized traffic and never passes through any flows. For more details on SNMP fake index, see "Virtual In-Path Deployments" on page 62.

For more details on out-of-path deployments, see the *Riverbed Deployment Guide*.

Configuring Steelhead Appliance for Flow Data Export

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

To enable flow data export on the Steelhead appliance

1. On the Steelhead appliance, choose Configure > Networking > Flow Export.

2. Select the Configure tab to display the Configuration menu.

3. Select Networking to expand the Networking menu.

4. Select Flow Export to display the Networking > Flow Export page.

5. Select Enable Flow Export.

6. Select Disable Top Talkers.

   Note: Riverbed recommends that you disable top talkers unless you have it enabled for reporting on the Steelhead appliance.

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

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

9. Click **Apply**.

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

Figure 5-3. Flow Export Page

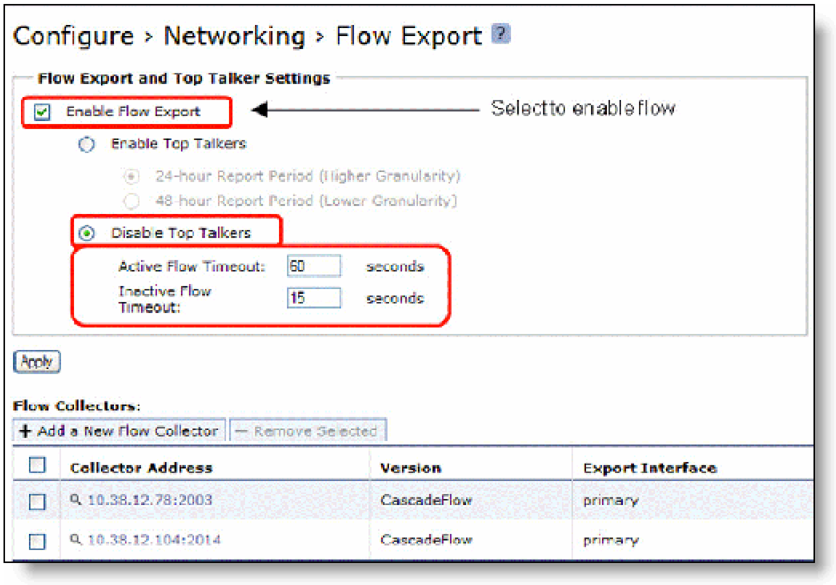

Configure > Networking > Flow Export

Flow Export and Top Talker Settings

☑ Enable Flow Export

Select to enable flow

○ Enable Top Talkers

● 24-hour Report Period (Higher Granularity)

○ 48-hour Report Period (Lower Granularity)

● Disable Top Talkers

Active Flow Timeout: 60 seconds

Inactive Flow Timeout: 15 seconds

Apply

Flow Collectors:

+ Add a New Flow Collector — Remove Selected

□	Collector Address	Version	Export Interface
□	10.38.12.78:2003	CascadeFlow	primary
□	10.38.12.104:2014	CascadeFlow	primary

To configure the Flow Collector and Exporting interfaces

1. On the Flow Export page, select the Add a New Flow Collector tab.

2. Enter the IP address and listening UDP port number of the Express or the Gateway.

3. Select the version of flow data to be exported:
   - For the Profiler v8.4 and later, use CascadeFlow.
   - For the Profiler v8.3.2, select CascadeFlow compatible and select the LAN Address check box.

4. For the desired interfaces, select All to export both optimized and non-optimized flow information.

5. Click **Add** to add the Express or the Gateway to the collector list.

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

Figure 5-4. Flow Collector and Exporting Interfaces

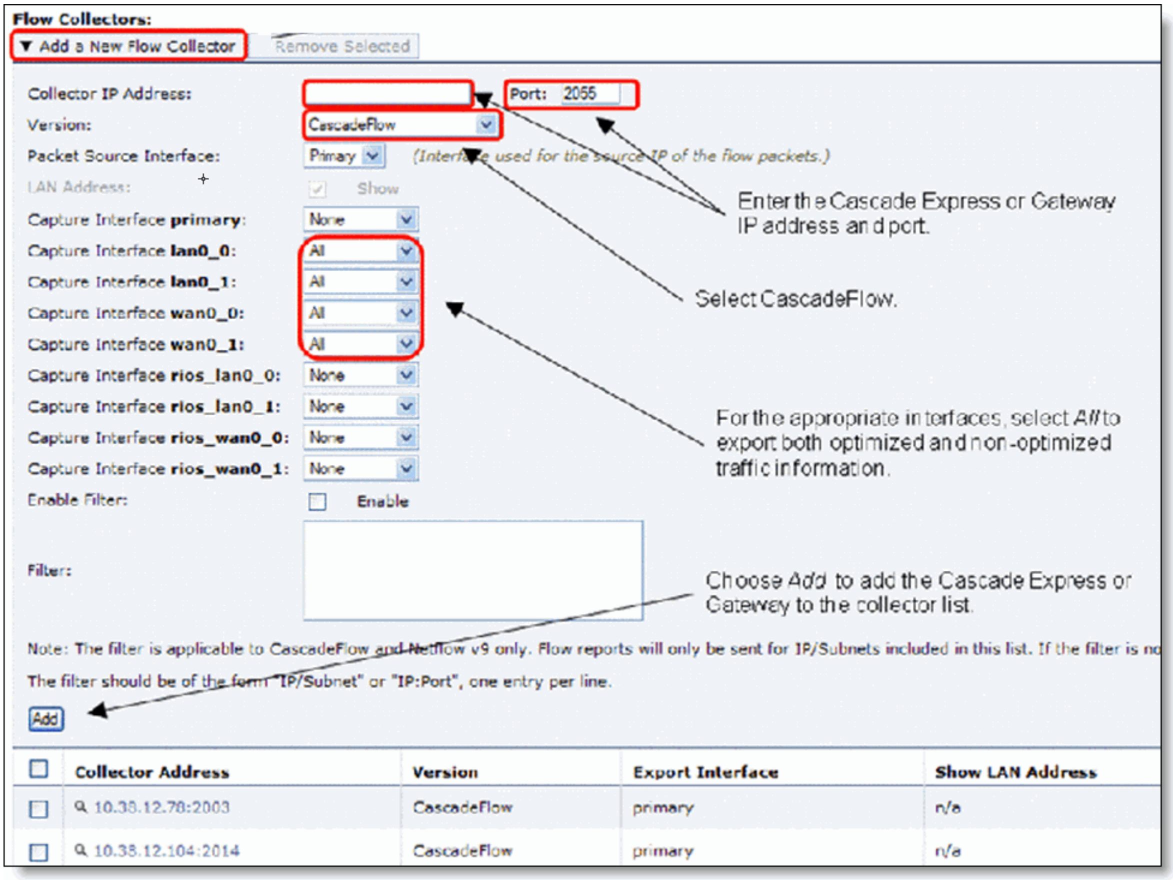

Flow Collectors:

▼ Add a New Flow Collector | Remove Selected

Collector IP Address:	Port: 2055
Version:	CascadeFlow
Packet Source Interface:	Primary (Interface used for the source IP of the flow packets.)
LAN Address:	☑ Show
Capture Interface **primary**:	None
Capture Interface **lan0_0**:	All
Capture Interface **lan0_1**:	All
Capture Interface **wan0_0**:	All
Capture Interface **wan0_1**:	All
Capture Interface **rios_lan0_0**:	None
Capture Interface **rios_lan0_1**:	None
Capture Interface **rios_wan0_0**:	None
Capture Interface **rios_wan0_1**:	None
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 no

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

Add

□	Collector Address	Version	Export Interface	Show LAN Address
□	10.38.12.78:2003	CascadeFlow	primary	n/a
□	10.38.12.104:2014	CascadeFlow	primary	n/a

Enter the Cascade Express or Gateway IP address and port.

Select CascadeFlow.

For the appropriate interfaces, select *All* to export both optimized and non-optimized traffic information.

Choose *Add* to add the Cascade Express or Gateway to the collector list.

Note: 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.

Configuring Cascade Sensor-VE on RSP

Sensor-VE captures packets and sends resulting Cascade Sensor flow to the Profiler. Sensor-VE can capture packets at the same time from both the in-path interfaces and the Steelhead appliance's auxiliary interface.

Note: In RiOS v7.0.1 and later, RSP was replaced with VSP. VSP comes preinstalled in the Steelhead EX appliance. For more information about VSP, see the *Steelhead EX Management Console User's Guide*. Your existing RSP packages work on VSP.

Due to the unique architecture of the RSP virtual network interface (VNI) switching platform, Sensor-VE is in the direct flow of data moving though the Steelhead appliance. A monitor VNI is inserted directly after the LAN VNI, which allows Sensor-VE to capture all traffic from the remote site to the data center. This VNI copies all of the traffic in the Steelhead appliance data path to Sensor-VE. For more details on VNI, see the *RSP User's Guide*.

When you capture packets from the auxiliary interface, you can configure SPAN or port mirror on a switch that is in the proximity of the Steelhead appliance. This allows you to capture traffic within a remote site that does not necessarily traverse the Steelhead appliance.

Note: Sensor-VE v8.5.1 or later is required to monitor in-path and auxiliary interfaces. Earlier versions of Sensor-VE allow only in-path or auxiliary interfaces to be configured.

For details about RSP, see the *RSP User's Guide*.

Requirements for Sensor-VE on RSP

You must meet the following requirements for Sensor-VE on RSP:

- A valid RSP license on the target Steelhead appliance
- RSP installed, configured, and running
- A minimum of 17.5 GB of free RSP disk space
- A minimum of 260 MB RAM on Steelhead appliance 250 models and 512 MB RAM on all other Steelhead appliance models

Installing and Configuring Sensor-VE

From the RSP Service page, verify that the Steelhead appliance has sufficient resources. Figure 5-5 shows a valid RSP installation, as indicated by RSP Installed and RSP Enabled, and shows that there is enough disk space and RAM.

Note: The auxiliary port is disabled by default. To enable the auxiliary port for monitoring, on the Steelhead appliance, choose Configure > Networking > Base Interfaces. Select Enable Aux Interface.

Figure 5-5. RSP Service Page

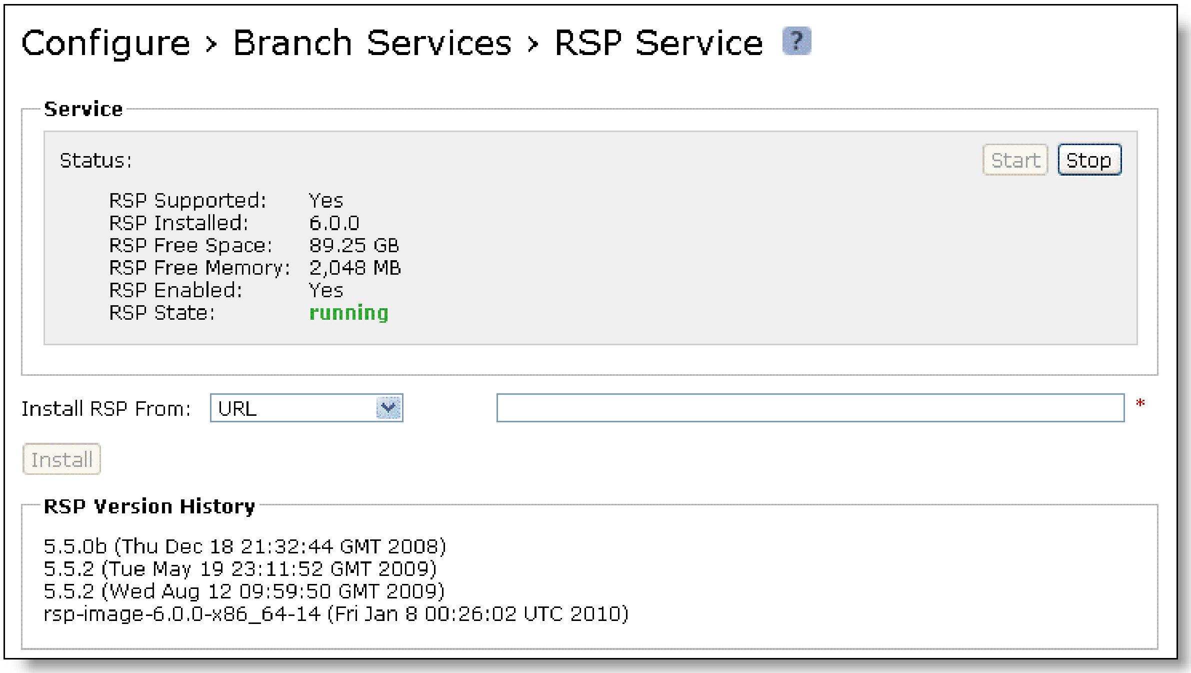

Configure > Branch Services > RSP Service

Service

Status:

RSP Supported:	Yes
RSP Installed:	6.0.0
RSP Free Space:	89.25 GB
RSP Free Memory:	2,048 MB
RSP Enabled:	Yes
RSP State:	running

Install RSP From: URL

RSP Version History

```
5.5.0b (Thu Dec 18 21:32:44 GMT 2008)
5.5.2 (Tue May 19 23:11:52 GMT 2009)
5.5.2 (Wed Aug 12 09:59:50 GMT 2009)
rsp-image-6.0.0-x86_64-14 (Fri Jan 8 00:26:02 UTC 2010)
```

To install the Sensor-VE RSP package

1. Download the latest package from https://support.riverbed.com.

2. Choose Configure > Branch Services > RSP Packages to display the RSP Packages page.

3. Select Local File.

4. Specify the path or click **Choose File** and select the Sensor-VE RSP package.

   Note: You cannot use this option to upload a package file that is larger than 2 GB. If the file is larger than 2 GB, use SCP or FTP to transfer it using the CLI. For details, see the *RiOS Services Platform User's Guide*.

5. Click **Add Package**.

   Do not go to another page until the package has finished uploading.

Figure 5-6. RSP Packages Page

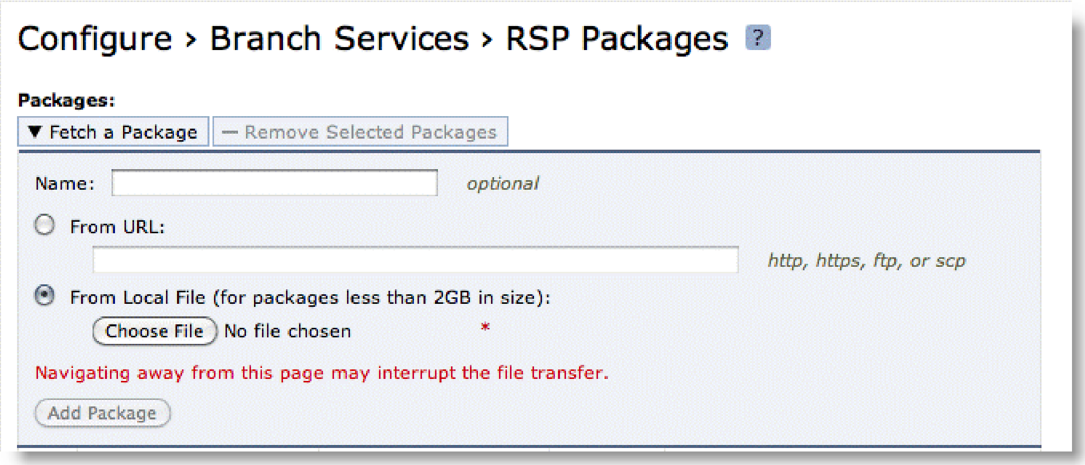

Configure › Branch Services › RSP Packages ?

Packages:

▼ Fetch a Package | — Remove Selected Packages

Name: [ ] *optional*

○ From URL:

[ ] *http, https, ftp, or scp*

● From Local File (for packages less than 2GB in size):

Choose File No file chosen *

Navigating away from this page may interrupt the file transfer.

Add Package

6. Configure the RSP slot to capture SPAN data and in-path flows by choosing Configure > Branch Services > RSP Packages.

7. Specify the correct Memory Footprint.

8. Assign Sensor-VE:aux to the auxiliary port and Sensor-VE:ctl to the primary port.

Figure 5-7. RSP Slots Page

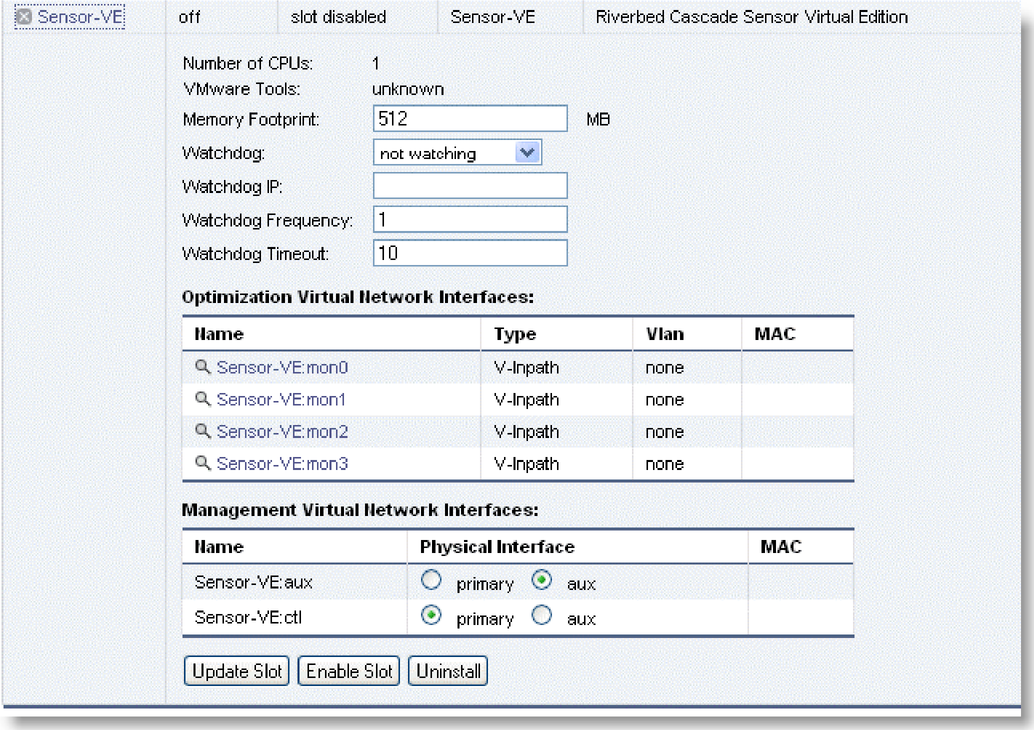

☒ Sensor-VE	off	slot disabled	Sensor-VE	Riverbed Cascade Sensor Virtual Edition

Number of CPUs:	1
VMware Tools:	unknown
Memory Footprint:	512 MB
Watchdog:	not watching
Watchdog IP:	
Watchdog Frequency:	1
Watchdog Timeout:	10

Optimization Virtual Network Interfaces:

Name	Type	Vlan	MAC
Sensor-VE:mon0	V-Inpath	none	
Sensor-VE:mon1	V-Inpath	none	
Sensor-VE:mon2	V-Inpath	none	
Sensor-VE:mon3	V-Inpath	none	

Management Virtual Network Interfaces:

Name	Physical Interface	MAC
Sensor-VE:aux	○ primary ● aux	
Sensor-VE:ctl	● primary ○ aux	

Update Slot | Enable Slot | Uninstall

9. Click **Enable Slot** to activate Sensor-VE.

10. Click **Save** to save your settings permanently.

Configuring In-Path Traffic Capture for Sensor-VE

You must configure the RSP virtual network interfaces (VNIs) for the Steelhead appliance to capture in-path flows and copy data to the Sensor-VE RSP Package. For more details on configuring VNIs, see the *RSP User's Guide*.

To add a VNI for Sensor-VE

1. On the Steelhead appliance, choose Configure > Branch Services > RSP Data Flow to display the RSP Data Flow page.

2. Select Add a VNI.

3. Select the first Sensor-VE monitor interface: Sensor-VE:mon0.

4. In the Data Flow Position drop-down list, select Start.

   This places the Sensor-VE interface closest to the LAN-side interface of the Steelhead appliance.

5. Click **Add**.

6. To edit the VNI properties, select the newly created VNI in the VNI list.

7. In the Default IP Policy drop-down list, select Copy Traffic to Slot.

8. In the Default Non-IP Policy drop-down list, select Copy Traffic to Slot.

9. Click **Apply**.

Figure 5-8. RSP Data Flow inpath0_0 Page

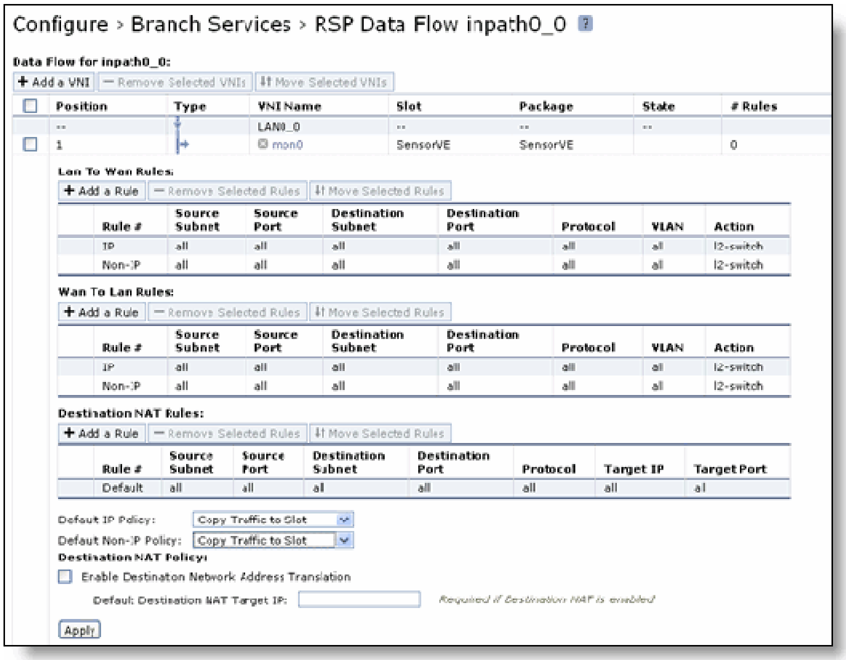

Configure > Branch Services > RSP Data Flow inpath0_0

Data Flow for inpath0_0:

□	Position	Type	VNI Name	Slot	Package	State	# Rules
	--		LAN0_0	--	--	--	
□	1		mon0	SensorVE	SensorVE		0

Lan To Wan Rules:

Rule #	Source Subnet	Source Port	Destination Subnet	Destination Port	Protocol	VLAN	Action
IP	all	all	all	all	all	all	l2-switch
Non-IP	all	all	all	all	all	all	l2-switch

Wan To Lan Rules:

Rule #	Source Subnet	Source Port	Destination Subnet	Destination Port	Protocol	VLAN	Action
IP	all	all	all	all	all	all	l2-switch
Non-IP	all	all	all	all	all	all	l2-switch

Destination NAT Rules:

Rule #	Source Subnet	Source Port	Destination Subnet	Destination Port	Protocol	Target IP	Target Port
Default	all	all	al	all	all	all	al

Defaut IP Policy: Copy Traffic to Slot

Defaut Non-IP Policy: Copy Traffic to Slot

Destination NAT Policy:

□ Erable Destinaton Network Address Translation

Defaul: Destination NAT Target IP: Required if Destination NAT is enabled

Sensor-VE Configuration

Sensor-VE requires very little configuration. At the initial startup, on the CLI, you are prompted to specify:

- the IP address.
- the subnet mask.
- the default gateway.

You can use your Web browser to go to the Sensor-VE IP address for additional configuration. On your initial visit to the GUI, you see a setup screen. If you are using the Steelhead appliance auxiliary port for monitoring, select the check box to monitor and collect traffic as shown in Figure 5-9.

For details on additional configuration, see the *Cascade Sensor and Cascade Gateway User's Guide*.

Figure 5-9. Monitor Port Configuration Page

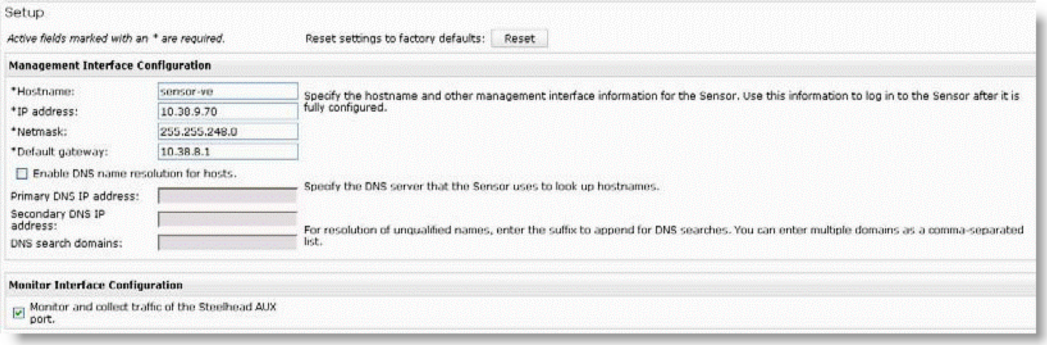

Setup

*Active fields marked with an * are required.*

Reset settings to factory defaults: Reset

Management Interface Configuration

*Hostname: sensor-ve

*IP address: 10.30.9.70

*Netmask: 255.255.248.0

*Default gateway: 10.38.8.1

Specify the hostname and other management interface information for the Sensor. Use this information to log in to the Sensor after it is fully configured.

□ Enable DNS name resolution for hosts.

Primary DNS IP address:

Secondary DNS IP address:

DNS search domains:

Specify the DNS server that the Sensor uses to look up hostnames.

For resolution of unqualified names, enter the suffix to append for DNS searches. You can enter multiple domains as a comma-separated list.

Monitor Interface Configuration

☑ Monitor and collect traffic of the Steelhead AUX port.

CHAPTER 6 Additional Cascade Integration

Cascade contains 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 Cascade integrations:

For additional assistance, contact Riverbed Professional Services.

SNMP Integration

This section describes the following SNMP integrations.

SNMP Integration for Flow Sources

When devices send flow, standard SNMP interface identifiers (ifindex values) indicate which interfaces the flow traverses. You must map the interface identifiers to names and descriptions to identify them on the Profiler. You must also obtain the link speed information so you can convert raw bandwidth numbers to link utilization percentages. Cascade gathers the following information using standard SNMP from all devices sending standard flow, or Steelhead CascadeFlow, to the Gateway or the Express:

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

Ensure that you configure firewalls between the Express (or the Gateway) and flow-reporting devices to allow SNMP access between the Cascade device 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 Express or the Gateway.

For more information about configuring of Cascade for SNMP collection of these items, see the *Cascade Profiler and Cascade Express User's Guide*.

SNMP Integration for Device Management of Cascade Components

You can monitor the status of Cascade components through SNMP from an external SNMP manager. Currently, the Profiler, the Gateway, and the Sensor generally support the industry-standard UCD-SNMP-MIB. For more information, see
http://www.oidview.com/mibs/2021/UCD-SNMP-MIB.html.

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

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

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

SNMP Integration for Sending Traps

The Profiler 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 Profiler UI. Some of the use cases for sending SNMP traps are as follows:

- Sending Cascade 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 Profiler MIB (labeled Mazu-MIB). This MIB is available from either the Riverbed Support site or Profiler help downloads page.

You can route the appropriate events to the appropriate devices by first configuring recipient groups within the Profiler and then configuring 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 *Cascade Profiler and Cascade Express User's Guide*.

SNMP for Switch Port Discovery

Standard flow records identify hosts (servers, desktops, printers, and so on) by their IP addresses. The Profiler supports discovery of MAC addresses and switch port locations of individual hosts based upon their IP addresses. This enables the Profiler to display complete information, as shown in Figure 6-1. This information appears in multiple places throughout the Profiler UI.

Figure 6-1. MAC and Host Switch Information Displayed as a Result of Switch Port Discovery

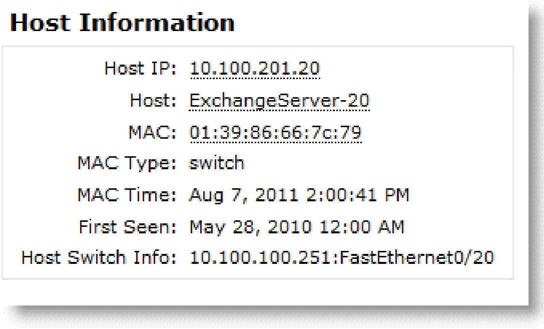

Host Information

Host IP: 10.100.201.20
Host: ExchangeServer-20
MAC: 01:39:86:66:7c:79
MAC Type: switch
MAC Time: Aug 7, 2011 2:00:41 PM
First Seen: May 28, 2010 12:00 AM
Host Switch Info: 10.100.100.251:FastEthernet0/20

Supported Routers and Switches for Switch Port Discovery

The following table shows a partial list of supported switches and routers for switch port discovery in Cascade Profiler v9.0 and later.

Vendor	Model	Lookup Router	Switch	Comments
Airespace wireless controllers	3500, 4101, 4102	No	Yes	APs appear as switch ports
Allied Telesyn	AT-8000	No	Yes	
Aruba wireless controllers	5000	No	Yes	APs appear as switch ports
Cisco 2500 series	2501, 2503, 2511, 2514, AS2509RJ, AS2511RJ	Yes	Yes	
Cisco 2600 series	2610, 2610XM, 2611, 2620, 2620XM, 2621, 2621XM, 2651XM, 2691	Yes	Yes	
Cisco 2800 series	2811, 2821, 2851	Yes	Yes	
Cisco Catalyst 2900 series	2908xl, 2912MfXL, 2924CXL, 2924CXLv, 2924 MXL	No	Yes	
Cisco 2940 and 2950 series	2940-8TT, 29500t24	No	Yes	
Cisco 2970 series	2960, 2970G-24T-E	No	Yes	
Cisco Catalyst 3500 series	3508GXL, 3524XL, 3548SL	No	Yes	Layer-2 devices—only when you run IOS
Cisco 3550 (partial)	3400 with MetroBase, 3550-12T	No	Yes	Running IOS

Vendor	Model	Lookup Router	Switch	Comments
Cisco Catalyst 3550 (partial)	3550, 3560, 3550-24, 3550-48	Yes	Yes	Running IOS
Cisco Catalyst 4000 series	4006, 4503, 4506, 4507, 4510, wsc4003, wsc4006, wsc4503, wsc4506, wsc4912g	No	Yes	
Cisco Catalyst 5000 series	Wsx5302	Yes	No	Most models not supported
Cisco Catalyst 6500 series	6503, 6509, sp72033, s3223, s32p3, s222, 6kMsfc, 6kMsfc2, wsc6509	Yes	Yes	
Cisco wireless controllers	2006, 4112, 4124, 4136, 4402, 4404	No	Yes	APs appear as switch ports
Dell PowerConnect 3000 and 5000 series	3348, 3448P, 3424, 3424P, 5324	No	Yes	
Dell PowerConnect 6000 series	6024F, 6224, 6248	Yes	Yes	
IBM BladeCenter Ethernet switch family	All	No	Yes	
Linksys 2048 family	All	No	Yes	
Enterasys Networks Matrix series	Matrix N-series DFE	Yes	Yes	
Enterasys Networks SuperStack C-series	C3G124-24, C3G124-48, C2G124-24, C2G124-48	Yes	Yes	
Extreme Network Alpine and Summit	Alpine 3808, Summit 7i, 48si	Yes	Yes	
Foundry EdgeIron series	EdgeIron 24G	No	Yes	
Foundry IronWare family	2512, 2524, 2626, 2650, 4000, 4104GL, 4108GL, 8000	No	Yes	ProCurve devices are widely supported (newer devices not in list likely supported)
Juniper M-series Router series	All	Yes	No	
Juniper Netscreen series	All	Yes	No	
Nortel Alteon AD series	180, 183, 184, AD2, AD3, AD4	Yes	Yes	
Nortel AP222x series	AP-2220, AP-2221	No	Yes	
Nortel BayStack Hub series	102, System 5000	No	Yes	Requires Advanced NMM
Nortel Business Switch series	-50, 110, 120, 210, 220, 1010, 1020	Yes	Yes	
Nortel Centillion series	5000BH, 5005BH, C50, C100	No	Yes	v4.x/5.x or later

Vendor	Model	Lookup Router	Switch	Comments
Nortel Ethernet Routing/ BaystackYes-Yes-	2526, 2550, 3510, 4524, 4526, 4548, 4550, 5510, 5520, 5530	Yes	Yes	
Nortel Passport 1600 series	1612, 1624, 1648	Yes	Yes	
Nortel Routing, Accelar Family	1050, 1100, 1150, 1200, 8106, 8110, 8603, 8606, 8610, 8610co	Yes	Yes	For 8600, code for switch support must be v3.2 and later
Nortel Baystack Switch series	303, 304, 350, 380, 410, 420, 425, 450, 460, 470, BPS	No	Yes	
Nortel Multiprotocol Router/BayRS	2430, 5430, AN, ARN, ASN, BLN, BCN	Yes	Yes	
Nortel Synoptics	281X, System3000	No	Yes	
Nortel VPN Router/ Contivity	100, 400, 600, 1000, 1010, 1050, 1500, 1600, 1700, 1740, 1750, 2500, 2600, 2700, 4500, 4600, 5000	Yes	No	
Nortel Wireless 2270	2270	No	Yes	APs appear as switch ports

Active Directory

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

- the users associated with an end station on the network.
- all the end stations that a user has logged into.

These two capabilities are shown in Figure 6-2 and Figure 6-3.

Figure 6-2. Users Logged into a Given Host for a Selected Time Period

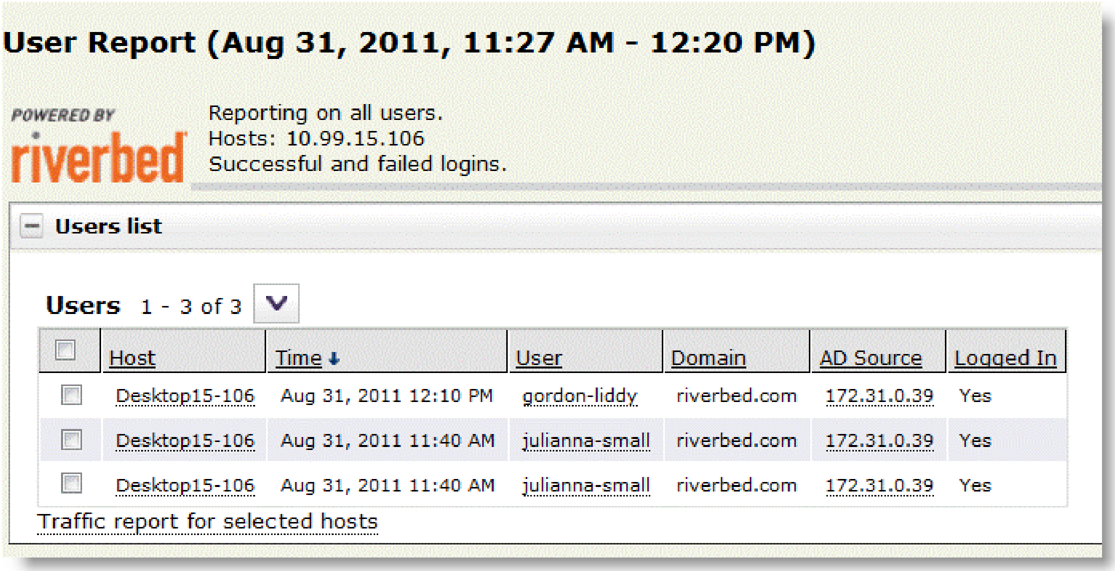

User Report (Aug 31, 2011, 11:27 AM - 12:20 PM)

Reporting on all users.
Hosts: 10.99.15.106
Successful and failed logins.

Users list

Users 1 - 3 of 3

□	Host	Time ↓	User	Domain	AD Source	Logged In
□	Desktop15-106	Aug 31, 2011 12:10 PM	gordon-liddy	riverbed.com	172.31.0.39	Yes
□	Desktop15-106	Aug 31, 2011 11:40 AM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop15-106	Aug 31, 2011 11:40 AM	julianna-small	riverbed.com	172.31.0.39	Yes

Traffic report for selected hosts

Figure 6-3. Hosts That a Given User Logged into for a Selected Time Period

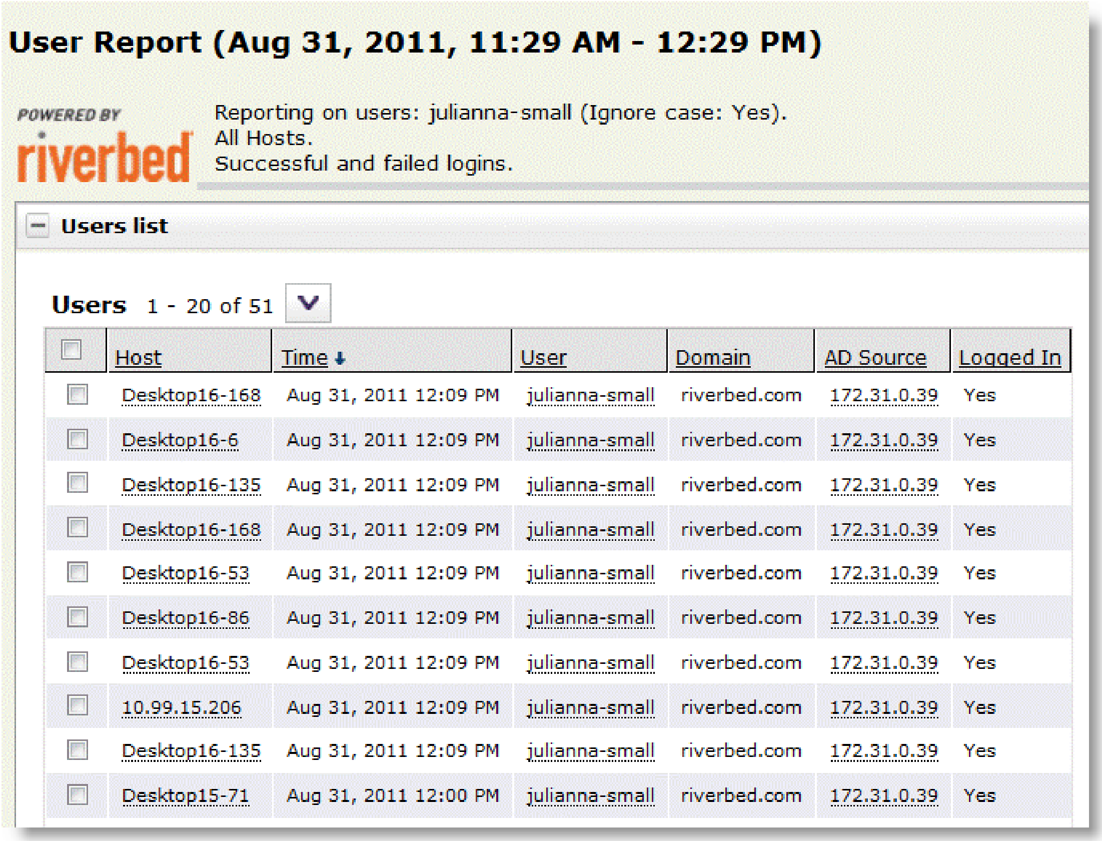

User Report (Aug 31, 2011, 11:29 AM - 12:29 PM)

Reporting on users: julianna-small (Ignore case: Yes).
All Hosts.
Successful and failed logins.

Users list

Users 1 - 20 of 51

□	Host	Time ↓	User	Domain	AD Source	Logged In
□	Desktop16-168	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-6	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-135	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-168	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-53	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-86	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-53	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	10.99.15.206	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop16-135	Aug 31, 2011 12:09 PM	julianna-small	riverbed.com	172.31.0.39	Yes
□	Desktop15-71	Aug 31, 2011 12:00 PM	julianna-small	riverbed.com	172.31.0.39	Yes

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 Profiler from a domain controller, or for AD-2008, an event collector host. Riverbed provides a service application named Cascade AD Connector that forwards the appropriate events from a domain controller, or event collector, to the Profiler.

Integration for Active Directory 2008

For AD-2008, you must use the Cascade AD Connector v2.0. 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 Cascade *AD Connector 2008 for Windows Server 2008:* Riverbed Cascade *AD 2.0 Release Notes*.

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

Integration for Active Directory 2003

For AD-2003, you must use the Cascade AD Connector v1.5. 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 Profiler help downloads page.

CHAPTER 7

Cascade Profiler Analytics and Service Monitoring

Cascade 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. It also provides best practices for how to stay within these limits and which specific metrics to use.

This chapter contains the following sections:

Before reading this chapter, you must also be familiar with the *Cascade Profiler and Cascade Express User's Guide*.

Analytic License Limits per Profiler Platform

If you are running the Cascade Profiler v9.0.5 or later with a deployed analytics license, the following table shows the number of analytics policy metrics allowed system wide.

Profiler Model	Available Analytic Policy Metrics
Express	Up to 5000
Standard	Up to 7500
Enterprise	Up to 10,000

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 Profiler 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 7-1. Simple SharePoint Service Configuration

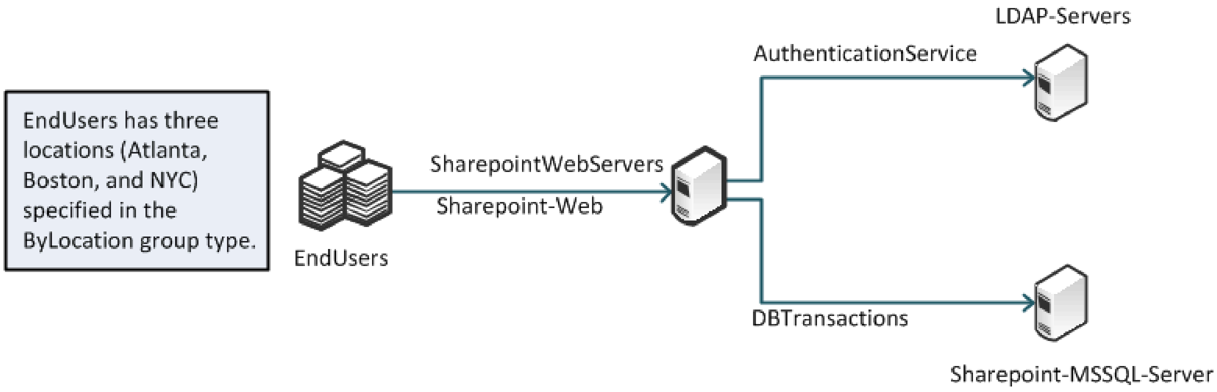

Figure 7-1 shows a simple SharePoint service configuration. There are three end-user locations (Atlanta, Boston, and New York) connected to the SharePoint Web servers. The SharePoint Web servers are supported by LDAP and MS-SQL on the back end.

Using the example shown in Figure 7-1, five metrics are selected (Figure 7-2) for each of the segments during service discovery for the SharePoint service.

Figure 7-2. Five Metrics for SharePoint Service

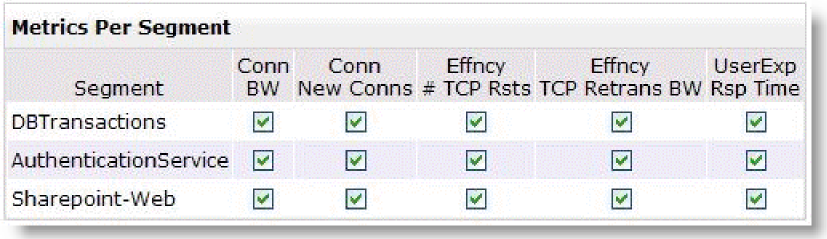

Metrics Per Segment

Segment	Conn BW	Conn New Conns	Effncy # TCP Rsts	Effncy TCP Retrans BW	UserExp Rsp Time
DBTransactions	☑	☑	☑	☑	☑
AuthenticationService	☑	☑	☑	☑	☑
Sharepoint-Web	☑	☑	☑	☑	☑

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 fifteen metrics (five metrics time three locations). The total number of metrics used to monitor the SharePoint service is twenty-five.

A location with a 100 end users is more representative of a larger enterprise network. If there are 100 end-user locations, the SharePoint service requires over 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 7-3. Metrics Shown on the Commit Changes to Service Page

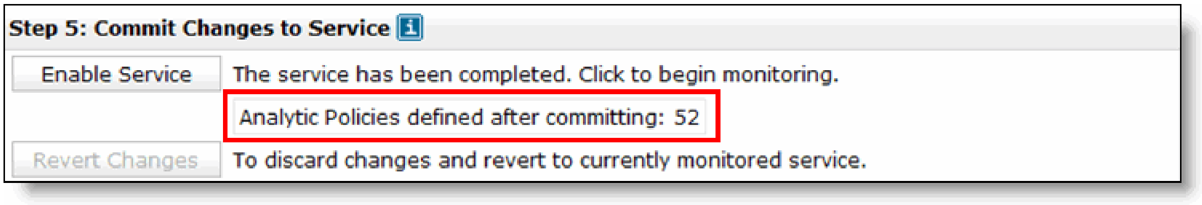

Step 5: Commit Changes to Service

The service has been completed. Click to begin monitoring.

Analytic Policies defined after committing: 52

To discard changes and revert to currently monitored service.

> **Note:** 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 7-4. Metrics Shown on the Information Page

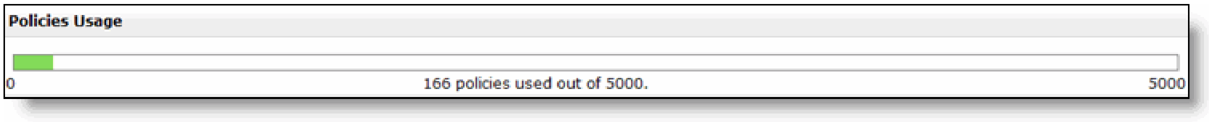

Policies Usage

166 policies used out of 5000.

Reducing the Number of Metrics

You can reduce the number of metrics used with services in the Profiler 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 7-5 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 7-5. Reducing the Number of Locations on the Customize Monitoring on the Locations Page

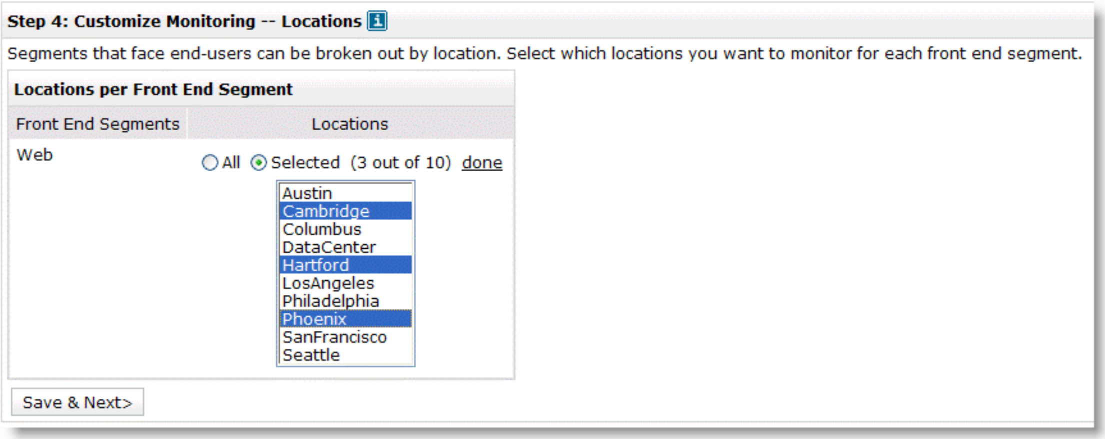

Step 4: Customize Monitoring -- Locations

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

Locations per Front End Segment

Front End Segments	Locations
Web	○ All ● Selected (3 out of 10) done Austin **Cambridge** Columbus DataCenter **Hartford** LosAngeles Philadelphia **Phoenix** SanFrancisco Seattle

---

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

- **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 deselect the Track By End-User checkbox on the front-end end-user component.

Note: You still 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 7-6 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 7-6. Default and Configurable Metrics on the Customize Monitoring on the Metrics Page

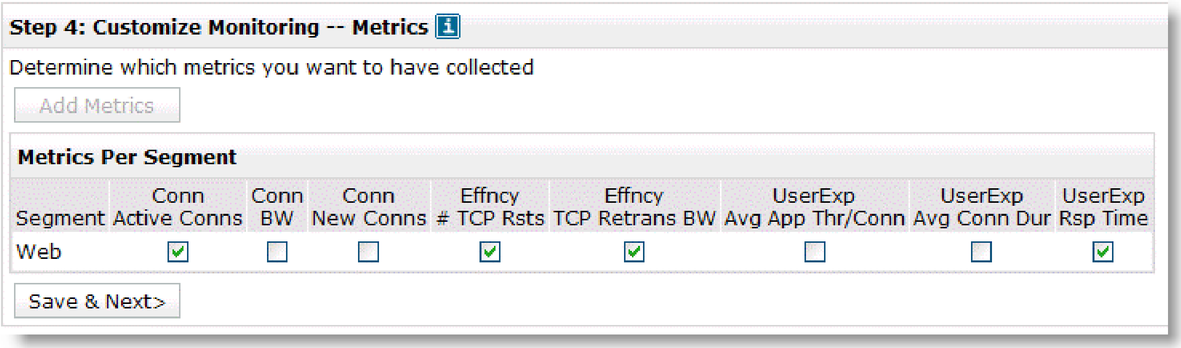

Step 4: Customize Monitoring -- Metrics

Determine which metrics you want to have collected

Metrics Per Segment

Segment	Conn Active Conns	Conn BW	Conn New Conns	Effncy # TCP Rsts	Effncy TCP Retrans BW	UserExp Avg App Thr/Conn	UserExp Avg Conn Dur	UserExp Rsp Time
Web	☑	□	□	☑	☑	□	□	☑

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% 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 Profiler, Step 4 of the configuration wizard allows 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 on characteristics to consider per metric, see "Best Practices for Choosing Metrics" on page 86.

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.

Best Practices for Choosing Metrics

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

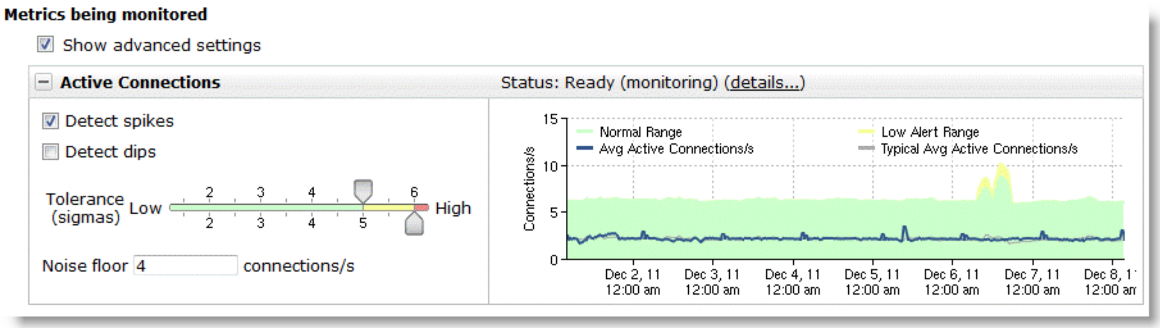

Metrics being monitored

☑ Show advanced settings

− Active Connections

Status: Ready (monitoring) (details...)

☑ Detect spikes

□ Detect dips

Tolerance (sigmas) Low 5 High 6

Noise floor 4 connections/s

- **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 (after 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 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 8 Troubleshooting

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

In its simplest form, the Profiler 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

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 Profiler calculates RTT information only for TCP-based flows. Any flow that is not TCP-based (ICMP, UDP, and so on) does display RTT information.
- **Flow not seen by the Shark, Sensor, or Sensor-VE** - The flow must be detected by a Cascade packet analysis device 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 Shark, Sensor, or Sensor-VE, the RTT information is not valid and is discarded.

- **Retransmits during the initial flow setup** - The Profiler does not show RTT information if there is retransmitted information during the initial flow setup. If the Profiler 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 Profiler 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 Cascade 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 allow connections from the Profiler and that it is able to relay messages to the appropriate destinations.

Note: The Profiler does not currently support encrypted SMTP.

DNS Names Not Being Resolved in Reports

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

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

The following are options for DNS resolution in UI Preferences:

- **Resolve host names using DNS** - Turns on the use of DNS to resolve an IP address into a host name.
- **Resolve host names for hosts managed by DHCP** - Uses the optional Cascade DHCP integration to use the information provided during DHCP imports for name resolution.
- **Suppress DHCP/DNS search domains from resolved host names** - 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 Profiler from overwhelming DNS servers, you can place the following limits on resolving hosts. You can control the:

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

Riverbed recommends that you set these values to 500 each to prevent overwhelming the DNS server. This mean that the Profiler does not attempt to resolve more then 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 Profiler waits only one second for responses from DNS servers. This is 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 Profiler 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 Profiler (for example, a branch office might not have a device reporting traffic internal to the branch). It is impossible to report on traffic for which there is no coverage.

  – **Too many devices reporting data** - The Profiler 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 statics are counted on an ingress interface. When Profiler 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, Profiler 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 Profiler has no way of knowing how much data actually went out the other interface, the numbers can be incorrect. If the 10Mbps-interface is receiving 5 Mbps, the Profiler reports 5 Mbps leaving the 1.5Mbp-interface because this is the best data Profiler received.

– **Routing issues** - The Profiler 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 Profiler can miss one side of the conversation. This results in inaccurate information from reports.

- **Different data resolution** - While the Profiler provides several different data resolutions (one minute, 15 minute, one hour, six hours, an so on) 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 Profiler and five minutes on a router) is very likely to result in differences in reported values.

Sensor Protocol Violations

The Sensor protocol violation (SPV) error is one of the more common errors that you can receive on the Profiler. The error appears as a system event indicating that the violation occurred between the Profiler and one or more reporting devices (Gateway, Shark, Sensor, or Sensor-VE). There are two potential causes of SPV - slow transfers or lack of time synchronization:

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

  If you detect SPV errors, send a file transfer between the Profiler 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.

  Note: 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 Shark, Sensor, or Gateway is unable to NTP synchronize with the Profiler (for example, due to a firewall or ACL), the time on that device can drift sufficiently from the Profiler's time. The Profiler then has difficulty ensuring that the data being sent is for the same time slice the Profiler is currently processing. You must ensure that no firewall or ACL are blocking NTP.

Communication Issues

The various Cascade devices communicate with each other on select ports. These ports must be open between the Profiler 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 Profiler and remote devices such as the Shark and Gateway. You must open this port bidirectionally between devices.

- **TCP/8443** - Used to facilitate the exchange of SSL keys between the Profiler and remote devices such as the Gateway and Sensor. You must open this port bidirectionally between devices.
- **UDP/123** - Used for NTP synchronization between the Profiler and the remote devices such as the Shark and Gateway. The Profiler acts as the NTP server for the remote devices.

If anything is blocking communications on the specified ports, that portion of the Profiler system does not work correctly. For example, if UDP/123 (NTP) is blocked between the Profiler and the Shark, the time on the Shark 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 Profiler or Express. Also check the device is in the list of supported devices in “Supported Routers and Switches for Switch Port Discovery” on page 75. If polling still fails, you can enter the following CLI commands to verify SNMP communication and support.

For a lookup router, enter the command:

```
Profiler# snmpinfo <RouterIP>  --dev   --fw  --version <1 or 2>  --comm <READonlyCommunityString>
```

For a switch, enter the command:

```
Profiler# snmpinfo <RouterIP>  --dev   --fw  --version <1 or 2>  --comm <READonlyCommunityString>
```

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 Cascade Integration” on page 73.

APPENDIX A Licensing

This appendix explains various aspects of licensing for your Cascade products. It contains the following sections:

Overview

The Cascade Profiler, Gateway, and Sensor systems have the following types of licenses:

- "Runtime Licenses" on page 95
- "Capacity Licenses" on page 95
- "Option Licenses" on page 96

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 allows 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 Gateway can accept no more than 100,000 flows per minute).

Different types of Cascade devices have different flow capacities:

- **The Gateway** - 100,000, 200,000, 500,000, 800,000, and 1,400,000 flows per minute
- **The Express** - 25,000, 50,000, and 75,000 flows per minute
- **The Profiler** - 100,000, 200,000, and 500,000 flows per minute
- **The Enterprise cluster** - 800,000 or more, depending on the number of analyzer or expansion modules you have installed

Some Cascade devices do not require a capacity license. The Sensor and the Shark allow data to be processed up to the capacity of the system without requiring additional capacity licenses.

Option Licenses

Option licenses enable you to add additional functionality to the Express, the Profiler, or an Enterprise cluster. Examples of functionality that requires an option license include:

- Cascade analytics
- Cascade security module
- SAN storage support

Only the security module license does not require you to purchase an additional part number. The Express, the Profiler, and the Enterprise cluster systems support the security module through the runtime license.

License Installation

License installation varies depending on the type of hardware the license is installed on. Cascade systems based on current-model hardware use an automatic license server that provides license keys directly to the Cascade device or through a Web portal that allows 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 Cascade system for xx60 hardware, a license key is generated and automatically assigned to the system in the Riverbed licensing portal.

If the Cascade system has Internet access, you can automatically download the new license to the device. If the Cascade system does not have Internet access, you can manually add the new license to the system by entering the provided key into the Cascade UI.

Previous-Generation Hardware License Installation

When you purchase a license for a Cascade system for earlier hardware (typically hardware purchased prior to February 2012), an email is sent from Riverbed that includes the command you run to enable the purchased license. You must run this command exactly as specified, and only on the system the license is intended for.

Other Device License Installation

The Pilot and the Shark have different license installations.

You purchase Pilot licenses one complete Pilot software installation per workstation. You do not need to purchase any options for the Pilot beyond the initial license. To have a fully functional software licence without expiration, you must activate the license with Riverbed through the Internet (if Internet access is not available the license can be activated by phone, fax, or mail). You can use Pilot software for 30 days without activating the license.

There is no software-based upgrade or license required for the Shark. The Shark does not support an upgradeable software license model (similar to the Sensor). Instead, the Shark is sold with a base chassis (1U, 2U, or 3U) that supports as many packets per second as the NIC cards that the Shark can support.

Assigning Licenses

When you purchase Cascade hardware, all appropriate licenses are assigned to each unit. The build-time license (indicating what type of system the unit is: for example, a Gateway) is installed during the manufacturing process and you can not 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 CAG-2260 with an F1 capacity license, a license is generated on the Gateway that enables it to support 100,000 flows per minute. This license is not activated, but it is available to you in the Riverbed licensing portal. The license automatically downloads to the Gateway when the 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, imagine if you are a customer that wants to deploy 10 Gateways at sites around the world. Each site requires a different capacity level. You purchase four F4 Gateways, two F3 Gateways, and four F1 Gateways. If each 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 Gateway to each location. When the Gateway is at the location and installed, you can use the licensing portal to assign capacity licenses to the 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 Cascade device 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 Cascade devices.

Automatic License Upgrades

When you purchase an upgrade for a specific device (for example, upgrading an Express from a 2260L to a 2260M), the licensing process is automatic. After the purchase is approved, the licensing portal receives a new, upgraded license. The next time that device 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 Cascade supports. 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 Cascade appliances.

Licensing the Express

The Express (CAX-2260) has three capacity licenses (L/M/H) and one optional license. The capacity licenses provide:

- 25,000 (L) flows per minute.
- 50,000 (M) flows per minute.
- 75,000 (H) flows per minute.

The optional license provides 5,000 Cascade analytics.

Licensing the Profiler

The Profiler (CAP-2260) has three capacity licenses (L/M/H) and two optional licenses. The capacity licenses provide:

- 100,000 (L) flows per minute.
- 200,000 (M) flows per minute.
- 500,000 (H) flows per minute.

The optional licenses provide:

- SAN support.
- 7,500 Cascade analytics.

Licensing the Enterprise Cluster

The Enterprise cluster (CAP-4260-UI, CAP-4260-DB, CAP-4260-AN) provides a base capacity license of 800,000 flows per minute. You can expand up to seven additional CAP-4260-EX modules, each providing 400,000 flows per minute of capacity. No additional capacity licenses are required for the Enterprise cluster.

The optional licenses available provide:

- SAN support
- 10,000 Cascade analytics

Licensing the Gateway

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

- 100,000 (F1) flows per minute.
- 200,000 (F2) flows per minute.
- 500,000 (F3) flows per minute.
- 800,000 (F4) flows per minute.
- 1,400,000 (F5) flows per minute.

Licensing the Sensor, Sensor-VE, and Shark

The Sensor, Sensor-VE, and Shark does not have any capacity or optional licenses available. All Sensors ship with the ability to forward their maximum number of flows per minute available.

Index