



Course Details

Steelhead™ Appliance Advanced Deployment & Troubleshooting

Course Information

Number of Days: 4 days

Instruction Method: Instructor-led

Lab Type: eLab (Riverbed Remote Labs)

Riverbed Software (RiOS) Version Used in Class: 5.x

This course uses RiOS software version 5.x and is also effective for students using version 4.x or earlier.

Course Description

Steelhead Appliance Advanced Deployment & Troubleshooting gives students the in-depth experience needed to deploy and troubleshoot Steelhead appliance in advanced network environments. Throughout the course, students use real-life examples and hands-on activities to learn advanced Steelhead deployment configurations. This course includes the following topics: QoS, SSL Optimization, WCCPv2 Deployment, PBR Deployment, Riverbed Interceptor, and Troubleshooting Logical In-Path Deployments.

Who should take this course?

Steelhead Appliance Advanced Deployment & Troubleshooting is targeted toward those individuals (system engineers, channel partners/resellers, customers) involved in the design, deployment, or administration of Steelhead appliances.

Prerequisites

- *Steelhead Appliance Deployment & Management* training course **required**
- Certification as a RCSP/CCNP/CCIP or the equivalent knowledge
- Familiarity with RIPv2, OSPF, IP QoS, network design and troubleshooting
- Familiarity with content networking terms and concepts
- Basic knowledge of the Linux operating system
- Participants are **required** to bring along their own laptops **AND**,
- If using our hardware, a serial cable is required or a USB-to-serial adaptor for the class labs (please view the confirmation email sent to you regarding version of course being delivered)

Detailed Course Outline and Objectives

After completing this course the student should be able to:

Module 1: QoS

- Understand DiffServ QoS model in relation to the Steelhead appliance
- Describe the various terminologies used in DiffServ QoS model

- Understand what is classification, marking or coloring, scheduling or queuing, bandwidth guarantee, and traffic conditioning and traffic shaping
- Understand the relationship between IP Precedence bits, DSCP bits, and ToS octet
- Understand Differentiated Services Code Point (DSCP) encodings and per-hop-behaviors (PHBs)
- Describe Expedited Forwarding (EF), Assured Forwarding (AF), and Class Selector (CS) PHBs
- Describe IP precedence values and names
- Describe well-known DSCP values and class names
- Configure DSCP markings on the Steelhead appliance for optimized and pass-through traffic
- Make valid changes to QoS policies on the WAN router to accommodate QoS markings from Steelhead appliances
- Convert DSCP to its corresponding ToS
- Describe the characteristics of Hierarchical Fair Service Curves (HFSC) and its classes of services
- Understand Riverbed DiffServ QoS parameters
- Understand how Riverbed QoS classification rules work
- Enable QoS and specify the default link rate on the WAN interfaces of the Steelhead appliance
- Create Riverbed QoS classes
- Define Riverbed QoS classification rules
- Verify Riverbed QoS implementation
- View a Riverbed QoS statistics report
- Give an overview on how Riverbed MX-TCP works
- Describe Riverbed MX-TCP performance
- Configure MX-TCP on the Steelhead appliance
- Verify MX-TCP implementation
- Describe some of the MX-TCP deployment constraints and caveats
- View MX-TCP statistics report

Module 2: SSL Optimization

- Give a brief overview of SSL
- Understand the Riverbed solution to SSL optimization
- Configure and deploy a Riverbed SSL optimization solution
- Troubleshoot Riverbed SSL optimization issues

Module 3: WCCPv2 Deployment

- Give a brief overview of logical in-path deployment
- Give a brief overview of WCCPv2 features and functionalities
- Understand WCCPv2 control messages
- Describe WCCPv2 load distribution with hash assignment
- Describe WCCPv2 failsafe mechanisms
- Deploy Connection Forwarding within a WCCPv2 cluster

- Implement Active / Active Data Store Synchronization within a WCCPv2 cluster
- Describe WCCPv2 ingress and egress interception and the performance impact of these interceptions
- Describe the two WCCPv2 packet return methods
- Configure WCCPv2 on the router
- Understand the caveats for WCCPv2 outbound redirection
- Configure WCCPv2 multicast, security, and group-list on the router
- Implement WCCPv2 with NBAR on the router
- Implement WCCPv2 with Tracking, EEM, and mHSRP on the router
- Verify WCCPv2 configurations on the router
- Troubleshooting WCCPv2 on the router
- Configure WCCPv2 on the Steelhead
- Configure WCCPv2 Multicast and Password on the Steelhead
- Configure WCCPv2 Priority, Weight and Encapsulation Scheme on the Steelhead
- Configure WCCPv2 Service Flags and Ports on the Steelhead
- Verify WCCPv2 configurations on the Steelhead

Module 4: PBR Deployment

- Give a brief overview of PBR
- Describe the advantages and disadvantages of using PBR over WCCPv2
- Define routing policies using a route-map
- Use the PBR match and set commands with a route-map
- Enable PBR on an inbound interface
- Verify PBR configurations
- Give a brief overview of PBR next-hop availability verification
- Configure CDP to verify PBR next-hop availability
- Configure IP SLA with ICMP echo to verify PBR next-hop availability
- Verify PBR object tracking configuration
- Configure PBR support on the Steelhead
- Enable PBR support on multiple in-path interfaces

Module 5: Riverbed Interceptor

- Describe the positioning of the Interceptor and its various functionalities
- Describe the different types of in-path rules on the Interceptor
- Configure the various in-path rule types on the Interceptor
- Describe the load balancing rules on the Interceptor
- Understand the load balancing scheme on the Interceptor
- Configure load balancing rules on the Interceptor
- Aware of the load balancing rules deployment considerations
- Configure peer neighbors on the Interceptor
- Illustrate a simple configuration example with a single Interceptor
- Describe Interceptor redirect

- Describe the basic packet flow during connection startup with client-side Interceptor deployment
- Describe the basic packet flow during connection startup with server-side Interceptor deployment
- Describe inter-Interceptor redirects
- Describe failover buddy pass-through during inter-Interceptor redirects
- Describe the use of connection forwarding after connection startup
- Configure peer Interceptors
- Illustrate a configuration example with two Interceptors in parallel
- Configure Interceptor failover buddies
- Illustrate a configuration example with two Interceptors as failover buddies
- Illustrate a configuration example with two pairs of Interceptors as failover buddies with one pair parallel to the other
- Verify configurations on the Interceptor
- Interpret the statistics on the Interceptor

Module 6: Troubleshooting Logical In-Path Deployments

- Troubleshooting duplex mismatch
- Select the recommended IOS versions for the WCCPv2 servers
- Select the recommended switching platforms as WCCPv2 servers
- Implement fair queuing and rate-limiting to reduce CPU utilization
- Implement adaptive Embedded Event Manager (EEM) applets to ease off the high CPU utilization caused by WCCPv2 redirection by disabling it altogether and re-enable it back automatically when the prolonged CPU spikes subsided
- Implement NetFlow switching to reduce the high CPU utilization caused by lengthy access-lists
- Place the Steelhead in a different network subnet or VLAN from that of the server to avoid potential looping scenarios
- Verify current connections on the Steelhead
- Verify Steelhead peer connectivity
- Apply the WCCPv2 troubleshooting checklist
- Verify WCCPv2 interception on the router and on the Steelhead
- Apply the PBR troubleshooting checklist
- Verify PBR interception on the router
- Apply the Interceptor troubleshooting checklist
- Identify the common issues encountered in Interceptor deployments
- Verify auto discovery

Lab Outline

- Lab 1: Implement DSCP Markings at the Steelhead
- Lab 2: Convert Steelhead to a DiffServ Node
- Lab 3: Configure MX-TCP on Steelhead
- Lab 4: Configure SSL Optimization on Steelhead
- Lab 5: Logical In-Path Deployment with WCCP

- Lab 6: WCCP Clustering
- Lab 7: Client-Side WCCP with PFS
- Lab 8: Advanced WCCP Deployment
- Lab 9: Logical In-Path Deployment with PBR
- Lab 10: Interconnecting Private Subnets
- Troubleshooting Lab

Registration

Sign up by sending an email to training@riverbed.com. Note that class size is limited, and classes fill up well in advance, so sign up soon!

Cancellation Policy

If attendee cancels any training session(s) upon written notice which shall be received by Riverbed (i) no less than fourteen (14) days prior to the first day of such cancelled training, attendee shall be entitled to a full refund of any enrollment fees paid for such cancelled training or such amount may be applied to another available scheduled training session which attendee selects; (ii) between seven (7) and thirteen (13) days prior to the first day of such cancelled training, attendee shall be entitled to a fifty percent (50%) refund of any enrollment fees paid for such cancelled training or such amount may be applied to another available scheduled training session which attendee selects; or (iii) less than seven (7) days prior to the first day of such cancelled training, attendee shall receive no refund for any fees paid for any such cancelled training (however a qualified replacement may be substituted in the place of the designated attendee at no additional charge).

If you have any additional questions, please contact your local Riverbed Sales Manager or email <mailto:training@riverbed.com>.