Virtual Steelhead® Appliance Installation Guide

RiOS Version 8.6

April 2014

riverbed®

Riverbed Technology
199 Fremont Street
San Francisco, CA 94105

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

Part Number
712-00069-06

Contents

Preface **1**

About This Guide 1
Audience 2
Document Conventions 2
Product Dependencies and Compatibility 2
Third-Party Software Dependencies 3
SNMP-Based Management Compatibility 3
Antivirus Compatibility 3
Additional Resources 4
Release Notes 4
Riverbed Documentation and the Support Knowledge Base 4
Contacting Riverbed 4
Internet 4
Technical Support 4
Professional Services 5
Documentation 5

Chapter 1 - Overview of Virtual Steelhead **7**

Understanding Virtual Steelhead 7
Virtual Steelhead Optimization 8
Configuring Optimization 10
New Features in Version 8.6 11
VMware ESX and ESXi 11
Microsoft Hyper-V 12
Virtual Steelhead Deployment Guidelines 12
Network Configuration 12
Network Performance 13
Deployment Options 13
In-Path Deployment 14
Virtual In-Path Deployment 14
Out-of-Path Deployment 14

Virtual Steelhead Platform Models 15
Flexible RiOS Data Store 16
Multiple RiOS Data Stores 17
In-Path Pairing for NIC Interfaces 17
NICs for Virtual Steelhead 17
Requirements for Virtual Steelhead Deployment with a NIC 18
Virtual Steelhead on the Cisco SRE 19

Chapter 2 - Setting Up Virtual Steelhead on ESX and ESXi 21
Basic Steps for Installing Virtual Steelhead on ESX/ESXi 21
Obtaining the Virtual Steelhead Software Package 22
Installing Virtual Steelhead with a Riverbed NIC 22
Completing the Preconfiguration Checklist 23
Installing Virtual Steelhead 23

Chapter 3 - Setting Up Virtual Steelhead on Hyper-V 35
Basic Steps for Installing and Configuring Virtual Steelhead 35
Obtaining the Virtual Steelhead Software Package 35
Completing the Preconfiguration Checklist 36
Installing Virtual Steelhead 36
Manual Installation on Hyper-V 39

Chapter 4 - Configuring Virtual Steelhead 41
Basic Steps for Configuring Virtual Steelhead 41
Completing the Initial Configuration 42
Logging In to the Virtual Steelhead Management Console 45
Purchasing the Token and Receiving the Licenses 46
Managing Licenses and Model Upgrades 47
Flexible Licensing Overview 47
Activating the Token and Installing the Licenses 48
Model Upgrade Overview 49
Upgrading RiOS to Version 8.6 54
Upgrading RiOS Software 54
Downgrading the Software 55
Rebooting and Shutting Down Virtual Steelhead 55
Verifying Your Connections 56
Verifying Your Configuration 56

Chapter 5 - Using Discovery Agent 57
Overview of the Discovery Agent 57
Discovery Agent Requirements 58
Installing the Discovery Agent on a Windows Server 59
Installing the Discovery Agent on a Linux Server 60
Configuring the Discovery Agent 60
Configuring the Discovery Agent on a Linux Server 60
Configuring the Discovery Agent on Windows 60
Configuring Transparency Modes 63
Enabling Optimization Using the Discovery Agent 63

Appendix A - Configuring a Riverbed NIC in ESX 4.1 65
Configuring the bpvm0 Interface (ESX/ESXi 4.1) 65
Configuring Riverbed NIC Interfaces (ESX/ESXi 4.1) 66

Appendix B - Troubleshooting 69
Duplex Mismatch 69
Possible Cause 70
Oplock Issues 70
Possible Causes 71
CIFS Overlapping Open Optimization Denies Multi-User Access 71
Solution 72
IP Address Configuration 73
Solutions 73
Asymmetric Routing 74
Possible Cause 74
Packet Ricochet 74
Possible Cause 75
Packet Ricochet—Internet Control Messaging Protocol (ICMP) Redirects 75
Possible Causes 75
Simplified Routing 76
Auto-Discovery Failure 77
Possible Causes 77
Protocol Optimization Errors 77
Solutions 77
Resetting a Lost Password 78
Bypass NIC Log Messages 79

Index 81

Preface

Welcome to the *Virtual Steelhead Appliance Installation Guide.* The Virtual Steelhead appliance is a software version of the Steelhead appliance that runs on the VMware ESX/ESXi and Microsoft Hyper-V hypervisors. Read this preface for an overview of the information provided in this guide and the documentation conventions used throughout, software dependencies, additional reading, and contact information. This preface includes the following sections:

- "About This Guide" on page 1
- "Product Dependencies and Compatibility" on page 2
- "SNMP-Based Management Compatibility" on page 3
- "Antivirus Compatibility" on page 3
- "Additional Resources" on page 4
- "Contacting Riverbed" on page 4

About This Guide

The *Virtual Steelhead Appliance Installation Guide* describes how to install and configure the Virtual Steelhead appliance. This guide includes information relevant to the following products:

- Riverbed Virtual Steelhead Appliance (Virtual Steelhead)
- Virtual Steelhead Management Console (Management Console)
- Riverbed Optimization System (RiOS)
- Riverbed Steelhead appliance (Steelhead appliance)
- Virtual Services Platform (VSP)
- Riverbed Services Platform (RSP)
- Central Management Console (CMC)
- Steelhead Mobile Controller (Mobile Controller)

Audience

This guide is written for administrators familiar with managing virtual environments, LANS, and WANs using common network protocols. You should also be familiar with using the Riverbed Command-Line Interface (CLI) as described in the *Riverbed Command-Line Interface Reference Manual*, and with Microsoft Hyper-V or VMware ESX/ESXi hypervisors.

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>}**

Product Dependencies and Compatibility

This section provides information about product dependencies and compatibility. It includes the following sections:

- "Third-Party Software Dependencies" on page 3
- "SNMP-Based Management Compatibility" on page 3
- "Antivirus Compatibility" on page 3

Third-Party Software Dependencies

The following table summarizes the software requirements for Virtual Steelhead.

Component	Software Requirements
Microsoft Hyper-V Hypervisor	Virtual Steelhead models VCX255 through VCX1555 support Hyper-V, available on Windows Server 2012 and Windows Hyper-V Server.
VMware ESX/ESXi Hypervisor	Virtual Steelhead supports ESX/ESXi 4.0, 4.1, 5.0 and 5.1. If you use the Riverbed network interface card (NIC), you must use ESXi 4.1 or later. For ESXi 5.0 and later, the method for supporting the card differs from the 4.1 method. For information, see the section "Completing the Preconfiguration Checklist" on page 23.
Virtual Steelhead Management Console	Any computer that supports a Web browser with a color image display. The Management Console has been tested with Mozilla Firefox v10.0, Mozilla Firefox Extended Support Release version 10.0, and Microsoft Internet Explorer v7.0 and v8.0. **Note:** JavaScript and cookies must be enabled in your Web browser.

SNMP-Based Management Compatibility

The Steelhead appliance supports a proprietary Riverbed MIB accessible through SNMP. SNMPv1 (RFCs 1155, 1157, 1212, and 1215), SNMPv2c (RFCs 1901, 2578, 2579, 2580, 3416, 3417, and 3418), and SNMPv3 are supported, although some MIB items might only be accessible through SNMPv2 and SNMPv3.

SNMP support enables the Steelhead appliance to be integrated into network management systems such as Hewlett-Packard OpenView Network Node Manager, BMC Patrol, and other SNMP-based network management tools.

Antivirus Compatibility

The Steelhead appliance has been tested on clients and file servers generating CIFS traffic with the following antivirus software with no impact on performance:

- Network Associates (McAfee) VirusScan v7.0.0 Enterprise on the server
- Network Associates (McAfee) VirusScan v7.1.0 Enterprise on the server
- Network Associates (McAfee) VirusScan v7.1.0 Enterprise on the client
- Symantec (Norton) AntiVirus Corporate Edition v8.1 on the server

The Steelhead appliance has been tested on clients and file servers generating CIFS traffic with moderate impact on performance:

- F-Secure Anti-Virus v5.43 on the client
- F-Secure Anti-Virus v5.5 on the server
- Network Associates (McAfee) NetShield v4.5 on the server
- Network Associates VirusScan v4.5 on the client
- Symantec (Norton) AntiVirus Corporate Edition v8.1 on the client

Additional Resources

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

- "Release Notes" on page 4
- "Riverbed Documentation and the Support Knowledge Base" on page 4

Release Notes

The online software release notes supplement the information in this manual. The release notes are available in the Software section of the Riverbed Support site at https://support.riverbed.com. The following table describes the release notes.

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.

Read this document before you begin the installation and configuration process. It contains important information about this release of Virtual Steelhead.

Riverbed Documentation and the Support Knowledge Base

For a complete list and the most current version of Riverbed documentation, go 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.

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/us/products/professional_services/.

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.

CHAPTER 1 Overview of Virtual Steelhead

This chapter provides an overview of Virtual Steelhead. It includes the following sections:

- "Understanding Virtual Steelhead" on page 7
- "Virtual Steelhead Optimization" on page 8
- "New Features in Version 8.6" on page 11
- "VMware ESX and ESXi" on page 11
- "Microsoft Hyper-V" on page 12
- "Virtual Steelhead Deployment Guidelines" on page 12
- "Deployment Options" on page 13
- "Virtual Steelhead Platform Models" on page 15
- "NICs for Virtual Steelhead" on page 17
- "Virtual Steelhead on the Cisco SRE" on page 19

Understanding Virtual Steelhead

Virtual Steelhead is software that delivers the benefits of WAN optimization, similar to those offered by the Steelhead appliance hardware, while also providing the flexibility of virtualization.

Built on the same RiOS technology as the Steelhead appliance, Virtual Steelhead reduces bandwidth utilization and speeds up application delivery and performance. Virtual Steelhead on VMware vSphere is certified for the Cisco SRE Service Module with Cisco Services Ready Engine Virtualization (Cisco SRE-V).

Virtual Steelhead runs on both the VMware vSphere and Microsoft Hyper-V hypervisors, installed on industry-standard hardware servers.

Figure 1-1. Virtual Steelhead and Hypervisor Architecture

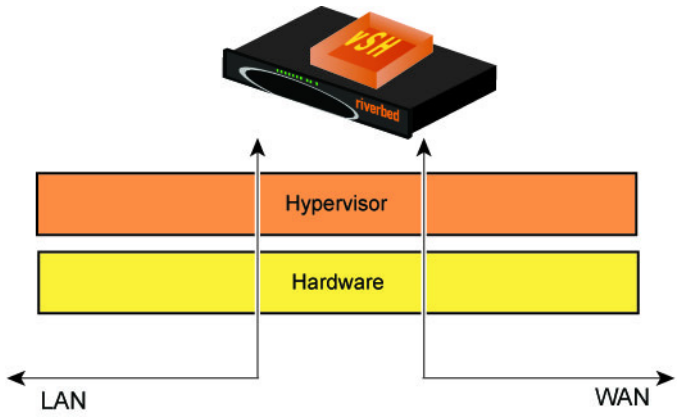

Virtual Steelhead enables consolidation and high availability while providing most of the functionality of the physical Steelhead appliance, with the following exceptions:

- Virtual Services Platform (VSP) or Riverbed Services Platform (RSP)
- Proxy File Service (PFS)
- Fail-to-wire (unless deployed with a Riverbed NIC card)

Note: Hyper-V does not currently support the Riverbed bypass NIC card.

- Hardware reports such as the Disk Status report
- Hardware-based alerts and notifications, such as a RAID alarm

You can integrate Virtual Steelhead into a wide range of networks. You can deploy Virtual Steelhead out-of-path, virtual in-path, or using the Discovery Agent. Virtual Steelhead supports both asymmetric route detection and connection forwarding features. You can make Virtual Steelhead highly available in active-active configurations, with data store synchronization as serial clusters.

After you license and obtain a serial number for Virtual Steelheads, you can manage them across the enterprise from a Riverbed Central Management Console (CMC) v8.0.0 or later.

Virtual Steelhead supports up to 24 virtual CPUs and ten interfaces.

Virtual Steelhead Optimization

With Virtual Steelhead, you can solve a range of problems affecting WANs and application performance, including:

- Insufficient WAN bandwidth
- Inefficient transport protocols in high-latency environments
- Inefficient application protocols in high-latency environments

RiOS intercepts client-server connections without interfering with normal client-server interactions, file semantics, or protocols. All client requests are passed through to the server normally, while relevant traffic is optimized to improve performance.

The optimization techniques RiOS uses are:

- **Data Streamlining** - Steelhead products (Virtual Steelhead, Steelhead appliances, and Steelhead Mobile) can reduce WAN bandwidth utilization by 65% to 98% for TCP-based applications using data streamlining. In addition to traditional techniques like data compression, RiOS uses a Riverbed proprietary algorithm called Scalable Data Referencing (SDR). SDR breaks up TCP data streams into *unique data chunks* that are stored in the hard disk (*data store*) of the device running RiOS (a Steelhead appliance or Steelhead Mobile host system). Each data chunk is assigned a unique integer label (*reference)* before it is sent to a peer RiOS device across the WAN. When the same byte sequence is seen again in future transmissions from clients or servers, the reference is sent across the WAN instead of the raw data chunk. The peer RiOS device (Virtual Steelhead software, Steelhead appliance, or Steelhead Mobile host system) uses this reference to find the original data chunk in its data store, and reconstruct the original TCP data stream.
- **Transport Streamlining** - Virtual Steelhead uses a generic latency optimization technique called transport streamlining. Transport streamlining uses a set of standards and proprietary techniques to optimize TCP traffic between Steelhead appliances. These techniques:
  - ensure that efficient retransmission methods, such as TCP selective acknowledgements, are used.
  - negotiate optimal TCP window sizes to minimize the impact of latency on throughput.
  - maximize throughput across a wide range of WAN links.
- **Application Streamlining** - In addition to data and transport streamlining optimizations, RiOS can apply application-specific optimizations for certain application protocols: for example, CIFS, MAPI, NFS, TDS, HTTP, and Oracle Forms.
- **Management Streamlining** - Management streamlining refers to the methods that Riverbed developed to simplify the deployment and management of RiOS devices. These methods include:
  - **Auto-Discovery Process** - Auto-discovery enables Virtual Steelhead, the Steelhead appliance, and Steelhead Mobile to automatically find remote Steelhead installations, and to optimize traffic using them. Auto-discovery relieves you from having to configure manually large amounts of network information. The auto-discovery process enables administrators to control and secure connections, specify which traffic is optimized, and specify peers for optimization.

    Enhanced auto-discovery automatically discovers the last Steelhead appliance in the network path of the TCP connection. In contrast, the original auto-discovery protocol automatically discovers the first Steelhead appliance in the path. The difference is only seen in environments where there are three or more Steelhead appliances in the network path for connections to be optimized.

    Enhanced auto-discovery works with Steelhead appliances running the original auto-discovery protocol, but it is not the default. When enhanced auto-discovery is enabled on a Steelhead appliance that is peering with other appliances using the original auto-discovery method in a “mixed” environment, the determining factor for peering is whether the next Steelhead appliance along the path uses original auto-discovery or enhanced auto-discovery (regardless of the setting on the first appliance).

    If the next Steelhead appliance along the path is using original auto-discovery, the peering terminates at that appliance (unless peering rules are configured to modify this behavior). Alternatively, if the Steelhead appliance along the path is using enhanced auto discovery, the enhanced probing for a peer continues a step further to the next appliance in the path. If probing reaches the final Steelhead appliance in the path, that appliance becomes the peer.
  - **CMC** - The CMC enables remote Steelhead appliances to be automatically configured and monitored. It also gives you a single view of the data reduction and health of the Steelhead network.

  - **Steelhead Mobile Controller** - The Mobile Controller is the management appliance you use to track the individual health and performance of each deployed software client, and to manage enterprise client licensing. The Mobile Controller enables you to see who is connected, view their data reduction statistics, and perform support operations such as resetting connections, pulling logs, and automatically generating traces for troubleshooting. You can perform all of these management tasks without end-user input.

Virtual Steelhead is typically deployed on a LAN, with communication between appliances occurring over a private WAN or VPN. Because optimization between Steelhead appliances typically occurs over a secure WAN, it is not necessary to configure company firewalls to support Steelhead-specific ports.

For detailed information about how Virtual Steelhead, the Steelhead appliance, or Steelhead Mobile works and deployment design principles, see the *Steelhead Appliance Deployment Guide*.

Configuring Optimization

You configure optimization of traffic using the Management Console or the Riverbed CLI. You configure the traffic that Virtual Steelhead optimizes and specify the type of action it performs using:

- **In-Path rules** - In-path rules determine the action that a Virtual Steelhead takes when a connection is *initiated*, usually by a client. In-path rules are used only when a connection is initiated. Because connections are usually initiated by clients, in-path rules are configured for the initiating, or client-side, Virtual Steelhead. In-path rules determine Virtual Steelhead behavior with SYN packets. You configure one of the following types of in-path rule actions:
  - **Auto** - Use the auto-discovery process to determine if a remote Steelhead appliance is able to optimize the connection attempted by this SYN packet.
  - **Pass-through** - Allow the SYN packet to pass through the Steelhead appliance. No optimization is performed on the TCP connection initiated by this SYN packet.
  - **Fixed-Target** - Skip the auto-discovery process and use a specified remote Steelhead appliance as an optimization peer. Fixed-target rules require the input of at least one remote target Steelhead appliance; an optional backup Steelhead appliance might also be specified.
  - **Deny** - Drop the SYN packet and send a message back to its source.
  - **Discard** - Drop the SYN packet silently.
- **Peering rules** - Peering rules determine how a Virtual Steelhead reacts to a probe query. Peering rules are in ordered lists of fields that a Virtual Steelhead uses to match with incoming SYN packet fields — for example, source or destination subnet, IP address, VLAN, or TCP port — as well as the IP address of the probing Virtual Steelhead. This is especially useful in complex networks. Following are the types of peering rule actions:
  - **Pass** - The receiving Steelhead appliance does not respond to the probing Steelhead appliance and allows the SYN+ probe packet to continue through the network.
  - **Accept** - The receiving Steelhead appliance responds to the probing Steelhead appliance and becomes the remote-side Steelhead appliance (the peer) for the optimized connection.
  - **Auto** - If the receiving Steelhead appliance is not using enhanced auto-discovery, Auto has the same effect as Accept. If enhanced auto-discovery is enabled, the Steelhead appliance becomes the optimization peer only if it is the last Steelhead appliance in the path to the server.

For detailed information about in-path and peering rules and how to configure them, see the *Steelhead Appliance Management Console User's Guide*.

New Features in Version 8.6

Virtual Steelhead v8.6 provides these new features:

- Support for the following new models:
  - VCX5055M, and VCX5055H
  - VCX7055M, and VCX7055L

  These models are only supported as ESXi 5.1-based Virtual Steelhead appliances.

  Only Virtual Steelhead appliances that are built on version 8.6 or greater can be upgraded to these newly supported, high-end models. For example, an v8.6 VCX1555H model appliance can be upgraded to VCX5055M (or any other of the above models), but an v8.5.2 VCX1555H model that has been upgraded to v8.6 cannot be upgraded to a higher capacity model.

  Following are the baseline CPU/RAM/Storage specifications used for Virtual Steelhead models VCX5055 and VCX7055:
  - Server: HPDL560 Gen8
  - CPU: 4 x Intel E5-4640
  - RAM: 8 x 8GB DDR3-1600
  - Storage: 5 x 600GB Intel DC S3500 SSD
- Support for multiple RiOS data stores. The newly supported models can have up to 14 data stores using FTS. Each data store disk can be of any size, but all data store disks on an appliance must be the same size.
- Support for a two-port 10GbE Riverbed multimode fiber bypass NIC.
- Support for multiqueue and RSS for paravirtualized devices. Riverbed now uses the most current vmxnet3 drivers.
- Support for paravirtualized storage controllers. Using paravirtualized storage controllers with multiple data store disks improves performance.

VMware ESX and ESXi

VMware ESX and ESXi are hypervisors that enable you to install and run the Steelhead appliance as a virtual appliance. For details about VMware ESX and ESXi, see http://www.vmware.com.

Your hardware must be compatible with VMware ESX or ESXi to deploy Virtual Steelhead. To ensure hardware compatibility, see http://www.vmware.com/resources/compatibility/search.php.

Virtual Steelhead supports ESX/ESXi 4.0 and later. If you use the NIC card, you must use ESXi 4.1 or later. For ESXi 5.0 and later, the method for supporting the card differs from the 4.1 method. For information, see "Completing the Preconfiguration Checklist" on page 23.

Note: For detailed information about ESX or any other VMware products, see the VMware documentation.

Microsoft Hyper-V

With RiOS v8.0.3 and later, some models of Virtual Steelhead run on the Microsoft Hyper-V hypervisor, which is available on Windows Server 2012 and Hyper-V Server 2012. For information on Hyper-V, see http://www.microsoft.com/en-us/server-cloud/hyper-v-server/.

In RiOS v8.0.3 and later, virtual in-path and out-of-path modes are supported, while direct in-path deployment with a Riverbed NIC card is not.

The underlying RiOS image is the same as the physical models, so that after you have installed the Virtual Steelhead package, you can upgrade using a standard image.

Virtual Steelhead Deployment Guidelines

Important: Riverbed requires that you follow these guidelines when deploying the Virtual Steelhead package on a hypervisor. If you do not follow the configuration guidelines, Virtual Steelhead might not function properly, or might cause outages in your network.

Network Configuration

When you deploy either hypervisor, follow this guideline:

- **Ensure that a network loop does not form** - An in-path interface is, essentially, a software connection between the lanX_Y and wanX_Y interfaces. Before deploying a Virtual Steelhead, Riverbed strongly recommends that you connect each LAN and WAN virtual interface to a distinct virtual switch and physical NIC (through the vSphere Networking tab). *Connecting LAN and WAN virtual NICs to the same vSwitch or physical NIC could create a loop in the system and might make your hypervisor unreachable.*

When you deploy Virtual Steelhead on ESX or ESXi, follow these guidelines:

- **Enable promiscuous mode for the LAN/WAN vSwitch** - Promiscuous mode allows the LAN/WAN Virtual Steelhead NICs to intercept traffic not destined for the Steelhead installation and is mandatory for traffic optimization on in-path deployments. You must accept promiscuous mode on each in-path virtual NIC. You can enable promiscuous mode through the vSwitch properties in vSphere. For details, see "Installing Virtual Steelhead" on page 23.
- **Use distinct port groups for each LAN or WAN virtual NIC connected to a vSwitch for each Virtual Steelhead** - If you are running multiple Virtual Steelhead virtual machines (VMs) on a single virtual host, you must add the LAN (or WAN) virtual NIC from each virtual machine (VM) into a different port group (on each vSwitch). This prevents the formation of network loops.

Network Performance

Follow these configuration tips to improve performance:

- **Use at least a Gigabit link for LAN/WAN** - For optimal performance, connect the LAN/WAN virtual interfaces to physical interfaces that are capable of at least 1 Gbps.
- **Do not share physical NICs** - For optimal performance, assign a physical NIC to a single LAN or WAN interface. Do not share physical NICs destined for LAN/WAN virtual interfaces with other VMs running on the hypervisor. Doing so can create performance bottlenecks.
- **Ensure that the host has resources for overhead** - In addition to reserving the CPU resources needed for the Virtual Steelhead model, verify that additional unclaimed resources are available. Due to hypervisor overhead, VMs can exceed their configured reservation. For details on hypervisor resource reservation and calculating overhead, see "Managing Licenses and Model Upgrades" on page 47.
- **Do not overprovision the physical CPUs** - Do not run more VMs than there are CPUs. For example, if a hypervisor is running off a 4-core CPU, all the VMs on the host should use no more than four vCPUs.
- **Use a server-grade CPU for the hypervisor** - For example, use a Xeon or Opteron CPU as opposed to an Intel Atom.
- **Always reserve RAM** - Memory is another very important factor in determining Virtual Steelhead performance. Reserve the RAM that is needed by the Virtual Steelhead model, but ensure there is extra RAM for overhead. This overhead can provide a performance boost if the hypervisor exceeds its reserved capacity.
- **Virtual RAM should not exceed physical RAM** - The total virtual RAM provisioned for all running VMs should not be greater than the physical RAM on the system.
- **Do not use low-quality storage for the RiOS data store disk** - Make sure that the Virtual Steelhead disk used for the data store VMDK (for ESX) or VHD (for Hyper-V) resides on a disk medium that supports a high number of Input/Output Operations Per Second (IOPS). For example, use NAS, SAN, or dedicated SATA disks.
- **Do not share host physical disks** - To achieve near-native disk I/O performance, do not share host physical disks (such as SCSI or SATA disks) between VMs. When you deploy Virtual Steelhead, allocate an unshared disk for the RiOS data store disk.
- **Do not use hyperthreading** - Hyperthreading can cause contention among the virtual cores, resulting in significant loss of performance.
- **BIOS Power Management Settings** - If configurable, power management settings in the BiOS should be set to maximize performance.

Deployment Options

Typically you deploy Virtual Steelhead on a LAN with communication between appliances taking place over a private WAN or VPN. Because optimization between Steelhead appliances typically takes place over a secure WAN, it is not necessary to configure company firewalls to support Steelhead appliance-specific ports.

For optimal performance, minimize latency between Virtual Steelheads and their respective clients and servers. Place the Virtual Steelheads as close as possible to your network end points: client-side Virtual Steelheads as close to your clients as possible, and server-side Virtual Steelheads as close to your servers as possible.

Ideally, Virtual Steelheads optimize only traffic that is initiated or terminated at their local sites. The best and easiest way to achieve this traffic pattern is to deploy the Virtual Steelheads where the LAN connects to the WAN, and not where any LAN-to-LAN or WAN-to-WAN traffic can pass through (or be redirected to) the Steelhead appliance.

For detailed information about deployment options and best practices for deploying Steelhead appliances, see the *Steelhead Appliance Deployment Guide*.

Before you begin the installation and configuration process, you must select a network deployment.

Note: You can also use the Discovery Agent to deploy the Virtual Steelhead. For information, see Chapter 5, "Using Discovery Agent."

In-Path Deployment

You can deploy Virtual Steelhead in the same scenarios as the Steelhead appliance, with the following exception: Virtual Steelhead software does not provide a failover mechanism like the Steelhead appliance fail-to-wire. For full failover functionality, you must install a Riverbed NIC with Virtual Steelhead.

Riverbed bypass cards come in four-port and two-port models. For more information on NICs and Virtual Steelhead, see "NICs for Virtual Steelhead" on page 17.

For deployments where a Riverbed bypass card is not an option (for example, in a Cisco SRE deployment) Riverbed recommends that you do not deploy your Virtual Steelhead in-path. If you are not using a bypass card, you can still have a failover mechanism, by employing either a virtual in-path or an out-of-path deployment. These deployments allow a router using WCCP or PBR to handle failover.

Promiscuous mode is required for in-path deployments.

Note: Hyper-V does not support promiscuous mode, or direct in-path deployment.

Virtual In-Path Deployment

In a virtual in-path deployment, Virtual Steelhead is virtually in the path between clients and servers.Traffic moves in and out of the same WAN interface, and the LAN interface is not used. This deployment differs from a physical in-path deployment in that a packet redirection mechanism, such as WCCP or PBR, directs packets to Steelhead appliances that are not in the physical path of the client or server. In this configuration, clients and servers continue to see client and server IP addresses.

On Virtual Steelhead models with multiple WAN ports, you can deploy WCCP and PBR with the same multiple interface options available on the Steelhead appliance.

For a virtual in-path deployment, attach only the WAN virtual NIC to the physical NIC, and configure the router using WCCP or PBR to forward traffic to the VM to optimize. You must also enable in-path OOP on Virtual Steelhead.

Out-of-Path Deployment

The Virtual Steelhead is not in the direct path between the client and the server. Servers see the IP address of the server-side Steelhead installation rather than the client IP address, which might have an impact on security policies.

For a virtual out-of-path (OOP) deployment, connect the primary interface to the physical in-path NIC and configure the router to forward traffic to this NIC. You must also enable OOP on Virtual Steelhead.

The following caveats apply to server-side OOP Virtual Steelhead configuration:

- OOP configuration does not support auto-discovery. You must create a fixed-target rule on the client-side Steelhead appliance.
- You must create an OOP connection from an in-path or logical in-path Steelhead appliance and direct it to port 7810 on the primary interface of the server-side Steelhead appliance. This setting is mandatory.
- Interception is not supported on the primary interface.
- An OOP configuration provides nontransparent optimization from the server perspective. Clients connect to servers, but servers treat it like a server-side Steelhead appliance connection. This affects log files, server-side ACLs, and bidirectional applications such as rsh.
- You can use OOP configurations along with in-path or logical in-path configurations.

Virtual Steelhead Platform Models

The tables in this section list the platform models available for Virtual Steelhead and Virtual Steelhead CX (VCX). Each Virtual Steelhead has a primary and an auxiliary interface. Confirm that you have the resources required for the Virtual Steelhead model you are installing before you download and install Virtual Steelhead.

The following table lists the Virtual Steelhead xx50 Models.

Virtual Steelhead Model	Virtual CPU	Min. CPU Speed	Memory	Management Disk (VMDK1)	RiOS Data Store Disk (VMDK2)	Optimized WAN Capacity	Max. Connections
V150M	1 CPU	1200 MHz	1 GB	30 GB	44 GB	1 Mbps	20
V250L	1 CPU	1200 MHz	1 GB	30 GB	44 GB	1 Mbps	30
V250M	1 CPU	1200 MHz	1 GB	30 GB	44 GB	4 Mbps	125
V250H	1 CPU	1200 MHz	1 GB	30 GB	44 GB	4 Mbps	200
V550M	2 CPUs	1200 MHz	2 GB	30 GB	80 GB	2 Mbps	300
V550H	2 CPUs	1200 MHz	2 GB	30 GB	80 GB	4 Mbps	600
V1050L	2 CPUs	1800 MHz	2 GB	30 GB	102 GB	8 Mbps	800
V1050M	2 CPUs	1800 MHz	2 GB	30 GB	102 GB	10 Mbps	1300
V1050H	2 CPUs	1800 MHz	4 GB	30 GB	202 GB	20 Mbps	2300
V2050L	4 CPUs	2000 MHz	6 GB	30 GB	400 GB	45 Mbps	2500
V2050M	4 CPUs	2000 MHz	6 GB	30 GB	400 GB	45 Mbps	4000
V2050H	4 CPUs	2000 MHz	6 GB	30 GB	400 GB	45 Mbps or 90 Mbps with a separate upgrade	6000

The following table lists the Virtual Steelhead CX xx55 Models.

Virtual Steelhead Model	Virtual CPU	Min. CPU Speed	Memory	Management Disk (VMDK1)	RiOS Data Store Disk (VMDK2+)	QoS Bandwidth	Optimized WAN Capacity	Max. Connections
VCX255U	1 CPU	1000 MHz	2 GB	38 GB	50 GB	4 Mbps	2 Mbps	50
VCX255L	1 CPU	1000 MHz	2 GB	38 GB	50 GB	12 Mbps	6 Mbps	75
VCX255M	1 CPU	1000 MHz	2 GB	38 GB	50 GB	12 Mbps	6 Mbps	150
VCX255H	1 CPU	1000 MHz	2 GB	38 GB	50 GB	12 Mbps	6 Mbps	230
VCX555L	1 CPU	1200 MHz	2 GB	38 GB	80 GB	12 Mbps	6 Mbps	250
VCX555M	1 CPU	1200 MHz	2 GB	38 GB	80 GB	20 Mbps	10 Mbps	400
VCX555H	1 CPU	1200 MHz	2 GB	38 GB	80 GB	20 Mbps	10 Mbps	650
VCX755L	2 CPUs	1200 MHz	2 GB	38 GB	102 GB	45 Mbps	10 Mbps	900
VCX755M	2 CPUs	1200 MHz	2 GB	38 GB	102 GB	45 Mbps	10 Mbps	1500
VCX755H	2 CPUs	1200 MHz	4 GB	38 GB	150 GB	45 Mbps	20 Mbps	2300
VCX1555L	4 CPUs	1200 MHz	8 GB	38 GB	400 GB	100 Mbps	50 Mbps	3000
VCX1555M	4 CPUs	1200 MHz	8 GB	38 GB	400 GB	100 Mbps	50 Mbps	4500
VCX1555H	4 CPUs	1200 MHz	8 GB	38 GB	400 GB	100 Mbps	100 Mbps	6000
VCX5055M	12 CPUs	2400 MHz	16 GB	82 GB	8 x 80 GB	no limit	200 Mbps	14,000
VCX5055H	12 CPUs	2400 MHz	16 GB	82 GB	8 x 80 GB	no limit	400 Mbps	25,000
VCX7055L	16 CPUs	2400 MHz	32 GB	178 GB	10 x 160 GB	no limit	622 Mbps	75,000
VCX7055M	24 CPUs	2400 MHz	48 GB	178 GB	14 x 160 GB	no limit	1 Gbs	100,000

The platform families are independent. You cannot upgrade a xx50 model to a xx55 model. The xx55 virtual models require RiOS v8.0 or later.

The data store size per model allocates extra disk space to accommodate hypervisor overhead. As of v8.5.1, the size of the management disk for new open virtualization appliance (OVA) deployments for the VCX models is 38 GB. Older models that upgrade still use a 50 GB management disk.

Flexible RiOS Data Store

As of RiOS v8.5.1, the flexible data store feature for VCX models supports a smaller data store size, down to a minimum 12 GB.

To change the disk size of a running Virtual Steelhead, you must first power off the VM. From the Settings section, you can expand or remove the RiOS data store (second) disk, and replace it with a smaller disk. (Reducing the disk size will not work.) Modifying the disk size causes the RiOS data store to automatically clear.

If you provide a disk larger than the configured RiOS data store for the model, the entire disk is partitioned but only the allotted amount for the model is used.

Memory and CPU requirements are a hard requirement for a model to run. Flexible RiOS data store is not supported for the older Vxx50 models.

Multiple RiOS Data Stores

As of RiOS v8.6, Virtual Steelhead models VCX5055, and VCX7055 support up to 14 RiOS data stores using FTS. Riverbed recommends that all RiOS data stores on an appliance are the same size.

To add additional data stores, you must power off the VM.

In-Path Pairing for NIC Interfaces

Virtual Steelhead models are not limited to a fixed number of NIC interfaces. However, the in-path pair limit is four (four LAN and four WAN interfaces), including bypass cards. If you want to use the Virtual Steelhead bypass feature, you are limited to the number of hardware bypass pairs the model can support.

Each Virtual Steelhead requires a primary and aux interface, which are the first two interfaces added. If you add additional interface pairs to the VM, they are added as in-path optimization interfaces. Total bandwidth and connection limits still apply.

NICs for Virtual Steelhead

Riverbed NICs provide hardware-based fail-to-wire and fail-to-block capabilities for Virtual Steelhead. The configured failure mode is triggered if the ESX or ESXi host loses power or is unable to run the Virtual Steelhead guest, if the Virtual Steelhead guest is powered off, or if the Virtual Steelhead guest experiences a significant fault (using the same logic as the physical Steelhead appliance).

Note: Physical fail-to-wire and fail-to-block NICs in Virtual Steelhead are not supported on Hyper-V.

Riverbed NICs are available in two-port and four-port configurations:

Riverbed NICs for Virtual Steelhead	Orderable Part Number	Virtual Steelhead Models
Two-Port 1GbE TX Copper NIC	NIC-001-2TX	All
Four-Port 1GbE TX Copper NIC	NIC-002-4TX	1050L, 1050M, 1050H, 2050L, 2050M, and 2050H VCX255, VCX555, VCX755, and VCX1555, VCX5055, VCX7055
Two-Port 10GbE Multimode Fiber NIC (direct I/O only)	NIC-008-2SR	VCX5055, and VCX7055

You must use Riverbed NICs for fail-to-wire or fail-to-block with Virtual Steelhead. NIC cards without a bypass feature from other vendors are supported for functionality other than fail-to-wire and fail-to-block, if supported by ESX or ESXi.

Requirements for Virtual Steelhead Deployment with a NIC

To successfully install a NIC in an ESXi host for Virtual Steelhead, you need the following:

- ESXi host with a PCIe slot.
- vSphere Client access to the ESXi host.
- VMware ESXi 5.0 and later and RiOS v8.0.3 or later.

  —*or*—

  VMware ESXi 4.1 and one of the following RiOS versions:

  – For V150, RiOS v7.0.3a or later.

  – For V250, V550, V1050, and V2050, RiOS v7.0.2 or later.

  – For VCX555, VCX755, and VCX1555, RiOS v8.0 or later.

For ESXi 4.1, you also need the following:

- ESXi bypass driver (a .VIB file) available from https://support.riverbed.com.
- Intel 82580 Gigabit network interface driver.
- By default, ESXi does not include the Intel 82580 Gigabit Ethernet network interface driver needed for the Riverbed bypass card. If you do not have this driver installed, you can download it from the VMware Web site.

  For ESX 4.1:
  http://downloads.vmware.com/d/details/dt_esxi4x_intel_10g_825xx/ZHcqYnQldypiZCVodw==
- SSH and SCP access to the ESXi host.

For more information on Riverbed NICs installation, see the *Network Interface Card Installation Guide*. The installation procedure in this manual assumes you have successfully installed a Riverbed NIC following the instructions in the *Network Interface Card Installation Guide*.

The number of hardware bypass pairs (that is, one LAN and one WAN port) supported is determined by the model of the Virtual Steelhead:

- models V150, V250, and V550: one bypass pair
- models V1050 and V2050: two bypass pairs (that is, two LAN and two WAN ports)
- models VCX555, VCX755, VCX1555, VCX 5055, and VCX 7055: two bypass pairs

---

Note: You can install a four-port card in an ESXi host for a Virtual Steelhead 150, 250, or 550. However, only one port pair is available because the Virtual Steelhead model type determines the number of pairs.

The following configurations have been tested:

- Two Virtual Steelhead guests, each using one physical pair on a single four-port Riverbed NIC card
- Two Virtual Steelhead guests connecting to separate cards
- One Virtual Steelhead guest connecting to bypass pairs on different NIC cards

For more information on installation and configuration of Virtual Steelhead with a Riverbed NIC, see "Completing the Preconfiguration Checklist" on page 23.

Virtual Steelhead on the Cisco SRE

In addition to standard ESX and ESXi, you can run Virtual Steelhead on a Cisco server blade, using the SRE platform, based on ESXi v5.0. The following table lists the Virtual Steelhead models supported on each supported Cisco SRE model, and the required version of RiOS, disk configuration, and RAM.

SRE Model	Virtual Steelhead Model	RiOS Version	Disk Configuration	RAM
910	V1050H, VCX755H	v6.5.4+, 7+, 8+	RAID1	8 GB
910	V1050M, VCX755M	v6.5.4+, 7+, 8+	RAID1	4 GB
900	V1050M, VCS755M	v6.5.4+, 7+, 8+	RAID1	4 or 8 GB
700/710	V250H	v6.5.4+, 7+, 8+	Single disk	4 GB
300	NOT SUPPORTED			

For more information on deploying Virtual Steelhead on a Cisco SRE blade, see the Riverbed deployment guide, “Virtual Steelhead on Cisco SRE/ISR G2” on the Riverbed Support site at https://support.riverbed.com/docs/technotes.htm.

You can find more information on the Cisco SRE platform on the Cisco Web site. For example: http://www.cisco.com/en/US/prod/collateral/modules/ps10598/data_sheet_c78-553913.html.

CHAPTER 2 Setting Up Virtual Steelhead on ESX and ESXi

This chapter describes how to install and configure Virtual Steelhead on VMware ESX and ESXi. It includes the following sections:

- "Basic Steps for Installing Virtual Steelhead on ESX/ESXi" on page 21
- "Obtaining the Virtual Steelhead Software Package" on page 22
- "Installing Virtual Steelhead with a Riverbed NIC" on page 22
- "Completing the Preconfiguration Checklist" on page 23
- "Installing Virtual Steelhead" on page 23

Basic Steps for Installing Virtual Steelhead on ESX/ESXi

This section provides an overview of the basic steps to install and configure Virtual Steelhead on ESX and ESXi, followed by detailed procedures.

If you are installing on Cisco SRE, see the section "Virtual Steelhead Deployment Guidelines" on page 12. If you have installed a Riverbed NIC, see the section "Completing the Preconfiguration Checklist" on page 23.

Task	Reference
1. Confirm that ESX/ESXi is provisioned to run the Virtual Steelhead model. Check to make sure the resources are available and configure ESX/ESXi before choosing your Virtual Steelhead model.	"Installing Virtual Steelhead" on page 23
2. Obtain the Virtual Steelhead package from Riverbed Support.	"Obtaining the Virtual Steelhead Software Package" on page 22
3. Gather network settings for the configuration wizard.	"Completing the Preconfiguration Checklist" on page 23
4. Deploy the Virtual Steelhead image, which automatically configures the Virtual Steelhead to the VCX255U model specifications.	"Installing Virtual Steelhead" on page 23

Task	Reference
5. In VMware vSphere Client or Open Virtual Machine Format (OVF) Tool, configure the VM to accommodate the correct target Virtual Steelhead model specifications. You can skip this step if you are installing a model VCX255U and have the appropriate hardware.	"Installing Virtual Steelhead" on page 23
6. Power on the VM, start the Virtual Steelhead, and log in.	

Obtaining the Virtual Steelhead Software Package

Virtual Steelhead is provided by Riverbed as a software image that contains the VMX and VMDK files necessary to create the VM.

The Virtual Steelhead image is an installable OVA package. OVA is a platform-independent, efficient, extensible, and open packaging distribution format. The OVA package provides the complete default specification for Virtual Steelhead, including its required virtual disks, CPU, memory, networking, and storage. To install a Virtual Steelhead model other than the default model, you first install the default and then upgrade it to a higher model.

The default model for the xx55 VCX models is a VCX255U.

The OVA is a compressed .tar.gz package that quickly creates a VM with predefined settings. It contains the following files:

- **OVF file** - Provides the XML description of Virtual Steelhead.
- **Two VMDK files** - One file contains the management system (the smaller VMDK) and the other contains the data store. The separate files let you resize the RiOS data store without losing any data in the management partitions.
- **Manifest file** - Contains the SHA1 checksum of the OVF and VMDK.
- **VMX file** - Contains the primary configuration.

You can download the OVA package from the Riverbed Support Web site at https://support.riverbed.com. Access to the software downloads requires registration.

Installing Virtual Steelhead with a Riverbed NIC

In RiOS v8.0.3 and later, you can configure bypass support using the VMware Direct Path feature on ESXi 5.0 and later virtual hosts. This feature allows Virtual Steelhead to directly control the physical bypass card. The procedure for configuring bypass support for ESXi 5.0 and later is documented in the *Network Interface Card Installation Guide*.

For instructions on how to configure bypass support on prior releases of ESX/ESXi, see Appendix A, "Configuring a Riverbed NIC in ESX 4.1."

In RiOS v8.6 and later, some Virtual Steelhead models support a Two-Port 10GbE Multimode Fiber NIC (direct I/O only).

Important: You must use a Riverbed-branded NIC. Virtual Steelhead does not support cards not provided by Riverbed. If you currently use a Riverbed-branded NIC with ESXi 4.1, you can use the same card if you want to upgrade the ESXi version. However, you must reconfigure the card to support the bypass method used in ESXi 5.0.

Note: Using passthrough devices requires that a memory reservation be made for the full amount of allocated memory. This reservation is done automatically initially, but if a model upgrade requires more memory, you must manually increase the reservation before powering on the VM.

Completing the Preconfiguration Checklist

This section lists the parameters you specify to complete the initial configuration of Virtual Steelhead.

Be prepared to provide values for the network settings listed in the following checklist when prompted by the configuration wizard.

Network Setting	Your Value
Hostname	
IP address	
Netmask	
Default gateway	
DNS Server	
Domain Name	

Installing Virtual Steelhead

This section describes the procedures for installing the VM OVA package obtained from Riverbed. You install the package using your VMware management tools, either OVF Tool or VMware vSphere Client. This section describes how to install and configure the default Virtual Steelhead model on a VMware ESX host using the vSphere Client.

The default model for the xx55 VCX models is a VCX255U.

To install a Virtual Steelhead model other than the default model, install the default and upgrade it to a higher model.

The installation package contains predefined hardware requirements and configuration for the default model Virtual Steelhead. Do not open or modify any of the files in the package. The package files use several gigabytes of disk space (the package itself is less than 1 GB).

Note: See the VMware Web site for documentation on OVF Tool and vSphere Client.

To install Virtual Steelhead

1. Obtain the VM package from https://support.riverbed.com and download it locally.
2. Open VMware vSphere, type the hostname IP address or name, type your user name, password, and click **Login**.

Figure 2-1. vSphere Client Login Page

3. Choose File > Deploy OVF template.

4. Select Deploy from file, click **Browse**, select the OVA file, and click **Open**.

Figure 2-2. Deploy OVF Template Page

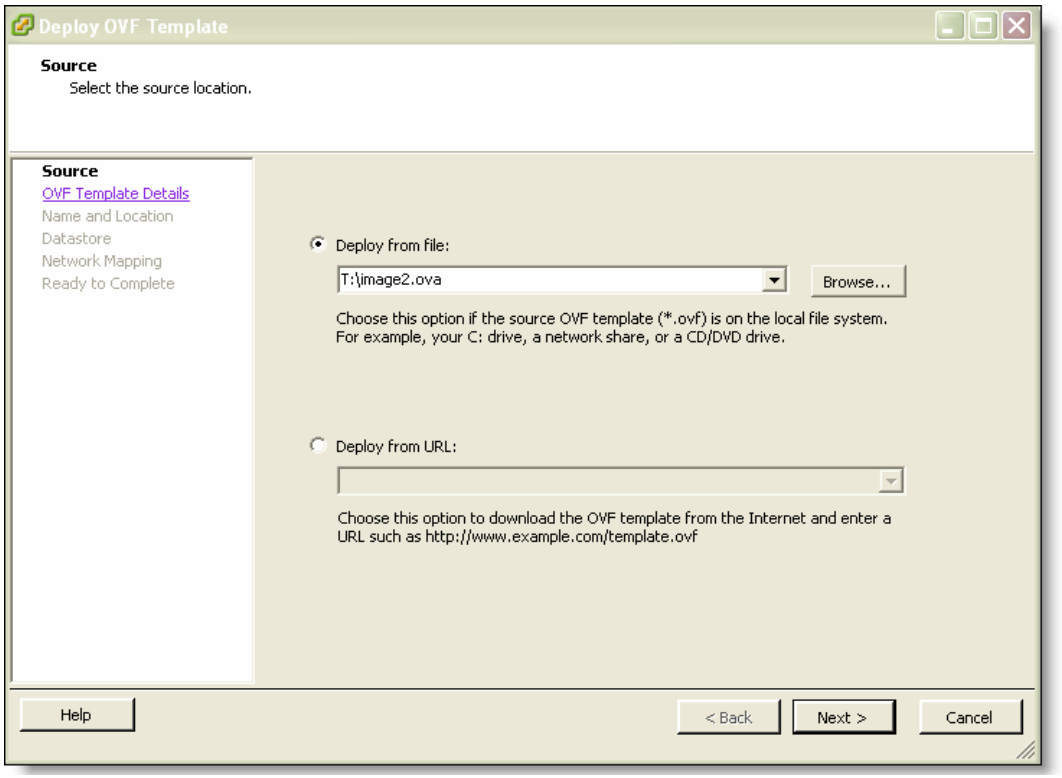

5. Click **Next**.

Figure 2-3. Deploy OVF Template Details Page

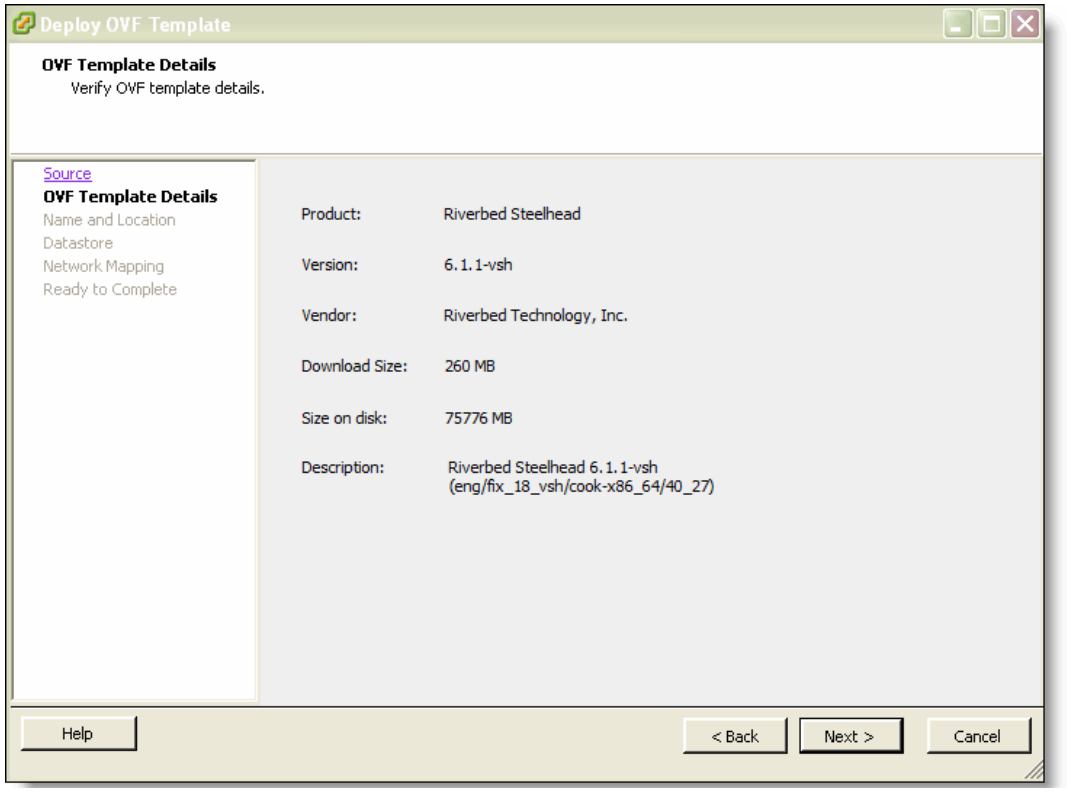

6. Verify that the OVA file is the one you want to deploy, and click **Next**.

Figure 2-4. Name and Location Page

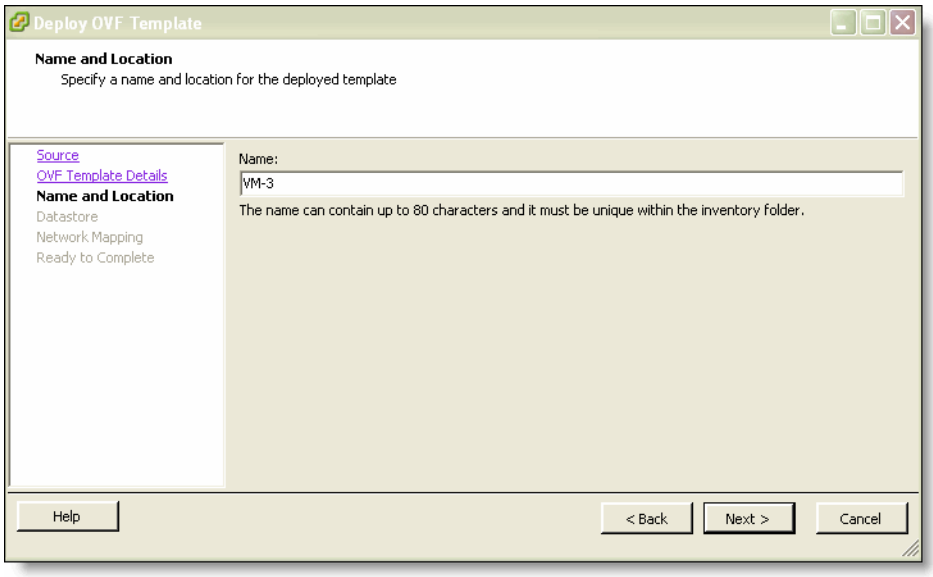

7. Specify a name for the VM.

8. Click **Next**.

Figure 2-5. VMware Data Store Page

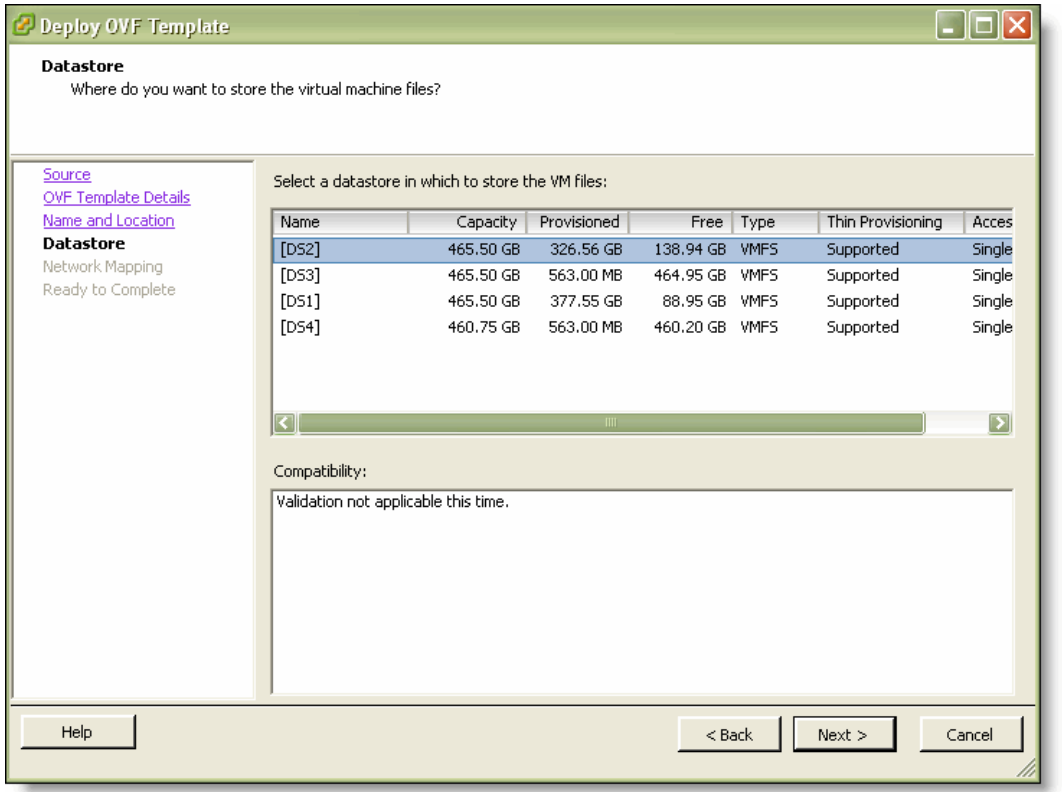

9. Select a host datastore in which to store the VM and its virtual disk files:
   - The standard installation puts both VMDKs on a single host datastore. The datastore that holds the VMDKs can be modified later in the install process.
   - Make sure that the host datastore you select has enough capacity for the OVA package to install. For example, for a VCX255U you need at least 88 GB. For a VCX555M you need at least 130 GB.
   - You can install the smaller VMDK containing the management disk on a datastore backed by any type of underlying storage media.
   - Riverbed recommends that you put the larger VMDK containing the RiOS data store on a host datastore backed by the fastest available storage media. That datastore should have enough room to expand to the required size of the Virtual Steelhead model.
   - Do not share host physical disks (such as SCSI or SATA disks) between VMs. Select an unshared disk for the data store disk.
   - Never delete the first VMDK, which contains the VM's operating system.

10. Click **Next**.

11. On the Disk Format page, select Thick provisioned format.

    Thick provisioning preallocates all storage.

12. Click **Next**.

Figure 2-6. Network Mapping Page

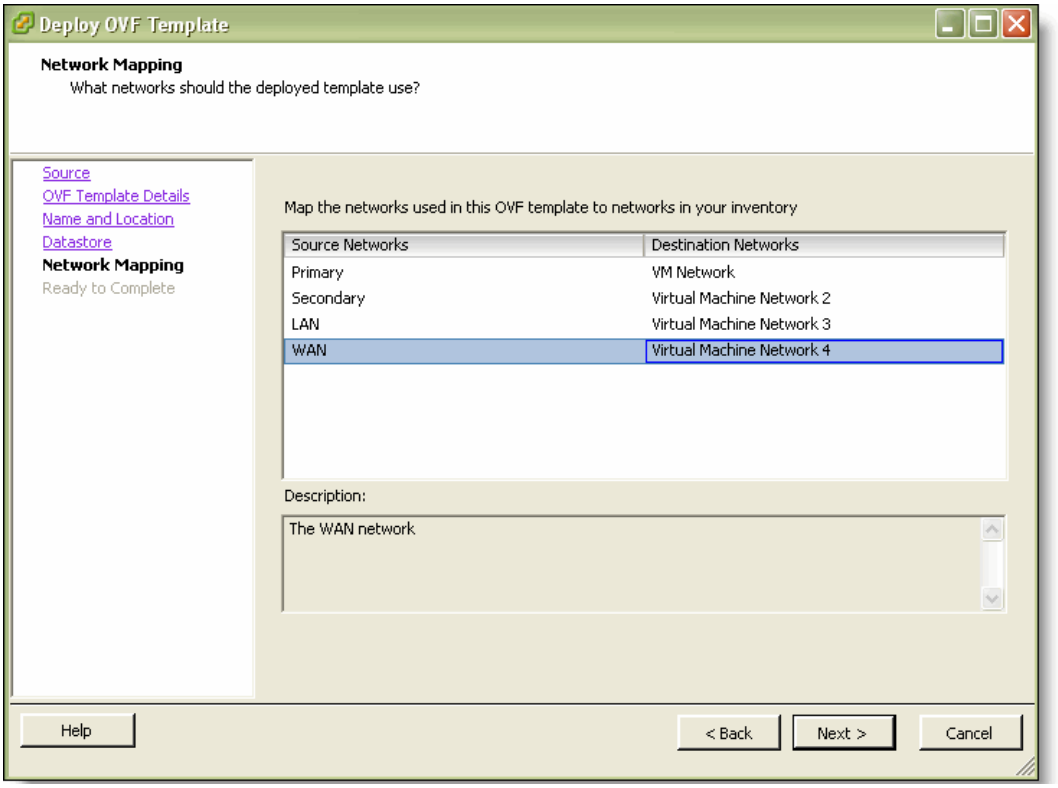

13. Select the destination network name and select a network from the drop-down list to map the source network to a destination network.

    If you have installed a Riverbed NIC, you must map the LAN source network to the pg-vmnic3 port label and the WAN source network to the pg-vmnic2 port label.

Important: Make sure that you map each source network to a unique destination network. If a source network is mapped to the same destination as another source, an error message appears. Mapping source networks to the same destination network can create a loop in the system and might make your ESX host unreachable. For details, see "Virtual Steelhead Deployment Guidelines" on page 12.

---

14. Click **Next**.

Figure 2-7. Ready to Complete Page

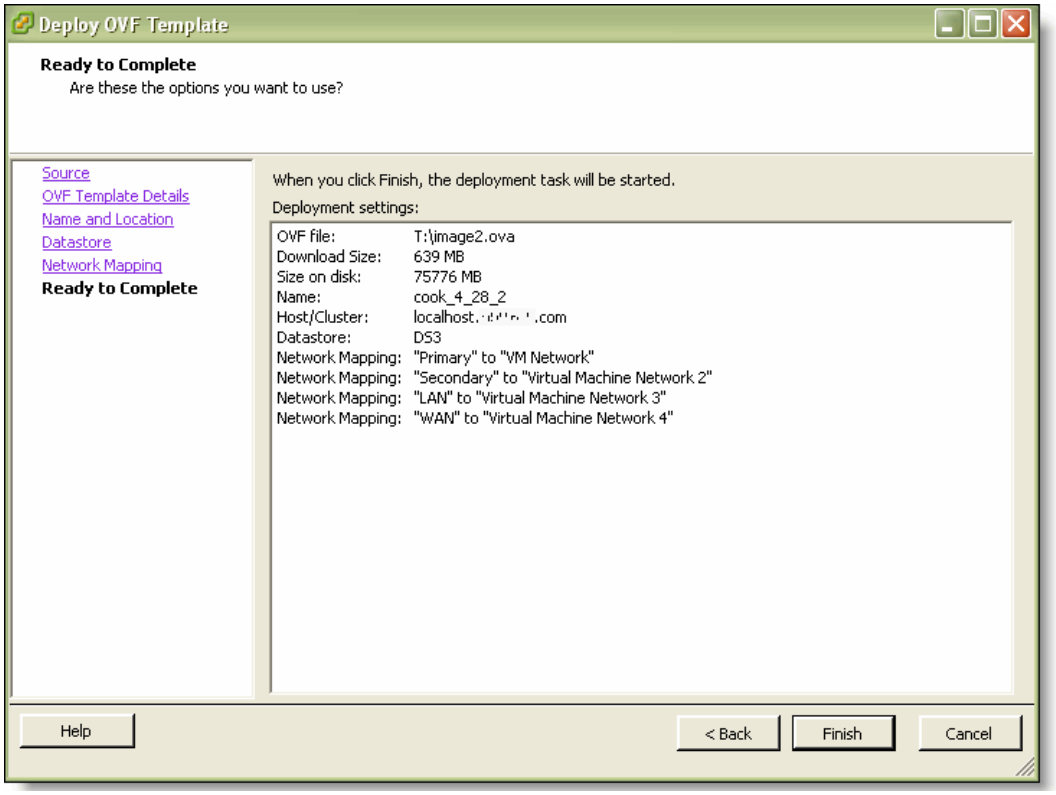

15. Verify the deployment settings and click **Finish**.

    A dialog box shows the amount of time remaining for the deployment.

    When the deployment finishes, a dialog box informs you that the deployment was successful.

16. Click **Close**.

    The new VM appears under the hostname or host IP address to the VM inventory.

    If you do not have a Riverbed NIC, skip the next section and go to, "To set Promiscuous Mode for in-path deployments" on page 31.

To set the adapter type for deployments with Riverbed NICs

Note: For ESXi 5.0 and greater, see the *Network Interface Card Installation Guide*.

1. In vSphere, select the ESX or ESXi host.
2. Under the host, right-click the Virtual Steelhead guest and choose Edit Settings.
3. Click **Add**.
4. Select **Ethernet adapter** and click **Next**.
5. From the adapter Type drop-down list, select **VMXNET3**.
6. From the Network Label drop-down list, select **pg-vmnic1**.
7. Click **Next**.
8. Click **Finish**.
9. Repeat the steps above, specifying network label **pg-vmnic0** instead of **pg-vnic-1**.
10. Repeat the steps in this procedure, specifying network label **pg-bpvm0** instead of **pg-vnic-1**.
11. Verify that the connected state of each of the four network adapters in the Virtual Steelhead guest are set to **Connect at power on**.

    To confirm this setting, choose Edit Settings > Hardware and select your adapter.

12. Under the host, select **Edit Settings**.

    Your Edit Settings window looks similar to Figure 2-8.

Figure 2-8. Edit Settings

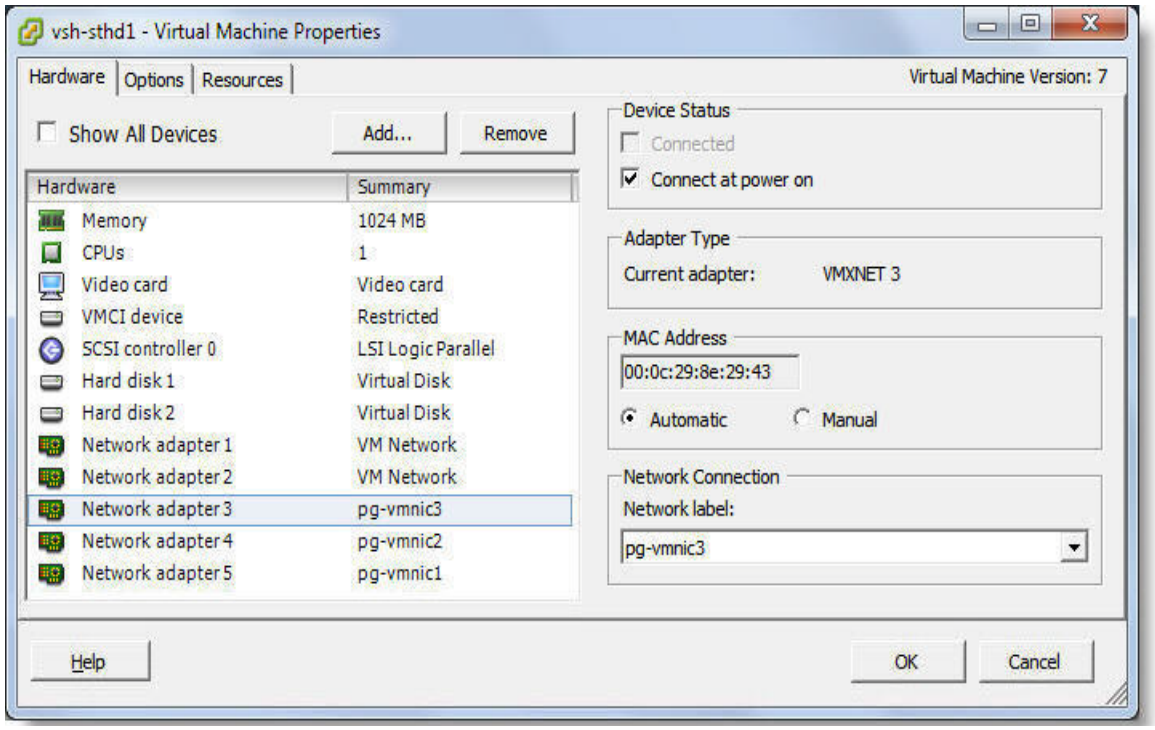

13. Click **OK**.

Next, you need to accept promiscuous mode on each in-path virtual NIC. Promiscuous mode allows the LAN/WAN Virtual NICs to intercept traffic not destined for the VM and is mandatory for traffic optimization for in-path deployments. If you are deploying Virtual Steelhead out-of-path or virtual in-path, skip this procedure and go to the next section, "To power on the VM" on page 34.

To set Promiscuous Mode for in-path deployments

1. Open the vSphere client.
2. In the left panel, select the hostname or IP address.
3. Select the Configuration tab.

4. In the Hardware section, select Networking.

   A list of virtual switches appears.

Figure 2-9. Configuration Tab

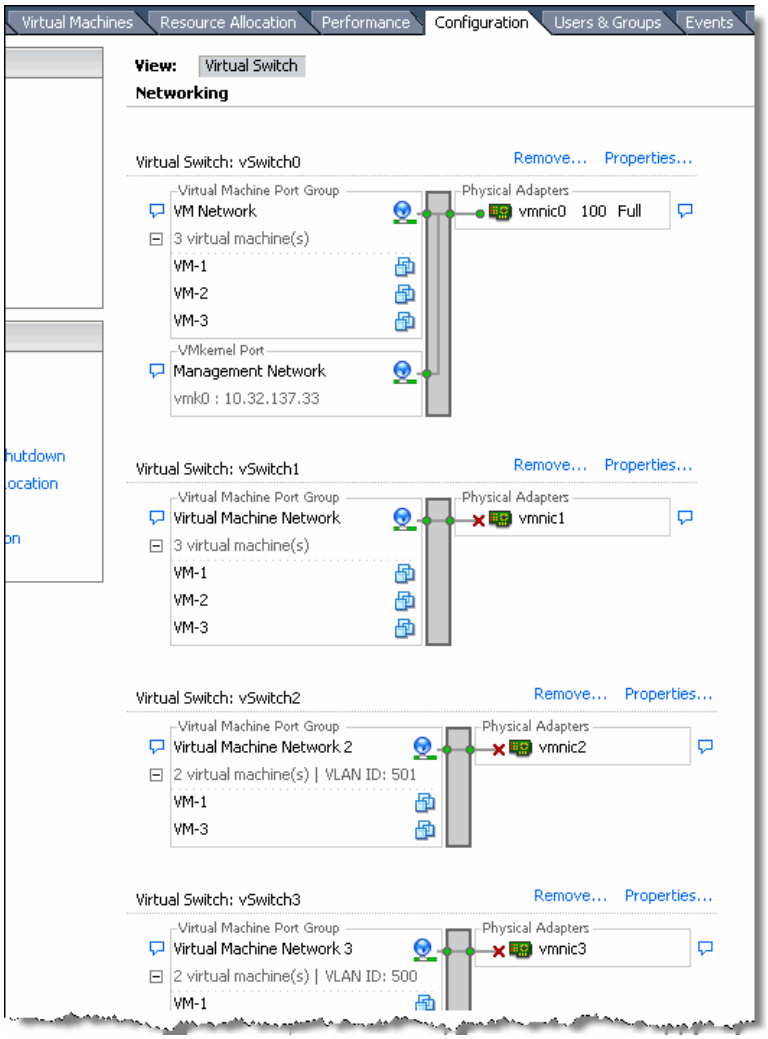

5. To the right of the first virtual switch on the tab, select Properties.

   A vSwitch Properties dialog box appears.

Figure 2-10. vSwitch Properties

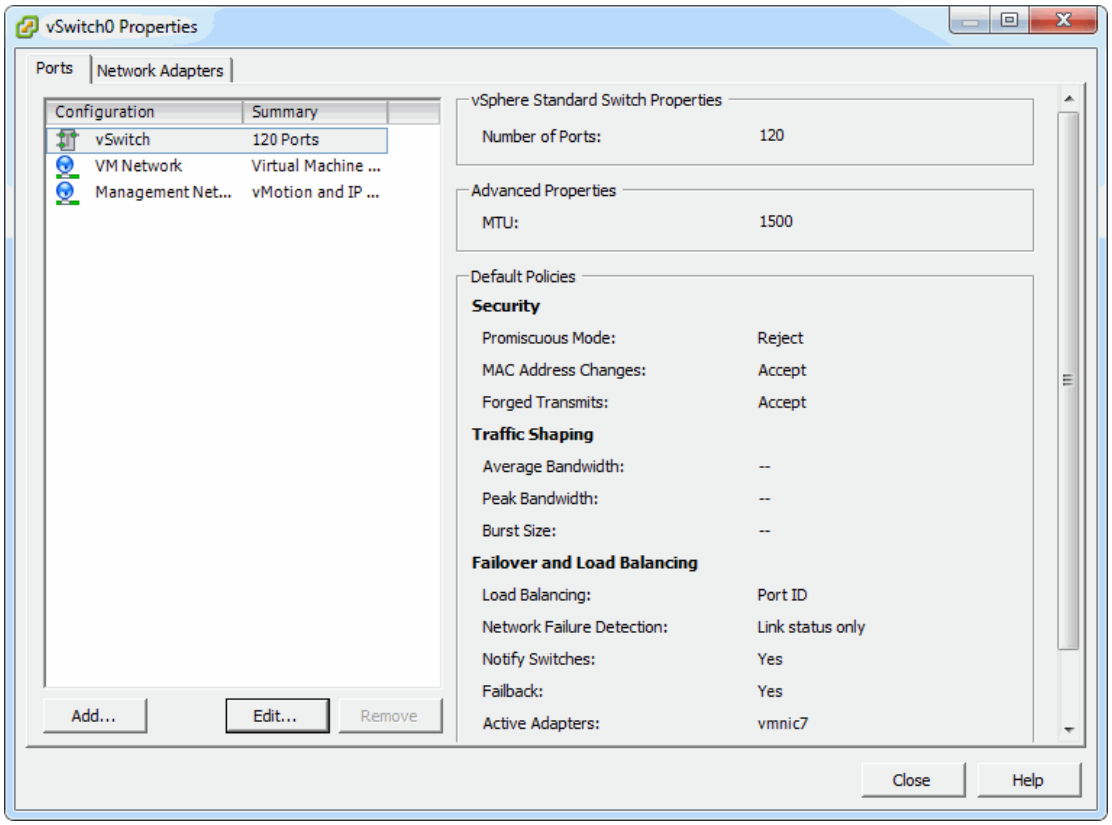

6. From the Configuration list, select the vSwitch configuration and click **Edit**.

7. Select the Security tab.

Figure 2-11. vSwitch Properties Security Tab

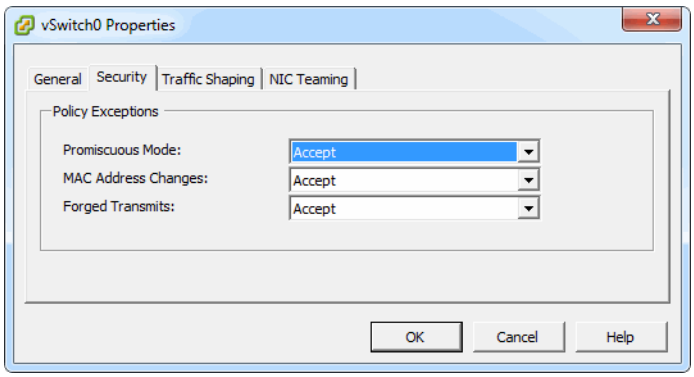

8. In the Promiscuous Mode list, select Accept and click **OK**.

   You return to the Properties dialog box.

9. From the Configuration list, select the Port Group configuration and click **Edit**.

10. Select the Security tab for the port group.

11. In the Promiscuous Mode list, select Accept and click **OK**.

    You return to the Properties dialog box.

12. Repeat steps 5 through 11 for each in-path virtual NIC.

To power on the VM

1. Right-click the VM you created, choose Power and choose Power On.

2. Select the Console tab.

3. Click the dark screen.

    Virtual Steelhead starts up, and the login prompt appears.

Figure 2-12. Log In to Virtual Steelhead

Tip: To release the cursor from the console, press Ctrl+Alt.

4. Log in to Virtual Steelhead as an administrator.

    The default administrator login is **admin** and the default password is **password**.

For information on completing the Virtual Steelhead configuration, see Chapter 4, "Configuring Virtual Steelhead."

Note: To monitor Virtual Steelhead resource use, you can configure alarms in vCenter. For example, you can configure an alarm when a VM experiences an unusually long wait time for CPU or other resources. For more information, see vCenter documentation from VMware.

CHAPTER 3 Setting Up Virtual Steelhead on Hyper-V

This chapter describes how to install and configure Virtual Steelhead on Hyper-V. It includes the following sections:

- "Basic Steps for Installing and Configuring Virtual Steelhead" on page 35
- "Obtaining the Virtual Steelhead Software Package" on page 35
- "Completing the Preconfiguration Checklist" on page 36
- "Installing Virtual Steelhead" on page 36

Basic Steps for Installing and Configuring Virtual Steelhead

This section provides an overview of the basic steps to install and configure Virtual Steelhead on Hyper-V, followed by detailed procedures.

Task	Reference
1. Confirm that Hyper-V is provisioned to run the Virtual Steelhead model. Check to make sure the resources are available before choosing your model.	"Virtual Steelhead Platform Models" on page 15
2. Obtain the Virtual Steelhead package from Riverbed Support.	"Obtaining the Virtual Steelhead Software Package" on page 35
3. Gather network settings for the configuration wizard.	"Completing the Preconfiguration Checklist" on page 36
4. Install and configure the Virtual Steelhead image.	"Installing Virtual Steelhead" on page 36
5. Power on the VM, start the Virtual Steelhead, and log in.	

Obtaining the Virtual Steelhead Software Package

The Hyper-V Virtual Steelhead package is a zip file containing the management virtual hard disk (VHD) and an install script. To download the zip package from the Riverbed Support Web site, go to https://support.riverbed.com. Access to software downloads requires registration.

During installation, you will unzip the package and run the RIVERBED_INSTALL.ps1 script from Windows Powershell. To run the script, you might need to configure the security policy to "Unrestricted."

To configure the security policy

1. Right-click the Windows Powershell program and select **Run as administrator**.
2. At the command prompt, enter the command: `Set-ExecutionPolicy Unrestricted`

Completing the Preconfiguration Checklist

This section lists the parameters you specify to complete the initial configuration of Virtual Steelhead. Be prepared to provide values for the network settings listed in the following checklist when prompted by the installation script.

Network Setting	Your Value
InstallLocation (required)	Path to the directory for the VM.
Model (required)	The hardware model to be configured. Choosing the model causes the installation to allocate the correct disk sizes, memory, and CPU cores.
VHDLocation (optional)	The default is the selected directory. The script looks for the mgmt VHD image at this location.
VMName (optional)	The default is Riverbed Steelhead.
ComputerName (optional)	The default is localhost. If you are installing to a remote computer, enter the name of that computer.
NumInpaths (optional)	The default is 1. Enter the number of in-path pairs to create.
SegstoreSize (optional)	The default is the allocated disk size for your model. Enter a value in bytes (B) or gigabytes (GB) to override the allocated size.
PowerOn (optional)	Include this setting if you want the Virtual Steelhead to start up after the install is complete.
PrimaryNetwork (optional)	Enter the name of the vSwitch to connect the primary NIC to.
AuxNetwork (optional)	Enter the name of the vSwitch to connect the auxiliary NIC to.
{WL}an{01234}_0Network (optional)	Enter the name of the vSwitch to connect the named network interface to.

Installing Virtual Steelhead

Follow these steps to install Virtual Steelhead on Hyper-V.

Note: If the install script in Step 6 displays a message about insufficient disk space, try using the parameter `-SegstoreSize <size> GB`. Set the size to an appropriate value. The management disk uses 38 GB in addition to this allocation.

To install Virtual Steelhead from the Hyper-V Manager

1. To obtain the VM package, and download it locally, go to https://support.riverbed.com.
2. Extract the zip file into the directory you want to use.
3. On your Windows desktop, open the Hyper-V Manager.

Figure 3-1. Hyper-V Manager Window

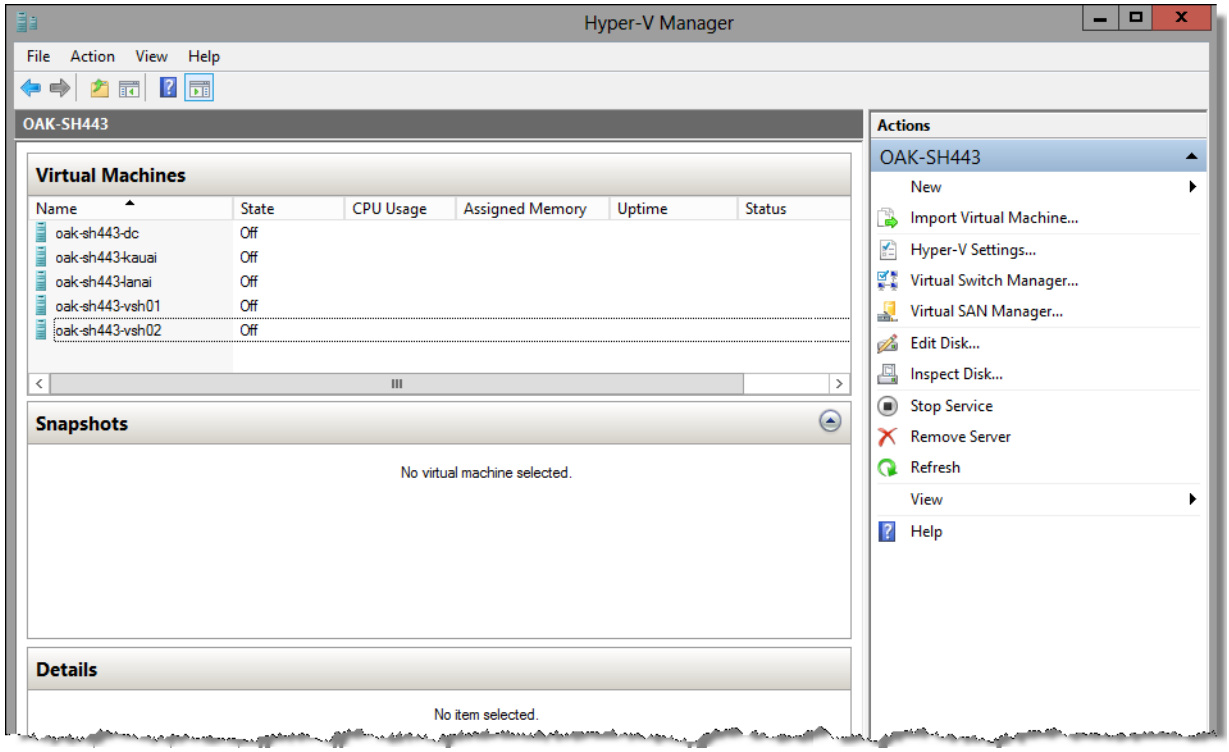

4. Open the Virtual Switch Manager in the right-hand Actions pane.
5. Use the Virtual Switch Manager to create a virtual switch for each Virtual Steelhead interface (for example, primary, aux, lan0_0, and wan0_0).

Figure 3-2. Virtual Switch Manager Window

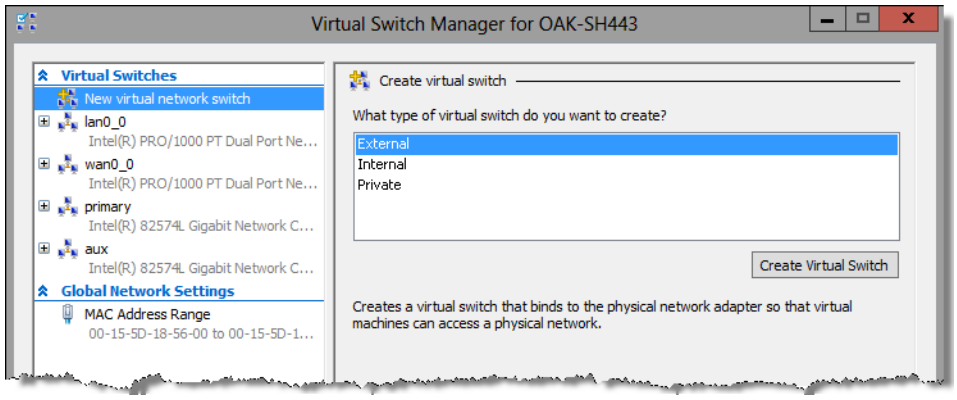

6. Open Windows Powershell.

7. Run the install script.

   You can enter all the script parameters as part of the run command. If you do not enter any parameters, you are prompted for the two required parameters in Steps 7 and 8.

8. Enter the install location (required).

9. Enter the Steelhead appliance hardware model (required).

   The message “Creating new VM” appears. VM creation can take 30 or more minutes to complete.

10. After the VM creation is complete, check all the VM settings in the Hyper-V Manager to verify they are correct.

11. Connect each virtual switch interface to the corresponding virtual switch.

Figure 3-3. VM Settings Window

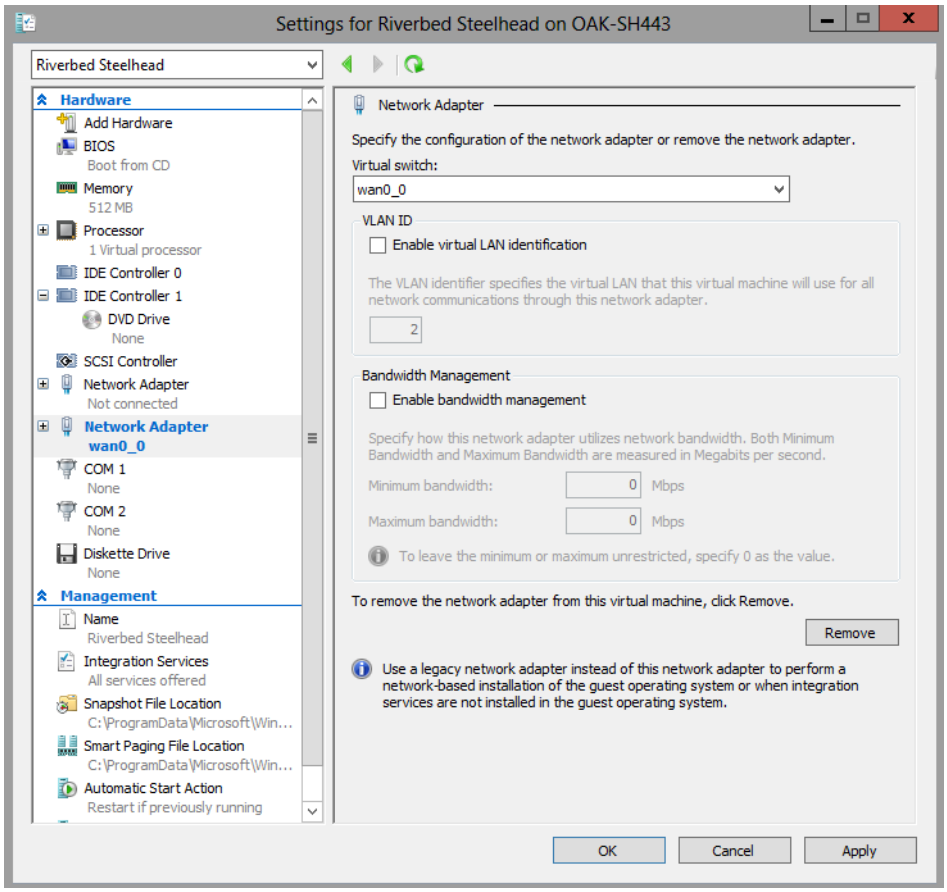

12. Select **Start**, in the right-hand Riverbed Steelhead pane, to power on the VM.

Figure 3-4. VM Options Pane

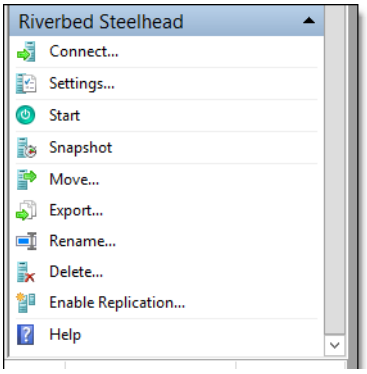

13. Select **Connect** to connect to the terminal window.

14. Log in to Virtual Steelhead as an administrator.

The default administrator login is **admin** and the default password is **password**.

After powering on the VM, if you see messages about missing interfaces or disks, check these troubleshooting tips:

- If there are missing interfaces on the Virtual Steelhead, check the VM settings and verify that you are using synthetic NICs (not legacy), and that the cards are connected.
- If RiOS logs messages about missing disks, ensure that the RiOS data store disk is present and is in slot 1 of controller 0.

For information on completing the Virtual Steelhead configuration, see Chapter 4, "Configuring Virtual Steelhead."

Note: After you deploy Virtual Steelhead, set the reserve weight for CPU to 100 and the memory weight to High.

Manual Installation on Hyper-V

This section describes how to manually install Virtual Steelhead from the Hyper-V Manager.

Note: Before you begin, see "Installing Virtual Steelhead" on page 36 to create and connect the virtual interfaces and switches.

To manually deploy Virtual Steelhead from the Hyper-V Manager

1. Create a new VM.

You need the correct amount of memory and CPU for the hardware model.

2. Remove the CD drive.

3. Create a fixed-size disk for the management VHD of the correct size for the model.
   You can perform this step from the Hyper-V Manager, or you can use the Convert-VHD script.
4. Add the management VHD as the disk in controller 0 slot 0.
5. Create a fixed-size disk for the RiOS data store of the correct size for the model.
6. Add this disk to controller 0 slot 1.
7. Create a synthetic NIC for primary and auxiliary, and two for each in-path pair you want.

CHAPTER 4 Configuring Virtual Steelhead

This chapter describes how to configure Virtual Steelhead after deploying it on a hypervisor. It includes the following sections:

- "Basic Steps for Configuring Virtual Steelhead" on page 41
- "Completing the Initial Configuration" on page 42
- "Logging In to the Virtual Steelhead Management Console" on page 45
- "Purchasing the Token and Receiving the Licenses" on page 46
- "Managing Licenses and Model Upgrades" on page 47
- "Upgrading RiOS to Version 8.6" on page 54
- "Rebooting and Shutting Down Virtual Steelhead" on page 55
- "Verifying Your Configuration" on page 56

The information in this chapter applies to both Hyper-V and ESX/ESXi hypervisors, except where otherwise noted.

Basic Steps for Configuring Virtual Steelhead

This section provides an overview of the basic steps to configure Virtual Steelhead, followed by detailed procedures.

Task	Reference
1. Complete the initial configuration.	"Completing the Initial Configuration" on page 42
2. Exit the configuration wizard.	
3. Purchase a token from Riverbed Sales.	"Purchasing the Token and Receiving the Licenses" on page 46
4. Go to Configure > Maintenance > Licenses and enter the token, which generates a license request string.	"Managing Licenses and Model Upgrades" on page 47
5. Go to the Riverbed Licensing Portal at https://licensing.riverbed.com and enter the license request string to generate your licenses.	"Activating the Token and Installing the Licenses" on page 48
6. Add the licenses.	"Activating the Token and Installing the Licenses" on page 48

Task	Reference
7. Select the new Virtual Steelhead model in the form below the license table and submit the form.	"Model Upgrade Overview" on page 49
8. Save the configuration and restart.	"Rebooting and Shutting Down Virtual Steelhead" on page 55
9. Power on the VM and log in to Virtual Steelhead.	
10. Verify your configuration—the Management Console appears, and Virtual Steelhead is healthy.	"Verifying Your Configuration" on page 56
11. Refer to the Riverbed product documentation for more information. **Note:** In the Riverbed product documentation, the term "Steelhead appliance" refers to the physical Steelhead appliance as well as Virtual Steelhead unless otherwise stated.	*Getting Started Guide* *Steelhead Appliance Management Console User's Guide* *Riverbed Command-Line Interface Reference Manual* *Steelhead Appliance Deployment Guide*

Completing the Initial Configuration

This section describes how to complete the initial configuration of Virtual Steelhead.

To configure Virtual Steelhead

1. After you log in to Virtual Steelhead as administrator, the system prompts you to start the configuration wizard.

   Enter **yes** at the system prompt. For example:

   ```
   Configuration wizard.
   Do you want to use the wizard for initial configuration? yes
   ```

Tip: Press Enter to accept the default value. If you mistakenly answer **no**, you can start the configuration wizard by specifying **configuration jump-start** at the system prompt.

Tip: Press **?** for help. Press Ctrl-B to go back to the previous step.

2. Complete the configuration wizard steps on client-side Virtual Steelhead as described in the following table.

Wizard Prompt	Description	Example
Step 1: Host name?	Enter the host name for Virtual Steelhead.	Step 1: `hostname? amnesiac`
Step 2: Use DHCP on the primary interface?	You are given the option to enable the DHCP to automatically assign an IP address to the primary interface for Virtual Steelhead. Riverbed recommends that you do not set DHCP. The default value is `no`.	Step 2: `Use DHCP? no`
Step 3: Primary IP address?	Enter the IP address for Virtual Steelhead.	Step 3: `Primary IP address? 10.10.10.6`
Step 4: Netmask?	Enter the netmask address.	Step 4: `Netmask? 255.255.0.0`
Step 5: Default gateway?	Enter the default gateway for the Steelhead appliance.	Step 5: `Default gateway? 10.0.0.1`
Step 6: Primary DNS server?	Enter the primary DNS server IP address.	Step 6: `Primary DNS server? 10.0.0.2`
Step 7: Domain name?	Enter the domain name for the network where Virtual Steelhead is to reside. If you set a domain name, you can enter hostnames in the system without the domain name.	Step 7: `Domain name? example.com`
Step 8: Admin password?	Riverbed strongly recommends that you change the default administrator password at this time. The password must be a minimum of 6 characters. The default administrator password is `password`.	Step 8: `Admin password? xxxyyy`
Step 9: SMTP server?	Enter the SMTP server. External DNS and external access for SMTP traffic is required for email notification of events and failures to function. **Important:** Make sure that you provide a valid SMTP server to ensure email notifications for events and failures.	Step 9: `SMTP server? natoma`
Step 10: Notification email address?	Enter a valid email address to which notification of events and failures are to be sent.	Step 10: `Notification email address? example@xample.com`
Step 11: Set the primary interface speed?	Enter the speed on the primary interface (that is,Virtual Steelhead). Make sure this value matches the settings on your router or switch. The default value is `auto` and Riverbed recommends this setting for Virtual Steelhead.	Step 11: `Set the primary interface speed? [auto] auto`

Wizard Prompt	Description	Example
Step 12: Set the primary interface duplex?	Enter the duplex mode on the primary interface. Make sure this value matches the settings on your router or switch. The default value is `auto` and Riverbed recommends this setting for Virtual Steelhead.	Step 12: `Set the primary interface duplex? [auto] auto`
Step 13: Would you like to activate the in-path configuration?	Enter `yes` at the system prompt to configure in-path support. An in-path configuration is a configuration in which the Steelhead appliance is in the direct path of the client and server. For detailed information about in-path configurations, see the *Steelhead Appliance Deployment Guide*.	Step 13: `Would you like to activate the in-path configuration? yes`
Step 14: In-Path IP address?	Enter the in-path IP address for Virtual Steelhead.	Step 14: `In-Path IP address? 10.11.11.6`
Step 15: In-Path Netmask?	Enter the in-path netmask address.	Step 15: `In-Path Netmask? 255.255.0.0`
Step 16: In-Path Default gateway?	Enter the in-path default gateway (the WAN gateway).	Step 16: `In-Path Default Gateway? 10.11.11.16`
Step 17: Set the in-path: LAN interface speed?	Accept the default value of `auto`. **Note:** If you have configured direct path with ESXi, you can enter the speed, matching the settings on your router or switch.	Step 17: `Set the in-path: LAN interface speed? [auto] auto`
Step 18: Set the in-path: LAN interface duplex?	Accept the default value of `auto`. **Note:** If you have configured direct path with ESXi, you can enter the speed, matching the settings on your router or switch.	Step 18: `Set the in-path: LAN interface duplex? [auto] auto`
Step 19: Set the in-path: WAN interface speed?	Accept the default value of `auto`. **Note:** If you have configured direct path with ESXi, you can enter the speed, matching the settings on your router or switch.	Step 19: `Set the in-path: WAN interface speed? [auto] auto`
Step 20: Set the in-path: WAN interface duplex?	Accept the default value of `auto`. **Note:** If you have configured direct path with ESXi, you can enter the speed, matching the settings on your router or switch.	Step 20: `Set the in-path: WAN interface duplex? [auto] auto`

3. The system confirms your settings:

```
You have entered the following information:
1. Hostname: amnesiac
2. Use DHCP on primary interface: no
3. Primary IP address: 10.10.10.6
4. Netmask: 255.255.0.0
5. Default gateway: 10.0.0.1
6. Primary DNS server: 10.0.0.2
```

```
   7. Domain name: example.com
   8. Admin password: xxxyyy
   9. SMTP server: natoma
   10. Notification email address: example@example.com
   11. Set the primary interface speed: auto
   12. Set the primary interface duplex: auto
   13. Would you like to activate the in-path configuration: yes
   14. In-Path IP address: 10.11.11.6
   15. In-Path Netmask: 255.255.0.0
   16. In-Path Default gateway: 10.11.11.16
   17. Set the in-path:LAN interface speed: auto
   18. Set the in-path:LAN interface duplex: auto
   19. Set the in-path:WAN interface speed: auto
   20. Set the in-path:WAN interface duplex: auto
To change an answer, enter the step number to return to.
Otherwise hit <enter> to save changes and exit.
Choice:
```

The Virtual Steelhead configuration wizard automatically saves your configuration settings.

4. To log out of the system, enter the following command at the system prompt:

```
amnesiac> exit
```

Logging In to the Virtual Steelhead Management Console

This section describes how to log in to the Virtual Steelhead Management Console. The Management Console provides a Web browser interface that facilitates managing Virtual Steelhead.

You can connect to Virtual Steelhead through any supported Web browser. To connect, you must know the host, domain, and administrator password that you assigned during the initial setup.

Note: Cookies and JavaScript must be enabled in your browser.

To log in to Virtual Steelhead

1. Enter the URL for Virtual Steelhead in the location box of your browser:

   protocol`://`*host.domain*

   protocol is `http` or `https`. The secure HTTPS uses the SSL protocol to ensure a secure environment. If you use HTTPS to connect, you are prompted to inspect and verify the SSL key.

   host is the IP address or hostname you assigned to Virtual Steelhead during the initial configuration. If your DNS server maps the IP address to a name, you can specify the DNS name.

Note: Alternatively, you can specify the IP address instead of the host and domain.

The Management Console Login page appears.

Figure 4-1. Login Page

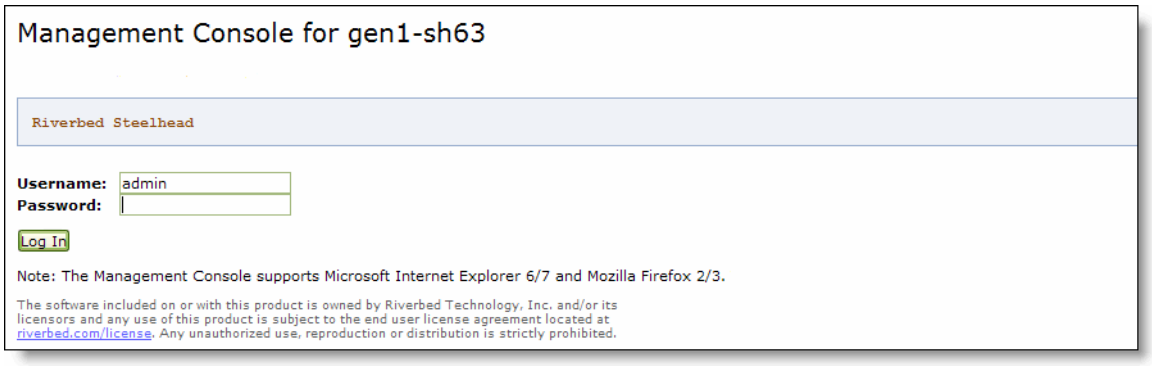

2. In the Username text box, type the user login: admin or monitor. The default login is admin.

   Users with administrator (admin) privileges can configure and administer the Steelhead appliance. Users with (monitor) privileges can view connected Steelhead appliances and reports. A monitor user cannot make configuration changes, modify private keys, view logs, or manage cryptographic modules in the system.

3. In the Password text box, type the password you assigned in the configuration wizard.

4. Click **Log In** to display the Home page.

   The Home page summarizes the current status of Virtual Steelhead.

Purchasing the Token and Receiving the Licenses

Before you can add licenses to Virtual Steelhead, you must purchase a token from Riverbed. The token is associated with a model number that is assigned to Virtual Steelhead during licensing.

To view your purchased tokens, log in to your account at https://support.riverbed.com.

After you receive a token, you are ready to install the licenses.

Starting in RiOS v8.0.2, you can delete an instance of Virtual Steelhead and deploy a new instance with the same token. When you reuse a token, the system indicates the reuse and recommends removing the old instances. Each new instance using the same token invalidates the previous instance. Requests from an old instance to get licenses result in an error message. The serial number remains the same for the new instance and licenses.

You can reuse a token up to 5 times. After reusing a token 5 times, contact Riverbed Support for a new token.

Managing Licenses and Model Upgrades

This section describes how to install, update, and remove a license. It also describes how to use flexible licensing to manage model configurations and upgrades. It includes the following topics:

- "Flexible Licensing Overview" on page 47
- "Activating the Token and Installing the Licenses" on page 48
- "Upgrading a Model That Requires No Additional Virtual Hardware" on page 50
- "Upgrading a Model That Requires Additional Virtual Hardware" on page 51
- "Removing a License" on page 54

You perform all license management and Virtual Steelhead model upgrades in the Configure > Maintenance > Licenses page.

Flexible Licensing Overview

RiOS provides a flexible way to manage Virtual Steelhead licenses, model configurations, and upgrades. Rather than performing an incremental model upgrade or replacing a virtual appliance with a new OVA image, RiOS provides *specification licenses* that configure specific performance characteristics of Virtual Steelhead A specification license points to a specific, validated model and includes the required licenses and the virtual hardware specification.

Licensing and Model Upgrade

Version 8.5 and later include improvements in the licensing and model upgrade areas.

Virtual Steelhead can run with a specification license for a higher model. For example, if there is only VCX755H hardware, but there is a specification license to run a VCX1555H, then the appliance can operate as the VCX755H model until more hardware is available.

If you are licensed for a higher model than you have hardware for, you can shut down Virtual Steelhead, add the hardware, and power it back on. When Virtual Steelhead comes up again it automatically upgrades to the highest runnable model.

In addition, when you add a specification license for the first time (or whenever RiOS is not running), Virtual Steelhead automatically upgrades to the highest runnable model based on the available hardware and license. No reboot is required.

By activating a specification license on Virtual Steelhead, you can transform its capabilities to meet performance characteristics for any model within a platform family.

Virtual Steelhead xx50 Model	License
V150M	MSPECV150M
V250L	MSPECV250L
V250M	MSPECV250M
V250H	MSPECV250H
V550M	MSPECV550M
V550H	MSPECV550H

Virtual Steelhead xx50 Model	License
V1050L	MSPECV1050L
V1050M	MSPECV1050M
V1050H	MSPECV1050H
V2050L	MSPECV2050L
V2050M	MSPECV2050M
V2050H	MSPECV2050H

Virtual Steelhead xx55 Model	License
VCX555M	MSPECVCX555M
VCX555H	MSPECVCX555H
VCX755L	MSPECVCX755L
VCX755M	MSPECVCX755M
VCX755H	MSPECVCX755H
VCX1555L	MSPECVCX1555L
VCX1555M	MSPECVCX1555M
VCX1555H	MSPECVCX1555H
VCX5055M	MSPECVCX5055M
VCX5055H	MSPECVCX5055H
VCX7055L	MSPECVCX7055L
VCX7055M	MSPECVCX7055M

Model downgrades are not supported, even for model evaluations. If you purchase a V250L and want to evaluate a V550M, you can install an evaluation license. When the trial period for the V550M expires, you cannot downgrade Virtual Steelhead back to the V250L. For this reason, you must create a new, separate VM for the evaluation. If you decide to upgrade to the new model, you purchase the full license for the V550M (in this example) and upgrade the Steelhead appliance. This requires a new token and hardware specification for the new model, and you must restart Virtual Steelhead with a clean data store.

You cannot upgrade between the Virtual Steelhead xx50 models and the VCX xx55 models.

Activating the Token and Installing the Licenses

This section describes how to activate a token, receive the license, and install the license on Virtual Steelhead. The procedures in this section describe both autolicensing and manual licensing.

To activate the token and install a license

1. Restart with a clean RiOS data store.

   For details, see "Rebooting and Shutting Down Virtual Steelhead" on page 55.

2. Log in to Virtual Steelhead and display the Management Console.

3. Choose Configure > Maintenance > Licenses to display the Licenses page.

4. Under License Request Token, type the token number and click **Generate License Request Key**.

   The console does not support pasting text directly into the field.

   Starting in RiOS v8.0.2, Virtual Steelhead performs autolicensing by default. With autolicensing, once you enter the token and click **Generate License Request Key**, Virtual Steelhead contacts the Riverbed licensing server and automatically downloads and installs the licensing keys.

   For more information about autolicensing, including commands to enable or disable autolicensing and to configure a licensing server, see the *Riverbed Command-Line Interface Reference Manual*.

   If you disable autolicensing or if Virtual Steelhead cannot connect to the licensing server, you can manually configure your licenses. With manual licensing, you enter the token and click **Generate License Request Key** and RiOS displays a license request key. Continue with the next step.

5. After you have obtained the license request key, go to the Riverbed Licensing Portal at https://licensing.riverbed.com (unregistered users) or to the Licenses tab on the Riverbed Support site at http://support.riverbed.com/ to generate your license keys. The license keys include the VBASE license as well as any other licenses needed for the Virtual Steelhead model.

   The Licensing Portal is a public site; the Riverbed Support site requires registration.

   After your licenses are generated, they appear online and are also emailed to you for reference.

6. Return to the Configure > Maintenance > Licenses page and click **Add a New License**.

7. Copy and paste the license key into the text box. Separate multiple license keys with a space, Tab, or Enter.

8. Click **Add.**

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

10. Choose Configure > Maintenance > Services and click **Restart** to restart the optimization service.

Model Upgrade Overview

You can use a hardware specification to upgrade a model. Some model upgrades require additional virtual hardware. When the appliance has the required virtual hardware, activating the hardware specification upgrades the appliance to the new model number. When the existing hardware is not adequate, the message `hardware required` appears after the hardware specification description.

For details on Virtual Steelhead model requirements, see "Virtual Steelhead Platform Models" on page 15.

Important: Upgrading Virtual Steelhead from a V1050L to a 1050H or from a V1050M to a V1050H clears all data in the RiOS data store and requires a reboot.

Model Upgrade and Flexible RiOS Data Store

With Virtual Steelhead v8.5 and later, you can configure the size of the RiOS data store disk for Virtual Steelhead on VCX models. This configuration is still fixed for other Steelhead models.

If you upgrade to a version later than v8.0.3, your RiOS data store configuration remains intact. However, due to changes in the RiOS data store disk layout, if you modify the disk size either manually or during a model upgrade, and later downgrade to the earlier image, data store corruption can result.

If you see errors related to the data store, you must clean the RiOS data store. If your disk is smaller than the Steelhead appliance model is expecting based on the values required in the earlier version, increase the disk size to the expected size, and then enter the command **restart clean**.

To avoid this situation, upgrade the Steelhead appliance model prior to upgrading to Virtual Steelhead v8.6 or later.

Note: When you upgrade Virtual Steelhead to a model that can use more of the available RiOS data store disk, the data store is automatically cleared.

Next Steps

After installing a license, the next steps to complete a model upgrade depend on whether the upgrade requires additional virtual hardware:

- If you do not need to add virtual hardware to Virtual Steelhead, see "Upgrading a Model That Requires No Additional Virtual Hardware" on page 50.
- If you are upgrading Virtual Steelhead to a model that requires new virtual hardware components, see "Upgrading a Model That Requires Additional Virtual Hardware" on page 51.

Upgrading a Model That Requires No Additional Virtual Hardware

This section describes how to activate a hardware specification that does not require additional virtual hardware on Virtual Steelhead.

Upgrades to models within a platform family (V250, V550, V1050, V2050) require only a new license, because the existing virtual hardware is sufficient. For example, an upgrade from a model V1050L to a model V1050M is a license-only upgrade. Downgrades are not permitted.

To activate a hardware specification

1. Stop the optimization service.
2. Choose Configure > Maintenance > Licenses to display the Licenses page.

The hardware specifications appear at the bottom of the page. The hardware specification description includes the potential bandwidth and connection counts. The current specification appears in bold.

You might see specifications listed but grayed out. After the license or required hardware for the model is installed, these specification become available.

3. Select the hardware specification that you want to activate.
4. Click **Apply**.
5. Click **Restart** to restart the optimization service.

Upgrading a Model That Requires Additional Virtual Hardware

This section describes how to activate a hardware specification that requires additional virtual hardware on Virtual Steelhead.

Upgrades to models from one platform family to another require additional virtual hardware. For example, to upgrade from a V250L to a V550M requires another CPU, more RAM, and additional RiOS data store disk space.

Hypervisor Reservations and Overhead

When you reserve CPU and other resources, reserve the full amount required for the Virtual Steelhead model. In addition, verify that additional unclaimed resources are available. Due to hypervisor overhead, VMs can exceed their configured reservation.

The overhead calculations for the ESX/ESXi and Hyper-V hypervisors differ:

- For ESX/ESXi, reserve the memory and CPU cycles needed for the Virtual Steelhead model and verify that the host has resources to accommodate the 5% VMware overhead.
- For Hyper-V, reserve the memory and CPU percentage needed for the Virtual Steelhead model and verify that the host has 1.5 GB and 15% CPU remaining, for overhead purposes.

Note: After you deploy Virtual Steelhead on Hyper-V, set the reserve weight for CPU to 100 and the memory weight to High.

To upgrade a model and add virtual hardware, perform the following steps after installing the license.

To upgrade a model

1. In the Management Console, choose Configure > Maintenance > Licenses.
2. Click **Stop** to stop the optimization service, and log out of Virtual Steelhead.
3. Right-click the name of the VM that you want to upgrade, choose Power, and choose Power Off.
4. In the vSphere Client, right-click the name of the VM that you want to upgrade and select Edit Settings.
5. Check the memory, CPU, and hard disk settings. Change any setting necessary to accommodate the model upgrade.

For more information on changing memory, CPU, and hard disk settings, see "To add memory" on page 53, "To increase the size of a hard disk" on page 53, and "To add CPU capacity" on page 53.

Figure 4-2. VM Properties Page - Hardware Tab

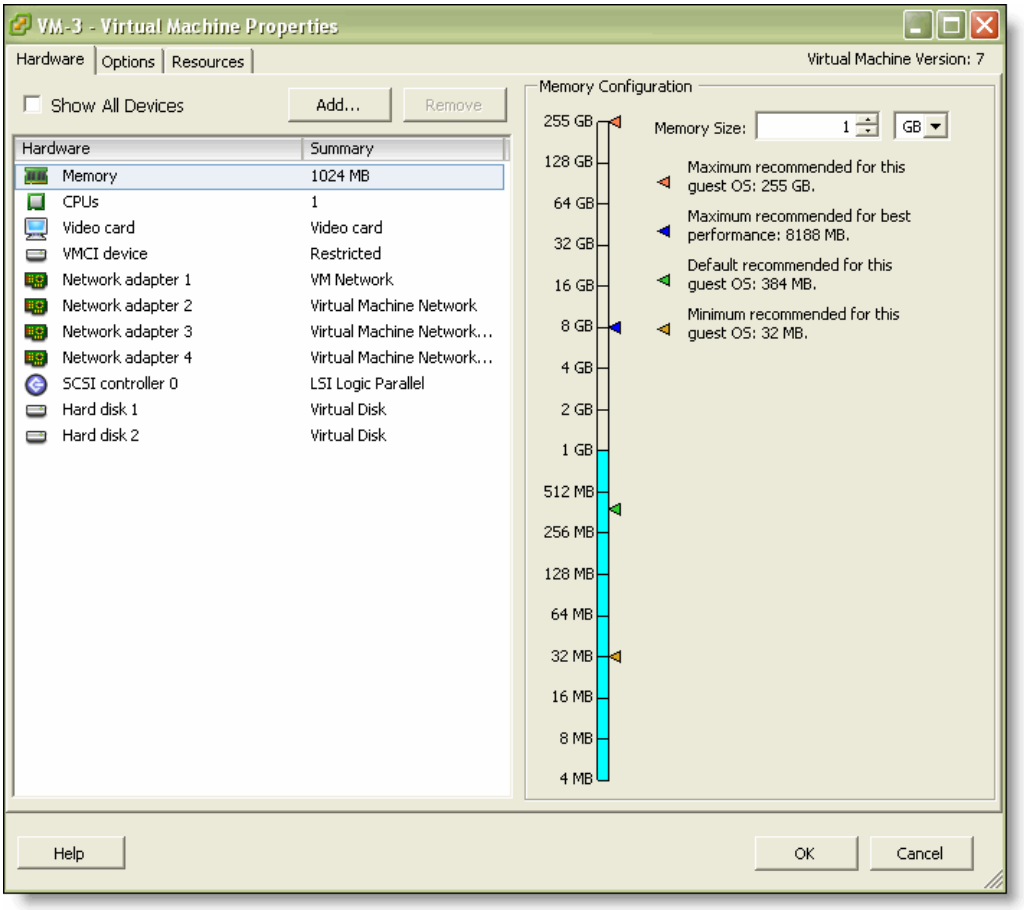

6. Right-click Virtual Steelhead, choose **Power**, and select **Power On**.

7. Log in to Virtual Steelhead and display the Management Console.

8. Choose Configure > Maintenance > Licenses to display the Licenses page.

   The bottom of the screen lists the available hardware specifications. The current specification appears in bold. The hardware specification description includes the potential bandwidth and connection counts. Hardware specifications that will be available after the license or required virtual hardware have been installed are included in the list but are dimmed.

9. Select the hardware specification that you want to activate.

   If a hardware specification requires a reboot after activation, you see the message `activation reboots appliance`.

10. Click **Apply**.

    Virtual Steelhead reboots and the optimization service restarts.

When the upgrade is complete, the new model number appears on the Virtual Steelhead banner in the upper-right corner of the screen.

To add memory

1. On the Hardware tab, click **Memory**.

   Reserve the RAM needed by the Virtual Steelhead model.

2. Under Memory Configuration, increase the memory by clicking a colored triangle (on the slider or in the legend), using the slider control, or selecting a number from the drop-down list.

   Only multiples of 4 MB are valid for memory settings. If you manually enter a value that is not a multiple of 4 MB, a warning message appears.

3. Click **OK**.

To increase the size of a hard disk

1. On the Hardware tab, select Hard Disk 2.

2. In the Disk Provisioning section, specify the disk size, in gigabytes.

3. Click **OK**.

To add a hard disk

1. On the Hardware tab, click Add.

2. Select Hard Disk.

3. Specify the disk size, in gigabytes.

4. Click **OK**.

To add CPU capacity

1. On the Hardware tab, click **CPUs**.

2. Increase the number of virtual CPUs to two or four, depending on the model upgrade.

   You can configure how the virtual CPUs are assigned in terms of sockets and cores. For example, you can configure a VM with four virtual CPUs in the following ways:

   – Four sockets with one core per socket

   – Two sockets with two cores per socket

   – One socket with four cores per socket

3. Click **OK**.

4. Select the **Resources** tab.

5. Use the slider control to reserve the number of clock cycles (in terms of CPU MHz).

   For example, for a model V550M requiring 2 two virtual CPUs running on a quad-core Xeon-based system running at 2.6 GHz on a ESX/ESXi host, reserve 2 virtual CPUs and 2 * 2.6 GHz CPU cycles.

6. Click **OK**.

Downgrade Limitation

After using flexible licensing to upgrade, you cannot return Virtual Steelhead to a lower model.

Removing a License

You can remove a Virtual Steelhead license.

To remove a license

1. Choose Configure > Maintenance > Licenses to display the Licenses page.
2. Select the license you want to delete.
3. Click **Remove Selected**.
4. Click **Save** to save your settings permanently.

Upgrading RiOS to Version 8.6

RiOS v8.6 is backward compatible with previous RiOS versions. However, to obtain the full benefits of the new features in RiOS 8.6, you must upgrade the client-side and server-side Steelhead appliances on any given WAN link. After you have upgraded all appliances, all the benefits of the RiOS v8.6 features and enhancements are available.

If you mix RiOS software versions in your network, the releases might support different optimization features and you cannot take full advantage of the features that are not part of the older software versions.

Upgrading RiOS Software

Follow these steps to upgrade your RiOS software; you must already be familiar with the Steelhead appliance, the CLI, and the Management Console.

To upgrade RiOS software

1. Download the software image from the Riverbed Support site to a location such as your desktop.
2. Log in to the Management Console using the Administrator account (admin).
3. Go to the Configure > Maintenance > Software Upgrade page and choose one of the following options:
   - **From URL** - Type the URL that points to the software image in the text box.
   - **From Local File** - Browse your file system and select the software image.
   - **Schedule Upgrade for Later** - Select this option to schedule an upgrade for a later time. Type the date and time in the Date and Time text boxes using the formats YYYY/MM/DD and HH:MM:SS.
   - Click **Install**.
4. Reboot the appliance.

The software image is large, and uploading the image takes a few minutes.

After the upload completes, the system reminds you to reboot the system to switch to the new version of the software. After reboot, the software version displays on the Home page of the Management Console.

Downgrading the Software

If you are downgrading to a previous version of the Steelhead appliance software, you must downgrade to a version of the software that has previously run on Virtual Steelhead.

Rebooting and Shutting Down Virtual Steelhead

You can reboot or shut down the system in the Maintenance > Reboot/Shutdown page.

Rebooting the system disrupts existing network connections that are currently proxied through it. Rebooting can take a few minutes.

When you shut down the system, connections are broken and optimization ceases. Shutdown can take several minutes.

To reboot or shut down the system

1. Choose Configure > Maintenance > Reboot/Shutdown to display the Reboot/Shutdown page.

Figure 4-3. Reboot/Shutdown Page

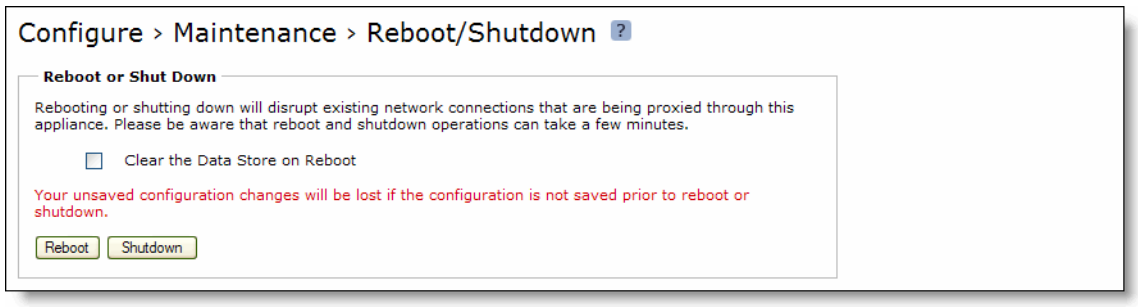

2. Click **Reboot**.

   After you click **Reboot**, you are logged out of the system and it reboots.

3. Click **Shutdown** to shut down the system.

   After you click **Shutdown**, the system turns off. To restart the system, you must manually turn on the Steelhead appliance.

Important: To remove data from the RiOS data store, select **Clear the Data Store on Reboot**.

Verifying Your Connections

This section describes how to verify that you have properly connected Virtual Steelhead.

To verify your connections

1. From a remote host, connect to the CLI. At the system prompt, enter one of the following commands:

   ```
   ssh admin@host.domain
   ```

   ```
   ssh admin@ipaddress
   ```

2. When you are prompted for a password, specify the administrator password you set when you ran the configuration wizard.

3. At the system prompt, enter ping commands to verify the connections. For example:

   ```
   ping -I <primary-IP-address> <primary-default-gateway>
   ```

   —or, to verify in-path connectivity—

   ```
   ping -I <inpath interface> <IP address>
   ```

Verifying Your Configuration

Perform the following tasks to verify that you have properly configured Virtual Steelhead.

To verify optimization in an in-path configuration

1. Navigate to the Reports > Optimization > Bandwidth Optimization page in the Management Console to verify optimization.

2. Map a remote drive on a client machine.

3. Drag and drop a 1 MB file from the client to the remote server.

   Ensure that the server is located across the WAN.

4. Drag and drop the 1 MB file again.

   Performance improves significantly.

Note: For details on managing and configuring Virtual Steelhead, see the *Steelhead Appliance Management Console User's Guide*, the *Steelhead Appliance Deployment Guide*, and the *Riverbed Command-Line Interface Reference Manual*. In the Riverbed product documentation, the term "Steelhead appliance" refers to the physical Steelhead appliance as well as Virtual Steelhead unless otherwise stated.

CHAPTER 5 Using Discovery Agent

This chapter describes how to use the Discovery Agent, an alternate method for deploying Virtual Steelhead. It includes the following sections:

- "Overview of the Discovery Agent" on page 57
- "Discovery Agent Requirements" on page 58
- "Installing the Discovery Agent on a Windows Server" on page 59
- "Installing the Discovery Agent on a Linux Server" on page 60
- "Configuring the Discovery Agent" on page 60
- "Configuring Transparency Modes" on page 63
- "Enabling Optimization Using the Discovery Agent" on page 63

Overview of the Discovery Agent

The Discovery Agent is a software package that you download from the Riverbed Support site and install on the client or server that is optimized.

In a server-side Discovery Agent deployment, when a client Steelhead appliance connects to a server with Discovery Agent installed, the Discovery Agent redirects any auto-discovery probe to a configured Virtual Steelhead from its list. The client Steelhead appliance then discovers and starts peering and optimizing with the server-side Virtual Steelhead. After the auto-discovery process completes, the connection is terminated locally between the Discovery Agent and server-side Steelhead appliance without going over the WAN.

In a client-side Discovery Agent deployment, when a client with Discovery Agent installed connects to a server, the Discovery Agent redirects any TCP connection to a configured Virtual Steelhead from its list. The client-side Virtual Steelhead next sends an auto-discovery probe, discovers the remote Steelhead, and starts peering and optimizing with it.

The Discovery Agent provides the following features:

- **Optimization** - Enables you to intercept (and optimize) inbound and outbound connections.
- **Auto-discovery** - Enables configured Steelheads to automatically find one another. Auto-discovery relieves you of having to manually configure the Steelhead appliances with fixed target rules to find the remote Virtual Steelheads and to optimize traffic through them.

- **Transparency** - Enables the application on the server to continue to send and receive data from the same client IP address (as if there was no Steelhead) so that logging, reporting, or any feature that uses the IP address continues to work the same as before you configured the Steelhead.

Note: In a Discovery Agent deployment, you cannot configure WAN transparency for the connection between Steelheads.

- **Failure detection** - Detects Virtual Steelhead failures and connectivity issues to the Virtual Steelhead so that traffic can be passed through instead of being redirected to the failed Virtual Steelhead.
- **Load balancing** - Redirects all traffic to the selected Steelhead. If there are multiple Steelheads in the group, the Discovery Agent uses the round-robin or priority load-balancing method to select a Steelhead. When the primary Steelhead is unavailable or overloaded, it redirects all new connections to the next Steelhead on the list.

Figure 5-1 shows a Discovery Agent deployment on the server side.

Figure 5-1. Discovery Agent Server Side Deployment

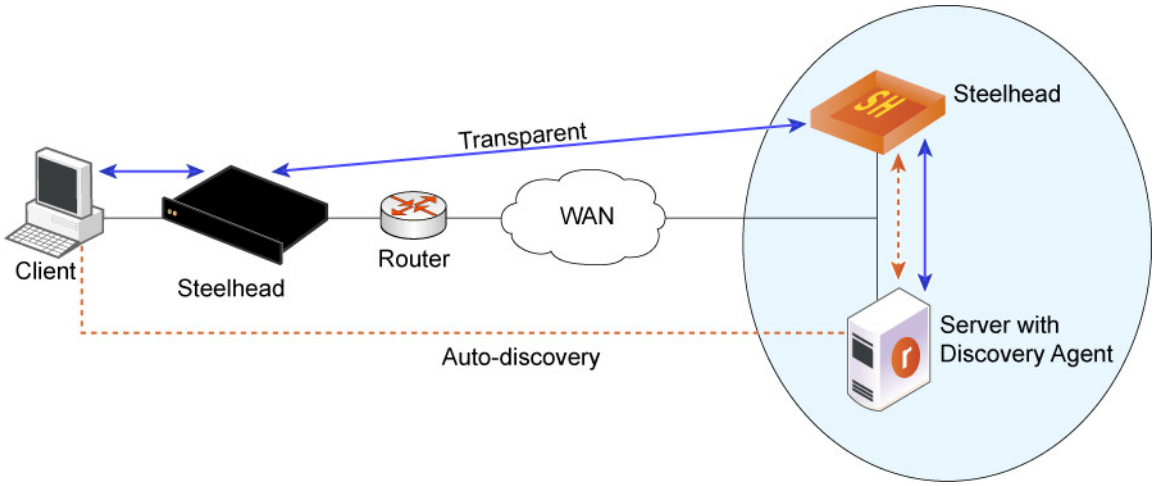

In Figure 5-1, the Discovery Agent enables the client-side Steelhead and the server-side Steelhead to discover each other.

When the client connects to the server, the client-side Steelhead sends an auto-discovery probe to the server. The Discovery Agent redirects the auto-discovery probe to the Virtual Steelhead. The Virtual Steelhead sends an auto-discovery probe response back to the Discovery Agent, which sends it to the client-side Steelhead. After the client-side Steelhead receives the probe response, it starts peering with the Virtual Steelhead to intercept and optimize the connection from the client to the server.

Discovery Agent Requirements

The Discovery Agent requires the following hardware:

- **Disk** - At least 160 MB on Windows and 120 MB on Linux. This space is mainly used to store binary files, configuration files, and log files.
- **RAM** - At least 110 MB for 20000 optimized connections (the current limit).
- **CPU** - Depends on the throughput. For example, the Discovery Agent uses 5-10% of a 2.66 GHz CPU to process 1 Gbps of optimized traffic.

Installing the Discovery Agent on a Windows Server

To install the Discovery Agent on a Windows server, you first download the package from the Riverbed Support Web site.

The Discovery Agent supports the following Windows servers:

- **Windows Server 2003 R2** - 32 bit and 64 bit
- **Windows Server 2008** - 32 bit and 64 bit
- **Windows client Windows 7** - 32 bit and 64 bit

Note: Riverbed does not support the Steelhead Mobile Client and the Discovery Agent on the same Windows computer.

To install the Discovery Agent on a Windows server

1. From the Riverbed Support Web site, click **Software & Documentation**.
2. In the Search text box, type **Discovery Agent** and click the arrow icon.
3. Click the link for the Discovery Agent package you want and save the file.
4. Log in to the Windows server and double-click the executable file to display the Discovery Agent Installation Wizard.
5. Click **Next** to display the Discovery Agent Installation Warning message.

   When you install, uninstall, or upgrade the Discovery Agent on a Windows server, there is a temporary loss of network connectivity. You should save your work and close any Windows program that might be affected by the disruption before you continue.
6. Click **Cancel** to quit the program, or click **Next** to continue with the installation.
7. Read and accept the license agreement and click **Next** to display the Riverbed Discovery Agent Configuration page.
8. Select the **Other** cloud type from the drop-down list.
9. Select the target folder for the installation and click **Install**.
10. Click **Finish** to complete the install process.

    The Discovery Agent starts automatically and the Riverbed icon appears on the system tray. If the icon appears gray, it signifies that the Discovery Agent service is just starting or has failed to start. If the Discovery Agent does not start, reboot the system and verify that the Discovery Agent starts.

Installing the Discovery Agent on a Linux Server

To install the Discovery Agent on a Linux server, you first download the package.

The Discovery Agent supports the follows Linux servers:

- **Centos 5.0, 5.2, 5.3, and 5.4** - 32 bit and 64 bit
- **Linux Ubuntu 8.04, 10.04, and 12.0.4** - 32 bit and 64 bit
- **Linux Fedora (Fedora core 8)** - 32 bit and 64 bit

To install the Discovery Agent on a Linux server

1. From the Riverbed Support Web site, click **Software & Documentation**.
2. In the Search text box, type **Discovery Agent** and click the arrow icon.
3. Click the link for the Discovery Agent package you want and save the file.
4. Copy the downloaded tar file (Discovery Agent package) to the Linux server and log in to the server as the root user.
5. Uncompress the tar file and extract its contents by entering the following command on the Linux command line:

   ```
   tar zxvf <filename>.tar.gz
   ```

6. Follow the steps in the README file to install and configure the Discovery Agent on the Linux server.

Configuring the Discovery Agent

The process for configuring the Discovery Agent differs depending on whether you are installing on Linux or Windows. In both cases, however, you deploy in manual mode.

Configuring the Discovery Agent on a Linux Server

Follow the instructions in the Linux Discovery Agent README file, obtained during download and installation, to configure the Discovery Agent on a Linux server.

Configuring the Discovery Agent on Windows

Follow these steps for a manual deployment on Windows.

To configure the Discovery Agent

1. Log in to your Windows server and double-click the Riverbed Discovery Agent icon in the system tray.
2. Select the Settings tab in the Discovery Agent to display the Settings page.

Figure 5-2. Discovery Agent Settings Page

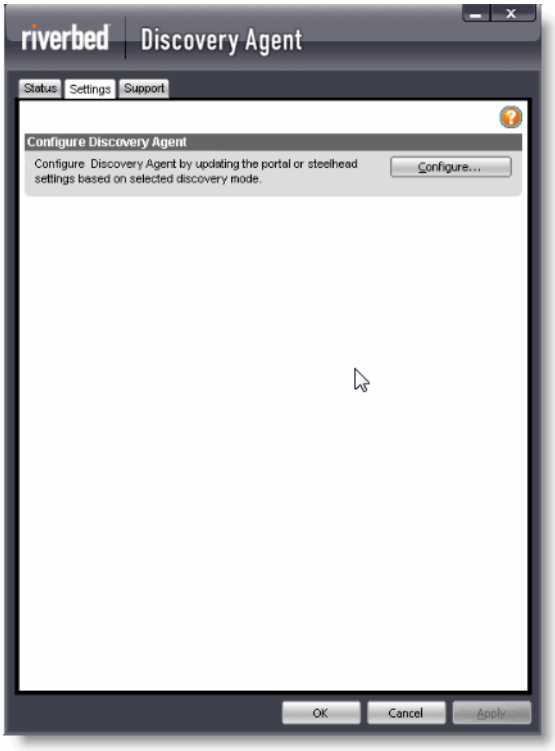

3. Click **Configure** to display the Configure Discovery Agent page.

Figure 5-3. Configure Discovery Agent - Use Riverbed Portal

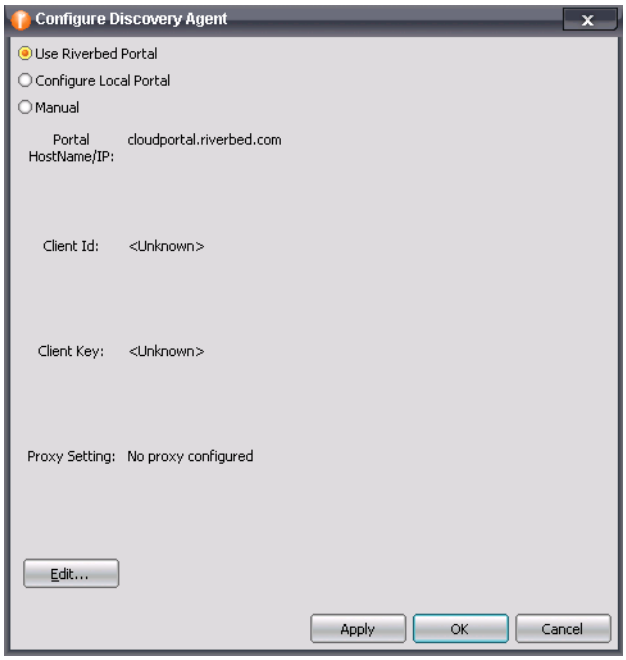

4. Select **Manual** to display the Manual Configuration page.

Figure 5-4. Configure Discovery Agent - Manual

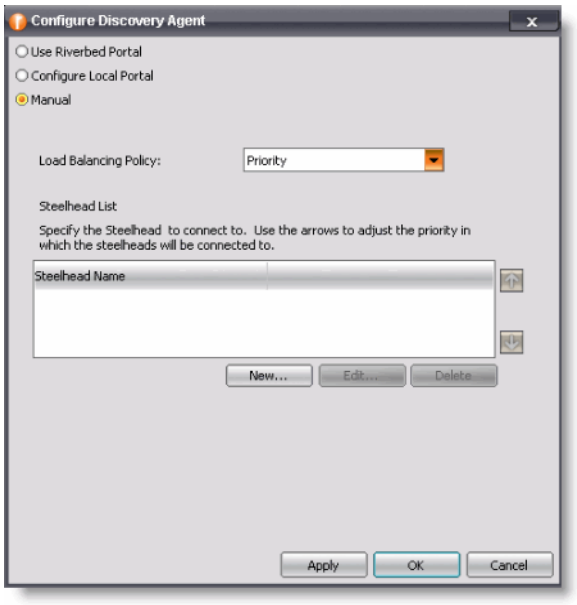

5. Choose one of the following load-balancing policies from the drop-down list:
   - **Priority** - Selects a Virtual Steelhead for load balancing until its connection count exceeds the maximum and then moves on to the next available Virtual Steelhead. When the load of the first Virtual Steelhead decreases below the maximum, it is available again. This is the default mode.
   - **Round Robin** - Selects a Virtual Steelhead and then another (using the round-robin method) for load balancing. Use the Round Robin mode only if the connection rate is high and you need more than one Virtual Steelhead to handle the load.
6. Click **New** to add a new Virtual Steelhead to connect to the Discovery Agent.

Figure 5-5. Add a New Steelhead

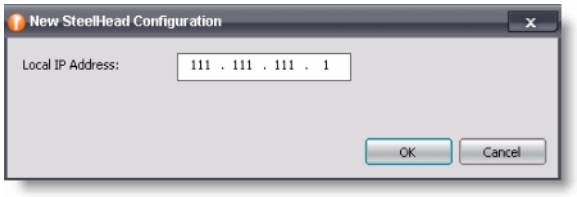

7. Type the IP address of the Virtual Steelhead and click **OK**.

   The Virtual Steelhead is added to the Steelhead List in the Configuration Dialog box.

   Use the arrows to adjust the priority in which the Discovery Agent connects to the Virtual Steelheads.

Configuring Transparency Modes

You configure the Discovery Agent transparency modes on the Steelhead appliance. You must use the CLI command:

```
in-path agent-intercept server-nat-mode ?
```

For information, see the *Riverbed Command-Line Interface Reference Manual*.

The Discovery Agent provides three transparency modes for connections between the client/server and the corresponding Steelhead. You configure the selected transparency mode in the Virtual Steelhead and it is transmitted to the Discovery Agent.

The transparency mode selected does not affect packets of the connection on the network. Packets viewed on the network are still addressed between the client/server and Steelhead appliance. However, the Discovery Agent performs NAT detection for these packets before sending them up the stack, so the transparency mode affects which IP address is visible to the application and client/server machine's network stack.

The three modes are:

- **Safe transparent** - If the client is behind a NAT device, the client connection to the application server is nontransparent—the application server detects the connection as a connection from the server-side Steelhead and not the client IP address. All connections from a client that is not behind a NAT device are transparent, which means that the server sees the connections from the client IP address instead of the Virtual Steelhead IP address.
- **Restricted transparent** - All client connections are transparent with the following restrictions:
  - If the client connection is from a NAT network, the application server sees the private IP address of the client.
  - You can use this mode only if there is no conflict between the private IP address ranges (there are no duplicate IP addresses) and ports.

  This is the default mode.
- **Non-transparent** - All client connections are nontransparent—the application server detects the connections from the server-side Steelhead IP address and not the client IP address. Riverbed recommends that you use this mode only if you cannot use one of the other two modes.

Enabling Optimization Using the Discovery Agent

After you configure the Discovery Agent, connect to the Virtual Steelhead CLI and enter the following commands to enable agent-intercept mode:

```
en
conf term
in-path agent-intercept enable
in-path enable
```

APPENDIX A Configuring a Riverbed NIC in ESX 4.1

The procedure for configuring bypass support for ESXi 5.0 and later is documented in the *Network Interface Card Installation Guide*.

Configuring the bpvm0 Interface (ESX/ESXi 4.1)

During the installation of your Riverbed NIC, you install an ESX driver of the file type .vib. Installation of the ESX driver is described in the *Network Interface Card Installation Guide*. After you install the ESX driver, a network adapter device named **bpvm0** appears in the vSphere client. If this device is not displayed after you install your ESX driver, reboot your machine. This device is not a real network adapter; it is used as a communication channel between the Virtual Steelhead guest and the ESX or ESXi host. The normal speed setting displayed for the bpvm0 adapter is 0 Half in the Network adapters tab.

To create virtual switches and port groups for the bpvm0 interface

1. In vSphere, connect to the ESX or ESXi host.
2. Select the ESX or ESXi host in the left Inventory panel.
3. Select the **Configuration** tab.
4. In the Hardware menu, select **Network adapters**.
5. In the Hardware menu, select **Networking**.
6. Click **Add Networking** in the upper right of the main panel.
7. Select the Connection Type **Virtual Machine**.
8. Click **Next**.
9. Select **Create a virtual switch**.
10. Select the box next to the bpvm0 network adapter.
11. Click **Next**.

12. Enter **pg-bpvm** for the name of the network label.

13. Click **Next**.

14. Verify that the diagram displayed in vSphere shows the pg-bpvm port group wired to the bpvm0 adapter.

15. Click **Finish**.

Configuring Riverbed NIC Interfaces (ESX/ESXi 4.1)

Before you install Virtual Steelhead, you must create a separate virtual switch and port label for each interface on the Riverbed NIC.

To create virtual switches and port groups for Riverbed NICs

1. In vSphere, connect to the ESX or ESXi host.

2. Select the ESX or ESXi host in the left Inventory panel.

3. Select the **Configuration** tab.

4. In the Hardware menu, select **Network Adapters**.

   In the vSphere client, the interfaces for the Riverbed NIC appear as two or four vmnic interfaces, under Intel Corporation 82575EB Gigabit Network Connection.

Figure 5-6. Interfaces in vSphere

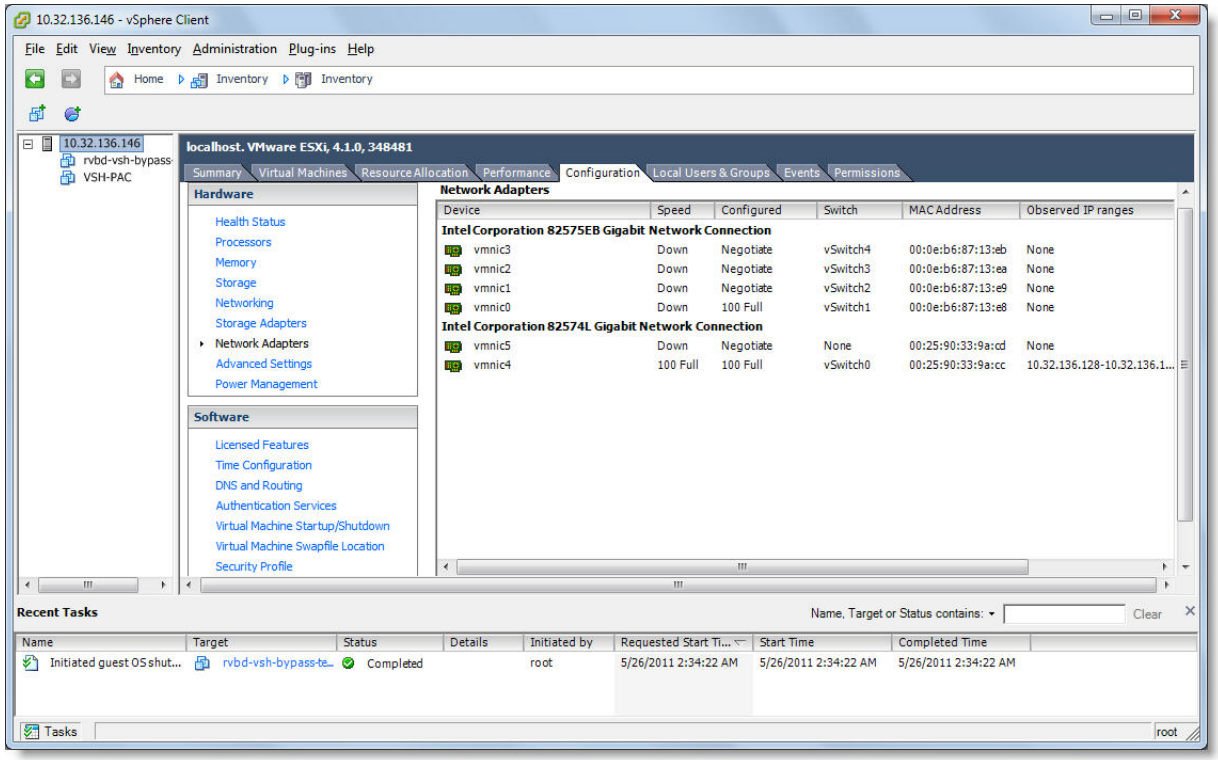

5. Under the heading "Intel Corporation 82575EB Gigabit Network Connection," find and write down the names of the two or four vmnic adapters.
6. In the Hardware menu, select **Networking**.
7. Click **Add Networking** in the upper right of the main panel.
8. Select the Connection Type **Virtual Machine**.
9. Click **Next**.
10. Select **Create a virtual switch**.
11. Select the box next to the adapter with the lowest number.

    Select only one of the adapters you noted in Step 5.
12. Click **Next**.
13. Use the following table to match your adapter to the correct network label and type it in.

ESX NIC Name	Network Label	Interface on Virtual Steelhead
vmnic0	pg-vmnic0	wan1_0
vmnic1	pg-vmnic1	lan1_0
vmnic2	pg-vmnic2	wan0_0
vmnic3	pg-vmnic3	lan0_0

14. Click **Next**.
15. Verify that the diagram displayed in vSphere shows the correct port group wired to the adapter.
16. Click **Finish**.
17. On the virtual switch you just created, click **Properties.**
18. On the Ports tab in the Properties dialog box, click the interface.
19. Click **Edit**.
20. Select the **Security** tab.
21. In the Security tab, select **Accept** for Promiscuous Mode.
22. Click **OK**.
23. Repeat this procedure for each of the remaining three interfaces.

Note: Deploying Riverbed NICs requires some additional steps, noted in the standard installation procedure.

APPENDIX B Troubleshooting

This chapter describes how to troubleshoot Virtual Steelhead. It includes the following sections:

- "Duplex Mismatch" on page 69
- "Oplock Issues" on page 70
- "CIFS Overlapping Open Optimization Denies Multi-User Access" on page 71
- "IP Address Configuration" on page 73
- "Asymmetric Routing" on page 74
- "Packet Ricochet" on page 74
- "Simplified Routing" on page 76
- "Auto-Discovery Failure" on page 77
- "Protocol Optimization Errors" on page 77
- "Resetting a Lost Password" on page 78
- "Bypass NIC Log Messages" on page 79

Duplex Mismatch

If the pass-through rule becomes ineffective, this indicates duplex mismatch. Duplex mismatch could also be the problem if you experience:

- Access is not faster after configuring Virtual Steelhead.
- The interface counters display error messages. An alarm or log message about error counts appears.
- There are many retransmissions in packet traces.
- You cannot connect to an attached device.
- You can connect to a device when you choose auto-negotiation, but you cannot connect to the same device when you manually set the speed or duplex.
- Good performance in one direction of data flow, but poor performance in the opposite direction.

Possible Cause

You have probably set the duplex value for your router to 100Full (fixed) and for the Virtual Steelhead to `Auto`.

Example

The following example shows applications that appear slower with Virtual Steelheads configured in an in-path deployment. The timed performance numbers to transfer a 20 MB file over FTP are:

- no Virtual Steelhead – 3:16.
- cold Virtual Steelhead – 5:08.
- warm Virtual Steelhead – 3:46.

Adding a pass-through rule for an application does not help. Slow connections appear as optimized in the Management Console on the Current Connections report page. However, stopping the Virtual Steelhead service while leaving the system powered on in an in-path configuration returns performance to original levels.

Solutions

You resolve duplex mismatch on the virtual host. You cannot configure speed and duplex settings through the Management Console of the Virtual Steelhead.

To resolve the duplex mismatch error:

- connect to the Virtual Steelhead CLI and enter the **ping** command with the flood (`-f`) option to check the duplex mismatch:

  ```
  ping -f -I >in-path-ip> -s 1400 <clientIP>
  ```

- ensure there is a speed and duplex match between each in-path interface and its peer network interface. If they do not match, you might have a large number of errors on the interface when it is in the bypass mode, because the switch and the router are not set with the same duplex settings. Also, ensure connectivity when service is down.

  If matching speed and duplex do not reduce collisions or errors, try hard-setting one end and auto-setting the other. Try the half-duplex mode.

If all combinations fail, as a last resort, add an intermediary hub or switch that is more compatible with both network interfaces.

Oplock Issues

The following symptoms occur due to opportunistic lock (oplock) issues:

- File access is not faster or tasks such as drag-and-drop are fast but applications might benefit from acceleration.
- The Current Connections report page in the Management Console (select Reports > Networking > Current Connections) displays slow connections as optimized.

Possible Causes

- The client is running an old antivirus software such as McAfee v4.5, the most common type, which competes with the application for an oplock instead of opening as read-only. The antivirus causes multiple file opens.
- The server has oplocks disabled.

Example

You can open a previously accessed file in 5 seconds on PC1, but you cannot open the same file in under 24 seconds on PC2. If you close the file on PC1, you can open it in 5 seconds on PC2. However, it takes you 24 seconds to open the same file on PC1.

Solutions

Windows Common Internet File System (CIFS) uses oplock to determine the level of safety the OS or the application has in working with a file. Oplock is a lock that a client requests on a file in a remote server.

An oplock controls the consistency of optimizations such as read-ahead. Oplock levels are reduced when you make conflicting opens to a file.

To prevent any compromise to data integrity, Virtual Steelhead optimizes data only when a client has exclusive access to the data.

When an oplock is not available, Virtual Steelhead does not perform application-level latency optimization but still performs Scalable Data Referencing (SDR) and data compression, as well as TCP optimization. Therefore, even without the benefits of latency optimization, Virtual Steelheads still increase WAN performance, but not as effectively as when application optimizations are available.

To resolve oplock issues:

- Upgrade your antivirus software to the latest version.
- Use FileMon (sysinternals) to check for file access, if available on your version of Windows.
- Enable CIFS Overlapping Opens (by default, this function is enabled). For details, see "CIFS Overlapping Open Optimization Denies Multi-User Access" on page 71.
- Ensure that the server has oplock enabled by verifying registry settings on Windows servers or the system configuration (for NetApp or EMC servers).
- Run a network analyzer such as Cascade Pilot, which is fully integrated with Wireshark, and determine that the server grants oplocks when the client opens a file.
- Check whether the client is running an antivirus software that is scanning the files over the WAN or that the antivirus software does not break the oplock.

CIFS Overlapping Open Optimization Denies Multi-User Access

The CIFS overlapping open optimization issue prevents a client from accessing a file when different clients access the file at the same time.

Solution

To resolve the CIFS overlapping open optimization issue, configure CIFS overlapping open optimization on the client-side Steelhead appliance as follows:

1. Connect to the Steelhead Management Console.

   For details, see the *Steelhead Appliance Management Console User's Guide*.

2. On the client-side Steelhead appliance, choose Configure > Optimization > CIFS (SMB1) to display the CIFS (SMB1) page.

Figure 5-7. CIFS (SMB1) Page

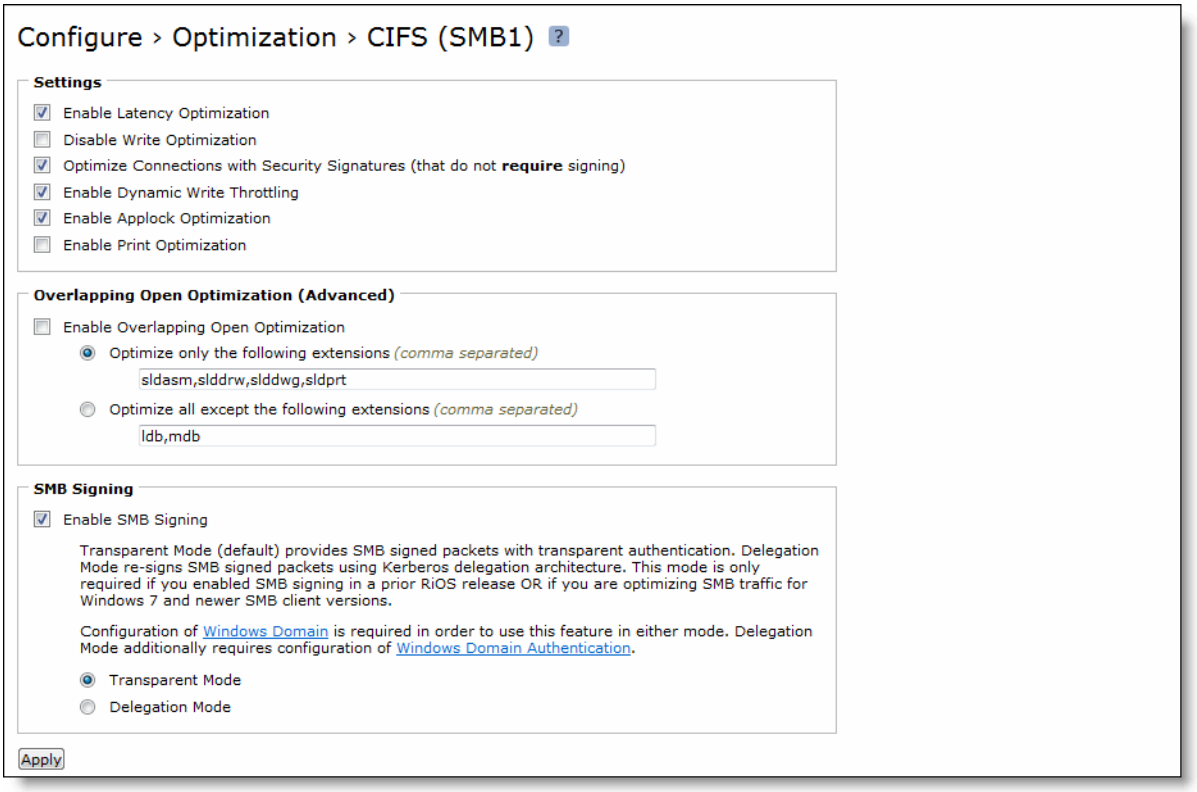

3. Under Overlapping Open Optimization (Advanced), complete the configuration as described in the following table.

Control	Description
Enable Overlapping Open Optimization	Enables overlapping opens to obtain better performance with applications that perform multiple opens on the same file: for example, CAD applications. By default, this setting is disabled. **Note**: Enable this setting on the client-side Steelhead appliance. With overlapping opens enabled, the Steelhead appliance optimizes data where exclusive access is available (when locks are granted). When an oplock is not available, the Steelhead appliancedoes not perform application-level latency optimizations but still performs SDR and compression on the data, as well as TCP optimizations. **Note**: If a remote user opens a file that is optimized using the overlapping opens feature on a v3.x.x or later Steelhead appliance, or if the file does not go through a Steelhead appliance and a second user opens the same file, the second user might receive an error message if the file fails to go through. For example, certain applications that are sent over the LAN may cause this problem. If this occurs, you should disable overlapping opens for such applications. Use the radio buttons to set either an include list or exclude list of file types subject to overlapping opens optimization.
Optimize only the following extensions	Specify a list of extensions that you want to include in overlapping opens optimization.
Optimize all except the following extensions	Specify a list of extensions that you do not want to include. For example, you should specify any file extensions that use Enable Applock Optimization.

4. Click **Apply** to apply your settings to the current configuration.

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

IP Address Configuration

If you have not configured IP addresses correctly, the Steelhead appliances cannot connect to each other or to your network.

Solutions

Follow the suggestions below to correctly configure IP addresses:

- Ensure that the Steelhead appliances are reachable through the IP address by pinging their primary and in-path interfaces.
- Ensure that the Steelhead appliances in the network can reach each other through their own interfaces.

  Connect to the Steelhead CLI. For details, see the *Riverbed Command-Line Interface Reference Manual*. Enter the following command to ping from a specific in-path interface on a Steelhead appliance to another in-path interface:

  ```
  ping -f -I {Local-Steelhead-Inpath-IP} -s 1400 {Remote-Steelhead-Inpath-IP}
  ```

- Ensure that the default gateways, both for the Steelhead appliance and for its in-path interfaces, are correct.

- For physical or virtual in-path installations, verify that the server-side Steelhead appliance can be auto-discovered by the client-side Steelhead appliance.

  Connect to the Steelhead CLI. For details, see the *Riverbed Command-Line Interface Reference Manual*. Enter the following command:

  ```
  tproxytrace -i inpath0_0 <example-server-IP-address>:<example-server-TCP-port>
  ```

  This command causes the Steelhead appliance to generate a fake TCP SYN packet, destined for the specified IP address and TCP port, and send it to the specified in-path interface. A remote Steelhead appliance should respond if it sees the SYN packet.

- Verify that the client-side Steelhead appliance is visible to the server-side Steelhead appliance.

  Connect to the Steelhead CLI. For details, see the *Riverbed Command-Line Interface Reference Manual*. Enter the following command:

  ```
  tproxytrace -i inpath0_0 <example-client-IP-address>: <example-client-TCP-port>
  ```

Asymmetric Routing

If there is an asymmetric routing issue, many connections fail during data transfer or they fail to start.

Possible Cause

Asymmetric routing occurs when a TCP connection takes one path to the destination and another when returning to the source. If the Virtual Steelhead sees only the LAN-to-WAN or only the WAN-to-LAN packets, it cannot optimize the data.

Solutions

To resolve the asymmetric routing issue, do one of the following:

- Rank the following solutions from most to least preferable, with respect to complexity and cost, and select one:
  - Configure a fixed-target rule.
  - Use a logical in-path configuration such as WCCP or PBR.
  - Configure connection-forwarding with two Steelhead appliances.
- Remove the asymmetry.

Packet Ricochet

The following symptoms occur due to packet ricochet:

- Performance is less than expected.
- The following log message appears:

  ```
  » [fionr taelrcreeapdt/y lnoactaltekde rnceoln/neiccotireo.n c:119426.316]
  8.n7a3t._1c5h:e1c6k1: 1 SYN ==> packet 192.168.208.12:80 ==> 192.168.72.9:7801
  ```

Possible Cause

Traffic to the LAN is traveling to the WAN router on the way to the LAN.

Solutions

To resolve packet ricochet issues:

- change the in-path gateway to the LAN router.
- add static routes to LAN subnets through the LAN router.
- enable in-path simplified routing.

Packet Ricochet—Internet Control Messaging Protocol (ICMP) Redirects

The following symptoms occur due to packet ricochet ICMP redirects:

- Connections fail on first attempt, but succeed on second attempt.
- On one or both sites, the in-path interface on the Steelhead appliance is on a different network than the local host.
- There are no in-path routes defined.

Possible Causes

- Traffic to the LAN is travelling to the WAN router on the way to the LAN, but the router drops the packet.
- Outer connections to clients or servers are routed through the WAN interface to the WAN gateway, and then routed through the Steelhead appliance to the next hop LAN gateway.
- The WAN router is probably dropping the SYN from the Steelhead appliance before issuing an ICMP redirect.

Solutions

To resolve the packet ricochet ICMP redirects issue, do one of the following:

- Change the router ICMP configuration to forward the packet or turn off ICMP redirect.
- Change the in-path gateway to the LAN router.
- Add static routes to LAN subnets through the LAN router.
- Enable in-path simplified routing. For details, see "Simplified Routing" on page 76.
- Add in-path routes to local destinations to prevent the ICMP redirect and subsequent drop.

Simplified Routing

Simplified routing changes the process used to select the destination Ethernet address for packets transmitted from in-path interfaces.

Simplified routing collects the IP address for the next-hop MAC address from each packet it receives to address traffic. With simplified routing, you can use either the WAN-or LAN-side device as a default gateway. Virtual Steelhead sets the correct gateway to use by based on where the switch or router sends the traffic, and by associating the next-hop Ethernet addresses with IP addresses. Enabling simplified routing eliminates the need to add static routes when the Virtual Steelhead is in a different subnet from the client and the server.

Without simplified routing, if Virtual Steelhead is installed in a different subnet from the client or server, you must define one router as the default gateway and static routes for the other routers so that traffic is not redirected back through the Steelhead appliance. In some cases, even with the static routes defined, the Access Control List (ACL) on the default gateway can still drop traffic that should have gone through the other router. Enabling simplified routing eliminates this issue.

Simplified routing has the following constraints:

- You cannot enable WCCP.
- The default route must exist on each Steelhead appliance in your network. (For detailed information, see the *Steelhead Appliance Deployment Guide*.)

To enable simplified routing

1. Choose Configure > Networking > Simplified Routing to display the Simplified Routing page.

Figure 5-8. Simplified Routing Page

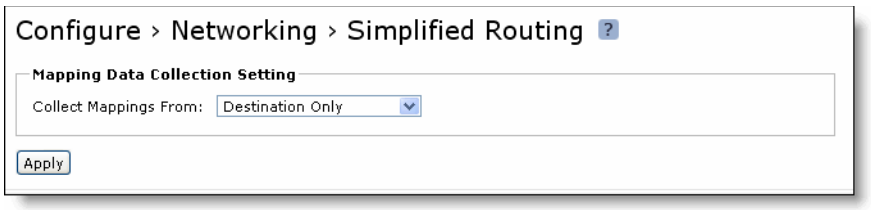

2. Under Mapping Data Collection Setting, complete the configuration as described in the following table.

Control	Description
Collect Mappings From	Select one of the following options from the drop-down list: • **None** - Do not collect mappings. • **Destination Only** - Collects destination MAC data. Use this option in connection-forwarding deployments. This is the default setting. • **Destination and Source** - Collect mappings from destination and source MAC data. Use this option in connection-forwarding deployments. • **All** - Collect mappings for destination, source, and inner MAC data. Also collect data for connections that are not translated using NAT. You cannot enable this option in connection-forwarding deployments. Riverbed recommends that you use this option to maximize the effects of simplified routing.

3. Click **Apply** to save your settings to the running configuration.

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

Auto-Discovery Failure

All traffic passes through the in-path (physical or logical) Virtual Steelhead due to auto-discovery failure.

Possible Causes

- Cisco PIX 7.x or Raptor firewalls
- Satellite
- Intrusion detection system (IDS) or intrusion prevention system (IPS)

Solutions

- Create a fixed-target rule on the client-side Steelhead appliance:
  - Specify the **Target Appliance IP Address** and its port as 7800 on the opposite Steelhead appliance (in-path without auto-discovery).
- Configure end nodes (firewalls) to allow your probe to pass through.
- Configure the Steelhead IP address as the friendly IP address for IDS or IPS.
- Cisco PIX Firewall IOS v7.0 might block the auto-discovery probe. Some firewall configurations strip TCP options or drop packets with these options. You can keep this configuration and switch to fixed-target rules or change the configuration on the firewall.

Protocol Optimization Errors

Virtual Steelhead fails to optimize expected protocols.

Solutions

To resolve protocol optimization errors, check:

- that connections have been successfully established.
- that Steelhead appliances on the other side of a connection are turned on.
- for secure or interactive ports that are preventing protocol optimization.
- for any pass-through rules that could be causing some protocols to pass through the Steelhead appliances unoptimized.

Resetting a Lost Password

To reset your password, you must have access to the serial console or monitor and must be able to see the entire boot process to perform these steps:

To reset a lost password

1. Start or reboot the Steelhead appliance.
2. When prompted, click any key to continue.
3. Immediately press **E**.

   A GNU GRUB menu appears.
4. Press the **V** or **^** keys to select the disk image to boot.
5. Press **E**.

   A GRUB menu appears, with options similar to the following:

   ```
   -----------------
   0: root (hd0,1)
   1: kernel /vmlinuz ro root=/dev/sda5 console=tty0 console=ttyS0,9600n8
   -----------------
   ```
6. Press **V** or **^** to select the kernel boot parameters entry.
7. Press **E** to edit the kernel boot parameters.

   The CLI displays a partially completed line of text similar to the following:

   ```
   kernel /vmlinuz ro root=/dev/sda5 console=tty0 console=ttyS0,9600n8
   ```
8. The line of text contains two console= entries. Modify this line as follows:
   - If you are accessing the Steelhead appliance remotely, delete the console=tty0 entry.
   - If you are accessing the Steelhead appliance directly (through a keyboard and monitor connected to the appliance), delete the console=ttyS0 entry.
   - At the end of the line, type a space and append with `single fastboot`. It is important to type a space before `single`.

Tip: Use the arrow keys to access the entire command line.

9. Press Enter.
10. Press the **B** key to continue booting.

    The system starts.
11. At the command prompt, enter **/sbin/resetpw.sh**.

    There is no password.

12. Type **reboot** and press Enter to reboot the appliance.

Bypass NIC Log Messages

If you have deployed Virtual Steelhead with a Riverbed NIC, you can view log messages related to NIC card performance.

Note: The Microsoft Hyper-V hypervisor does not currently support deployment with a Riverbed NIC.

To view log messages regarding the Riverbed NIC

1. Log in to the Management Console using an account with permission to read logs, such as root.

2. Select the **Logging** tab.

3. On the System Log page, enter **vsh_bypass** in the Filter dialog box.

4. Click the **Filter** button.

Only log messages regarding the NIC appear. For example:

```
Feb 17 00:56:51 localhost rbtkmod: vsh_bypass_init: probing for hardware bypass
Feb 17 00:56:54 localhost rbtkmod: vsh_bypass_init: control interface lan0_0 mac 00:0C:29:05:5A:3F
devnum 5 slave 6 esx_nic vmnic6 esx_mac 00:0E:B6:87:13:E8
Feb 17 00:56:54 localhost rbtkmod: vsh_bypass_init: slave interface wan0_0 mac 00:0C:29:05:5A:49
devnum 6 esx_nic vmnic7 esx_mac 00:0E:B6:87:13:E9
Feb 17 00:56:54 localhost rbtkmod: vsh_bypass_init: management interface bpvm0 mac 00:0C:29:05:5A:67
driver_version 2.0.1.12 device_count 8
Feb 17 00:56:54 localhost rbtkmod: vsh_bypass_init: success, hardware bypass enabled on [ inpath0_0
: lan0_0 ESX nic vmnic6 ESX mac 00:0E:B6:87:13:E8 ; wan0_0 ESX nic vmnic7 ESX mac 00:0E:B6:87:13:E9 ]

Feb 17 00:57:37 amnesiac wdt[6317]: [wdt.NOTICE]: vsh_bypass watchdog init: control interface
"lan0_0" (guest mac 00:0C:29:05:5A:3F ESX nic vmnic6 ESX mac 00:0E:B6:87:13:E8 devnum 5) slave
interface "wan0_0" (guest mac 00:0C:29:05:5A:49 ESX nic vmnic7 ESX mac 00:0E:B6:87:13:E9 devnum 6),
interval "1000", bypass_interval 7000, disable_on_timeout "false”
```

Index

A
Antivirus compatibility 3
Appliance
  installing 69
Application Streamlining, overview of 9
Auto-discovery 57
Auto-discovery process, overview of 9
Auto-discovery rule, overview of 10

C
CMC 9
Compatibility
  hardware 11
Configuration checklist
  ESXi 23
  Hyper-V 36
Configuration wizard, completing 42
Configuration, initial 42
Configuration, verifying 56
Configuring
  Virtual Steelhead on Hyper-V 35
CPU
  allocating cycles 53
  physical 13
  virtual 13
CX platform models 16

D
Data store
  disk storage 13
  see RiOS data store
Data Streamlining, overview of 9
Default gateway 23
Deny rules, overview of 10
Deploy OVF 24
Discard rules, overview of 10
Discovery Agent
  configuring 60
  installing on Linux server 60
  installing on Windows server 59
  overview 57
DNS Server 23
Document conventions, overview of 2
Documentation, contacting 5
Domain Name 23
Downgrading an appliance model 48
Downloading the OVA package 22

E
ESX 4.1
  NIC deployment 18
ESX/ESXi deployment guidelines 12, 14
ESX/ESXi overhead and reservations 51
ESXi 5.0
  NIC deployment 18
ESXi direct path configuration settings 44
Evaluation license 48

F
Failure detection 58
Fixed-target rules, overview of 10
Flexible licensing 47

H
Hardware dependencies, overview of 3
Hardware required message 49
Hardware specifications 50
Host server 13
Hostname 23
Hyper-V
  manual install 39
  overhead and reservations 51
  parameters 36
  troubleshooting install 39
  v8.0.3 limitations 12
  Virtual Steelhead package file 35
Hyper-V Manager 37, 39
Hypervisor overhead and reservations 51

I
Initial configuration 42
In-Path rules, overview of 10

Installing
Discovery Agent on Linux server 60
Discovery Agent on Windows server 59
Hyper-V manually 39
Virtual Steelhead on Hyper-V 36
IP address 23

K
Known issues 4

L
LAN/WAN
virtual interface link size 13
License Request Key 49
Licenses, managing 47
Licenses, removing 54
Licensing
for RiOS v8.0.3 47
Load balancing 58
Logging in 45
Login page 46

M
Management Console
logging in to 45
Management disk, v8.0.3 size 16
Management Streamlining, overview of 9
Manifest file 22
Model upgrade
for v8.0.3 47
RiOS data store 50

N
Netmask 23
Network loops, preventing 12
Network mapping 28
NIC
assignment 13
connecting 12
deployment on ESXi 18
physical 13
Non-transparent mode 63

O
Online documentation 4
Optimization 57
OVA package 23
Overhead, hypervisor 51
OVF file 22
OVF Tool 23

P
Package 23
installing 23
obtaining from Riverbed 23
Parameters
Hyper-V 36
Pass-through rules, overview of 10
Peering rules, overview of 10
Performance 13
Platform models 15
Port groups 12
Professional services, contacting 5
Promiscuous mode 12, 31

R
RAM 13, 53
Reboot 55
Related reading 4
Release notes 4
Reservations
hypervisor 51
Restricted transparent mode 63
RiOS data store
downgrade errors 50
flexible size 16
RiOS v8.0.3
licensing and model upgrade 47
Riverbed, contacting 4
RIVERBED_INSTALL.ps1 script 36

S
Safe transparent mode 63
Scalable Data Referencing, overview of 9
SDR, overview of 9
Setting up
Virtual Steelhead on Hyper-V 35
Shut down 55
SNMP
compatibility 3
Specification license 47
SSL protocol 45
Steelhead Mobile Controller, overview of 10

T
Technical Publications, contacting 5
Technical specifications 15
Technical support, contacting 4
Token ID number 49
Token key 48
Transparency 58
Transparency modes, configuring 63
Transport Streamlining, overview of 9
Troubleshooting
Hyper-V install 39

U
Updating
license 47
Upgrading
model 47

V
vCenter
  alarms 34
VCX models
  flexible data store 16
Virtual Machine
  image 22
  naming 26
Virtual NICs, connecting 12
Virtual RAM 13
Virtual Steelhead 12
  Hyper-V install 36
  platform models 15
  running multiple on a host 12
  setting up 35
Virtual switch
  Hyper-V configuration 38
Virtual Switch Manager 37
VMDK files 22
VMware
  vSphere client 23
vSwitch
  configuring 12

W
Windows Powershell 36, 37

X
xx50 platform models 15